CentER

Non-standard Preferences in Asset Pricing and Household Finance

JORGO T.G. GOOSSENS

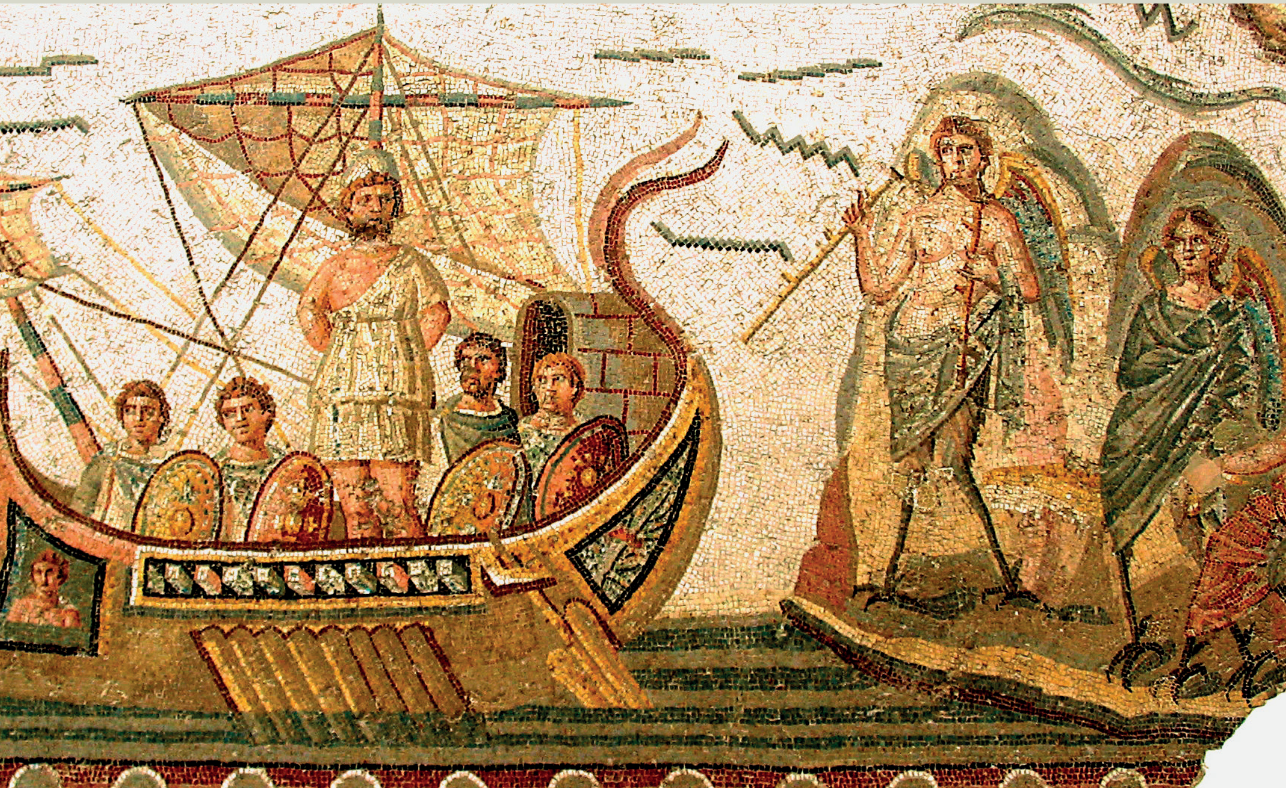

Tilburg University

Non-standard Preferences in Asset Pricing and Household Finance

Jorgo T.G. Goossens

Non-standard Preferences in Asset Pricing and Household Finance

Proefschrift ter verkrijging van de graad van doctor aan Tilburg University op gezag van de rector magnificus, prof. dr. W.B.H.J. van de Donk, in het openbaar te verdedigen ten overstaan van een door het college voor promoties aangewezen commissie in de Aula van de Universiteit op vrijdag 17 februari 2023 om 13.30 uur door

door

Jorgo Theodorus Godefridus Goossens,

geboren te Breda

Promotores:
prof. dr. B.J.M. Werker (Tilburg University)
prof. dr. M.G. Knoef (Universiteit Leiden)
prof. dr. E.H.M. Ponds (Tilburg University)

Copromotor:
dr. R. van den Goorbergh PhD (APG Asset Management)

Leden promotiecommissie:
prof. dr. H. Bateman (University New South Wales)
prof. dr. S. Zeisberger (Radboud Universiteit)
prof. dr. F.C.J.M. de Jong (Tilburg University)
dr. S. Cassella PhD (Tilburg University)
prof. dr. J.J.M. Potters (Tilburg University)

This thesis was supported by NWO (grant number NWA.ID.17.003).

Acknowledgments

This dissertation is the result of a satisfying and enjoyable period of my life. It started with my quantitative bachelors and masters in Econometrics at Tilburg University, followed by a teaching degree in Mathematics at Eindhoven University, and then a research master in Finance at Tilburg University. Finally, I started my PhD at Tilburg University together with APG. I owe a lot to the many people that I have met during these times, as they helped, supported, and guided me. I would like to take the opportunity to thank them now.

First and foremost, I thank my supervisors Bas Werker, Marike Knoef, Eduard Ponds, and Rob van den Goorbergh. I met Eduard and Rob when I wrote my master thesis in Econometrics as an intern at APG. Both, amongst others, stimulated me to do a PhD, gave me the opportunity to do so, and advised me to reach out to Bas Werker. As I aspired to combine micro-level behavior with asset pricing, Bas asked Marike to join as a supervisor. With the help and advice of all four, I wrote a research proposal for a NWO Industrial Doctorate grant, which we received and funded my PhD. This marked the beginning of a fruitful collaboration between academia and the industry.

I am indebted much to my four supervisors. Bas is a great supervisor, teacher, and academic, and he will be an important role model for the rest of my life. In general, when I struggled with a difficult question, I went to Bas and he always had an answer or made me find the answer myself. I also enjoyed our informal discussions about pensions in The Netherlands. Over the years, I think I learned something which I call the three 'laws of Bas'. First, when you enter his room, you leave the room with a new question. Second, intuition and mathematics must

be in line with each other, otherwise your result does not hold. Third, for each paper, you will reach a point in your research where your results become seemingly simple to you; this is the point where you understand what you have done and you can explain it to others. Finally, there is also a principle that I liked and which I call ‘efficiency’: on average, a meeting with Bas lasts for twenty minutes. These laws, and the principle, form fundamental buildings blocks for how he learned me to be a successful academic.

Marike taught me valuable knowledge and guidance about two matters: doing empirical research, and the connection between academia and society. She is a kind, approachable, and helpful supervisor. I have always enjoyed your genuine interest and our small talks, also about personal matters. Myself having a theoretical quantitative background, but also desiring to do behavioral empirical work, Marike had the patience to teach me how to create and field surveys, and how to handle such data. I appreciate that Marike helped and stimulated me with reaching out to society and the industry, already early on. Over the years, we started multiple projects and we have received multiple grants with co-authors, for which I thank you. I owe many industry connections, and especially my precious relation with Netspar, to you.

Eduard and Rob gave me the helpful and cherishable opportunity to see how research and academia work in practice. You allowed me every week to work very closely with APG, to develop my soft skills, and to steer my research within “the industry”. It goes without saying how valuable my time was, and I hope continues to be, at APG. You offered me the opportunity to do research at APG and its clients, which is invaluable for my own research, but also for the facilitation of knowledge transfers between academia and the sector. Both of you always pushed me to think about the relevance and practical implications of my academic work, which is a valuable skill that I have acquired thanks to both of you.

Eduard, I owe you much in terms of the opportunities you gave me: connections with the industry, access to populations and data for research, and teaching together at EMAS (which I hope to continue for a long time). You made me enthusiastic me with your skills of quickly generating ideas and connecting different

topics. I always looked forward talking to you, as you gave me inspiration, and you cared about how I was doing. Rob, you have made many things possible for me during my PhD. From the very start offering the opportunity to do a PhD with APG, till the current opportunity to continue collaborating with APG. Our quantitative discussions learned me a lot about industry- and policy-relevant research. It was always a pleasure to meet together with the three of us and also talk about life outside academia, especially during COVID-19.

Second, I thank Hazel Bateman, Stefano Cassella, Frank de Jong, Jan Potters, and Stefan Zeisberger for accepting to be in my dissertation committee. I appreciate the time and effort they have spent on my chapters, and for their comments on my work. Stefano provided me with a useful list of detailed suggestions to improve my papers. Frank and Jan have been engaged in my own education and research early on, and I appreciate their constructive help throughout all the years. I want to thank Hazel for the opportunity of a short study visit at CEPAR. Finally, Stefan thank you for your trust in me and your willingness to offer me a job at Radboud University in the direction of behavioral finance.

Third, I want to thank Ralph Koijen. In my fourth year, during my job market year, I had the opportunity to visit Chicago Booth. This was a period of hard work, but also a rewarding experience. Ralph was my host and he improved my job market paper with multiple suggestions. Thanks to Ralph I got exposure to the academic life in the US, and the opportunity to see and meet great researchers and PhD students in the field.

Furthermore, I would like to thank all faculty members of Tilburg University, especially those in the Econometrics & OR and the Finance departments. Ramon van den Akker and Nikolaus Schweizer thank you for the cooperation during the course we taught together. Anja de Waegenaere thank you for your advice during the thesis supervision we did together, and I also want to thank you for the enjoyable master Open Days together. Susan van Soest and Marieke Quant thank you for all educational matters and the opportunity to get my University Teaching Qualification (BKO). Finally a big thanks to the secretaries for all administrative help!

I also like to express my gratitude to Netspar for all their support, from the beginning of my studies till the end of my PhD, and hopefully in the future. You provided assistance during my masters, during process for the NWO Industrial Doctorate grant, and during the topicality grants. Specifically, I want to thank Theo Nijman, Anne Balter, Bertrand Melenberg, Arthur van Soest, and Rogier Potter van Loon for all their support and guidance throughout the years. Regarding administrative support over the years, I want to thank Gaby, Sylvia, Femke, Silvie, Peter, Sheila, and Diana.

My gratitude also goes to the colleagues at APG. I started at the department SBOC, where Michiel van Leuvensteijn gave me the opportunity to write my master thesis as an intern. Amongst others, I greatly benefited from discussions with Peter Gortzak, Jurre de Haan, Tinka den Arend, Youri van Dorssen, Franck Erkens, and Patroesjka Zuurhout. I also owe a thank you to all the colleagues during my years at SPA, and specifically the departments Research & Analytics, ESA, and ALM. I cannot thank all of you on paper, but I especially thank: Bart Kuijpers, Zina Lekniute, Pieter Kasse, Jan Bonenkamp, Patrick Lighaam, Peter Vlaar, Onno Steenbeek, Hindrik Angerman, Remco Rijntjes, Niek Hensen, Joyce Augustus Vonken, Henk Hoek, Johannes Krotje, and Jordy Duits. We had wonderful lunches, presentations, and discussions together.

Of course, I want to thank all my great fellow PhD students. Foremost I want to thank my office-mates throughout the years: Hanan, Xiaoyue, and Ernst. Even though our research interests were quite different, I appreciated the laughs, discussions, and serious conversations we had on so many diverse topics. Again, I cannot thank all others on paper, but I specifically like to thank: Bart, Lennart, Martijn, Pepijn, Henk, Albert, Junhee, Nora, Jop, Riley, Lieke, Emily, Johan, Lieske, Mario, Kadircan, Tomas, Daniel, Lkhagvaa, Zihao, and Miao. We have shared great times together with personal experiences and academic discussions.

Tenslotte wil ik mijn familie en vrienden bedanken. Lenny, mijn vriendin, bedankt voor jouw onvoorwaardelijke liefde en ruimte die je mij hebt geboden sinds het begin van mijn studietijd. Zonder jou, jouw steun, jouw flexibiliteit en onze leuke momenten samen had ik niet kunnen bereiken waar ik nu ben. Je bent een enorme belangrijke comfortabele toeverlaat voor mij tijdens drukke periodes en ik realiseer me dat ik dan soms afwezig overkom. Volgens mij gaan we een erg mooie toekomst tegemoet. Michel Goossens en Ingrid Goossens-Buermans, mijn ouders, en Judy Goossens, mijn zus, bedankt voor alle mogelijkheden en steun die jullie mij hebben gegeven om mij te ontwikkelen. Jullie relativeringsvermogen, wijze raad en enthousiasme hebben in een bijzondere mate bijgedragen aan hetgeen ik heb bereikt. Op de juiste momenten wisten jullie mij altijd de juiste adviezen te bieden. Papa en mama, jullie hebben mij gemaakt tot wie ik ben en jullie hebben mij altijd geïnspireerd. Zonder jullie steun, trots en enthousiasme was het onmogelijk geweest om mij te storten op mijn opleiding. Judy, jij bent de beste zus die iemand kan wensen en je hebt altijd gezorgd voor een goede dosis relativeringsvermogen en de juiste afleiding om leuke dingen samen te doen; ik hoop dat dit in de toekomst ook zo zal blijven. Uiteindelijk dank ik de rest van mijn familie en vrienden, in het bijzonder Bart, Jeffrey, Sander, Stuart, Cheu Thai, en Willem, voor hun steun en interesse in alle afgelopen jaren.

Thank you all, bedankt allemaal!

Jorgo T.G. Goossens
Breda, November 2022

Contents

Introduction

This dissertation is a collection of four chapters using non-standard preferences to better understand the behavior of asset prices and households. Part I of this dissertation studies the effects of non-standard preferences on equilibrium stock and bond prices. I show that non-standard preferences can be a useful ingredient to explain empirical regularities in the financial markets. Part II of this dissertation studies the relationships between non-standard preferences and household financial decision making. I show that non-standard preferences, sometimes even time varying, can be a useful ingredient to explain observed household behavior. Overall, this dissertation shows that non-standard preferences have explanatory power for macro- and micro-economic phenomena.

Part I of this dissertation concerns "Asset Pricing", and encompasses the first two chapters. The chapters focus on aggregate macro-finance decision making. The non-standard preferences that I study are present bias and regret aversion. In these chapters, I model a representative agent that is either present biased or regret averse, in an otherwise standard model of financial markets. I show, mostly in closed form, that in equilibrium these preferences influence asset prices. Consequently, I bring the models to the data and I present evidence that the models also have an empirical value.

In the first chapter, "Present Bias, Asset Allocation, and Bond Behavior", which is joint work with Bas Werker, we present a present-biased general equilibrium model that explains multiple features of bond behavior. Present-biased investors increase short-term hedge demands to satisfy short-term needs, compared to standard time-consistent preferences. Hence, a present-biased investor drives down short-term yields and requires a premium on long-term bonds, leading to an upward sloping yield curve. Observed bond behavior is best explained using a short-term orientation of at most 1 year, providing an estimate for the investor's "duration of the present".

In the second chapter, my job market paper, "Regret and Asset Pricing", I investigate the consequences of regret aversion for asset prices in an otherwise standard model of financial markets. Accounting for investors' regret aversion can help explain the risk-free rate puzzle, excess stock return volatility, the downward sloping term structure of equity risk premiums, and the predictability of stock returns both in the time series and in the cross section. The model also evaluates bond behavior and predicts a downward sloping real yield curve. I provide an empirical measure of regret which confirms empirically the main model's testable predictions. This research is the first to document the linkage between regret aversion and many stylized facts concerning asset prices. I pursue to understand regret in more detail in future work.

Part II of this dissertation concerns "Household Finance", and encompasses the next two chapters. The chapters focus on micro individual-level decision making. The non-standard preferences that I study in these chapters are present bias, patience, and risk aversion. I measure these preferences among pension fund participants and a representative group of The Netherlands. In each chapter, I relate the measured preferences with one of the following two financial-economic decisions: retirement decisions and investment decisions. In general, I find evidence that non-standard preferences explain such financial-economic decisions.

In the third chapter, "Can Estimated Risk and Time Preferences Explain Real-life Financial Choices?", which is joint work with Marike Knoef and Eduard Ponds, we combine experimentally elicited preferences with administrative microdata to study actual financial decision making. We simultaneously elicit and estimate risk and time preferences in a real-life context, using the convex time budgets method. Within an expected utility framework, we show that individually estimated preferences explain actual retirement decisions up to 82% of our sample for a utility indifference of at most 2% annual certainty equivalent consumption. Freedom of choice by means of a front-loaded annuity creates annual potential welfare gains up to 2.77%, but realized welfare gains are lower or even negative due to suboptimal choices.

In the fourth chapter, "Time-varying Risk and Time Preferences: Relation with Trading Behavior", which is joint work with Marike Knoef, we show that risk and time preferences are time varying, and are related to trading behavior. Using simultaneous elicitation and estimation of risk and time preferences for 2240 individuals, we find that risk aversion and patience correlate positively with daily changes in national COVID-19 hospitalizations. Daily hospital changes are temporarily uncorrelated and a two standard deviation increase in the *daily* change in COVID-19 hospitalizations decreases the *annual* discount rate from 4.3% to 2.6%, increases the *annual* present-bias factor by 0.05, and increases risk aversion by 0.11. At the same time, the disposition effect is time varying: it declines when COVID-19 hospitalizations increase, as investors hold on to winning stocks relatively more. This observation is in line with the time-varying investor's risk aversion and patience, as predicted by intertemporal realization utility models.

In summary, this dissertation has two main messages. First, non-standard preferences matter in the aggregate for asset prices in financial markets. Second, non-standard preferences drive parts of individual-level behavior of households.

Part I

Asset Pricing

Chapter 1

Present Bias, Asset Allocation, and Bond Behavior*

*We thank Anne Balter, Hazel Bateman, Servaas van Bilsen, Stefano Cassella, Rob van den Goorbergh, Frank de Jong, Marike Knoef, Eduard Ponds, Stefan Zeisberger, and seminar participants at the American Finance Association (2021, Poster), World Finance Conference (2021), Econometrics of Option of Markets (2021), Netspar International Pension Workshop (2020) and Netspar Pension Day (2018, 2019) as well as participants of the APG Asset Management and ABP Pension Fund Seminars for useful comments.

Do present-biased preferences shape bond behavior? And if so, what is the investor's duration of the present? Economic agents are subject to present bias (Laibson, 1997; Frederick et al., 2002). Present bias is an old idea and forms a pillar of the modern behavioral economics literature, having added generally to our understanding of financial-economic intertemporal decision making.

We show that multiple features of observed nominal bond behavior can be understood with a simple modification of time preferences in the standard representative-agent equilibrium model. Standard time-consistent preferences have difficulties matching observed bond behavior. The central ingredient in our model is present bias, or time inconsistency, added to basic time-separable power utility. As we will see, a present-biased investor has a higher demand for short-term bonds than long-term bonds, as she is more oriented on the short term than the long term. As a result, a present-biased investor requires a higher premium on long-term bonds. Therefore, compared to standard time-consistent behavior, excess returns on long-term bonds are higher and the slope of nominal yield curve is steeper in equilibrium, which matches the observed data.

Present bias distinguishes explicitly between the short term and the long term. However, the "duration of the present" for present-biased preferences remains an open empirical question (Ericson and Laibson, 2019). How soon is "now" in an investment context? We find that investors using a short-term discounting duration of at most one year matches observed nominal bond behavior. Thus, we establish a connection between asset-pricing models and the experimental literature.

Our results show that present bias matches the observed sign, magnitude and Sharpe ratio of excess nominal bond returns. In historical U.S. data, the excess return on ten-year bonds is 4.10% and the Sharpe ratio is 0.38. Present bias produces an excess return of 4.45% and a Sharpe ratio of 0.42, while standard time-consistent behavior produces an excess return of only 2.80% with a Sharpe of only 0.26. Second, present bias produces yield spreads that match the historical unconditional averages, in line with Liu and Wu (2021), but also the time-series dynamics of yield spreads are closer to the data for present-biased behavior than time-consistent behavior. The average observed yield spread for five-year and ten-year bonds is 1.33% and 1.78% respectively. Assuming present bias yields a five-year yield spread of 1.09% and a ten-year yield spread of 1.95%, while time-consistent behavior produces a flatter term structure of interest rates by slopes of 0.67% and 1.11% respectively. Finally, present bias fits other moments of nominal bond behavior as well, such as term premia (Backus et al., 1989) and predictability (Campbell and Shiller, 1991).

The mechanism that drives our results is due to a preference for higher short-term *hedge* demands and lower long-term *hedge* demands for bonds, compared to a standard time-consistent investor. A present-biased investor puts less value to the distant future and as such cares less about hedging risks for the long run. A naïve present-biased investor prefers short-term bond investments over long-term bond investments to hedge for short-term consumption satisfaction rather than long-term consumption. Consequently, short-term bond prices rise and long-term bond prices decrease, leading to a positive yield spread in our model. The positive yield spread from the present-bias model matches the data, while time-consistent behavior does not. We assume that risk premia are unpredictable such that speculative demand is independent of time preferences, because speculative demands are independent of current time and the investor's horizon (Merton, 1969; Brennan and Xia, 2002). For this reason, equity and cash holdings are not directly affected by present bias in our model. We derive this mechanism in closed form.

Present bias has been extensively documented in diverse groups of the population using financial events, actual consumption events, field studies, and lab studies (see the overview papers of Ericson and Laibson, 2019; Cohen et al., 2020).[1] Evidence for overvaluing the present compared to the future for specifically investors and institutions comes from Porter (1992), Laverty (1996), Bushee (1998), Bushee (2001), Van Binsbergen, Brandt, et al. (2008), and Dikolli et al. (2009). Overall, they report that institutional investors exhibit preferences for the short term due to (i) compensation horizons of managers, (ii) competitiveness among managers, and (iii) pressure to deliver superior short-term performance rather than long-term firm value (i.e., the "short-term pressure" hypothesis, "managerial myopia" and "in-

[1] McClure et al. (2007) provides neuro-scientific evidence for present bias via a neuro-imaging laboratory experiment.

stitutional myopia"). Practitioners and academics have repeatedly documented the excessive focus on short-term performance at the expense of long-term firm value (Lowenstein, 1988; Lehmann, 2004; Haldane, 2010; Stout, 2012; Chung and Low, 2017).[2] We can capture such behavior that overvalues the present, while putting less value to the future, naturally using present bias.

The most popular way to model present bias is by a simple generalization of the standard time-consistent expected utility model. Quasi-hyperbolic discounting introduces an additional present-bias factor $\beta < 1$ (Strotz, 1956; Phelps and Pollak, 1968; Laibson, 1997) and the quasi-hyperbolic model collapses to standard exponential discounting with only a long-term discount factor δ if $\beta = 1$.[3] Time-consistent discounting does not permit any noticeable discounting over short horizons, because short-term orientation in the exponential discounting model would compound to predict counterfactually severe discounting over longer horizons (O'Donoghue and Rabin, 2015).

The time-consistent investor, one that discounts exponentially, makes the same choice no matter when asked. However, this assumption of time consistency has been challenged empirically (see for example Thaler, 1981, Ainslie, 1992, Loewenstein and Prelec, 1992). The naïve present-biased investor holds a belief (that proves incorrect) that her current self can commit future selves to act in a time-consistent manner. Strotz (1956) suggested the assumption of naïvete, and more and more research suggests that naïvete regarding time-inconsistency seems to explain behavior (O'Donoghue and Rabin, 2015).[4]

The investor can invest in stocks, bonds, and cash in the arbitrage-free complete financial market. The risk-free rate is driven by a two-factor Vasicek process, whereas stocks follow a geometric Brownian motion. Our derivations hold for any N-factor Vasicek model, but we confine our empirical analysis to two factors. The

[2]Yan and Zhang (2009) find that institutions with short investment horizons are better informed than long-term institutions.

[3]Financial economists typically call the standard long-term discount factor β, while (experimental) economists call the long-term discount factor typically δ. Quasi-hyperbolic discounting is also known as "beta-delta" discounting and we follow this convention.

[4]An alternative is (partial) sophistication about present bias, which to some extent is a higher form of rationality.

investor solves a dynamic multi-period consumption problem, in combination with her optimal investment demands, under present bias. Our derivations are general and hold for any deterministic discount structure, in which we consider present bias as a special case. In equilibrium, investment demands of the representative investor are set equal to the exogenous supply of U.S. bonds.

Our paper contributes to the literature on modeling, estimating, and explaining bond behavior. Unlike our simple approach of time-separable utility, studies analyzing bond risk premia often use recursive preferences, see for example Piazzesi and Schneider (2006a) and Van Binsbergen, Fernández-Villaverde, et al. (2012). Bansal and Shaliastovich (2013) highlight the preference for early resolution of uncertainty with respect to long-run risks, and Augustin and Tédongap (2020) combine disappointment aversion with a preference for early resolution of uncertainty to explain term structures. Wachter (2006) uses habit formation to explain nominal and real bond behavior, and Gallmeyer et al. (2017) use a preference shock with a monetary policy rule.

Similar to ours, Creal and Wu (2020) and Gomez-Cram and Yaron (2021) study how time preferences shape bond behavior, but they use recursive preferences. They assume shocks to the rate of time preference, while we keep stable time preferences. Relatively little research exists on the stability of time preferences, but using a large two-year longitudinal study, Meier and Sprenger (2015) find that aggregate discount rates are unchanged, which supports the assumption of stable time preferences. Gao et al. (2018) study the term structure of interest rates with stable time preferences in the context of financial intermediaries. Financial intermediaries offer contracts to time-inconsistent hyperbolic discounting investors. The authors show that arbitrage-free linear contracts allow for a unique term structure of interest rates which includes a premium for naïvete, such that more naïve investors in the population create a larger yield spread. Related to their results, we find that the naïve present-biased investor demands a premium on long-term bonds and this premium is higher if the amount of present bias increases such that the yield spread increases.[5]

[5]Riedel (2004) shows that heterogeneity in (stable) time preferences, but the same degree of

Our main finding of an average upward sloping yield curve relates to the preferred habitat model of Vayanos and Villa (2021). They model the term structure of interest rates as an interaction between investors with preferences for specific maturities and risk-averse arbitrageurs, and also find an upward sloping term structure of interest rates. According to the preferred habitat theory, there are investor clienteles for specific maturity segments. The interest rate for a given maturity is mainly driven by shocks affecting the demand of the corresponding clientele. Our approach is different as it relates to the consumption-based approach of macro-finance models, which does not require unobservable 'preferred habitat' parameters. Our results are interpretable in a preferred habitat context since present-biased investors prefer to invest in the habitat of short-term bonds compared to long-term bonds and, thereby, create an average upward sloping yield curve. Stated differently, present bias might provide a micro foundation for preferred habitats.

Our theoretical model can be both given a nominal and a real interpretation, similar to Vayanos and Villa (2021). However, we calibrate our model using nominal yields, so the empirical results should be understood in a nominal context. A motivation is that economic agents have a bias to evaluate outcomes in nominal terms rather than real terms, because economic agents are often found to be subject to money illusion (Shafir et al., 1997). Note that if agents indeed overestimate future inflation, then the increase in nominal term structures does not necessarily lead to increasing real term structures.

We also contribute to the literature on optimal consumption and portfolio decisions for time-inconsistent investors. Marín-Solano and Navas (2010) study optimal consumption and investment for naive and sophisticated individuals in the classical Merton (1969) case, without interest rate risk.[6] They find that the optimal share invested in the stock is independent of time preferences for CRRA utility. We confirm that speculative demands and, therefore, the share invested

risk aversion, can produce additional humps in the yield curve if there is a sharp 'habitat' for a large group of investors.

[6]Zhao et al. (2014) study the same consumption-investment problem as in Marín-Solano and Navas (2010), but with a general discount function, logarithmic utility, and stochastic coefficients.

in stocks is independent of time inconsistency. However, we show that hedge demands and, thereby, the share invested in bonds depends on time preferences if one considers interest rate risk.

1.1 The model

The representative present-biased investor maximizes utility over consumption to find her optimal asset allocation. The investment options are a risk-free asset, a stock, and bonds. Equilibrium prices adjust such that the supply equals demand for these assets.

1.1.1 Preferences

The present-biased investor solves

$$\max_{\{W_{t,T_1},\ldots,W_{t,T_n}\}} \mathbb{E}_t \left[\sum_{j=1}^{n} D(T_j - t) \frac{W_{t,T_j}^{1-\gamma}}{1-\gamma} \right]. \tag{1.1}$$

Here γ is the investor's risk aversion and W_{t,T_j} denotes optimal planned consumption at decision time $t \in [0, T_j]$ for consumption moments $T_j \geq t$. $D(x)$ reflects the discounting associated with delay x

$$D(x) = \begin{cases} \delta^x & \text{if } x = [0, T_S] \\ \beta\delta^x & \text{if } x \in (T_S, T_L]. \end{cases} \tag{1.2}$$

Here $0 < \delta < 1$ is the standard long-term discount factor and $0 < \beta \leq 1$ is the present-bias factor. T_S denotes the duration of the present (i.e., short-term horizon), and after the end of the present the future starts which lasts until the terminal long-term horizon T_L.[7] Consumption during the short term is only discounted by the standard discount factor δ, while consumption in the long term is additionally discounted by the present-bias factor β.

[7] We choose to work with the typically used deterministic horizons, but the end of the present T_S might be stochastic (Harris and Laibson, 2013).

Many variants exist for modeling time inconsistency, but we use present bias captured by the generalized quasi-hyperbolic discount function (Harris and Laibson, 2013).[8] With this functional form, the present-bias factor $\beta = 1$ corresponds to standard exponential discounting (Samuelson, 1937), while $\beta \in (0, 1)$ reflects present bias. Thus, a present-biased investor values consumption after the present T_S less than a standard time-consistent investor. Only in case the investor has time-consistent preferences, which is mathematically equivalent to exponential discounting, consumption W_{t,T_j} is independent of decision time t because the time-consistent investor makes the same choice no matter when asked.[9]

Figure 1.1 presents our main empirical specification, in which the present-biased investor has a short-term orientation of T_S in years and a long-term orientation of T_L in years. Essentially, we group consumption in the present and in the future, such that the investor only has two consumption moments. One aggregated consumption moment during the present $[0, T_S]$ and one aggregated consumption moment during the future $(T_S, T_l]$.[10] The representative investor is infinitely lived such that at each decision time t the investor discounts consumption during the present by the standard discount factor δ and consumption during the future by the additional present-bias factor β. Thus, at each time t, the investor decides how to optimally smooth consumption over the present and the future, and how to optimally invest for the present and the future.

The 2-period model with grouped aggregate consumption and investment provides a helpful discipline. It cleanly separates behavior in durations of the present and the future, revealing the key intuition for present-biased investing along with the key mechanism for the upward sloping yield curve. Nevertheless, we solve our model theoretically for multiple consumption moments as well. Moreover, our the-

[8]The instantaneous quasi-hyperbolic discount function as proposed by Laibson (1997) follows if $T_S = 0$.

[9]Optimal *planned* consumption for in-between dates $t \leq t + h \leq T_j$ follows from $W_{t,T_j,t+h}$, please see Appendix C.

[10]Mathematically, consumption during the present coincides with time T_S and consumption during the future coincides with time T_L. For interpretability and tractability, we let them coincide. The model yields qualitatively the same predictions if consumption differs from the bounds, as long as there is consumption somewhere in the present interval and somewhere in the future interval.

Figure 1.1: Time-line present-bias model. This timeline presents the discount structure for a present-biased investor with quasi-hyperbolic discounting, from the perspective of time $t = 0$. The quasi-hyperbolic discount function discretely drops with size β when the present ends and the future begins. The present-to-future transition occurs at $T_S = 1$ year and the future ends at year $T_L = 10$, respectively the short-term and long-term planning horizons.

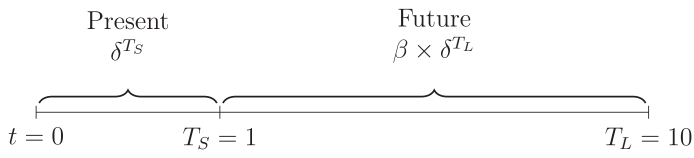

oretical analysis holds for any general discount structure $D(x)$, such that modeling present bias is just a special case.

1.1.2 Financial market

The investor has access to an arbitrage-free complete financial market consisting of a stock, constant-maturity bonds and cash. The short rate r_t is assumed to be affine in an N-dimensional factor of state variables $\boldsymbol{F}_t$

$$r_t = A_0 + \boldsymbol{\iota}' \boldsymbol{F}_t, \tag{1.3}$$

where $\boldsymbol{\iota}$ denotes an N-dimensional vector of ones.[11] The factors $\boldsymbol{F}_t$ follow an N-dimensional multivariate Ornstein-Uhlenbeck process:

$$d\boldsymbol{F}_t = \boldsymbol{\kappa} \left(\boldsymbol{\theta} - \boldsymbol{F}_t \right) dt + \boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,t}, \tag{1.4}$$

where $\boldsymbol{\kappa}$ is a $N \times N$ diagonal mean-reversion speed matrix, $\boldsymbol{\theta}$ is an N-dimensional column vector of long-run averages, $\boldsymbol{\sigma_F}$ is a $N \times N$ lower triangular covariance matrix with strictly positive elements on its diagonal and $\boldsymbol{Z}_{F,t}$ is an N-dimensional column vector of independent standard Brownian motions.

The investment opportunities depend on the pricing kernel in the economy, which determines the expected returns on all securities in the financial market.

[11] We assume that the instantaneous interest rate equals a constant plus the sum of the factors, see for example De Jong (2000).

We assume absence of arbitrage and, thus, the existence of a stochastic discount factor process M_t with $M_0 = 1$:

$$\frac{dM_t}{M_t} = -r_t dt - \boldsymbol{\lambda}' d\boldsymbol{Z}_t, \tag{1.5}$$

where $\boldsymbol{\lambda} = [\lambda_S; \boldsymbol{\lambda}_F]$ and $\boldsymbol{Z}_t = [Z_{S,t}; \boldsymbol{Z}_{F,t}]$ are $(N+1)$-dimensional vectors. λ_S is the constant price-of-risk for the stock and $\boldsymbol{\lambda}_F$ is the N-dimensional vector with the constant prices-of-risk for the bonds. $Z_{S,t}$ is a standard Brownian motion representing shocks to the stock, and it is independent of the shocks $\boldsymbol{Z}_{F,t}$ to the state variables.

The dynamics of the stock price and the dynamics of the N-dimensional vector of (constant-maturity) bond prices follow from

$$\begin{aligned} \frac{dS_t}{S_t} &= (r_t + \boldsymbol{\sigma}'\boldsymbol{\lambda})dt + \boldsymbol{\sigma}' d\boldsymbol{Z}_t \\ \frac{d\boldsymbol{P}_t}{\boldsymbol{P}_t} &= (r_t - \mathcal{B}(\tau)'\boldsymbol{\sigma}_F \boldsymbol{\lambda}_F)dt - \mathcal{B}(\tau)'\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,t}, \end{aligned} \tag{1.6}$$

where $\boldsymbol{\sigma} = [\sigma_S; \boldsymbol{\sigma}_{FS}]$ is an $N+1$-dimensional vector. σ_S is the volatility parameter of the stock and $\boldsymbol{\sigma}_{FS}$ is an N-dimensional vector governing the covariance between stock and bond returns.[12] $\mathcal{B}(\tau)$ follows from

$$\mathcal{B}(\tau) = [\boldsymbol{B}(\tau_1)\iota, ..., \boldsymbol{B}(\tau_N)\iota],$$

with

$$\boldsymbol{B}(t) = (\boldsymbol{I} - \exp(\boldsymbol{\kappa} t))\,\boldsymbol{\kappa}^{-1}. \tag{1.7}$$

in which τ_j for $j = 1, ..., N$ denotes the maturity of bond j. $\boldsymbol{B}(.)$ is an $N \times N$-dimensional matrix, so $\mathcal{B}(.)$ has the same dimensions.

[12] The dynamics of the bond prices follow from the explicit bond prices in the financial market obtained by Lemma 1 in Appendix B for $\alpha = 1$.

1.1.3 Consumption, investment, and the yield curve

We present closed-form solutions for optimal consumption and investment with a general discount structure $D(x)$. The investor determines both the optimal allocation of wealth over time for consumption and the optimal investment strategy to finance her consumption. Since we assume that risk premia are unpredictable, speculative demands are independent of time and the investor's horizon (Merton, 1969; Brennan and Xia, 2002). However, hedge demands depend on time and the investor's horizon, such that hedge demands depend on the discount structure through the channel of discounted consumption. As a result, time preferences of an investor shape the yield curve through hedge demands.

We solve the investor's problem (1.1) with n consumption moments using the martingale method of Cox and Huang (1989) subject to her budget constraint at each time t

$$\mathbb{E}_t\left[\sum_{j=1}^{n} W_{t,T_j} M_{T_j}\right] = W_t M_t, \tag{1.8}$$

with W_t total available wealth at time t and W_{t,T_j} the to be planned investor's consumption at time t with investment horizon is T_j.

Consider the 2-period model with consumption grouped in the present and the future. The investor has to split total available wealth at each time t in an amount to finance her first consumption moment during the present and in an amount to finance her second consumption moment during the future. She uses a part of her total wealth to finance each consumption moment, which one can think of as a money pot (Balter and Werker, 2021). Thus, total wealth at each time t equals the present value of the first money pot and the second money pot. Optimizing the allocation for the first money pot is a stand-alone terminal wealth problem and the second money pot is a stand-alone terminal wealth problem as well. What eventually matters, is the optimal division of total wealth over each of the money pots, i.e., the short-term present and the long-term future.

This approach provides a convenient way to think about intermediate allocation and smoothing of consumption over time by sequentially solving multiple terminal consumption problems, which works thanks to the completeness of the market. The intuition carries easily to the more general n-period model. The theorem below characterizes optimal consumption and investment, as indicated by the asterisks.

Theorem 1. *For an investor that solves (1.1) subject to her budget constraint (1.8):*

1. *The optimal consumption at time t for investment horizon T_j equals W^*_{t,T_j} and is (among others) a function of the discount structure $D(x)$.*

2. *The optimal fraction of wealth invested in bonds, a stock, and cash for a single terminal consumption problem equals $\pi^*_i(t,T) = (\boldsymbol{\pi}^*_B(t,T), \pi^*_S, \pi^*_C)$.*

   - *The optimal fraction of wealth invested in constant-maturity τ-year bonds at time t for investment horizon T is an N-dimensional vector*

   $$\boldsymbol{\pi}^*_B(t,T) = \frac{1}{\gamma}\left(\frac{\lambda_S}{\sigma_S}\boldsymbol{\sigma}'_{FS} - \boldsymbol{\lambda}'_F\right)(\boldsymbol{\sigma}_F)^{-1}(\mathcal{B}(\tau)')^{-1} + (1-1/\gamma)\boldsymbol{\iota}'\boldsymbol{B}(T-t)(\mathcal{B}(\tau)')^{-1}.$$

   - *The optimal fraction of wealth invested in stocks is*

   $$\pi^*_S = \frac{\lambda_S}{\gamma\sigma_S}.$$

   - *The remainder is invested in cash.*

3. *The optimal fraction of total wealth invested in asset i = {stock, constant-maturity τ-year bond, cash} at time t for all consumption moments T_j is*

$$\omega^*_i(t,T_j) = \frac{\sum_{j,j>t}^{n}\pi^*_i(t,T_j)W^*_{t,T_j}}{\sum_{j,j>t}^{n}W^*_{t,T_j}} \tag{1.9}$$

We prove this formally in Appendix C. At a less formal level, the investor's problem (1.1) clearly depends on the discount structure $D(x)$. Because discounting influences optimal consumption W^*_{t,T_j}, the optimal fraction of total invested wealth $\omega^*_i(t,T_j)$ reacts to discounting accordingly since the distribution of total available wealth over the money pots changes. Because risk premia are unpredictable, speculative demands are constant such that the key mechanism that shapes bond behavior is through the hedge demands.

More specifically, Statement 3 shows that the optimal proportions $\omega_i^*(t, T_j)$ for each asset i invested over all money pots T_j depend on two things: (i) the optimal fraction of wealth invested per terminal consumption problem $\pi_i^*(t, T_j)$, and (ii) the actual allocation of total wealth W^*_{t,T_j} over all consumption moments.

Starting with consumption, Statement 1 implies that the actual allocation of total wealth W^*_{t,T_j} over all consumption moments depends on the investor's time preferences through the discounting structure $D(x)$. The consumption strategy W^*_{t,T_j} describes the investor's *actual* consumption at time t if the investment horizon is T_j. The consumption rule $W^*_{t,T_j,t+h}$ describes the investor's *planned* consumption at each time t for all in-between dates $t \leq t + h \leq T_j$ with investment horizon T_j.[13] Consumption for both actual and planned strategies depend on the discount structure.

For a present-biased investor the actual consumption path differs from her planned consumption path. Because a present-biased investor is time inconsistent, she deviates from her plans. At time t, the present-biased investor plans how much to consume at future date $t + h$. But, if she actually arrives at date $t + h$, she decides to consume more than initially planned at time t.

For a standard time-consistent investor, the actual consumption path is identical to her planned consumption path, i.e., $W^*_{t,T_j,t+h} = W^*_{t+h,T_j}$ for $t + h = t, ..., T_j$. Because the investor is time consistent, she sticks to her plan. The time-consistent investor consumes at each date $t + h$, what she initially planned at time t to consume at time $t + h$ given the state of the world. As a result, the ratio of current consumption to total wealth is higher for a present-biased investor than a time-consistent investor at every time t.[14] So, similar to Marín-Solano and Navas (2010) and Zou et al. (2014), we find that naïve present bias increases the consumption rate compared to standard time-consistent behavior, however we generalize to a model with interest rate risk.

[13] See Appendix C for the derivation.

[14] This consumption rate is concave upward sloping and reaches 1 at the terminal horizon.

Regarding investment, Statement 2 present the optimal fraction of wealth invested per single money pot. The allocation to bonds $\boldsymbol{\pi}_B^*(t,T)$ depends on time t and the investor's horizon T, while the allocation to stocks π_S^* is constant (Brennan and Xia, 2002). The investment strategy for bonds has two components: speculative demand and hedge demand. The investment strategy for stocks has only speculative demand, because risk premia are unpredictable (Merton, 1969). The bond and stock speculative demands for a single terminal wealth problem depend on risk aversion, volatility, and prices-of-risk but are independent of current time t and the investment horizon T.[15] Consequently, the speculative demand for total invested wealth $\omega_i^*(t,T_j)$ is independent of time preferences.[16]

However, the bond hedge demands for a single terminal wealth problem depend on current time t and the investment horizon T, such that hedge demands for total invested wealth $\omega_i^*(t,T_j)$ depend on the investor's time preferences through consumption W_{t,T_j}^* which depends on t and T_j. Similar to planned consumption, one can also easily derive the planned optimal investment by using $W_{t,T_j,t+h}^*$ rather than W_{t,T_j}^* in Theorem 1. Investment for both actual and planned strategies depend on the discount structure.

Overall, a present-biased investor allocates more wealth to consumption moments during the present, while allocating less wealth to future consumption moments. Consequently, the fraction of wealth allocated to the present over total wealth increases, while the fraction of wealth allocated to the future over total wealth decreases. Since hedge demands depend on the investor's discounting through the delay between T_j and t and because the present-biased investor has a higher fraction of total wealth invested for the short-term present, the fraction of short-term hedge demand over total hedge demand increases. Thus, a present-biased investor has a relatively larger demand for short-term bonds, such that short-term bond prices rise and short-term yields decrease. The investor uses this

[15] Although stock market returns are influenced by all $N+1$ sources of risk, only stock market shocks matter for the optimal investment fraction to stocks.

[16] Mathematically, since speculative demand is independent of time t and horizon T_j you can take the speculative part of $\pi_i^*(t,T)$ out of the summation in $\omega_i^*(t,T_j)$, such that $\omega_i^*(t,T_j)$ is constant as well because there is no relation anymore with the discount structure through optimal consumption W_{t,T_j}^*.

short-term hedge component to hedge against short-term unfavorable changes in the state variables. Vice versa, the fraction of long-term hedge demand over total hedge demand decreases, such that long-term bond prices drop and long-term yields increase. So, a present-biased investor requires a premium to hold long-term bonds. This is the key mechanism that drives the main result of an upward sloping nominal yield curve.

Theorem 1 generalizes the results of Marín-Solano and Navas (2010) and Zou et al. (2014). First, we extend their models to one with interest rate risk through a N-factor Vasicek model. Consequently, time preferences *do* influence the asset allocation through hedge demands and, therefore, shape the yield curve. Second, our result holds for any discount structure $D(x)$, in which we treat present bias as a helpful special case. Additionally, we bring our model to the data using an equilibrium approach to study the model's descriptive abilities.

1.1.4 Equilibrium bond behavior

Here, we extend the partial equilibrium to the concept of general equilibrium such that we can compute equilibrium bond behavior. The investor's optimal investment demands equal the exogenous historical U.S. supply of bonds. The definition below characterizes the general equilibrium.

Definition 1. *The market is in general equilibrium if both of the following conditions are satisfied:*

1. *The representative investor solves her consumption problem (1.1) subject to her budget constraint (1.8).*

2. *Bond markets clear continuously, such that for all t we have:*

$$\boldsymbol{\omega}_B^*(t, T_j) = \hat{\boldsymbol{w}}_B(t)$$

where $\boldsymbol{\omega}_B^*(t, T_j)$ is the optimal bond demand from Theorem 1 and $\hat{\boldsymbol{w}}_B(t)$ is the exogenous supply of bonds in the economy, both being N-dimensional vectors.

The first condition determines the demand for the stock, bonds, and cash in the economy by Theorem 1. The representative investor is infinitely lived by definition. Thus, at all time t, the representative investor solves her consumption problem with discount structure $D(x)$.

The second condition states that demand and supply for bonds is equal to each other, implying the general equilibrium. The market is complete, so the number of bonds with different maturities equals the number N of state variable factors. Hence, to match the supply and demand of N bonds, we need N free parameters for an exactly identified system. Standard macroeconomics uses prices to match demand and supply. We follow this approach, and we use the N prices-of-risk $\boldsymbol{\lambda}_F$ as free parameters to match demand with supply in equilibrium for all N bonds. The estimated equilibrium prices-of-risk are $\hat{\boldsymbol{\lambda}}_F$, which we solve for numerically.

The estimated prices-of-risk identify the equilibrium yields and equilibrium bond returns. In our economy, the investor may still invest in the stock market, but we do not impose equilibrium in that market. Since the optimal stock allocation is independent of the discount structure, there is no need to examine how discounting influences equilibrium stock prices.

1.2 Calibration

In this section we calibrate the model to U.S. historical bond data and we calibrate preferences.

1.2.1 Data

On the one hand, we require data and Kalman estimation to calibrate the financial market. On the other hand, we require data to determine the historical supply of U.S. bonds in the economy.

Starting with the latter, to determine the supply of bonds in the economy we use monthly data from 1 October 1976 to 1 January 2019 on U.S. government debt. The source is Datastream's U.S. Maturity Distribution, Interest Bearing Public Debt. The ratio of short-term debt (i.e., maturities smaller than 1 year, and between 1 and 5 years) to total debt declines during our sample period from roughly 84% to 69%. The ratio of long-term debt(i.e., maturities exceeding 5 years) to total debt increases from 15% at October 1976 to 30% at January 2019. During our sample period, on average 73% of the debt has a maturity lower than 5 years, while 17% of the debt has a maturity higher than 5 years.[17]

Empirically, not all the parameters of the affine model can be identified and certain assumptions are necessary. For identification purposes, we normalize the long-run means of the factors $\boldsymbol{\theta}$ to zero in (1.4), in line with De Jong (2000). We assume that there are two factors driving the short rate r_t in (1.3). Hence, to ensure market completeness, we need two constant-maturity bonds of different maturities. To determine the maturities of these two bonds, we rely on historical U.S. government debt data. We assume that all debt with maturity lower than 5 years reflects a U.S. government bond with a maturity of 3 years, while all debt with maturity higher than 5 years reflects a U.S. government bond with a maturity of 10 years.[18] So, in equilibrium, we match the supply of three-year and ten-year U.S. government bonds with the model-implied demand for both three-year and ten-year bonds respectively on a monthly basis during 1 October 1976 to 1 January 2019.

To calibrate the financial market, we use monthly zero-coupon yields with multiple maturities and monthly stock returns from 1 October 1976 to 1 January 2019.[19] Regarding the yield data, we take the yields with maturities 2 years, 3 years, 5 years, 7 years, and 10 years from the Treasury Constant Maturity Rates, while the yields with a maturity of 3 months and 1 year are from Treasury Bills: Secondary market rates, because these series contain less missing values. Source of the yield data is the Fed database at the St. Louis Federal Reserve. All yield data is reported in percent per annum, and annualized using a 360-day year or bank interest. We use the 3 month treasury bill to proxy for the instantaneous

[17]Table 1.6 in Appendix D provides additional information on U.S. government debt.

[18]Horvath et al. (2017) follow a similar approach.

[19]Regarding the choice of the sample period, we did not go further back than October 1976, because yield curve estimations and debt data from earlier years contain relatively high standard errors and missing values.

risk-free rate.[20] As for the stock data, we use Kenneth R. French's Website, which is a value-weighted index of all CRSP firms incorporated in the U.S. and listed on the NYSE, AMEX or NASDAQ.

We calibrate our model to the monthly market data by using a standard Kalman filter with maximum likelihood estimation. Exact discretization of our economy is possible by writing the financial market processes as a multivariate Ornstein-Uhlenbeck process. Regarding the choice of maturities for the yields in the estimation, we follow De Jong (2000). For maturities over 10 years, bond data is somewhat scarce, so interpolation is less accurate. Very short-term interest rates of one and two months exhibit sometimes exceptionally large one-period changes. We feel more confident using interest rates of 3 months and longer, and of 10 years and shorter. To keep the estimation feasible, we confine ourselves to four maturities: 3 months, 1 year, 5 years, and 10 years.

The maximum likelihood estimation starts from multiple initial values to prevent that the optimizer finds a local optimum. We estimate 12 model parameters and 4 measurement errors for each maturity. We assume that each maturity has its own measurement error, such that the variance of the errors depends on maturity (Geyer and Pichler, 1999). Each error is drawn from a uniform distribution, and both serially and cross-sectionally uncorrelated.[21] Table 1.5 in Appendix D shows the estimates for the financial market.[22]

1.2.2 Preferences

For our benchmark model, we use standard time-separable power utility with a typical CRRA risk-aversion parameter of $\gamma = 10$ (Mehra and Prescott, 1985),

[20] Empirically using shorter maturities, e.g., one or two months, suffers from data issues such as large changes within one period and fewer observations. So, we feel most confident approximating the theoretical instantaneous risk-free rate by the empirical 3 month rate.

[21] Assuming only one error variance for all maturities is a possible and convenient simplification.

[22] Factor 2 exhibits stronger mean reversion than factor 1. Factor 1 (level factor) is very highly correlated with the ten-year yield and factor 2 is very closely related to the spread between the three-month yield and the ten-year yield (slope factor). The long-term mean of the short-rate A_0 is estimated at 3.59%. Both prices-of-risk for the two factors are negative. The negative sign of $\sigma_{F,21}$ implies that there is negative correlation between the two factors. The price-of-risk for the stock λ_S equals 0.4095, and the volatility of the stock σ_S is 15.04% per annum.

which we keep fixed throughout the analysis. Regarding time preferences, we need durations of the present T_S and the future T_L as well as values for the present-bias factor β and the long-term discount factor δ. Institutions typically perform decentralized investment with different investment horizons (Van Binsbergen, Brandt, et al., 2008). Managers are usually compensated on an annual basis, such that their investment horizon is generally relatively short. In contrast, a Chief Investment Officer, who employs multiple managers to implement investment strategies, generally has a much larger horizon because of long-term mandates from the institution. Hence, as institutions are exemplary for our representative agent, in our benchmark model we set the short-term investment orientation T_S to 1 year (i.e., duration present), while we set the long-term investment orientation to $T_L = 10$ years (i.e., duration future). Benartzi and Thaler (1995) also provide evidence that investors tend to evaluate their portfolios on a one-year basis.

We set the annual long-term discount factor to $\delta = 0.97$, which is a typical value in the asset pricing literature. For a standard time-consistent investor with exponential discounting, this value implies an annual discount rate of 3% no matter when the investor is asked. Laibson (1997) argues that the annual present-bias factor β is in the interval $(0, 2/3)$ if δ is close to unity. In our analysis, we take roughly the mean value of this interval and set $\beta = 0.35$ on an annual basis for our naïve representative present-biased investor. We feel comfortable using these values, since Laibson, Maxted, et al. (2015) estimate a present-bias factor of $\beta = 0.35$ and a long-term discount factor of $\delta = 0.97$ on an annual basis for specifically a naive agent.[23]

1.3 Empirical findings

In this section we characterize the model's behavior. We compare the present-bias model with U.S. historical data. We evaluate the model's behavior by studying model-implied equilibrium nominal bond returns, yields, term premia, and pre-

[23]Their estimates are particularly helpful because the authors distinguish between time-preference estimates for sophisticated and naive agents.

dictability (Campbell and Shiller, 1991). As we will show, a present-biased investor with a short-term focus of at most 1 year (i.e., "duration of the present") matches many features of nominal bond behavior. The mechanism is through changes in asset allocation compared to time-consistent behavior.

1.3.1 Excess returns

Our first measure to evaluate the model's behavior is the excess return on bonds. The mean and standard deviation of excess bond returns are a popular measure for characterizing bond behavior (Rudebusch and Swanson, 2008). Model-implied excess bond returns are defined as the expected instantaneous τ-year bond return in excess of the the short rate r_t

$$r_B(\tau) = -(B(\tau)\iota)'\boldsymbol{\sigma}_F\hat{\boldsymbol{\lambda}}_F, \tag{1.10}$$

along with its standard deviation

$$s_B(\tau) = \sqrt{(B(\tau) * \iota)'\boldsymbol{\sigma}_F\boldsymbol{\sigma}_F'(B(\tau) * \iota)}. \tag{1.11}$$

The excess bond returns depend on the discount structure $D(x)$ through the equilibrium estimated prices-of-risk $\hat{\boldsymbol{\lambda}}_F$, such that (1.10) expresses the equilibrium excess bond returns. The standard deviation of bonds is independent of the deterministic prices-of-risk and, hence, independent of discounting. For this reason, we study the Sharpe ratio — the mean of the excess return (1.10) divided by its standard deviation (1.11) — which is a comprehensive measure of bond behavior. In a similar vein, we compute the excess return and Sharpe ratio of the stock in the financial market.

Table 1.1 shows the excess returns and Sharpe ratios on nominal bonds and stocks, on average over time. We report moments as found in the data, as in the present-bias model with several durations of the present and as in the time-consistent model. We first describe our benchmark present-bias model which assumes a duration of the present of 1 year and, afterwards, we describe the other

Table 1.1: Excess returns and duration of the present. This table reports average returns in excess of the risk-free rate and Sharpe ratios of a nominal 3-year bond, a nominal 10-year bond and a stock. The Sharpe ratio is the mean of the excess return divided by the standard deviation. The column "Data" gives statistics for excess returns as realized in the data. The column "Present bias" gives statistics as implied by the present-bias model (present-bias factor $\beta = 0.35$). The column "Time consistency" gives statistics as implied by standard time-consistent discounting (present-bias factor $\beta = 1$). The column "Duration present" gives statistics as implied by the present-bias model for a duration of the present equal to 3 months with present-bias factor $\beta = 0.7$, and 3 years with present-bias factor $\beta = 0.05$. Values are annualized. Data are monthly and run from 1 October 1976 to 1 January 2019.

		Duration present			
	Data	3 months	1 year	3 years	Time consistency
3-year bond					
Mean	1.90	1.67	1.60	1.05	1.06
Sharpe	0.48	0.42	0.41	0.27	0.27
10-year bond					
Mean	4.10	4.40	4.45	3.11	2.80
Sharpe	0.38	0.41	0.42	0.29	0.26
Stock					
Mean	7.27	7.51	7.48	7.05	7.02
Sharpe	0.48	0.49	0.49	0.46	0.46

durations.[24]

The table confirms that our benchmark present-bias one-year model supports the data in terms of sign, magnitude, and Sharpe ratio. Empirically, excess bond returns have a positive relation with maturities up to 10 years (Boudoukh et al., 1999). Indeed, model-implied excess returns are increasing in maturity. In the present-bias one-year model, the excess return on a ten-year bond is 4.45%, somewhat higher than its mean in the data. The excess return on a three-year bond is 1.60%, somewhat lower than its mean in the data. Although, by construction, we do not fit the standard deviation of bonds well, the present-bias one-year model produces Sharpe ratios fitting the observed data closely. Clearly, a time-consistent investor produces excess returns and Sharpe ratios that are too low compared to

[24] Table 1.8 in Appendix D shows excess returns and Sharpe ratios on nominal asset for an impatient time-consistent investor. The impatient time-consistent investor produces results similar to the present-bias model, however an impatient time-consistent investor produces counterfactually high annual discount rates over long-term horizons and as such is undesirable.

the data, such that present bias yields a large but simple improvement over time consistency. Time consistency produces a three-year excess bond return of only 1.06% and a ten-year excess bond return of only 2.80%.

Besides, the present-bias one-year model fits excess stock returns reasonably well, such that the equity premium puzzle is absent in our model. By definition, excess stock returns depend on both the equilibrium prices-of-risk $\hat{\boldsymbol{\lambda}}_F$ and the stock price-of-risk λ_S, so the discount structure $D(x)$ influences equilibrium excess stock returns partially. Unlike others that find distorted predictions in the stock market, we find reasonable equilibrium stock behavior in a model tailored to capture bond behavior. The remaining part of the predictive power for stocks comes from the Kalman estimated financial market parameters.

We now first discuss the mechanism for our main result and, then, we discuss the several durations of the present.

Asset allocation

The relatively higher excess return on a ten-year bond under present bias is a result of a higher hedge demand for short-term bonds and a lower demand for long-term bonds compared to standard time-consistent investing. Table 1.2 summarizes the optimal hedge and speculative holdings, as fractions of total wealth, under present bias and time consistency for a short-term three-year bond, a long-term ten-year bond, a stock, and cash, when the agent has two consumption moments according to the timeline in Figure 1.1. Thus, the optimal asset allocation for the four assets is weighted by the wealth allocated to both consumption moments, i.e., during the present and the future.

The short-term hedge demand for a present-biased investor is 9 percentage points higher than for a time-consistent investor, while the long-term hedge demand for a present-biased investor is 17 percentage points lower than for a time-consistent investor. The key mechanism is a higher consumption rate for a present-biased investor: a present-biased investor is less concerned with the future and, therefore, cares less about hedging risks for the long term but she demands to hedge short-term risks. If the short-term focus increases (i.e. a lower present-bias

Table 1.2: Optimal asset allocation with consumption in the present and future. Optimal fractions of total wealth — divided in hedge and speculative demands — invested in a 3-year bond, a 10-year bond, a stock index, and cash (total demand adds up to 1) according to Theorem 1. The agent has two consumption moments according to the timeline in Figure 1.1. The column "Present bias" presents results as implied by the present-bias model (present-bias factor $\beta = 0.35$). The column "Time consistency" presents results as implied by standard time-consistent discounting (present-bias factor $\beta = 1$).

Asset	Present bias			Time consistency		
	Hedge	Speculative	Total	Hedge	Speculative	Total
3-year bond	0.48	2.44	2.92	0.39	2.44	2.83
10-year bond	0.04	-0.63	-0.59	0.21	-0.63	-0.42
Stock	-	0.27	0.27	-	0.27	0.27
Cash	-	-	-1.61	-	-	-1.68
Total	-	-	1.00	-	-	1.00

factor β), then the hedge demand for short-term bonds rises.

Note that present bias and impatience are different from myopic behavior. A myopic investor ex-ante completely neglects the future, while a present-biased investor puts more weight to the present than the future, but still foresees that there is a future consumption moment. Myopic behavior would lead to no hedge demands (Van Binsbergen, Brandt, et al., 2008), such that myopic behavior would be virtually unaffected by discounting as hedge demands are virtually absent.

Table 1.7, Panel A, in Appendix D shows the optimal asset allocation for a present-biased investor for several present-bias factors. The main observation is that a lower present-bias factor drives up short-term hedge demands but lowers long-term hedge demands, because the less the investor cares about future consumption opportunities. An extremely low present-bias factor β of 0.05 causes the investor to go short in terms of ten-year bond hedge demands. Panel B shows that a standard time-consistent investor with a long-term annual discount factor of $\delta = 0.86$ replicates the optimal asset allocation decisions for a present-biased investor with $\beta = 0.35, \delta = 0.97$, however such exponential discounting implies counterfactually high discounting over long-term horizons inconsistent with the experimental evidence (O'Donoghue and Rabin, 2015).

In line with Theorem 1, the speculative demands are independent of the discount structure. Thus, in a financial market with interest rate risk and no predictability of Sharpe ratios, the optimal stock holdings are equal for present-biased and time-consistent investors. This result extends the finding of Marín-Solano and Navas (2010) in the classical Merton model.

Duration of the present

Coming back to Table 1.1, what is the investor's duration of the present, or what is the investor's short-term orientation? This is an open empirical question in the behavioral literature (see Ericson and Laibson, 2019). Using a discounting model based on present bias requires a distinction between now and later, i.e., the short term and the long term. Experimental evidence over consumption (e.g., juice, water, and effort) finds a duration of the present ranging from a few minutes to a few weeks (McClure et al., 2007; Augenblick, Niederle, et al., 2015). However, a different picture emerges if we study structural models with annual periods, in which models treat consumption anytime this year as immediate (Angeletos et al., 2001). For a stock-bond portfolio, Benartzi and Thaler (1995) report that investors are myopic with an investment evaluation period of 1 year. Van Binsbergen, Brandt, et al. (2008) argue that asset managers are compensated on an annual basis such that their investment horizon is relatively short. To the extent of our knowledge, there are no papers measuring the duration of the present for explicitly bond investment decisions and we try to fill this gap here.

We repeat our earlier analysis, however we estimate equilibrium excess returns for a short-term horizon of $T_S = 3$ months and for a short-term horizon of $T_S = 3$ years in Figure 1.1. For $T_S = 3$ months, the investor has one aggregated consumption moment during the present $[0, 3 \text{ months}]$ and one aggregated consumption moment during the future $(3 \text{ months}, 10 \text{ years}]$. Similarly, the time-line for a short-term horizon of 3 years follows. Because a short-term horizon of 3 months makes the investor by definition more short-sighted — the first consumption moment arrives earlier than the one-year benchmark — the investor uses a higher present-bias factor (i.e., less present biased). A short-term horizon of 3 years makes the investor more far-sighted — the first consumption moment arrives later than the one-year benchmark — such that the investor uses a lower present-bias factor (i.e., more present biased).

We do not allow ourselves the luxury of selecting the present-bias factors that would fit the data best, and we set the present-bias factor $\beta = 0.7$ for $T_S = 3$ months and we set the present-bias factor $\beta = 0.05$ for $T_S = 3$ years. The value $\beta = 0.7$ is in line with estimated present-bias factors in the experimental literature that focuses on eliciting present bias on a daily, weekly or monthly basis (see for example Ericson and Laibson (2019)). $\beta = 0.05$ indicates very strong present bias, but Laibson, Maxted, et al. (2015) do estimate even lower values of present bias for a sophisticated agent in a long-run life-cycle model with a CRRA risk aversion parameter of 3. All other parameters remain identical to the benchmark present-bias one-year model.

Table 1.1 shows that present bias with a duration of at most 1 year fits the cross section of excess returns as found in the data well, while a duration of 3 years fails to fit the data. In terms of magnitude, a present duration of 3 months even yields a slightly better match with the data than our benchmark present duration of 1 year. For a duration of the present of 3 years, excess returns and Sharpe ratios fail to quantitatively match the data although excess returns are still increasing in maturity. A present-biased investor with a duration of 3 years comes close to the standard time-consistent investor in terms of bond return behavior.

So, the main take-away is that observed bond behavior is best explained if the investor's short-term orientation is at most 1 year. We conclude that the present for bond investors has a duration of at most 1 year. We feel most confident with reporting a duration of the present of at most 1 year, compared to 3 months, as the 1 year duration is intuitive and in line with the annual myopic portfolio evaluation period of Benartzi and Thaler (1995). They are successful in explaining the equity premium puzzle by means of an evaluation period of 1 year. It may not come as a surprise that a similar short-term orientation is able to explain asset-pricing puzzles in the bond market.

1.3.2 Yield curve

Our second measure to evaluate the model's behavior is the yield curve. We study the shape of the yield curve by nominal yields and yield spreads. The model-implied yield on zero-coupon bonds is a linear transformation of the factors, where the intercept and factor loadings are time-invariant functions of time to maturity (vector) τ

$$\boldsymbol{Y}_t(\tau; \hat{\boldsymbol{\lambda}}_F) \equiv -\ln \boldsymbol{P}_t(\tau)/\tau = -A(\tau)/\tau + \boldsymbol{\iota}'\boldsymbol{B}(\tau)\boldsymbol{F}_t/\tau. \tag{1.12}$$

Here $\boldsymbol{P}_t(\tau)$ is the price of zero-coupon bonds with time-to-maturity τ, as given by (1.22) in Appendix C. The intercept $A(\tau)$ is a function of the discount structure $D(x)$ through the equilibrium estimated prices-of-risk $\hat{\boldsymbol{\lambda}}_F$, such that (1.12) expresses the equilibrium yields. The yield spread, also known as the slope of the yield curve, is simply the difference between the yield to maturity on a long-term bond and the instantaneous risk-free rate.

Table 1.3 shows that present bias matches the cross section of nominal yields as found in the data, averaged over time. A well-established feature of the unconditional nominal yield curve is that it slopes upward as maturity increases (Piazzesi and Schneider, 2006a). Panel A confirms that the average yield curve on five-year and ten-year nominal bonds is upward sloping as the yield spreads are positive. Clearly, the present-bias model does a better job in matching the observed yield spreads than the time-consistent model. Present bias produces a yield curve that matches on average the sign and magnitude of the slope of the yield curve in the data. A time-consistent investor (i.e., present-bias factor $\beta = 1$) produces a flatter relationship between yields and maturity as the yield spreads are lower than observed in the data and in the present-bias model. An implication of the present-bias model is that bond term premia are increasing in maturity for which we find support below in line with the data of Boudoukh et al. (1999).

Panel B shows the yield curves in the data, in the present-bias model, and in the time-consistent model. The average yield on the five-year nominal bond in the present-bias model is equal to 5.82%, similar to the data mean of 5.85%. The

Table 1.3: Yield curves and yield spreads of nominal bonds. Panel A presents the τ-year yield spread, or slope, defined as the difference in nominal yields between the τ-year bond and the 3-month bond. Panel B presents nominal yields for maturities of 3 months to 10 years. The column "Data" gives statistics for nominal yields on nominal bonds as realized in the data. The column "Present bias" gives statistics for nominal yields on nominal bonds as implied by the present-bias model (present-bias factor $\beta = 0.35$). The column "Time consistency" gives statistics for nominal yields on nominal bonds as implied by standard time-consistent discounting (present-bias factor $\beta = 1$). Values are annualized. Data are monthly and run from 1 October 1976 to 1 January 2019.

Panel A: Yield spreads

Maturity	Data		Present bias		Time consistency	
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
5 years	1.33	0.97	1.09	1.02	0.67	1.02
10 years	1.78	1.22	1.95	1.38	1.11	1.35

Panel B: Yield curves

Maturity	Data		Present bias		Time consistency	
	Mean	Std. dev.	Mean	Std. dev.	Mean	Std. dev.
3 months	4.52	3.59	4.73	3.91	4.71	3.91
1 year	5.03	3.83	4.93	3.82	4.84	3.83
5 years	5.85	3.50	5.82	3.53	5.38	3.60
10 years	6.29	3.23	6.68	3.26	5.82	3.40

average yield of the risk-free rate is 4.73%, about its mean in the data. Observe that the risk-free rate in the present-bias model and time-consistent model are nearly equal. The difference in yield spreads between both types of discounters is mainly driven through relatively higher yields for ten-year bonds in the present-bias model. Because the present-biased investor prefers to hedge her consumption in the short-run more than the time-consistent agent, the present-biased investor requires a higher yield on the long-term ten-year bond. The present-bias model also produces standard deviations of bond yields that are closer to the data than the time-consistent model. Long-maturity yields are less volatile than short-maturity yields. For the five-year yield, the standard deviation implied by the model is 3.53%, while in the data it is 3.50%. However, much of the explained standard deviation comes from fitting the Kalman filter on the data.

Rather than aggregate market moments of yields, a somewhat more challenging task is to see if the present-bias model fits the time series as well. In the top panel of Figure 1.2, we plot the level of the risk-free rate in percentage points over time, defined as the yield on three-month bonds. The solid line depicts the estimates from our present-bias model, and the dotted line is the data. In the bottom panels of Figure 1.2, we plot the slope of the yield curve defined as the five-year yield minus the three-month yield and the ten-year yield minus the three-month yield. The panel titles also display the mean absolute errors.

Our model-implied risk-free rate traces the data well. Both yield spreads, or slopes of the yield curve, match the time-series dynamics of the data as well. The model matches many of the short- and long-run fluctuations in the nominal data, considering that our model uses deterministic prices-of-risk. A potential reason is that we combine structural modeling of time preferences with a Gaussian affine two-factor term structure model. The mean absolute errors (MAE) for the risk-free rate, five-year yield spread, and ten-year yield spread are respectively 0.41, 0.46, and 0.45 percentage points.[25] For the time-consistent model, the mean absolute errors for the yield spreads are higher: 0.71 for the five-year yield spread, and 0.69 for the ten-year yield spread. While the mean absolute error 0.40 for the risk-free

[25]These pricing errors are two to three times smaller than the mean absolute errors reported by Creal and Wu (2020), who use a sample period from 1959 to 2014.

Figure 1.2: Time series of the risk-free rate and the yield spread implied by the present-bias model and in the data. The solid lines show the implied risk-free rate or n-year yield spread in equilibrium when monthly debt data is fed into the present-bias model. The dotted lines show the realized time series of the monthly data. The risk-free rate is the nominal yield on a 3-month bond. The n-year yield spread, or slope of the yield curve, is the difference in nominal yields between the n-year bond and the 3-month bond. Values are annualized. Data run from 1 October 1976 to 1 January 2019.

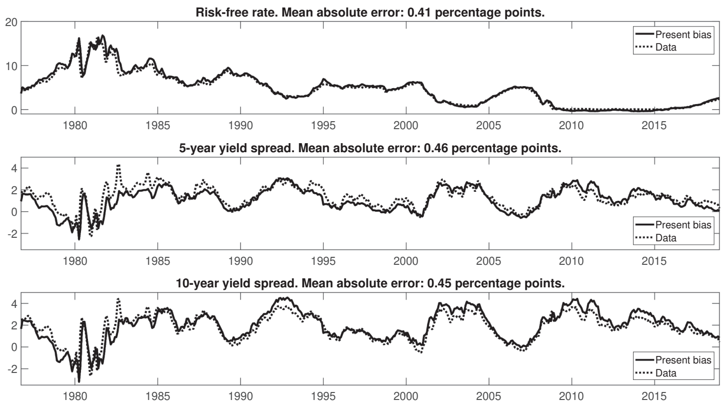

rate is similar. In terms of root mean squared errors (RMSE), the present-bias model also produces lower errors for the yield spreads than the time-consistent model. The root mean squared errors for the five-year and ten-year yield spreads in the present-bias model are 0.61 and 0.59 respectively, while for the time-consistent model they are 0.85 and 0.84 respectively.

1.3.3 Term premia and predictability

Our final measures to evaluate the model's behavior are bond term premia and predictability. Arguably a clean conceptual measure of long-term bond risk is the term premium, or the bond risk premium. The term premium measures the additional compensation a risk-averse investor needs to choose a long-term bond over a short-term bond. The term premium is typically expressed as the difference between the yield on the bond and the unobserved risk-neutral yield for that same bond (Rudebusch and Swanson, 2008). Thus, the term premium is not directly observed in the data and must be inferred using term-structure models (or other methods). The model-implied term premium equals

$$tp_t(\tau) = y_t(\tau) - \tilde{y}_t(\tau). \tag{1.13}$$

Here $y_t(\tau)$ is the model-implied yield from (1.12) at time t on a τ-year bond, which depends on the discount structure through the estimated equilibrium prices-of-risk. $\tilde{y}_t(\tau)$ is the risk-neutral yield on that same bond, which we obtain by setting both prices-of-risk for the stock λ_S and for the factors $\boldsymbol{\lambda}_F$ to zero. By construction, the risk-neutral yields are independent of the discount structure.

We measure the model's predictability by the "long-rate regressions" of Campbell and Shiller (1991)

$$y_{t+1}(\tau-1) - y_t(\tau) = \text{constant}_\tau + \alpha_\tau \frac{1}{\tau-1}(y_t(\tau) - y_t(3\text{ month})) + \varepsilon_{\tau,t+1}. \tag{1.14}$$

Here the dependent variable is the change in the τ-period zero-coupon yield from period t to $t+1$, and the independent variable is the slope of the yield curve at time

t divided by $\tau - 1$. The intercept 'constant$_\tau$' and slope coefficient α_τ are maturity specific. Thus, the slope coefficient from a Campbell-Shiller (1991) predictability regression assesses how the slope of the yield curve predicts changes in long-term yields.

Under the expectation hypothesis, the slope of the yield curve is the optimal forecast of future changes in long-term bond yields. In other words, α_τ should be equal to one. Instead, Campbell and Shiller (1991) empirically find a coefficient that is negative at all maturities and significantly different from one. Time variation in the yield spread pushes α_τ away from unity. Moreover, they observe that the higher the maturity, the lower is α_τ.

Table 1.4, Panel A, shows that the term premia and the Campbell-Shiller regression slopes for the present-bias model are in line with the observed data. Present bias matches the mean observed term premium over time well. The term premium on a five-year bond in the present-bias model equals 1.32%, nearly identical to its mean in the data of 1.36%, while the time-consistent model clearly yields a too low five-year term premium of 0.88%. The term premium on a ten-year bond is overstated by the present-bias model, but understated by the time-consistent model. Our main goal is to match historical bond behavior by a simple extension of standard time-separable CRRA utility, such that we have modeled constant prices-of-risk and, therefore, miss out on the variability of term premia.

The Campbell-Shiller regression slopes show the coefficients α_τ when we run regression equation (1.14) on the model-implied yields and on the observed yields. The present-bias model produces negative slope coefficients close to the data and they are decreasing in maturity. Our estimated five-year coefficient is similar to Wachter (2006), who reports a value in the range of -1.4 and -1.6. The model-implied ten-year slope coefficient equals -2.76, close to its estimate in the data -2.68. The five-year slope coefficient from the time-consistent model matches the data well, but the time-consistent model performs poorly compared to the present-bias model when matching the ten-year slope coefficient in the data.

Table 1.4: Term premia and predictability. Panel A presents average term premia and "long-rate" regression slope coefficients in the data, in the present-bias model, and in the time-consistent model. The τ-year term premium, or bond risk premium, is the difference between the yield on the τ-year bond and the unobserved risk-neutral yield for that same bond. The Campell-Shiller regression slope is the coefficient α_τ from the regression $y_{t+12}(\tau-12) - y_t(\tau) = \text{constant}_\tau + \alpha_\tau \frac{12}{\tau-12}(y_t(\tau) - y_t(3)) + \varepsilon_{t+12}$ using monthly periods t and bond yields with maturity $\tau \in \{60, 120\}$.

	5 years	10 years
Term premium		
Data	1.36	2.03
Present bias	1.32	2.42
Time consistency	0.88	1.56
Campbell-Shiller regression slope		
Data	-1.64	-2.68
Present bias	-1.41	-2.76
Time consistency	-1.52	-3.07

1.4 Conclusion

In this paper, we offer an explanation for several empirical features of nominal bond behavior based on present-biased preferences. Our solution is a simple extension of the standard exponential discounting time-separable CRRA utility framework by means of an additional present-bias factor β (Phelps and Pollak, 1968; Laibson, 1997). A substantial body of experimental literature suggests that present bias plays an important role in making intertemporal decisions (Frederick et al., 2002). Present-biased investors value the short term more than the long term, such that they focus more on hedging risks in the short term rather than hedging risks in the long term. For this reason, present bias increases the (hedge) demand for short-term bonds and creates a premium for long-term bonds. These effects can match a panoply of empirical stylized facts.

Present-biased preferences require a distinction between the short term and the long term. We answer the open empirical question about the duration of the present in an investment context (Ericson and Laibson, 2019). Bond behavior is best explained if bond investors use a present duration of at most 1 year. As such, we connect asset pricing models with the experimental literature. We are the first to show that present bias shapes bond behavior. Our present-bias model matches observed excess bonds returns and observed yield spreads especially well compared to a time-consistent model. Besides, the present-bias model matches term premia and bond predictability to a somewhat better extent than a time-consistent model.

The present-biased investor solves a dynamic investment problem with a general discount function in a financial market with bonds and stocks, in which interest rates are driven by a 2-factor Gaussian affine term structure model (such as Vasicek, 1977). We present explicit analytical solutions by using the martingale method (Cox and Huang, 1989) for the investor's optimal consumption path and investment strategy. These optimal investment demands imply equilibrium bond behavior.

1.5 Appendix

A. Consumption and investment in a N-factor model

The result is fairly standard and is present for completeness only. We solve firstly for the optimal consumption path and, then, for the optimal investment strategy. The representative investor maximizes expected CRRA utility of wealth at a terminal horizon. Specifically, the representative investor maximizes expected utility of future wealth W for investment horizon T:

$$\max_{W_T} E_0 \left[\frac{W_T^{1-\gamma}}{1-\gamma} \right], \tag{1.15}$$

subject to her budget constraint

$$E_0 \left[W_T M_T \right] = W_0, \tag{1.16}$$

where W_0 is initial total available wealth. Obviously, such a terminal consumption formulation is independent of any discount structure. For this reason, discounting models in a terminal horizon problem do not influence optimal consumption and optimal investment. We formalize this observation in the theorem below.

Standard calculations using the Lagrange method, lead to implicit optimal consumption

$$W_t^* = \frac{1}{M_t} E_t \left[W_T^* M_T \right] = \frac{W_0}{M_t} \frac{E_t \left[M_T^{1-1/\gamma} \right]}{E_0 \left[M_T^{1-1/\gamma} \right]}, \tag{1.17}$$

where we use the first fundamental theorem of asset pricing.

Lemma 1. *For positive α, the conditional expectation at time t of the stochastic* *discount factor follows from*

$$E_t\left[\left(\frac{M_T}{M_t}\right)^{\alpha}\right] = \exp\left(\alpha m(\boldsymbol{F}_t, T-t) + \frac{1}{2}\alpha^2 v^2(T-t)\right) \tag{1.18}$$

$$= e^{\frac{1}{2}(\alpha-1)\alpha v^2(T-t)} P_t(T-t)^{\alpha} \tag{1.19}$$

where ($\tau = T - t$)

$$m(\boldsymbol{F}_t, \tau) = -\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(\tau) - \boldsymbol{\iota}'\boldsymbol{B}(\tau)\left(\boldsymbol{F}_t - \boldsymbol{\theta}\right) - \frac{1}{2}\boldsymbol{\lambda}'\boldsymbol{\lambda}(\tau) \tag{1.20}$$

$$v^2(\tau) = \int_0^{\tau} ||\boldsymbol{\iota}'\boldsymbol{B}(\tau - v)\boldsymbol{\sigma}_F||^2 dv + \boldsymbol{\lambda}'\boldsymbol{\lambda}(\tau) + 2\int_0^{\tau} \langle \boldsymbol{\iota}'\boldsymbol{B}(\tau - v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F\rangle dv, \tag{1.21}$$

and

$$P_t(T-t) = \exp\left(A(T-t) - \boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{F}_t\right) \tag{1.22}$$

with deterministic functions

$$A(T-t) = -\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-t) + \boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{\theta}$$
$$+ \frac{1}{2}\int_t^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \int_t^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}'_F\rangle dv \tag{1.23}$$

Appendix B provides the proof of Lemma 1. Now, the explicit optimal consumption path follows directly from (1.17) by plugging in the result of Lemma 1.

Mathematically,

$$
\begin{aligned}
W_t^* = {} & W_0 \exp\left(-\int_0^t r_s ds - \int_0^t \boldsymbol{\lambda}' d\boldsymbol{Z}_s - \frac{1}{2}\int_0^t \boldsymbol{\lambda}'\boldsymbol{\lambda} ds\right)^{\alpha-1} \\
& \times \exp\left(\frac{1}{2}(\alpha-1)\alpha v^2(T-t) - \frac{1}{2}(\alpha-1)\alpha v^2(T)\right) \\
& \times \{\exp[-\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-t) - \boldsymbol{\iota}'\boldsymbol{B}(T-t)\left(\boldsymbol{F}_t - \boldsymbol{\theta}\right) \\
& + \frac{1}{2}\int_t^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \int_t^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F' \rangle dv]^\alpha \\
& - \exp[-\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T) - \boldsymbol{\iota}'\boldsymbol{B}(T)\left(\boldsymbol{F}_0 - \boldsymbol{\theta}\right) \\
& + \frac{1}{2}\int_0^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \int_0^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F' \rangle dv]^\alpha\}. \qquad (1.24)
\end{aligned}
$$

Note that the stochasticity in (1.24) comes from

$$-(\alpha-1)\left(\boldsymbol{\lambda}_F'\boldsymbol{Z}_{F,t} + \lambda_S Z_{S,t}\right)$$

and from

$$-\alpha\boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{F}_t.$$

So, rewriting (1.24) into a stochastic differential equation, with $\alpha = 1 - 1/\gamma$ due to CRRA utility, yields

$$d\log W_t^* = g_1(r, T-t)dt + \frac{\lambda_S}{\gamma}dZ_{S,t} + \left(\frac{\boldsymbol{\lambda}_F'}{\gamma} - (1-1/\gamma)\boldsymbol{\iota}'\boldsymbol{B}(T-t)\boldsymbol{\sigma}_F\right)d\boldsymbol{Z}_{F,t}, \qquad (1.25)$$

where the drift term $g_1(r, T-t)$ is a function of the interest rate r_t and the remaining investment horizon $\tau = T - t$.

Finally, consider the return on a portfolio that consists of the three available assets. Let $\boldsymbol{\pi^*} = (\pi_S^*(t), \boldsymbol{\pi_B^*}(t,\tau), \pi_M^*(t))$ be the optimal proportion of wealth invested at time t in a stock, in a vector of bonds with maturities τ and in cash. If the individual invests in such a portfolio, then total wealth $A(t)$ evolves according to

$$\frac{dA_t}{A_t} = \left(\pi_S(t)\frac{dS_t}{S_t} + \boldsymbol{\pi}_{\boldsymbol{B}}(t,\tau)'\frac{d\boldsymbol{P}_t}{\boldsymbol{P}_t} + \pi_M(t)\frac{dB_t}{B_t} \right), \tag{1.26}$$

where B_t is the process for the risk-free asset, or cash. Substituting the dynamics of the assets and taking the log, yields

$$d\log A_t = g_2(.)dt + \pi_S(t)\sigma_S dZ_{S,t} + \left(\pi_S(t)\boldsymbol{\sigma}'_{FS} - \boldsymbol{\pi}_{\boldsymbol{B}}(t,\tau)'\mathcal{B}(\tau)'\boldsymbol{\sigma}_F\right) d\boldsymbol{Z}_{F,t}, \tag{1.27}$$

where $g_2(.)$ is the drift term.

Then, the optimal investment demands $\boldsymbol{\pi}^*$ follow by simply equating the coefficients of the diffusion terms in (1.25) and (1.27). $\square$

B. Proof of Lemma 1

Using (1.3) and (1.4), we have

$$\begin{aligned} r_t &= A_0 + \boldsymbol{\iota}' \left[\boldsymbol{\theta} + \exp(-\boldsymbol{\kappa} t)(\boldsymbol{F}_0 - \boldsymbol{\theta})\right] + \int_0^t \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa}(t-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s} \\ &= A_0 + \boldsymbol{\iota}'\boldsymbol{\theta} + \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa} t)(\boldsymbol{F}_0 - \boldsymbol{\theta}) + \int_0^t \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa}(t-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s}, \end{aligned} \tag{1.28}$$

which follows directly from the solution of the stochastic differential equation to the Orstein-Uhlenbeck process (see, e.g., Chin et al. (2014)).

From the above, we find

$$\begin{aligned} \int_0^t r_v dv &= \int_0^t \left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta} + \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa} v)(\boldsymbol{F}_0 - \boldsymbol{\theta}) + \int_0^v \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa}(v-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s} \right) dv \\ &= \left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right) t + \boldsymbol{\iota}'\boldsymbol{B}(t)\left(\boldsymbol{F}_0 - \boldsymbol{\theta}\right) + \int_0^t \left(\int_0^v \boldsymbol{\iota}' \exp(-\boldsymbol{\kappa}(v-s))\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,s} \right) dv \\ &= \left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right) t + \boldsymbol{\iota}'\boldsymbol{B}(t)\left(\boldsymbol{F}_0 - \boldsymbol{\theta}\right) + \int_0^t \boldsymbol{\iota}'\boldsymbol{B}(t-v)\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,v} \end{aligned} \tag{1.29}$$

where in the second equality we use (1.7) and in the third equality we use stochastic integration by parts. More general

$$\int_t^T r_v dv = \left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-t) + \boldsymbol{\iota}'\boldsymbol{B}(T-t)\left(\boldsymbol{F}_t - \boldsymbol{\theta}\right) + \int_t^T \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F d\boldsymbol{Z}_{F,v}, \tag{1.30}$$

which has a normal distribution with (conditional) mean and variance:

$$E_t\left[\int_s^T r_v dv\right] = \left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-s) + \boldsymbol{\iota}'\boldsymbol{B}(T-s)\left(\boldsymbol{F}_s - \boldsymbol{\theta}\right), \tag{1.31}$$

$$V_t\left[\int_s^T r_v dv\right] = \int_t^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv, \tag{1.32}$$

where we use Itô isometry to obtain the variance.[26] Hence, the stochastic discount

[26] $||.||$ denotes the Euclidean norm of a vector (i.e., the square root of inner product of the vector and itself).

factor follows a log-normal distribution with mean

$$
\begin{aligned}
m(r_t, T-t) &= E_t\left[\ln\frac{M_T}{M_t}\right] \\
&= E_t\left[-\int_s^T r_v dv - \int_s^T \boldsymbol{\lambda}' d\boldsymbol{Z}_v - \frac{1}{2}\int_s^T \boldsymbol{\lambda}'\boldsymbol{\lambda} dv\right] \\
&= -\left(A_0 + \boldsymbol{\iota}'\boldsymbol{\theta}\right)(T-s) - \boldsymbol{\iota}'\boldsymbol{B}(T-s)\left(\boldsymbol{F}_s - \boldsymbol{\theta}\right) - \frac{1}{2}\boldsymbol{\lambda}'\boldsymbol{\lambda}(T-s), \qquad (1.33)
\end{aligned}
$$

and variance

$$
\begin{aligned}
v^2(T-t) &= V_t\left[\ln\frac{M_T}{M_t}\right] \\
&= V_t\left[-\int_s^T r_v dv - \int_s^T \boldsymbol{\lambda}' d\boldsymbol{Z}_v - \frac{1}{2}\int_s^T \boldsymbol{\lambda}'\boldsymbol{\lambda} dv\right] \\
&= \int_s^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \boldsymbol{\lambda}'\boldsymbol{\lambda}(T-s) + 2E_t\left[\int_s^T r_v dv \int_s^T \boldsymbol{\lambda}_F' d\boldsymbol{Z}_{F,v}\right] \\
&= \int_s^T ||\boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F||^2 dv + \boldsymbol{\lambda}'\boldsymbol{\lambda}(T-s) + 2\int_s^T \langle \boldsymbol{\iota}'\boldsymbol{B}(T-v)\boldsymbol{\sigma}_F, \boldsymbol{\lambda}_F'\rangle dv. \qquad (1.34)
\end{aligned}
$$

We use that the prices-of-risk are constant, Z_S is independent of $\boldsymbol{Z}_F$ and Itô isometry. $\square$

C. Proof of Theorem 1

We firstly solve for the optimal consumption path and, then, for the optimal investment strategy. The representative investor maximizes CRRA utility of intermediate consumption. She solves (1.1) subject to her budget constraint (1.8). The derivations are general in the sense that they hold for $W_{t,T_j,t+h}$ with $t \leq t+h \leq T_j$, such that *actual* and *planned* consumption may be separated.

Using η as Lagrange multiplier for the budget constraint, we find

$$\mathcal{L} = \mathbb{E}_t \left[\sum_{j=1,t\leq T_j}^{n} D(T_j - t) \frac{W_{t,T_j,T_j}^{1-\gamma}}{1-\gamma} \right] + \eta \left(W_t M_t - \mathbb{E}_t \left[\sum_{j=1,t\leq T_j}^{n} W_{t,T_j,T_j} M_{T_j} \right] \right). \tag{1.35}$$

Taking the first order conditions with respect to terminal wealth for every state of the world, yields

$$\frac{\partial \mathcal{L}}{\partial W_{t,T_j,T_j}} = D(T_j - t) W_{t,T_j,T_j}^{-\gamma} - \eta M_{T_j} = 0. \tag{1.36}$$

This implies *implicit* optimal *terminal* wealth

$$W_{t,T_j,T_j}^{*} = D^{1/\gamma}(T_j - t) \left(\eta M_{T_j} \right)^{-1/\gamma}. \tag{1.37}$$

and isolating the Lagrange multiplier with the budget constraint yields

$$\eta^{-1/\gamma} = \frac{W_t M_t}{\mathbb{E}_t \left[\sum_{j=1,t\leq T_j}^{n} D^{1/\gamma}(T_j - t) M_{T_j}^{1-1/\gamma} \right]}. \tag{1.38}$$

Substituting the expression for the Lagrange multiplier in implicit optimal terminal wealth yields *explicit* optimal *terminal* wealth

$$W_{t,T_j,T_j}^{*} = D^{1/\gamma}(T_j - t) M_t M_{T_j}^{-1/\gamma} \frac{W_t}{\sum_{j=1,t\leq T_j}^{n} D^{1/\gamma}(T_j - t) \mathbb{E}_t \left[M_{T_j}^{1-1/\gamma} \right]}. \tag{1.39}$$

Now, using the first the fundamental theorem of asset pricing, the *actual* optimal consumption *path* at time t with horizon $T_j \geq t$ for in-between dates $t+h = t, ..., T_j$ for every money pot $j = 1, ..., n$ equals

$$
\begin{aligned}
&W^*_{t,T_j,t+h}\left(D(T_j - t)\right) = \frac{1}{M_{t+h}} \mathbb{E}_{t+h}\left[W^*_{t,T_j,T_j-t} M_{T_j}\right] \\
&= D(T_j - t)^{1/\gamma} W_t \frac{M_t}{M_{t+h}} \frac{\mathbb{E}_{t+h}\left[M_{T_j}^{1-1/\gamma}\right]}{\sum_{j=1,t\leq T_j}^{n} D(T_j - t)^{1/\gamma} \mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]} \\
&= D(T_j - t)^{1/\gamma} W_t M_t M_{t+h}^{-1/\gamma} \frac{\exp\left(\frac{1}{2}\frac{1-\gamma}{\gamma^2} v^2 (T_j - (t+h))\right) P_{t+h}(T_j - (t+h))^{1-1/\gamma}}{\sum_{j=1,t\leq T_j}^{n} D(T_j - t)^{1/\gamma} \mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]} \\
&= D(T_j - t)^{1/\gamma} W_t M_t^{1-1/\gamma} \exp\left(-\int_t^{t+h} r_s ds - \boldsymbol{\lambda}'\left(\boldsymbol{Z}_{t+h} - \boldsymbol{Z}_t\right) - \frac{1}{2}\boldsymbol{\lambda}'\boldsymbol{\lambda}((t+h) - t)\right) \\
&\quad \frac{\exp\left(\frac{1}{2}\frac{1-\gamma}{\gamma^2} v^2 (T_j - (t+h))\right) P_{t+h}(T_j - (t+h))^{1-1/\gamma}}{\sum_{j=1,t\leq T_j}^{n} D(T_j - t)^{1/\gamma} \mathbb{E}_t\left[M_{T_j}^{1-1/\gamma}\right]},
\end{aligned}
\tag{1.40}
$$

where we use Lemma 1 in the third equation, and in the fourth equation we use the explicit expression of the stochastic discount factor.

The optimal investment strategy $\pi_i^*(t, T_j)$ for asset i = $\{$stock, constant-maturity τ-year bonds, cash$\}$ at time t for investment horizon T_j follows from Appendix A. So, the optimal fraction of actual total invested wealth at time t for each money pot j follows from the investor's *planned* optimal consumption path $W^*_{t,T_j,t+h}$ with $h = 0$, leading to the investor's *actual* optimal consumption path $W^*_{t,T_j,t}$. Both consumption paths are a function of the discount structure $D(T_j - t)$. This proves Theorem 1.[27] □

[27] Note that $j > t$ because investments are for future dates, i.e., the investor does not invest for a payment at a payment date

D. Additional Data & Sensitivities

Table 1.5: Estimates of model parameters. Maximum likelihood parameter estimates and standard errors for the joint process of nominal yields, stock returns and state-factors in (1.3) - (1.5) by implementing a standard Kalman filter using monthly observations. We observe the the stock market index and four points on the U.S. zero-coupon yield curve, corresponding with maturities of 3 months, 1 year, 5 years, and 10 years. The data runs from 1 October 1976 to 1 January 2019. The standard errors follow from the square root of the diagonal elements of the inverted Hessian matrix $\sqrt{H_{ii}^{-1}}$ for $i = 1, .., k$ where k is the number of estimated parameters.

Parameter	Estimate	Standard error
$\hat{\kappa}_1$	0.0398	0.0015
$\hat{\kappa}_2$	0.4623	0.0119
$\hat{A}_0$	0.0359	0.0015
$\hat{\lambda}_{F,1}$	-0.1378	0.0044
$\hat{\lambda}_{F,2}$	-0.4785	0.0383
$\hat{\lambda}_S$	0.4095	0.1381
$\hat{\sigma}_{F,11}$	0.0139	0.0004
$\hat{\sigma}_{F,21}$	-0.0065	0.0007
$\hat{\sigma}_{F,22}$	0.0166	0.0005
$\hat{\sigma}_{FS,1}$	-0.0264	0.0259
$\hat{\sigma}_{FS,2}$	-0.0171	0.0101
$\hat{\sigma}_S$	0.1504	0.0020

Table 1.6: U.S. government debt by maturity. Panel A presents the composition of U.S. government debt at 1 October 1976. Panel B presents the composition of U.S. government debt at 1 January 2019. Panel C shows the mean debt composition during the observation period 1 October 1976 to 1 January 2019.

	Debt outstanding (million USD)	Fraction of total debt outstanding
Panel A: Composition at 1 October 1976		
Total	294,595	1.00
< 1 year	153,302	0.52
1-5 years	94,845	0.32
5-10 years	31,247	0.11
10-20 years	7,939	0.03
> 20 years	7,262	0.02
Panel B: Composition at 1 January 2019		
	Debt outstanding (million USD)	Fraction of total debt outstanding
Total	13,385,359	1.00
< 1 year	3,927,279	0.29
1-5 years	5,426,079	0.41
5-10 years	2,524,238	0.19
10-20 years	113,097	0.01
> 20 years	1,394,666	0.10
Panel C: Composition from 1 October 1976 to 1 January 2019		
	Debt outstanding (million USD)	Fraction of total debt outstanding
Total	3,798,442	1.00
< 1 year	1,263,723	0.37
1-5 years	1,460,573	0.36
5-10 years	623,445	0.14
10-20 years	147,308	0.05
> 20 years	303,393	0.08

Table 1.7: Sensitivity optimal asset allocation. Optimal fraction of total wealth — divided in hedge and speculative demands — invested in a 3-year bond, a 10-year bond, a stock index and cash (total demand adds up to 1). Panel A presents optimal demands for the present-bias model with varying present-bias factors β. Panel B presents optimal demands under impatient time consistency.

	3-year bond	10-year bond	Stock	Cash
	Panel A: Present bias			
$\beta = 0.05, \delta = 0.97$				
Hedge demand	0.54	-0.06	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.98	-0.69	0.27	-1.56
$\beta = 0.5, \delta = 0.97$				
Hedge demand	0.46	0.09	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.90	-0.54	0.27	-1.63
$\beta = 0.8, \delta = 0.97$				
Hedge demand	0.42	0.16	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.85	-0.47	0.27	-1.66
	Panel B: Impatient time consistency			
$\beta = 1, \delta = 0.86$				
Hedge demand	0.49	0.04	-	-
Speculative demand	2.44	-0.63	0.27	-
Total demand	2.92	-0.59	0.27	-1.60

Table 1.8: Excess returns and impatient time consistency. This table reports average returns in excess of the risk-free rate and Sharpe ratios of a nominal 3-year bond, a nominal 10-year bond and a stock. The Sharpe ratio is the mean of the excess return divided by the standard deviation. The column "Data" gives statistics for excess returns as realized in the data. The column "Present bias" gives statistics as implied by the present-bias model (present-bias factor $\beta = 0.35$). The column "Impatient time consistency" gives statistics as implied by standard time-consistent discounting (present-bias factor $\beta = 1$) with impatience (discount factor $\delta = 0.86$). Values are annualized. Data are monthly and run from 1 October 1976 to 1 January 2019.

	Data	Present bias	Impatient time consistency
3-year bond			
Mean	1.90	1.60	1.61
Sharpe	0.48	0.41	0.41
10-year bond			
Mean	4.10	4.45	4.48
Sharpe	0.38	0.42	0.42
Stock			
Mean	7.27	7.48	7.49
Sharpe	0.48	0.49	0.49

Chapter 2

Regret and Asset Pricing*

*I am particularly grateful to Bas Werker and Ralph Koijen. I thank Hazel Bateman, Paul Bokern, Fabio Braggion, Stefano Cassella, Aditya Chaudry, Alex Clyde, Joost Driessen, Rik Frehen, Rob van den Goorbergh, Frank de Jong, Marike Knoef, Betrand Melenberg, Stefan Nagel, Antoon Pelsser, Eduard Ponds, Jan Potters, Nikolaus Schweizer, Martijn de Vries, Ole Wilms, Marcel Zeelenberg, Stefan Zeisberger, and seminar participants at the Spring Meeting of Young Economists (2022), the Bachelier Finance Society (2022), the University of Amsterdam (2022), the NEOMA Business School (2022), the UC3M (2022), the Radboud University (2022), the European Winter Meeting of the Econometric Society (2021), the European Economic Association (2021), the RCEA Money-Macro-Finance Conference (2021), the Free University of Brussels (ULB, 2021), the Netspar Pension Day (2021), the ENTER Jamboree (2021), the Royal Dutch Economic Association KVS New Paper Sessions (2021), the Netspar International Pension Workshop (2021), and Tilburg University (2022, 2021, 2020) for useful comments.

It appears to be somewhat of a challenge to provide a unifying explanation for stylized facts that have been uncovered in the asset pricing literature. In the time series, typical stock returns are excessively volatile and predictable using lagged prices scaled by fundamentals such as dividends (Shiller, 1981; Campbell and Shiller, 1988). The risk-free rate is low and stable (Weil, 1989), with unpredictable consumption and dividend growth. The unconditional term structure of equity risk premiums is downward sloping (van Binsbergen, Brandt, et al., 2012; van Binsbergen and Koijen, 2017).[1] In the cross section, there is a value premium (Basu, 1983; Fama and French, 1992) and "long-term reversal" (De Bondt and Thaler, 1985). Bond yields produce an unconditional downward sloping real yield curve (Piazzesi and Schneider, 2006b).[2]

I present a model that helps explaining these stylized facts in a unifying way. The central and only ingredient is regret and the aversion to it, added to an otherwise standard power utility function and standard financial market. Regret-averse investors are concerned not only about the returns they receive, but also about the foregone returns they could have received, had they invested differently. Investors anticipate disutility from a state of the world where they could have had higher consumption, weighted by a regret-aversion parameter. In case the foregone return is large, regret and marginal utility in that state of the world are high, such that risky asset prices are high and expected returns fall.

Why is regret aversion relevant for asset prices? Extensive psychological, experimental, and neuroscientific research provides abundant evidence for the role of regret in investment decisions and in trading behaviour.[3] Psychological evidence

[1]Bansal, Miller, et al. (2021) argue that the unconditional equity term structure is upward sloping.

[2]This is consistent with the evidence in the U.K., but it is more ambiguous in U.S. data. For the latter, we only have a small sample (Piazzesi and Schneider, 2006b).

[3]Evidence for regret in decision making has been extensively investigated by Larrick (1993), Gilovich and Medvec (1995), Larrick and Bowles (1995), Zeelenberg (1999), Connolly and Zeelenberg (2002), Connolly and Butler (2006), and Zeelenberg and Pieters (2007). Zeelenberg and Pieters (2004) document regret in a real-life non-student sample. According to Connolly and Zeelenberg (2002), regret is the emotion that has received the most attention from decision theorists. Saffrey et al. (2008) investigate the intensity and frequency of twelve emotions, and their participants rate regret as being the most intense negative emotion. Zeelenberg (2020) argues that regret fulfills all conditions for being classified as a basic emotion.

comes from Lin, Huang, et al. (2006), who study actual stock investors' behavior and document that regret influences investors' investment decisions through counterfactual thinking. Lohrenz et al. (2007) and Frydman and Camerer (2016) use neural scientific data, gathered simultaneously with actual investment behavior, to show that investors exhibit and experience regret while trading. More generally, using a neuro-psychological experiment, Camille et al. (2004) and Bourgeois-Gironde (2010) find that their respondents think counterfactually, anticipate regret and consider regret while making risky decisions.[4] Using detailed trading data, Strahilevitz et al. (2011) and Magron and Merli (2015) emphasize the important role of regret in financial decisions and they relate regret to the disposition and repurchase effects.[5]

I model regret, and the aversion to it, in line with three observations from the literature. First, at the moment of investment decision making, investors anticipate that they experience the feeling of regret in the future. Neuroscientific evidence (Camille et al., 2004) and the review of Zeelenberg (2018) show that anticipating future regret influences current decision making and, thus, current investors' holdings. In the model, regret enters in an otherwise standard power utility function as a multiplicative component that yields disutility, weighted by a regret-aversion parameter (Quiggin, 1994). Regret follows from counterfactual thinking about foregone returns. The representative investor invests in a portfolio and anticipates regret by a comparison (ex-post) of her realized consumption with the best unchosen alternative (i.e., "if only I made another investment decision") and the inaction alternative (i.e., "if only I did not invest").

Second, emotions, and thereby regret, follow laws (Frijda, 1988; Frijda, 2007). Regret is time varying and reverts over time around a mean (Wilson and Gilbert, 2005), but only slowly and gradually. Regret is persistent, extending up to years. Especially negative feelings and emotions, to which regret belongs, are persistent

[4]Bourgeois-Gironde (2010) state that regret helps to optimize decision behaviour. They define regret as a rational emotion.

[5]When an investor sells a stock, she is less likely to repurchase this same stock if the price has increased since the sale, compared with when the price has decreased since the sale. Strahilevitz et al. (2011) call this the repurchase effect. Fioretti et al. (2021) use a stock market experiment to study the influence of regret aversion on the decision to sell an asset if prices change over time.

phenomena (Coricelli et al., 2005; Wilson and Gilbert, 2005; Hajcak and Olvet, 2008). This feature produces mean-reversion in prices and return predictability in the time series and cross section. Consumption growth and dividend growth follow simple white noise lognormal processes (Campbell and Cochrane, 1999), with means and standard deviations consistent with the empirical asset pricing literature. Because of the law of emotional control (Frijda, 1988; Frijda, 2007), regret is not too volatile such that the risk-free rate in the economy remains stable.

Third, the main premise of regret theory (Bell, 1982; Loomes and Sugden, 1982), being an alternative to expected utility theory (Neumann and Morgenstern, 1947), is that we are averse to regret. Bleichrodt, Cillo, et al. (2010) introduce the first quantitative measurement of a regret aversion parameter. They estimate a regret-aversion parameter which implies more disutility when foregone consumption is high. Their evidence confirms regret aversion at the individual and aggregate level.

My results show that regret aversion is a helpful ingredient to understand behavior of assets in the time series and in the cross section, not only in terms of sign, but also in terms of magnitude consistent with the empirical asset pricing literature. I find a low stable risk-free rate with unpredictable, low, and stable consumption growth and dividend growth. Stocks are more volatile than the underlying dividends, and returns are predictable by the lagged price-dividend ratio and lagged returns. Regret produces an unconditional term structure of risk premiums that is downward sloping. In the cross section, I document long-term reversal and a value premium: stocks with low price-dividend ratios (i.e., value stocks) yield higher subsequent returns than stocks with high price-dividend ratios (i.e., growth stocks). The analysis on bonds shows that regret produces an unconditional downward sloping real yield curve, and bond returns are predictable by regret.

To understand the mechanisms, consider first the case of predictability in the time series and the cross section. If foregone returns on the risky asset are high, then regret is high. Investors regret having invested too little and demand more of the risky asset, which pushes up prices today. Prices relative to dividends become overvalued such that future returns fall. In the cross section, regret is asset specific. Stocks with high regret are typically growth stocks, or winner stocks, that yield subsequent lower returns. The mispricing of these stocks and the regret on growth stocks are highly persistent, consistent with (Arisoy et al., 2021; van Binsbergen, Boons, et al., 2021). Regret volatility, amplified by regret aversion, makes returns more volatile than the underlying cash flows. Because regret-averse investors are concerned with regret in the short run, as they confront their performance annually, investors require a premium to hold short-term assets such that the term structures of equity risk premiums and interest rates is downward sloping.

The risk-free rate in the economy is low and stable, since regret is not too volatile. When regret is high, the representative investor feels poor such that she starts to save more which drives down the equilibrium risk-free rate. Finally, regret-averse investors theoretically require a regret risk premium to hold risky assets through negative correlations between consumption and regret, and dividend and regret. If regret is high, then the foregone return on risky assets is high, such that marginal utility is also high. Regret-averse investors do not like such states of the world, as the could have been better off. Therefore, investors require a premium to hold risky assets. Theoretically this holds true, but in my calibration the regret premium is small and, thus, regret does not resolve the equity premium puzzle. Though, I study the equity premium on unlevered claims, which is typically lower than the equity premium on firms with leverage included (Abel, 1999).

An empirical measure of regret confirms the main model's predictions. In line with my regret model and the measure of Arisoy et al. (2021), I empirically measure annual regret by the highest return in the cross section per year. This regret measure is highly persistent and behaves historically similar to the price-dividend ratio, as predicted by the model that prices relative to dividends are a function of investors' regret. Moreover, the regret measure predicts future returns with a negative sign in the time series, specifically when the forecasting horizon is long. Arisoy et al. (2021) empirically study the implications of regret-sorted portfolios in the cross section and they find that growth stocks are stocks with high regret, as predicted by my theoretical model.

Regret-averse investors are concerned with the positive skewness of returns. A simple exercise shows that regret, or foregone returns, are on average high when the skewness of the underlying returns is positive and high. The more dispersion cross-sectionally in returns, the higher the probability of a foregone missed opportunity. Drerup et al. (2022) provide direct evidence for my main mechanism of predictability of returns, as the authors find a positive correlation between skewness expectations and investment decisions. That is, investors indeed increase portfolio allocations when skewness expectations are high, i.e., high expected skewness possibly yields high foregone returns which implies high regret. Eeckhoudt et al. (2007) and Gollier (2018) find that regret-averse agents have a preference for positively skewed risks and longshots. These upside risk concerns of regret-averse investors contrast with the downside risk capital asset pricing model of Lettau, Maggiori, et al. (2014), in which investors are concerned with downside risks.

My paper is the first to study regret in a consumption-based model and the first to show that regret also explains stylized facts in terms of magnitude. I contribute to the asset pricing literature on regret, which contains a few studies with explicit regret-theoretic models. Dodonova and Khoroshilov (2005) study regret aversion in a one-period terminal wealth asset pricing model with two firms and two types of agents. They theoretically find excess volatility and long-run negative autocorrelations of stock prices, but their regret-utility specification deviates from the original regret theory of Bell (1982), Loomes and Sugden (1982), and Quiggin (1994). Muermann et al. (2006) study optimal portfolio choice between a risky and risk-free asset in DC schemes when investors are regret averse, by following additive regret. Qin (2020) presents a regret-CAPM model, and indicates that a regret-related beta can help explain cross-sectional returns and possibly the high equity premium. Arisoy et al. (2021) empirically document a regret premium in the cross section.

Compared to leading asset pricing models, regret is able to match the downward sloping term structure of equity risk premiums and the cross-sectional stylized facts. Habit formation (Campbell and Cochrane, 1999) needs high risk aversion to explain the equity-premium risk-free rate puzzle, it produces a growth premium rather than a value premium (Santos and Veronesi, 2010) and the term structure of equity risk premiums is upward sloping (van Binsbergen, Brandt, et al., 2012). The long-run risk model has difficulties with the absence of predictability in consumption and dividend growth (Beeler and Campbell, 2012), needs to rely on time-varying consumption volatility for predictability, and produces an unconditionally upward sloping term structure of equity risk premiums. Disaster risk (Barro, 2006) requires time variation in disasters to match predictability, and disaster risk produces a flat term structure for equity and bonds (Gabaix, 2012). Prospect theory and loss aversion models (Barberis, Huang, and Santos, 2001; Pagel, 2016) have difficulties solving the risk-free rate puzzle and produce upward sloping equity term structures, but Barberis and Huang (2001) can math the value premium and long-term reversal.

Regret relates to long-term experience effects caused by the persistence and strength of emotions (Malmendier, 2021). Malmendier (2021) states that we as economists typically pay little attention to emotions, while we might want to reconsider choice behavior and beliefs as a function of emotional inputs rather than (only) informational inputs. The implications of regret on asset prices relate to return extrapolation (Atmaz, 2021) and optimism (Brunnermeier et al., 2007) as regret-averse investors chase good foregone returns. Moreover, regret relates to ex-ante rational expectations-based reference-dependent models (Köszegi and Rabin, 2006; Pagel, 2016) in which references points are (fixed) forward looking beliefs over all possible outcomes (i.e., good and bad news) rather than backward looking as in habit formation or prospect theory. Finally, regret aversion relates to disappointment aversion (Gul, 1991), but regret is neurologically different and significantly more intensely felt than disappointment (Camille et al., 2004).

Regret appears implicitly in the early behavioral finance literature (Shefrin and Statman, 1984; Shefrin and Statman, 1985). Taking regret more explicitly into account, Muermann et al. (2006) and Baule et al. (2019) study optimal portfolio choice for regret-averse investors with additive regret, whereas Gollier and Salanié (2012) study risk-sharing and portfolio allocation in a complete market for a general bivariate regret-utility function. Qin (2015) studies bubbles, herding, and market turbulence by regret over action and inaction, while Fogel and Berry (2006) find that regret aversion explains the disposition effect. Solnik and Zuo (2012) present a global equilibrium asset pricing model to explain the home bias. Regret also appears in the literature of insurance and pensions (Braun and Muermann, 2004; Frehen et al., 2008), as well as currency hedging (Michenaud and Solnik, 2008).

A. Related concepts

In this section, I discuss how the predictions of the regret model compare to leading asset pricing models. The section ends with economic and psychological related concepts to regret.

1. Asset pricing models

This section provides a comparison between the regret model and other leading asset pricing models such as habit formation, long-run risks, disaster risk and prospect theory. To the extent of my knowledge, current leading asset pricing models have difficulty explaining the behavior of equity and bonds in the time series and the cross section. Table 2.1 presents an overview of the models, and I explain it below.

Table 2.1: Overview of asset pricing models. This table presents an overview of the stylized facts that asset pricing models can explain, together with the number of parameters. It compares the regret model to four leading asset pricing models: habit formation, long-run risks, disaster risk and prospect theory. [a] Required to add economic uncertainty (i.e., time-varying consumption volatility). [b] Required to add that disaster risk changes over time. [c] Narrowing framing required.

	Regret aversion	Habit formation	Long-run risk	Disaster risk	Prospect theory
Risk-free rate puzzle	Yes	Partial	Yes	Yes	No
Equity premium puzzle	Partial	Partial	Partial	Yes	Yes
Excess volatility	Yes	Yes	Yes	Yes	Yes
St. Dev. Price-dividend ratio	Yes	Yes	Partial	Yes[b]	No
Time-varying risk premium	No	Yes	Yes[a]	Yes[b]	
AC(1) returns	Yes	Yes	Yes[a]		Yes
AC(1) price-dividend ratio	Yes	Yes	Yes[a]		Yes
Stock predictability	Yes	Yes	Partial[a]	Yes[b]	Partial
Unpredictable Δc	Yes	Yes	No[a]	Yes[b]	Yes
Unpredictable Δd	Yes	Yes	No[a]	Yes[b]	Yes
Value premium	Yes	No	Yes		Partial[c]
Long-term reversal	Yes	No	No		Yes[c]
Equity term structure	Downward	Upward	Upward	Flat	Upward
Real yield curve	Downward	Upward	Downward	Flat	
St. dev. yields	Decreasing	Increasing			
Total parameters	12	11	13	16	16

Habit formation - Campbell and Cochrane (1999) present an asset pricing model that explains many asset pricing phenomena by including one simple ingredient in an otherwise standard model: external habit formation. The habit is slow moving and yields time-varying risk aversion, whereas the regret model implies time-independent risk aversion. The external habit model delivers a high equity premium, excess volatility, with low mean consumption growth, and volatility, unpredictable consumption and dividend growth, and a low and slowly varying risk-free rate. But, the habit model does not always have low risk aversion and, as such, does not resolve the equity-premium risk-free rate puzzle.[6]

Also, the habit model delivers the observed return predictability, the countercyclical variation of stock market volatility, and time-varying risk premia. The regret model has all of the above, including low stable risk aversion, and as such has the potential to explain the equity premium and risk-free rate puzzles. However, in the current setup, regret lacks the time-varying stock market volatility (i.e., volatility is higher after a price drop) and time-varying risk premia.

The external habit formation model yields an upward sloping yield curve (Campbell and Cochrane, 1995), which contradicts the evidence in the U.K. (Piazzesi and Schneider, 2006b). Yields on long-term bonds vary more than yields on short bonds, which contradicts the data. The habit model implies an unconditional upward sloping term structure of risk premiums and volatility (Campbell and Cochrane, 1999), which slope upward indefinitely through the non-stationary state variable, inconsistent with the empirical findings of van Binsbergen, Brandt, et al. (2012).

In the cross section, the habit model produces a growth premium rather than a value premium (Santos and Veronesi, 2010), which is at odds with the data (Basu, 1983; Fama and French, 1992). Overall, the habit model requires eleven parameters in total.[7] Note that the authors set mean consumption growth and mean dividend growth equal.

Long-run risks - Bansal and Yaron (2004) present an asset pricing model that

[6]The risk aversion of investors varies over time to more than 100 when the probability of a recession becomes larger (Mehra, 2012).

[7]Based on Table I in Campbell and Cochrane (1999).

explains many stylized facts by using two ingredients: (i) recursive Epstein-Zin preferences, and (ii) small persistent shocks (i.e., news) to consumption and dividend growth. News about long-run future consumption growth is the state variable. However, the data seems to suggest that consumption growth is closer to a random walk than the assumed persistence (Beeler and Campbell, 2012). The model produces a low and stable risk-free rate, and excess stock return volatility. It can produce the high equity premium with high risk aversion and low consumption volatility, or low risk risk aversion and high consumption volatility.

When adding economic uncertainty (i.e., time-varying volatility of consumption growth) to the long-run risk model, it produces time-varying risk premia and return predictability, but the latter is less than observed in the data (Beeler and Campbell, 2012). Also, the model-implied volatility of the log price-dividend ratio appears low compared to the data (Beeler and Campbell, 2012). The long-run risk model produces a downward sloping real yield curve, with real yields below zero for a maturity of ten years or higher. Risks for the long-run imply an unconditional upward sloping term structure of risk premiums and volatilities (van Binsbergen, Brandt, et al., 2012). The model produces a value and size premium in the cross section, as well as momentum (Bansal, Dittmar, et al., 2005). Overall, the long-run risks model (including time-varying volatility of consumption growth) needs thirteen parameters in total.[8]

Rare disasters - Barro (2006) presents an asset pricing model that explains the basic asset pricing moments by including disaster risk in an otherwise standard model. Disasters reflect huge market crashes; one objection to the disaster model is that we might have seen too few disasters (Cochrane, 2017). The model resolves the equity premium and risk-free rate puzzles with low risk aversion, and explains excess stock return volatility.

To get rare disasters to account for return predictability, one needs to specify that the risk of a rare event changes over time (Gabaix, 2012; Cochrane, 2017). Gabaix (2012) shows that time-varying disaster risk produces the observed volatility of the log price-dividend ratio, the time-varying risk premia, and the return

[8]Based on the long-run risks parameters in Table I in Beeler and Campbell (2012).

predictability as observed in the data. As shown by van Binsbergen, Brandt, et al. (2012), the unconditional term structure of equity risk premiums is flat, which appears to be at odds with the data (Bansal, Miller, et al., 2021).

The yield curve on real bonds is flat: all yields are equal to the risk-free rate (Gabaix, 2012). Disasters with recoveries produce a downward sloping yield curve. To the extent of my knowledge, it is unknown whether disaster risks produce a value premium or long-term reversal in the cross section. Overall, the disaster model requires sixteen parameters.[9] Similar to Campbell and Cochrane (1999), the authors set the growth rates of consumption and dividends equal.

Prospect theory - Barberis and Huang (2001) and Barberis, Huang, and Santos (2001) present an asset pricing model that explains several stylized facts by using four ingredients of prospect theory. First, investors derive direct utility not only from consumption, but also from gains and losses itself. Second, investors are loss averse such that agents are more sensitive to losses than to gains. Third, loss-averse investors are risk-seeking over (large) losses. Finally, investors use narrowing framing and mental accounting (Barberis and Huang, 2001), or they distort probabilities (Barberis, Huang, and Santos, 2001).

Barberis, Huang, and Santos (2001) explain the equity premium puzzle and excess volatility, with low and stable consumption growth that is not predictable like their dividend growth. However, the model produces a too high risk-free rate and a too low log price-dividend ratio volatility. Their model yields return predictability, with an R^2 increasing with the return horizon, but the R^2 are lower than found in the data. It is unclear whether the model produces time-varying risk premiums. de Vries (2021) argues that the model of Barberis, Huang, and Santos (2001) implies an upward-sloping term structure of equity premiums and risks, but quantitative predictions are absent. I am unaware of a model with these ingredients that studies (real) bond yields.

Barberis and Huang (2001) use similar ingredients as Barberis, Huang, and Santos (2001), but the authors include narrow framing and mental accounting. This model has some success in the cross section as well, as it creates a value

[9]Based on Table I in Gabaix (2012).

premium and De Bondt-Thaler premium. Though, the value premium appears too high. On the other hand, Barberis, Jin, et al. (2021) also apply prospect theory to the cross section, but they find that prospect theory works especially poor in explaining the value premium. Overall, the model of Barberis and Huang (2001) needs sixteen parameters.[10]

In a similar vein, Pagel (2016) presents an asset pricing model by using the ex-ante rational expectations-based reference-dependent model of Köszegi and Rabin (2006) together with loss aversion. The reference point is forward looking, rather than backward looking as in habit formation (Campbell and Cochrane, 1999) or prospect theory (Barberis, Huang, and Santos, 2001). Specifically, the reference point is based on fully probabilistic rational beliefs about current and future consumption that the agent formed in the previous period. Investors receive gain-loss utility from unexpected changes in present consumption and from revisions in expectations over future consumption, such that gain-loss utility can be interpreted as utility over good and bad news. In contrast, the "reference point" in regret theory is simply based on the true best ex-post realization.

The model of Pagel (2016) explains the equity premium puzzle, the excess volatility and predictability of returns. The model fails to explain the risk-free rate puzzle and autocorrelation of returns. To match these empirical findings, the author needs to introduce long-run risks with time-varying consumption volatility (Bansal and Yaron, 2004) as well as time-varying disaster risk (Barro, 2006) and sluggish belief updating.

2. Economics and psychology

I end this section by discussing how regret relates to other economic and psychological concepts. Barberis and Huang (2001) and Barberis, Huang, and Thaler (2006) study implicitly the notion of regret by arguing that consumption is not the only carrier of utility and that regret is a possible interpretation for narrow framing and loss aversion. However, anticipated regret already occurs before any losses actually materialize (Janis and Mann, 1997), such that the basic emotion of regret

[10]Based on Table I in Barberis and Huang (2001).

(Zeelenberg, 2020) could be more primary than the aversion of realizing losses (Frydman and Camerer, 2016). The concept that losses loom larger than gains in prospect theory bears similarities with the finding that regret (after a negative experience) is more strongly felt than rejoicing (after a positive experience).

Proceeding on prospect theory, regret runs over final wealth levels as in expected utility theory, rather than gains or losses in cumulative prospect theory (Tversky and Kahneman, 1992). Loss aversion requires the specification of a reference point, similar to the counterfactual specification of regret aversion, but the psychological literature has found well-defined counterfactuals for investors (Lin, Huang, et al., 2006). Furthermore, regret enters convex in the utility function (i.e., the larger the foregone alternative, the higher regret, the more disutility) and is symmetric if one would consider rejoicing, rather than the asymmetric S-shaped concave-convex utility function of gains and losses in prospect theory.

Regret relates to long-term experience effects caused by the persistence and strength of emotions (Malmendier, 2021). The implications of regret on asset prices relate to return extrapolation (Atmaz, 2021) and optimism (Brunnermeier et al., 2007) as regret-averse investors have a preference for positively skewed risks (Gollier, 2018) and chase good returns. Frydman and Camerer (2016) argue that regret itself can provide a microfoundation for realization utility (Barberis and Xiong, 2012). Other interpretations of regret relate to cognitive dissonance (Chang et al., 2015) and belief-based explanations (Frydman and Camerer, 2016).

Regret is not a unique emotion when making decisions, because other emotions, such as disappointment, relief, anger, envy, satisfaction, and pride, are also often felt in a decision making context. However, all these other emotions can also be felt without one having made a decision (e.g., one can be disappointed in the weather, and proud of one's children), but regret is always linked to a decision (Zeelenberg, 2020). Thus, regret theory does not require an ex-ante fixed (probabilistic) reference point like the rational expectations-based model of Köszegi and Rabin (2006) and Pagel (2016), but regret uses the ex-post realized alternatives (see Lin, Huang, et al. (2006)).

Camille et al. (2004) provide neurological evidence that regret, based on ex-post alternative realizations, is different from disappointment, based on ex-ante (rational) expectations. Contrary to disappointment, which is experienced when a negative outcome happens relative to prior ex-ante expectations, regret is strongly associated with a feeling of responsibility for the ex-post outcome of the decision that has been made. The authors also report that disappointment is insignificant in the decision making process, while regret is significant and more intensely felt neurologically.

2.1 The Model

This section describes the preferences of the investors and how they set prices for equity and bonds in the market.

2.1.1 Preferences

Identical agents maximize expected discounted utility today t, with subjective discount factor δ, over the fraction ξ_t invested

$$\max_{\{\xi_t\}} u\left(C_t(\xi_t), X_t\right) + \delta \mathbb{E}_t\left[u\left(C_{t+1}(\xi_t), X_{t+1}\right)\right]. \tag{2.1}$$

Here, C_t is the realized consumption level, based on realized returns, while $X_t \geq C_t$ is the foregone consumption level, based on foregone returns. I follow the multiplicative regret theory of Quiggin (1994), which is based on the additive regret theory as originally formalized by Bell (1982) and Loomes and Sugden (1982). Regret theory is an alternative to expected utility theory (Neumann and Morgenstern, 1947).

The multiplicative regret-utility function is defined as

$$u(C_t, X_t) = \frac{C_t^{1-\gamma}}{1-\gamma} X_t^{\kappa}, \quad \gamma > 1 \tag{2.2}$$

with γ the risk-aversion parameter and κ the regret-aversion parameter.[11] The

[11] γ retains the classical interpretation of the Arrow-Pratt definition of risk aversion (Gollier,

(a) $u(C, X)$ for varying consumption levels C with fixed foregone consumption level X

(b) $u(C, X)$ for varying foregone consumption levels X with fixed consumption level C

Figure 2.1: Regret-utility function.

first term in the regret-utility function is standard CRRA utility. The second multiplicative term reflects the disutility of regret. Because $\gamma > 1$, the regret-utility function has a negative range and a positive domain. When the foregone consumption level X_t is larger than the realized consumption level C_t, investors feel regret as they could have been better off. If there is no regret aversion, i.e., $\kappa = 0$, then the regret-utility function is a standard CRRA utility function.

Figure 2.1 provides intuition for the behavior of the regret-utility function. For varying consumption levels C and a fixed foregone consumption level X, Figure 2.1a graphs the standard well-known behavior of CRRA utility: increasing and concave. A low consumption level corresponds to a bad state of the world. For varying foregone consumption levels X and a fixed consumption level C, Figure 2.1b shows that the regret-utility function is decreasing in foregone consumption. A large foregone consumption level yields high disutility, which corresponds to a bad state of the world for investors, since marginal utility is high. The strength of disutility depends on the regret-aversion parameter κ. To ensure that (i) marginal utility of consumption increases as foregone consumption increases and (ii) utility exhibits aversion to the foregone alternative, we need the condition $\gamma - 1 \geq \kappa \geq 1$.[12] The conditions are essential properties for modelling regret (Gollier, 2018).

Instead of multiplicative regret (Quiggin, 1994), one could also model addi-

2018).

[12]In line with Gabillon (2011), see properties P3 and P4b in Table 2.10 in the Appendix. The theorem of Diecidue and Somasundaram (2017) shows that regret theory holds in line with their behavioral foundation if the inequalities are strict, which I can easily assume as well (implying $\gamma - 1 > \kappa > 1, C > 0, X > 0$.

tive regret as originally and independently developed in the regret theory of Bell (1982) and Loomes and Sugden (1982). However, the multiplicative specification simplifies our calculations and leads to closed-form tractable results, as opposed to the additive specification which yields a nested utility function in a convex regret function. Conceptually, multiplicative and additive regret should lead to the same results. But, multiplicative regret, compared to additive regret, excludes rejoicing.

Rejoicing is an additional emotion, which is the opposite feeling of regret and results from downward counterfactual thinking. Rejoicing is felt when realized consumption turns out to be the more desirable result than the foregone consumption X_t, i.e., $C_t > X_t$: "the extra pleasure associated with knowing that, as matters have turned out, the agent has taken the best decision" (Loomes and Sugden, 1982). However, the emotional impact of regret is greater than rejoicing (Larrick and Bowles, 1995; Zeelenberg, Beattie, et al., 1996; Humphrey, 2004; Zeelenberg, 2020) since counterfactual thinking is primarily triggered after upward counterfactual thinking and negative experiences rather than downward counterfactual thinking and positive experiences (Kahneman and A.Tversky, 1979; Roese and Olson, 1995; Roese and Olson, 1997), as generally negative information exerts a greater influence on choices than positive information (Beattie et al., 1994) and pains persist longer than joys (Frijda, 1988; Frijda, 2007). For this reason, I only model regret, which is in line with the multiplicative regret theory of Quiggin (1994) which excludes the feeling of rejoicing compared to additive regret.

To ease interpretation, think of the unit of time as a year, so that consumption and foregone consumption are measured annually. Investors might check their portfolio more often than that, but I assume that it is only once a year that investors confront their performance in a serious way.[13] Without loss of generality of results, decision problem (2.1) can be extended to a multi-period model with a finite or an infinite horizon.

[13]This annual review period is in line with the yearly review periods of Benartzi and Thaler (1995) and Barberis and Huang (2001).

Consumption and foregone consumption

Having described regret utility, we now turn towards realized consumption and foregone consumption. Investors can buy or sell a risky asset with gross return R_{t+1}. We can think of the risky asset as the market portfolio. The representative investor invests a fraction ξ_t of her wealth w_t. e_t denotes the consumption level if the agent decides to invest all of her wealth, which ensures that consumption remains non negative. I exclude short selling such that $0 \leq \xi_t \leq 1$. Then, realized consumption today C_t and realized consumption next year C_{t+1} equal

$$\begin{aligned} C_t &= e_t + w_t (1 - \xi_t), \\ C_{t+1} &= e_{t+1} + w_t \xi_t R_{t+1}. \end{aligned} \tag{2.3}$$

Foregone consumption is defined as the largest level of consumption that would have been attainable if another decision would have been made (Bell, 1982; Loomes and Sugden, 1982; Quiggin, 1994). If there is a foregone alternative that yields higher consumption than realized consumption, then regret is felt. In case the foregone alternative equals realized consumption, the investor feels no regret as her decision is the best she could have made ex post.

Thus, regret follows from upward counterfactual thinking using foregone alternatives. Upward counterfactuals follow after a negative experience and take the form of “if only...” statements (Lin, Huang, et al., 2006). Investors consider a state of the world where they would have been better off in terms of consumption levels, i.e., $X_t \geq C_t$ for all t: “If only I had made another decision, I would have had a higher consumption level”. Lin, Huang, et al. (2006) study real investors and they find that investors base their foregone alternatives on two counterfactuals: the inaction alternative and the best unchosen alternative.[14] In my model, investors use these two counterfactuals to determine their foregone consumption levels today X_t

[14] The authors also report the expected outcome as a potential reference point. However, as also stated by Lin, Huang, et al. (2006), if an outcome does not match the investor’s expectation, then disappointment is felt. Camille et al. (2004) provide neurological evidence that regret is distinct from disappointment, such that disappointment falls outside the scope of the current analysis. Moreover, disappointment differs from the foregone and inaction alternative in the sense that the latter two are evoked after a negative experience only.

and next year X_{t+1}

$$\begin{aligned} X_t &= e_t + w_t, \\ X_{t+1} &= e_{t+1} + w_t \tilde{R}_{t+1}. \end{aligned} \tag{2.4}$$

First, the foregone consumption level today X_t originates from the inaction alternative. This is the counterfactual thought of not having invested anything in the risky asset, i.e., $\xi_t = 0$. The representative investor considers a state of the world where she would have been better off by just holding her wealth: "if only I did not invest." Second, the foregone consumption level next year X_{t+1} originates from the best unchosen alternative, which is the counterfactual thought of having invested all wealth, $\xi_t = 1$, in the risky asset (i.e., portfolio of assets) with a higher return than realized returns, i.e., $\tilde{R}_{t+1} \geq R$. We can think of $\tilde{R}$ as the foregone return on the risky asset. The representative investor anticipates the experience of a state of the world where she would have been better off by any alternative investment: "if only I made another investment decision." Hence, the foregone consumption level today X_t follows from the foregone inaction alternative, while the foregone consumption level next year X_{t+1} follows from the foregone best unchosen alternative. So, investors experience utility from realized consumption, but also experience disutility regarding the consumption they could have received, had they made a different decision. The foregone consumption levels are the main mechanism that drive investors' regret.

It is convenient to think of regret as the ratio of the foregone consumption levels, i.e., X_{t+1}/X_t, such that we can interpret regret in terms of foregone returns. Namely, the market-wide foregone return that ex-post higher alternative risky returns would have yielded over not investing. Thus, regret is high when the market-wide ex-post unrealized risky returns could have been high.

Regret

I model regret in line with three observations from the literature. First, at the moment of investment decision making investors anticipate the feeling of regret, as shown in the representative investor's problem (2.1). Neuroscientific evidence (Camille et al., 2004) and the review of Zeelenberg (2018) show that anticipating future regret influences current decision making under uncertainty. Although regret is only felt when consumption is realized and foregone consumption is known, the investor anticipates and takes into account this emotion when making her investment decision today such that counterfactual thinking influences the investor's holdings today. If we think of the risky return R as the market return, then the anticipation of the market-wide foregone return $\tilde{R}$ directly influences the investors' decisions today and, thereby, the composition and prices of the market today and consequently future to be realized market returns. Anticipation and the feeling of regret are possible, because asset prices are available after any investment decision such that feedback is always received.

Second, emotions, and thereby regret, follow laws. Regret varies over time around a mean, and is persistent, and is not too volatile. The laws of change and habituation (Frijda, 1988; Frijda, 2007) state that regret is time varying and reverts over time to a steady-state mean (Wilson and Gilbert, 2005), but only slowly and gradually. Additionally, the laws of hedonic asymmetry and conservation of emotional momentum (Frijda, 1988; Frijda, 2007) state that regret is persistent, which extends up to years. Especially negative feelings and emotions, to which regret belongs, are persistent phenomena (Coricelli et al., 2005; Hajcak and Olvet, 2008). People have the tendency to overestimate the anticipated intensity and duration of their emotional feelings (Wilson and Gilbert, 2005). Arisoy et al. (2021) find that regret extends up to years by empirically showing that regret-sorted portfolios are highly persistent.[15] The law of care for consequence manifests the presence of emotional control (Frijda, 1988; Frijda, 2007), such that the emotion of regret is not too volatile.

Third, investors are averse to regret. The main premise of regret theory (Bell, 1982; Loomes and Sugden, 1982) is that we are averse to regret. Bleichrodt, Cillo,

[15]Using financial market data, Arisoy et al. (2021) show that regret is a highly persistent phenomenon, extending up to years. The authors form portfolios based on a regret measure (i.e., the maximum return in the same industry), and they show that a regret stock remains in the highest regret quintile for five months up to years with a probability up to 60%.

et al. (2010) introduce the first quantitative measurement of regret aversion, and their evidence confirms regret aversion at the individual and aggregate level. To the best of my knowledge, they are the first and only ones estimating a utility curvature parameter and a regret-aversion parameter based on power utility forms of regret. They estimate a regret-aversion parameter $\kappa = 2$, which implies increasing disutility in the foregone alternative because $\kappa > 1$. $\kappa < 1$ implies decreasing disutility in the foregone alternative, and $\kappa = 1$ corresponds to linear regret. Regret aversion (Bell, 1982; Loomes and Sugden, 1982; Quiggin, 1994), which generates the distinctive predictions of regret theory, implies that regret should enter convex, i.e., $\kappa > 1$.

Dynamics

To complete the description of preferences, I specify how they develop over time. First, given the laws of emotions, log regret $x_t \equiv \log\left(\frac{X_{t+1}}{X_t}\right)$ evolves as an AR(1) process

$$x_{t+1} \equiv \log\left(\frac{X_{t+1}}{X_t}\right) = \phi x_t + \mu_x + \varepsilon_{x,t+1}, \tag{2.5}$$

in which μ_x, σ_x, and $0 < \phi < 1$ are parameters such that regret is stationary and positively autocorrelated. In line with the aforementioned psychological, neuroscientific and experimental evidence, regret evolves as a time-varying and persistent process with coefficient ϕ, gradually and slowly mean reverting around μ_x with a low emotional volatility of regret σ_x. The persistence coefficient ϕ ensures that regret slowly varies over time. Since regret equals the foregone return, we can think of long-run mean of regret, $\mu_x/(1-\phi)$, as the steady premium that the foregone risky return would have offered. Regret is subject to shocks, but emotions behave in a controlled and stable manner such that volatility of regret σ_x is low. Later, I specify the calibration of all these parameters.

Second, we close the model by describing the processes for consumption and dividends. It is convenient to introduce dividend growth already here, since we need dividend growth for computing returns. Following Campbell and Cochrane

(1999), I model realized consumption growth and dividend growth as independently and identically distributed (i.i.d.) lognormal processes. Thus,

$$\Delta c_{t+1} \equiv \log\left(\frac{C_{t+1}}{C_t}\right) = \mu_c + \varepsilon_{c,t+1}, \tag{2.6}$$

in which the parameter μ_c is mean consumption growth and the parameter σ_c denotes the volatility of consumption growth. And dividend growth is white noise

$$\Delta d_{t+1} \equiv \log\left(\frac{D_{t+1}}{D_t}\right) = \mu_d + \varepsilon_{d,t+1}, \tag{2.7}$$

in which the parameter μ_d is mean dividend growth and the parameter σ_d denotes the volatility of dividend growth.

All three processes have the following correlation structure

$$\begin{bmatrix} \varepsilon_{c,t} \\ \varepsilon_{x,t} \\ \varepsilon_{d,t} \end{bmatrix} \sim N\left(0, \begin{bmatrix} \sigma_c^2 & \sigma_{c,x} & \sigma_{c,d} \\ \sigma_{c,x} & \sigma_x^2 & \sigma_{d,x} \\ \sigma_{c,d} & \sigma_{d,x} & \sigma_d^2 \end{bmatrix}\right) \tag{2.8}$$

and i.i.d. over time, with covariances $\sigma_{c,d} = \sigma_c\sigma_d\rho_{c,d}$, $\sigma_{c,x} = \sigma_c\sigma_x\rho_{c,x}$ and $\sigma_{d,x} = \sigma_d\sigma_x\rho_{d,x}$. $\rho_{c,d}$, $\rho_{c,x}$ and $\rho_{d,x}$ respectively denote the correlation between consumption and dividends, the correlation between consumption and regret, and the correlation between dividends and regret. The growth rates of consumption and dividends are weakly correlated (Campbell and Cochrane, 1999). Consumption growth and regret (i.e., foregone returns) correlate negatively, and dividend growth and regret correlate negatively as well. Intuitively, when consumption or dividends are low, regret is high as the investor missed out on a good opportunity.[16]

[16] One may wonder if the specified dynamics for regret (2.5) and consumption growth (2.6) fulfil the condition $X_t \geq C_t$. A simulation exercise, with 7000 simulations and 500 years of data, shows that $X_t \geq C_t$ holds true for all dates t. The simulation for regret starts in the steady state. Initial values for both consumption and foregone consumption are normalized to one.

2.1.2 Equity

I now compute equilibrium equity asset prices and returns to show how regret influences these. As evident from the representative agent's problem (2.1), regret is independent of the invested fraction ξ since regret follows from the counterfactuals of having invested all wealth differently. Therefore, marginal utility with respect to consumption, its first argument, is

$$u_1(C_t, X_t) = C_t^{-\gamma} X_t^{\kappa}, \tag{2.9}$$

Intuitively, when realized consumption C_t is low, regret X_t is high due to their negative correlation such that marginal utility of consumption is relatively high as investors find themselves in a bad state of the world.

Taking the first-order conditions in this economy yields the stochastic discount factor

$$M_{t,t+1} = \delta \frac{u_1\left(C_{t+1}, X_{t+1}\right)}{u_1\left(C_t, X_t\right)} = \delta \left(\frac{C_{t+1}}{C_t}\right)^{-\gamma} \left(\frac{X_{t+1}}{X_t}\right)^{\kappa}. \tag{2.10}$$

It is related to the time-discount factor, innovations in consumption and regret, and the aversion to risk and regret. Since consumption growth and regret are negatively correlated, the stochastic discount factor is more volatile than under standard CRRA utility (i.e., $\kappa = 0$). We can now compute moments of the stochastic discount factor and find equity asset prices. To do so, I follow the approach of Lettau and Wachter (2007) who price zero-coupon equity as long-lived assets in the economy. With this approach, we can also easily price zero-coupon bonds in the economy, as I do later below.

Let $P_{n,t}$ be the value of a dividend paid n periods from now. Absence of arbitrage implies

$$P_{n,t} = \mathbb{E}_t[M_{t,t+n} D_{t+n}], \tag{2.11}$$

with boundary condition $P_{0,t} = D_t$, because equity maturing today must be worth aggregate dividend. The law of iterated expectations, as in a standard Lucas-tree model, yields the recursion

$$P_{n,t} = \mathbb{E}_t[M_{t,t+1}P_{n-1,t+1}], \quad n \geq 1. \tag{2.12}$$

Following standard practice, I guess that a solution to the recursion depends on the state variable of regret x_t

$$\frac{P_{n,t}}{D_t} = F_n(x_t), \tag{2.13}$$

such that, after dividing both sides of (2.12) by D_t,

$$F_n(x_t) = \mathbb{E}_t\left[M_{t,t+1}F_{n-1}(x_{t+1})\frac{D_{t+1}}{D_t}\right]. \tag{2.14}$$

Assuming joint lognormality of consumption growth, dividend growth and regret, we find a closed-form solution for the price-dividend ratio

$$\frac{P_{n,t}}{D_t} = F_n(x_t) = \delta^n e^{a_n + b_n x_t}, \tag{2.15}$$

with the coefficients recursively defined as

$$\begin{aligned} a_n &= a_{n-1} - \gamma\mu_c + \kappa\mu_x + \mu_d + b_{n-1}\mu_x + \frac{1}{2}\left(\gamma^2\sigma_c^2 + (\kappa + b_{n-1})^2\sigma_x^2 + \sigma_d^2\right) \\ &\quad - \gamma(\kappa + b_{n-1})\sigma_c\sigma_x\rho_{c,x} - \gamma\sigma_c\sigma_d\rho_{c,d} + (\kappa + b_{n-1})\sigma_d\sigma_x\rho_{d,x}, \\ b_n &= \kappa\phi + b_{n-1}\phi, \end{aligned} \tag{2.16}$$

and initial conditions $a_0 = b_0 = 0$. The Online Appendix gives a derivation with intermediate steps. Hence, prices relative to dividends depend on the time-discount factor, the time-invariant constants a_n and b_n, and the time-varying, but persistent, state variable of regret. Please notice that the recursive parameter b_n is an increasing function of the regret-aversion parameter κ and the persistence coefficient of regret ϕ.

Returns, predictability and volatility

To find returns, let $R_{n,t+1}$ denote the one-period return on zero-coupon equity that matures in n periods

$$R_{n,t+1} = \frac{P_{n-1,t+1}}{P_{n,t}} = \frac{F_{n-1}(x_{t+1})}{F_n(x_t)} \frac{D_{t+1}}{D_t}. \tag{2.17}$$

The price-dividend ratio on the aggregate stock market is a claim to all future dividends such that the aggregate price-dividend ratio is the sum of the price to aggregate dividend ratios for all n-period claims

$$\frac{P_t^m}{D_t} = \sum_{n=1}^{\infty} \frac{P_{n,t}}{D_t} = \sum_{n=1}^{\infty} F_n(x_t). \tag{2.18}$$

Then, the return on the aggregate stock market equals

$$R_{t+1}^m = \frac{P_{t+1}^m + D_{t+1}}{P_t^m} = \frac{P_{t+1}^m / D_{t+1} + 1}{P_t^m / D_t} \frac{D_{t+1}}{D_t}. \tag{2.19}$$

For the simulated results later below, I use the aggregate price-dividend ratio and I calculate aggregated results as expected returns, volatility, predictability and other interesting quantities. However, I present intuition for my results based on analyzing zero-coupon equity returns $R_{n,t+1}$. The intuition derived from zero-coupon equity extends to the aggregate market.

Starting with Shiller (1981) and Campbell and Shiller (1988), the literature shows that returns are excessively volatile and predictable using lagged prices scaled by fundamentals, such as dividends. Fama and French (1988) and Poterba and Summers (1988) find that returns have small negative autocorrelations due to slow mean reversion. To illustrate the intuition for predictability, we start by computing the returns $R_{n,t+1}$. Substituting the price-dividend ratio and the processes for dividend growth and regret, we find

$$\begin{aligned}\log(1+R_{n,t+1}) &= \log\left(\frac{F_{n-1}(x_{t+1})}{F_n(x_t)}\right) + \log\left(\frac{D_{t+1}}{D_t}\right), \\ &= -\log(\delta) + a_{n-1} - a_n + b_{n-1}x_{t+1} - b_n x_t + \mu_d + \varepsilon_{d,t+1}, \\ &= -\log(\delta) + a_{n-1} - a_n + (b_{n-1}\phi - b_n)x_t \\ &+ b_{n-1}\mu_x + b_{n-1}\varepsilon_{x,t+1} + \mu_d + \varepsilon_{d,t+1},\end{aligned} \tag{2.20}$$

in which the final equality follows from the process for regret x_{t+1}. Substitution of the recursion for b_n from (2.16) implies the log one-period holding return on the n-period dividend strip

$$\log(1+R_{n,t+1}) = -\log(\delta) + a_{n-1} - a_n - \kappa\phi x_t + b_{n-1}\mu_x + b_{n-1}\varepsilon_{x,t+1} + \mu_d + \varepsilon_{d,t+1}, \tag{2.21}$$

such that log conditional expected returns are given by

$$\begin{aligned}\log \mathbb{E}_t\left[1+R_{n,t+1}\right] = &-\log(\delta) + a_{n-1} - a_n - \kappa\phi x_t + b_{n-1}\mu_x + \mu_d \\ &+ \frac{1}{2}\left(b_{n-1}^2\sigma_x^2 + \sigma_d^2\right) + b_{n-1}\sigma_d\sigma_x\rho_{d,x}.\end{aligned} \tag{2.22}$$

Thus, returns are predictable and affected by the investor's regret through counterfactual thinking. If regret x_t is high today, prices relative to dividends are high, see (2.15), and future returns are low. Intuitively, if regret today is high, then the foregone return on the risky asset is high. The representative investor regrets having invested too little and demands more of the risky asset, which pushes up risky prices today. Consequently, prices relative to fundamentals are high today such that expected future returns must fall. For this reason, the lagged price-dividend ratio predicts future returns.

The emotion of regret also delivers the small negative autocorrelation of returns as observed in the data. High returns today forecast low returns in the future. Note that standard CRRA utility (i.e., $\kappa = 0$) or white noise regret (i.e., $\phi = 0$) yield no predictability or autocorrelation of returns, as returns are not time varying.

Regret also creates excessive volatility of prices relative to dividends and, consequently, high volatility of returns. The unconditional variance of returns (2.21) equals

$$V\left[\log(1+R_{n,t+1})\right] = \kappa^2\phi^2\frac{\sigma_x^2}{1-\phi^2} + b_{n-1}^2\sigma_x^2 + 2b_{n-1}\sigma_d\sigma_x\rho_{d,x} + \sigma_d^2. \tag{2.23}$$

Returns are more volatile than the volatility of the underlying dividends σ_d alone, as the stochastic discount factor is volatile. The volatility in returns arises from the emotional volatility of regret σ_x, which by itself is not too large, but the regret-aversion parameter κ and the persistence coefficient ϕ enlarge the impact of regret volatility through b_{n-1}. Similar to Guo and Wachter (2021), most excess volatility on the market comes from the second term, $b_{n-1}^2\sigma_x^2$, as it is an order of magnitude larger than the first and third terms. Intuitively, the return on the market is best represented for long-maturity equity strips, (i.e., large n) which implies that the recursion coefficient b_{n-1} is large. Returns in the economy are more volatile than standard CRRA utility or when the emotion of regret would be white noise, as in both cases return volatility equals dividend volatility.

Risk-free interest rate and equity premium

The equity-premium risk-free rate puzzle (Mehra and Prescott, 1985; Weil, 1989) is the stylized fact that stock returns are high and the risk-free rate is low, compared to the implications of a standard CRRA model. Empirically, the model needs low mean consumption growth and volatility, unpredictable consumption and dividend growth, matching market volatility, a slowly varying risk-free rate, with low risk aversion and a low subjective discount rate. To provide intuition for these stylized facts, we compute the risk-free rate and equity premium.

A risk-free asset exists in the economy as well, and is assumed to be in zero-net supply. Then, the real one-period risk-free rate is given by the reciprocal of the conditionally expected stochastic discount factor

$$R_{f,t} = 1/\mathbb{E}_t\left[M_{t,t+1}\right]. \tag{2.24}$$

The risk-free asset is priced with the investor's Euler equation as

$$1 = \mathbb{E}_t \left[\delta \left(\frac{C_{t+1}}{C_t} \right)^{-\gamma} \left(\frac{X_{t+1}}{X_t} \right)^{\kappa} R_{f,t} \right], \tag{2.25}$$

implying the log risk-free rate

$$\log(1 + R_{f,t}) = -\log\delta + \gamma\mu_c - \kappa(\mu_x + \phi x_t) - \frac{1}{2}\gamma^2\sigma_c^2 - \frac{1}{2}\kappa^2\sigma_x^2 + \gamma\kappa\sigma_c\sigma_x\rho_{c,x}. \tag{2.26}$$

The risk-free rate is low and stable. The term $-\kappa(\mu_x + \phi x_t)$ reflects an intertemporal substitution effect. Intuitively, when regret x_t is high, the investor feels poor as she missed out on a high alternative return. So, the agent is willing to save more, which drives down the equilibrium interest rate. The risk-free rate is stable, because the volatility of consumption σ_c and regret σ_x are not too large. If regret aversion is absent, then we fall back in the standard class of CRRA utility models and find the classical risk-free rate puzzle.[17]

Theoretically, regret could explain the high mean excess return that we observe empirically in the stock market. The equity premium equals the difference between the log conditional expected one-period return on the n-period dividend strip (2.22) and the log one-period risk-free rate (2.26) such that

$$\begin{aligned} \log \mathbb{E}_t \left[(1 + R_{n,t+1})/(1 + R_{f,t}) \right] = {} & a_{n-1} - a_n + b_{n-1}\mu_x + \mu_d + \frac{1}{2}\left(b_{n-1}^2\sigma_x^2 + \sigma_d^2 \right) \\ & + b_{n-1}\sigma_d\sigma_x\rho_{d,x} - \gamma\mu_c + \kappa\mu_x + \frac{1}{2}\left(\gamma^2\sigma_c^2 + \kappa^2\sigma_x^2 \right) \\ & - \gamma\kappa\sigma_c\sigma_x\rho_{c,x}. \end{aligned} \tag{2.27}$$

[17] Stronger regret aversion implies a lower and more volatile risk-free rate. The intertemporal substitution effect becomes stronger but also the emotional volatility.

Substituting the recursion for a_n from (2.16) yields

$$\begin{aligned}\log \mathbb{E}_t\left[(1+R_{n,t+1})/(1+R_{f,t})\right] &= \log \mathbb{E}\left[(1+R_{n,t+1})/(1+R_{f,t})\right] \\ &= \gamma\sigma_c\sigma_d\rho_{c,d} - \kappa\sigma_d\sigma_x\rho_{d,x} - \kappa b_{n-1}\sigma_x^2 + \gamma b_{n-1}\sigma_c\sigma_x\rho_{c,x},\end{aligned} \tag{2.28}$$

which could produce a higher equity premium compared to standard CRRA utility.

The first term represents the standard consumption risk premium as in CRRA utility models, i.e., the risk aversion multiplied by the covariance between consumption and dividends. The second term represents a regret risk premium for regret-averse investors, in line with the finding of Qin (2020). Intuitively, if dividends are low, regret is high such that marginal utility is high as investors do not like such a state of the world. However, as we will see below, the regret risk premium is small in my calibration and, thus, the equity risk premium remains similar to the standard consumption risk premium. The reason is that $\kappa\sigma_d\sigma_x\rho_{d,x}$ influences the volatility of returns (2.23) and the risk premium with different signs, such that there exists a trade-off. If $\kappa\sigma_d\sigma_x\rho_{d,x}$ is strongly negative, then the volatility of returns is smaller, while the risk premium will be larger. Vice versa, a small negative $\kappa\sigma_d\sigma_x\rho_{d,x}$ produces higher return volatility, but a lower risk premium.

The last two terms reflect a bond risk premium, which is decreasing in the maturity of the claim n, since $\rho_{c,x}$ is negative. Therefore, regret produces a downward sloping term structure of equity risk premiums. Regret-averse investors demand a higher premium on short-term assets than long-term assets. Intuitively, investors are concerned with regret in the short run as they confront their performance annually, such that investors require a higher premium for holding risky assets in order to compensate for potential regret in the short term. If regret aversion is absent $\kappa = 0$, then we find a flat equity term structure.

Observe that the equity risk premium is not time varying as the conditional and unconditional equity premium are equal to each other. The reason is that the state variable of regret x_t enters the risk-free rate and the risky return, such that it cancels in the equity risk premium. Though, introducing time-varying volatility of regret would produce time-varying risk premiums. However, I am unaware of behavioral studies regarding time-varying volatility of emotions, such that I cannot support or oppose this claim.[18]

2.1.3 Bonds

Analogous to zero-coupon equity, we can price zero-coupon bonds. Let $P^B_{n,t}$ denote the real price of a real bond maturing in n periods. I highlight the differences with equity by adding an additional B to the expressions for bonds. Bond prices are determined recursively by the investor's Euler equation

$$P^B_{n,t} = \mathbb{E}_t \left[M_{t,t+1} P^B_{n-1,t+1} \right], \quad n \geq 1, \tag{2.29}$$

in which the stochastic discount factor $M_{t,t+1}$ is the same as before and given by (2.10). We work with a real stochastic discount factor such that bond prices are in real terms. When $n = 0$, the bond is worth one unit of consumption good, implying the boundary condition $P^B_{0,t} = 1$. Similar to zero-coupon equity, regret x_{t+1} is a state variable and bond prices are an affine function of time-varying regret x_t up to some time-invariant constants. The solution to the recursion of the bond prices takes the form

$$F_n(x_t) = P^B_{n,t}, \tag{2.30}$$

such that we can write (2.29) as

$$F_n(x_t) = \mathbb{E}_t \left[M_{t,t+1} F_{n-1}(x_{t+1}) \right]. \tag{2.31}$$

Since the processes for consumption growth and regret are jointly lognormal, we find a closed-form analytical solution for the price of an n-period real bond at time t

$$P_{n,t} = F_n(x_t) = \delta^n e^{a_{B,n} + b_{B,n} x_t}, \tag{2.32}$$

[18]Likewise, the conditional variance of returns is not time-varying, but an extension of the model with time-varying regret volatility would do so.

with the coefficients recursively defined as

$$\begin{aligned} a_{B,n} &= a_{B,n-1} - \gamma\mu_c + \kappa\mu_x + b_{B,n-1}\mu_x \\ &\quad + \frac{1}{2}\left(\gamma^2\sigma_c^2 + (\kappa + b_{B,n-1})^2\sigma_x^2\right) - \gamma(\kappa + b_{B,n-1})\sigma_c\sigma_x\rho_{c,x}, \\ b_{B,n} &= \kappa\phi + b_{B,n-1}\phi, \end{aligned} \tag{2.33}$$

and initial values $a_{B,0} = b_{B,0} = 0$. The Online Appendix gives a derivation with intermediate steps. Similar to zero-coupon equity, notice that the recursive parameter $b_{B,n}$ is an increasing function of the regret-aversion parameter κ and the persistence coefficient of regret ϕ.

Returns and yields

Piazzesi and Schneider (2006b) find that real bond yield curves are unconditionally downward sloping in U.K. data, while U.S. data suggests an unconditional upward sloping curve. To illustrate intuition for the real yield curve in the regret model, we want to find bond yields and bond returns. In line with Wachter (2006), I define the real return on an n-period bond as

$$R_{n,t}^B = \frac{P_{n-1,t+1}^B}{P_{n,t}^B}, \tag{2.34}$$

and the (continuously compounded) yield on the n-period real bond as

$$y_{n,t} = -\frac{1}{n}\log P_{n,t}^B. \tag{2.35}$$

Substituting the bond prices from (2.32), we find the one-period holding return on the n-period bond

$$\log(1 + R_{n,t+1}^B) = -\log\delta + a_{B,n-1} - a_{B,n} - \kappa\phi x_t + b_{B,n-1}\mu_x + b_{B,n-1}\varepsilon_{x,t+1} \tag{2.36}$$

and the bond yield

$$y_{n,t} = -\log\delta - \frac{a_{B,n}}{n} - \frac{b_{B,n}}{n}x_t, \tag{2.37}$$

with the yield on an one-period bond as

$$\begin{aligned} y_{1,t} &= -\log\delta - a_{B,1} - b_{B,1}x_t, \\ &= \log(1+R_{f,t}). \end{aligned} \tag{2.38}$$

Thus, bond returns and bond yields are functions of regret x_t and the time-invariant recursion coefficients, which depend on the maturity n. As such, bond returns and yields would theoretically be predictable by regret. The one-period yield is identical to the risk-free rate in equation (2.26).

The unconditional mean yield curve follows from

$$\mathbb{E}\left[y_{n,t}\right] = -\log\delta - \frac{a_{B,n}}{n} - \frac{b_{B,n}}{n}\frac{\mu_x}{1-\phi} \tag{2.39}$$

and the unconditional mean yield spread equals

$$\mathbb{E}\left[y_{n,t} - y_{1,t}\right] = \left(a_{B,1} - \frac{a_{B,n}}{n}\right) + \left(b_{B,1} - \frac{b_{B,n}}{n}\right)\frac{\mu_x}{1-\phi}. \tag{2.40}$$

Bond yields and the yield spread are decreasing functions of maturity n, producing an downward sloping real term structure of interest rates. Intuitively, the model's bond risk premium predicts that long-term bonds have a lower risk premium than short-term bonds and, thus, the bond yield curve slopes downward. If the long-run mean of regret $\frac{\mu_x}{1-\phi}$ is large, or regret aversion κ (implicit in $b_{B,n}$) is strong, then the yield curve becomes steeper. In case of CRRA utility (i.e., no regret), the real yield curve is exactly flat as there is no bond risk premium.

2.1.4 Cross section

We have seen that regret can explain stylized facts in the time series. In this part, I show that regret also explains two cross sectional features of equity: the value premium (Basu, 1983; Fama and French, 1992) and the De Bondt-Thaler premium (De Bondt and Thaler, 1985), also known as long-term reversal.

Basu (1983) and Fama and French (1992) show that stocks with low price-to-fundamentals ratios (value stocks), such as price-dividend ratios, exhibit significantly higher subsequent returns than stocks with high price-to-fundamentals ratios (growth stocks). A stock's price-dividend ratio predicts the stock's subsequent return with a negative sign. The difference in returns earned by "value" stocks with low price-dividend ratios and "growth" stocks with high price-dividend ratios is known as the value premium.

De Bondt and Thaler (1985) show that a stock's return over the past three to five years (i.e., portfolio formation period) predicts the stock's subsequent return (i.e., portfolio evaluation period) with a negative sign in the cross section. The difference in the future returns earned by losing stocks and winning stocks is known as the De Bondt-Thaler premium. A slight change in interpretation of the model presented above accounts naturally for both stylized facts. The results above rely on the interpretation of regret by counterfactually considering the investor's realized *portfolio* returns with the inaction alternative and the foregone best unchosen alternative (Lin, Huang, et al., 2006). For the cross section, the two counterfactuals remain the same, but we are interested in the behavior of an individual risky asset itself. As such, we study the asset pricing implications if the investor counterfactually compares the return on a chosen individual risky *asset* with the two foregone alternatives of Lin, Huang, et al. (2006).

Consider a market with two risky assets k and l. The representative agent invests in both risky assets. However, rather than being concerned with the foregone return on her portfolio, the representative investor is concerned with the foregone return on each individual asset. That is, the investor holds two separate regret processes for each asset. An interpretation is that the agent holds an individual mental account for each asset. Thus, the investor considers the performance and feedback on stock k and stock l separately.

Consistent with the finding of investors' mental individual stock accounting (Thaler, 1985; Barberis, Huang, and Santos, 2001), agents maximize utility over allocations $\xi_{j,t}$ to individual assets $j = l, k$

$$u(C_t, X_t) \equiv \sum_j u(C_{j,t}, X_{j,t}), \quad \forall t \tag{2.41}$$

with realized and foregone consumption such that $\xi_{j,t} = [0, 1]$

$$C_{j,t} = e_{j,t} + w_{j,t}(1 - \xi_{j,t}) \quad C_{j,t+1} = e_{j,t+1} + w_{j,t}\xi_{j,t}R_{j,t+1} \tag{2.42}$$

$$X_{j,t} = e_{j,t} + w_{j,t} \qquad\qquad X_{j,t+1} = e_{j,t+1} + w_{j,t}\tilde{R}_{j,t+1} \tag{2.43}$$

The representative investor's foregone consumption today on asset j, $X_{j,t}$, reflects the inaction alternative (i.e.,"if only I did not invest in asset j"). Foregone consumption next year on asset j, $X_{j,t+1}$, follows from the (anticipated) foregone return $\tilde{R}_{j,t+1} \geq R_{j,t+1}$ on risky asset j.[19]

Thus, rather than a single process for regret, the investor holds two separate processes for regret based on each risky asset j such that log regret evolves as

$$x_{j,t+1} \equiv \log\left(\frac{X_{j,t+1}}{X_t}\right) = \phi x_{j,t} + \mu_x + \varepsilon_{x,j,t+1}, \text{ for } j = \{l, k\}, \tag{2.44}$$

in which μ_x, σ_x, and $0 < \phi < 1$ are the same parameters as before in equation (2.5). Intuitively, investors' regret is subject to same laws of emotion (Frijda, 1988; Frijda, 2007) such that regret inhibits the same persistence, mean and volatility, but the innovations to regret are different for each asset j. I assume that regret shocks and dividend shocks are uncorrelated across assets, such that the correlation structure per asset j in equation (2.8) remains.

Intuition for the value premium and long-term reversal now follows easily. Since asset j is subject to the asset-specific process of regret $x_{j,t+1}$, prices relative to dividends for asset j depend on asset-specific regret

$$\frac{P_{j,t}}{D_{j,t}} = \sum_{n=1}^{\infty} F_n(x_{j,t}). \tag{2.45}$$

[19] In case the realized return $R_{j,t+1}$ is negative, the representative investor imagines a situation where she would have invested in a risk-free asset.

Consequently, the log conditional expected one-period return on the n-period dividend strip on stock j equals

$$\log \mathbb{E}_t \left[1 + R_{j,n,t+1}\right] = -\log(\delta) + a_{n-1} - a_n - \kappa\phi x_{j,t} + b_{n-1}\mu_x + \mu_d + \frac{1}{2}\left(b_{n-1}^2\sigma_x^2 + \sigma_d^2\right) + b_{n-1}\sigma_d\sigma_x\rho_{d,x}. \tag{2.46}$$

The return on stock j is identical to the market return in equation (2.22), however returns now are asset specific through the asset specific regret $x_{j,t}$. Namely, the investor experiences regret on each asset separately rather than on her portfolio. The asset-specific regret drives the cross-sectional findings.

To fix ideas, suppose that regret of holding stock k is higher than regret of holding stock l, i.e., $x_{k,t} > x_{l,t}$. It follows from the price-dividend ratio that stock k has a high price-to-fundamentals ratio, whereas stock l has a low one. Thus, stock k is overpriced relative to stock l. So, stock k identifies as a growth stock and stock l as a value stock. The value premium predicts that stock k has lower subsequent returns, while stock l has higher subsequent returns.

Actually, we can compute the differential expected excess return on the value-minus-growth stock. It equals

$$\log \mathbb{E}_t \left[1 + R_{l,n,t+1}\right] - \log \mathbb{E}_t \left[1 + R_{k,n,t+1}\right] = \kappa\phi(x_{k,t} - x_{l,t}), \tag{2.47}$$

such that, value stocks, identified as stock l, indeed offer a premium over growth stocks, identified as stock k. The value premium equals the difference in regret multiplied by the regret-aversion parameter κ and the persistence parameter ϕ. Clearly, a value premium is absent in standard CRRA utility models (i.e., $\kappa = 0$) or when regret is white noise (i.e., $\phi = 0$).

Intuitively, if regret $x_{j,t}$ on stock j is high today, then this stock typically generated good returns in the past (a winner stock). The investor regrets having invested too little in stock j and demands more of stock j, pushing up the price-dividend ratio today. As a result, prices compared to fundamentals today are overvalued (a growth stock) such that expected future returns are worse. Overall, this mechanism explains the value premium, i.e., growth stocks earn lower future returns than value stocks. Also, it explains the De Bondt-Thaler premium: winner stocks, formed on past performance, earn lower future returns than loser stocks. One mechanism, namely regret, explains both stylized facts in the cross section, which is plausible since long-term reversal is an alternative proxy for value (Fama and French, 1996; Gerakos and Linnainmaa, 2018).

2.2 Model evaluation

In this section, I evaluate the implications of the model in magnitudes. I simulate data by drawing the normally distributed shocks for consumption growth, dividend growth and regret. The simulated data for equity and bonds replicates many interesting statistics consistent with the empirical asset pricing literature.

2.2.1 Parameter values

Table 2.2 summarizes the parameter choices for my simulations, which are done at an annual frequency. I simulate 7000 samples, each consisting of 500 years of data from the model, to calculate population values for a variety of statistics. The amount of samples and the length of each sample is chosen such that the population moments are close to the theoretical moments.

I choose the parameters in the model in three ways. First, I choose several parameters to match certain moments as found in the empirical asset pricing literature. The parameters for dividend growth and consumption growth are chosen in line with Barberis and Huang (2001) and Lettau and Wachter (2007). Dividend growth is a white noise process with mean zero and a standard deviation σ_d of 14%, while consumption grows annually with a mean μ_c of 1.84% and a standard deviation of σ_c 3.82%. Correlation between dividend and consumption growth is imperfect (Campbell and Cochrane, 1999) and set to 0.30. The persistence coefficient of regret is set as $\phi = 0.81$ to closely match the serial correlation of the log price-dividend ratio as found in the empirical asset pricing literature (Campbell and Cochrane, 1999).

Second, I choose the preference parameters in line with the experimental literature. Using the estimated value of Bleichrodt, Cillo, et al. (2010), I set the regret-aversion parameter to $\kappa = 2$. This value implies that marginal utility of consumption is increasing in the foregone consumption level. The subjective annual time-discount factor equals $\delta = 0.97$ (Frederick et al., 2002), which yields a very plausible annual discount rate of 3%. I set the risk-aversion parameter γ to 10, which is considered to be plausible by Mehra and Prescott (1985), Bansal and Yaron (2004), and Beeler and Campbell (2012).

Third, I need to calibrate the four parameters in the regret process to match moments as found in the empirical asset pricing literature. The mean μ_x and standard deviation σ_x of regret are chosen such that the risk-free rate is low and stable in the economy, and such that we can interpret the long-run mean of log regret as the foregone premium that risky assets offer over risk-free assets. I set $\mu_x = 0.0125$, such that the long-run mean of regret $\mu_x/(1-\phi)$ equals 6.51%, which yields a plausible interpretation of the premium that foregone risky assets offer. I set $\sigma_x = 0.015$, so emotional volatility of regret is not too large. We have seen that the correlations between consumption and regret must be negative, as well as between dividends and regret. I set $\rho_{c,x} = -0.1$ and $\rho_{d,x} = -0.1$ such that regret produces excess volatility and a downward sloping term structure of equity risk premiums.

2.2.2 Results aggregate market

Table 2.3 shows the basic moments of the aggregate market from the simulations, namely: the risk-free rate, the equity premium, the market and risk-free rate volatilities, the volatility of the log price-dividend ratio, and the mean and volatility of regret.

The results confirm the intuition from the earlier theoretical analysis. The model produces a low risk-free rate of 1.02%, which is also stable with a volatility of 4.18%. As such, regret solves the risk-free rate puzzle (Weil, 1989) since consumption growth is unpredictable with a realistic mean and low volatility, and the model has low risk aversion with positive time discounting. Notwithstanding

Table 2.2: Parameter choices and calibration. This table shows the parameters used in the simulations. The values are based on the asset pricing literature, experimental literature and a calibration. All values are annual.

Parameter	Variable	Value
Asset pricing literature:		
Mean consumption growth	μ_c	1.84%
Standard deviation consumption growth	σ_c	3.82%
Standard deviation dividend growth	σ_d	14%
Correlation consumption growth and dividend growth	$\rho_{c,d}$	0.30
Persistence coefficient	ϕ	0.81
Experimental literature:		
Regret aversion	κ	2
Time-discount factor	δ	0.97
Risk aversion	γ	10
Calibration:		
Mean regret	μ_x	1.25%
Standard deviation regret	σ_x	1.50%
Correlation consumption growth and regret	$\rho_{c,x}$	-0.1
Correlation dividend growth and regret	$\rho_{d,x}$	-0.1

these resolved stylized facts, the regret model has difficulty with explaining the equity premium puzzle (Mehra and Prescott, 1985). The equity premium is small, 1.03%, and mostly driven by the standard consumption CAPM. One might argue that the equity premium is low, because I compute the equity premium on unlevered claims. Empirically observed equity returns include firms with leverage and Abel (1999) shows that a correction for levered equity produces a higher equity premium in general equilibrium models.

Regret solves the excess volatility puzzle (Campbell and Shiller, 1988). Risky assets have a return volatility of nearly 19% and the volaility of the equity premium is about 20%, which is larger than the underlying dividend volatility. The volatility of the log price-dividend ratio is 0.22. All reported values of these moments are similar to the values reported in Bansal and Yaron (2004). The simulated mean and volatility of log regret equal their theoretical counterparts given by its AR(1) property, respectively $\mu_x/(1-\phi)$ and $\sqrt{(\sigma_x^2/(1-\phi^2))}$. Regret is not too volatile and the steady-state mean of regret can be interpreted as the foregone return.

Table 2.3: Simulated moments for the regret model. This table reports the means and standard deviations of simulated moments on the aggregate market. R^m denotes the return on the market, R_f the real risk-free rate, $R^m - R_f$ the equity premium, $p^m - d$ the log price-dividend ratio on the market, x regret and $\sigma(.)$ the standard deviation. The model is simulated at an annual frequency. All values are annualized.

Statistic	Model	Statistic	Model	Statistic	Model
$\mathbb{E}[R_f]$	1.02%	$\sigma[R_f]$	4.18%	$\sigma[p^m - d]$	0.22
$\mathbb{E}[R^m]$	2.05%	$\sigma[R^m]$	18.90%	$\mathbb{E}[x]$	6.51%
$\mathbb{E}[R^m - R_f]$	1.03%	$\sigma[R^m - R_f]$	20.04%	$\sigma[x]$	2.57%

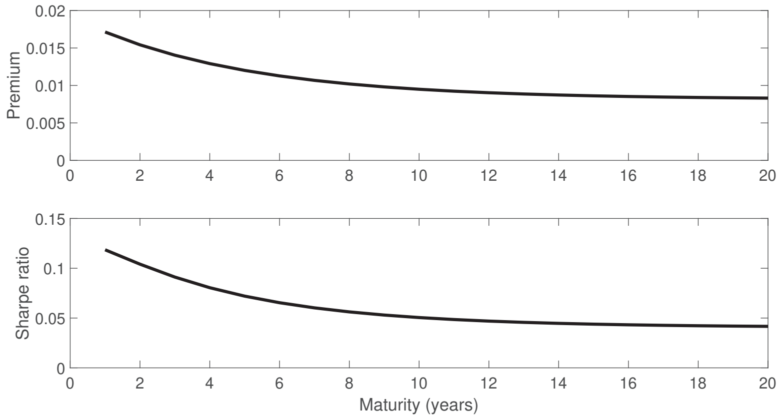

Figure 2.2: Term structures of equity risk premiums and Sharpe ratios. The graph shows the term structures of equity risk premiums and Sharpe ratios implied by the regret model. The graph plots the first 20 years of dividend strips.

Figure 2.2 shows the term structures of equity risk premiums and Sharpe ratios. The graphs show that the term structures of excess returns and Sharpe ratios are unconditionally downward sloping, consistent with the empirical evidence of (van Binsbergen, Brandt, et al., 2012; van Binsbergen and Koijen, 2017). In terms of shape, the regret model matches this new stylized fact in the asset pricing literature. The earliest dividend strips have an annual average excess return equal to 1.8 percent, while the latest dividend strips earn a annual average excess return of somewhat less than one percent. Note however that Bansal, Miller, et al. (2021) find an unconditionally upward sloping term structure of equity risk premiums.

2.2.3 Predictability

Here, I document autocorrelations, long-horizon predictability and cross-sectional predictability, consistent with the empirical asset pricing literature.

Autocorrelations

Table 2.4 presents autocorrelations for the excess returns, the price-dividend ratio and regret from the simulated data. The model fits the slight negative autocorrelation of returns (Fama and French, 1988; Poterba and Summers, 1988; Guo and Wachter, 2021). This produces mean reversion and time-series predictability: high returns today, predict low returns tomorrow. The price-dividend ratio and regret have identical autocorrelations as the price-dividend ratio is a function of regret and the persistence coefficient ϕ is chosen to generate the first-order annual autocorrelation of the log price-dividend ratio, which is close to the value of 0.78 as found in the data by Campbell and Cochrane (1999).

Table 2.4: Autocorrelations of simulated data. This table reports for several yearly lags the autocorrelations of the return on the market R^m, the equity premium $R^m - R_f$, the log price-dividend ratio on the market $p^m - d$, and regret x. The model values are based on time-aggregated annual values with a yearly simulation interval.

	Lag (years)				
Variable	1	2	3	5	7
R^m	-0.04	-0.03	-0.03	-0.02	-0.01
$R^m - R_f$	-0.02	-0.01	-0.01	-0.01	-0.01
$p^m - d$	0.80	0.64	0.51	0.33	0.20
x	0.80	0.64	0.51	0.33	0.20

Long-horizon regressions and the cross section

Table 2.5 documents time-series and cross-sectional predictability. Panel A reports regressions of lagged price-dividend ratios on future returns, consumption and dividends. The column with 'Returns' replicates our earlier intuition that the price-dividend ratio predicts subsequent returns with a negative sign. R^2 statistics are increasing in the horizon and range between 5% to 16%. The coefficients and R^2

are very similar to the reported values by Campbell and Cochrane (1999) and Guo and Wachter (2021). Columns 'Consumption' and 'Dividend' show, in contrast, that consumption growth and dividend growth are unpredictable, because they are simply white noise in the economy. Cochrane (2008), Beeler and Campbell (2012), and Cochrane (2017) argue that dividend growth and consumption growth are in fact unpredictable. The unpredictability is a necessary ingredient for a successful explanation of the equity-premium risk-free rate puzzle (Cochrane, 2017).

Table 2.5: Time-series and cross-sectional predictability. This table reports predictability results for the time series and the cross section. Panel A reports predictive coefficients and R^2-statistics from annual long-horizon regressions of cumulative log risky returns, consumption growth and dividend growth on the log price-dividend ratio: $\sum_{j=1}^{H} \log(1+R_{t+j}^m) = \beta_0 + \beta_1(p_t^m - d_t) + \varepsilon_{t+H}$, $\sum_{j=1}^{H} \Delta c_{t+j} = \beta_0 + \beta_1(p_t^m - d_t) + \varepsilon_{t+H}$, $\sum_{j=1}^{H} \Delta d_{t+j} = \beta_0 + \beta_1(p_t^m - d_t) + \varepsilon_{t+H}$. Panel B reports the mean annual value premium and its standard deviation (in percentage points). Panel C reports the mean annual De Bondt-Thaler premium for winner-loser formation periods n and for winner-loser evaluation periods N.

Panel A: Time series						
	Returns		Consumption		Dividend	
Horizon (years)	Coefficient	R^2	Coefficient	R^2	Coefficient	R^2
1	-0.20	0.05	0.00	0.00	0.00	0.00
2	-0.36	0.09	0.00	0.00	0.00	0.00
3	-0.49	0.12	0.00	0.01	0.00	0.01
5	-0.67	0.15	0.00	0.01	0.00	0.01
7	-0.80	0.16	0.00	0.01	0.00	0.01
Panel B: Price-to-fundamentals ratio						
Value premium	4.73%					
Standard deviation	23.66 pp					
Panel C: Long-term reversal						
Formation period n	3 years	5 years	5 years			
Evaluation period N	3 years	3 years	5 years			
De Bondt-Thaler premium	4.48%	5.01%	7.25%			

Panel B confirms our earlier analysis of the cross section. Each year, the two stocks in the economy are sorted based on their price-dividend ratio, and the returns of both stocks over the next year are measured. The value premium is the time-series mean of the difference between the returns. Regret produces a value premium of 4.73% with a standard deviation of the value-minus-growth portfolio of 23.66 percentage points. The magnitude of the value premium and its standard deviation are similar to the values of 5.42% and 20.39 percentage points as observed in the data by Guo and Wachter (2021). Additionally, the value premium is highly persistent in the model. On average, a growth (value) stock remains overpriced (underpriced) up to five years after portfolio formation, consistent with the finding of van Binsbergen, Boons, et al. (2021). As such, the 'mispricing' resolves gradually. Stated differently, a stock on average remains a low (i.e., value) or high (i.e., growth) regret stock for five years, in line with the observation from Arisoy et al. (2021) that regret on stocks is persistent.

In Panel C we see that regret produces a De Bondt-Thaler premium between 4.48% to 7.25%, depending on the portfolio formation and evaluation periods. The portfolio formation period sorts stocks in winners and losers, and the evaluation period tracks their subsequent returns. Every n years (formation period), the two stocks in the economy are sorted based on their n-year cumulative prior return, and the returns of both the winner and loser stocks over the next N years (evaluation period) are measured. The De Bondt-Thaler premium is the time-series mean of the difference between the returns of the loser and winner stocks over all non-overlapping periods. The premiums are smaller than observed by De Bondt and Thaler (1985), but do indicate that regret creates long-term reversal in asset prices.[20]

2.2.4 Bonds

Table 2.6 shows the implications of the model for means and standard deviations of real bond yields and returns. The maturities demonstrate that the average unconditional yield curve on real bonds is downward sloping, which is consistent with the empirical findings of Piazzesi and Schneider (2006b) on U.K. indexed bonds. However, U.S. data on indexed bonds suggests that the yield curve is unconditionally upward sloping, but the data series is very small. The magnitude of the mean yields on the one-year to five-years real bonds is similar to the model's average real yields of Piazzesi and Schneider (2006b). Volatilities on real yields are

[20] In the current exposition, regret cannot explain momentum. As regret is concerned with long-term persistence, it does not mean revert quickly enough in the short run to create momentum. Capturing both long-term reversal and momentum is a long-standing challenge (Barberis, 2018).

decreasing in maturity. Short-term yields are more volatile than long-term yields, a finding which Piazzesi and Schneider (2006b) empirically support with data.

The regret model implies that bond returns decrease in maturity, which contradicts the empirical asset pricing literature, while their volatilities rise with maturity, consistent with the empirical findings. The one-year bond return in logs behaves identical to the one-year bond yield, i.e., the risk-free rate in the economy, by construction. For this reason, the one-year bond return volatility appears high. The other maturities show that long-term bonds have higher volatility than short-term bonds.

Table 2.6: Moments for real bond yields and bond returns. This table reports the means and standard deviations for yields and returns on real zero-coupon bonds (i.e., bonds that pay off in units of aggregate consumption) in the simulated model for an annual holding period. Yields and returns are in annual percentages. Maturity is in years.

Bond yields	Mean	St. dev.	Bond returns	Mean	St. dev.
1 Year	0.96	4.12	1 Year	1.05	4.17
2 "	0.85	3.73	2 "	0.82	3.36
3 "	0.74	3.39	3 "	0.64	3.90
4 "	0.63	3.09	4 "	0.49	4.91
5 "	0.53	2.83	5 "	0.37	5.92
6 "	0.44	2.59	6 "	0.27	6.81
7 "	0.35	2.39	7 "	0.19	7.56
8 "	0.27	2.21	8 "	0.13	8.17
9 "	0.19	2.05	9 "	0.08	8.68
10 "	0.13	1.91	10 "	0.03	9.10
15 "	-0.13	1.39	15 "	-0.08	10.27
20 "	-0.29	1.07	20 "	-0.12	10.68

2.3 Empirical findings

In this section, I propose an empirical measure of regret based on stock returns. Consequently, I empirically show that the behavior of regret is in line with the model's predictions. The section concludes with the statistical properties of regret, which shows that regret-averse investors are concerned with positively skewed returns.

Testable predictions

Empirically, I focus on the following three main testable predictions regarding aggregate stock returns. First, regret is a persistent phenomenon. Second, the price-dividend ratio is a function of regret. High (anticipated) regret today implies high prices relative to dividends today. Third, (anticipated) regret today predicts future returns with a negative sign in the time series. High (anticipated) regret today predicts lower future returns.

My model also predicts a value premium in the cross section when considering regret. Stocks with high (low) regret are growth (value) stocks. Arisoy et al. (2021) extensively study empirically the behavior of regret sorted portfolios in the cross section. Besides their finding that regret is a persistent phenomenon, they also find that stocks with low (high) regret have low (high) price-dividend ratios, statistically significant at any reasonable level.[21] Thus, low regret stocks are value stocks and high regret stocks are growth stocks. This finding precisely matches my regret model's prediction in the cross section.

2.3.1 Regret measure

The basic premise of regret utility is that high regret corresponds to a high foregone alternative, which equals the maximum return that an investor could have achieved by an alternative foregone investment. I propose a regret measure based on counterfactual thoughts that approximate investors' regret as

$$REG_t^a = \max_j (R_{j,t}) \,. \tag{2.48}$$

Intuitively, in line with the definition of regret in the theoretical model, REG_t^a is the annual foregone return during year t. The measure resembles the market-wide regret investors experience during year t. When REG_t is high, the foregone return is high such that regret is high, which decreases investor's utility in (2.2).

Specifically, $R_{j,t}$ is the return during year t for stock j trading on the market. The representative agent considers all stocks j on the market as counterfactuals

[21] See Table 2, which displays the average stock characteristics of regret portfolios.

and compares her realized market portfolio return to the maximum return that could have been attained across these counterfactuals over the same period, which creates the measure REG_t. My proposed market-wide regret measure is in line with the individual-stock regret measure of Arisoy et al. (2021), who posit a regret measure for individual stocks by taking the maximum return over all stocks in the same industry. As shown later, regret relates to skewness, so my empirical regret measure also relates to the cross-sectional measures of Bali et al. (2011) and Lin and Liu (2018) who proxy expected skewness by the maximum return of a stock within a month.

As a robustness check, I also consider two other specifications of the regret measure. First, regret defined as the maximum foregone monthly return within a year, which I measure as

$$REG_t^m = \max_j \left(R_{j,t,m}\right) \text{ for each } t = 1, ..., T. \tag{2.49}$$

Here, $R_{j,t,m}$ equals the return of stock j during month m within year t. Second, regret defined as annual regret by averaging all monthly maximum foregone returns within year t, which I measure as

$$REG_t^{\bar{a}} = \frac{1}{m} \sum_{m=1}^{12} REG_{t,m} \quad \text{with} \quad REG_{t,m} = \max_j \left(R_{j,t,m}\right). \tag{2.50}$$

Here, the representative agent forms annual regret during year t by averaging all maximum foregone monthly returns within a year. To ensure comparability with model's regret x_{t+1}, I consider the log of the regret measures. Intuitively, x_{t+1} measures the foregone return in logs and $\log(REG_t)$ does likewise.

2.3.2 Data

To construct the regret measures, I download stock prices. Individual monthly stock prices come from the Center for Research in Security Prices (CRSP). I take all individual stocks that trade on the New York Stock Exchange (NYSE), American Stock Exchange (AMEX), and NASDAQ with share codes 10 and 11. I exclude stocks with share prices less than \$1 and more than \$1,000, such that the results are not driven by small and illiquid stocks. All these individual stocks arguably form the market.

To test the model's predictions, I download annual variables from Professor Shiller's website. The main interest is in the real price-dividend ratio and the market return, represented by the S&P500. I consider a postwar sample from 1947 to 2012, and a long sample from 1926 to 2012.[22]

2.3.3 Results

Persistence

First, I show that regret is a persistent phenomenon. Figure 2.3 presents the history of the log real price-dividend ratio in the postwar sample with the annual log regret measure $\log(REG_t^a)$. The graphs indicate that regret is a time-varying phenomenon, but gradually and slowly moving. In fact, regret is highly persistent.

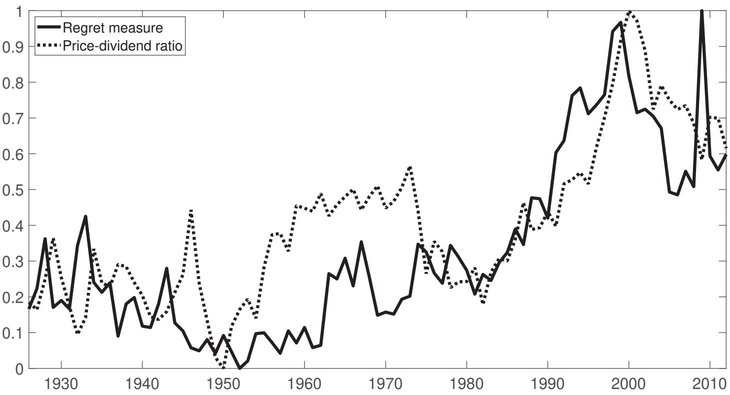

Figure 2.3: Historical price-dividend ratio and regret. The graph shows the normalized time-series of the real log price-dividend ratio and log regret measure $REG_t^{\bar{a}}$, which measures annual regret by averaging all monthly maximum foregone returns within a year. The values are based on annual data in the long sample (1926-2012).

[22]The sample stops at 2012 as the annual dataset of Professor Shiller's website ends in 2012.

Table 2.7 presents the first-order autocorrelation coefficients for the three regret measures, in the postwar and long samples. In the postwar sample the serial correlation coefficients of regret are somewhat higher than in the long sample. Overall, the persistence coefficients range from 0.68 to 0.92. So, this is in line with the dynamics of regret as prescribed by the laws of emotions (Frijda, 1988; Frijda, 2007) as found in the psychological, experimental and neuroscientific literature.

Table 2.7: Autocorrelations of regret measures. This table reports first-order autocorrelation, AR(1), coefficients of the regret measures in the postwar sample (1947-2012) and long sample (1926-2012). The values are based on annual data.

	$\log(REG_t^a)$	$\log(REG_t^m)$	$\log(REG_t^{\bar{a}})$
Postwar sample	0.73	0.81	0.92
Long sample	0.68	0.75	0.91

Price-dividend ratio

Second, the price-dividend ratio is a function of regret. The behavior of regret is similar to the behavior of the price-dividend ratio. Especially from the 1970s onward. The regret model predicts that high regret today implies high prices relative to dividends today. Table 2.8 confirms our eyeballing from Figure 2.3 and shows the regression results of regressing the price-dividend ratio on regret. In both samples and for all regret measures, the coefficients on regret are positive and highly statistically significant.[23] The results imply that high regret today yields high prices relative to dividends today. The R^2 in the postwar sample range from 42% to 56%. The relation between regret and prices relative to dividends is somewhat stronger in the postwar sample than in the long sample.

Return predictability

The third prediction to test is that high regret today implies lower future aggregate stock returns. Table 2.9, in a similar spirit to the earlier long-horizon regressions, regresses future returns on lagged regret. The main takeaway of the regression

[23]Throughout this section, we calculate standard errors using the Newey-West procedure with a lag of $T^{1/4}$, where T equals the number of observations.

Table 2.8: Price-dividend ratio and regret. This table reports OLS regression of the real price-dividend ratio on the regret measure i: $\log(P_t/D_t) = a + b\log(REG_t^i) + \varepsilon_t$. Newey-West standard errors $\sigma(b)$ to correct for heteroskedasticity and serial correlation. The values are based on annual data.

	Postwar sample			Long sample		
Measure	Coefficient	$\sigma(b)$	R^2	Coefficient	$\sigma(b)$	R^2
$\log(REG_t^a)$	0.41	0.09	0.42	0.37	0.09	0.31
$\log(REG_t^m)$	0.65	0.13	0.48	0.61	0.13	0.38
$\log(REG_t^{\bar{a}})$	1.13	0.19	0.56	1.18	0.18	0.54

results is that the point estimates for the coefficients on regret are negative: high regret today predicts lower future returns. The relation between regret and future returns becomes less statistically insignificant when the horizon increases. At a seven year horizon, the R^2 ranges from 5% to 8% in the postwar sample.

Table 2.9: Future returns and regret. This table reports predictive coefficients, standard errors and R^2-statistics from long-horizon regressions of cumulative log real risky returns on regret measure i: $\sum_{j=1}^{H}\log(R_{t+j}^m) = a + b\log(REG_t^i) + \varepsilon_{t+H}$. Newey-West standard errors $\sigma(b)$ to correct for heteroskedasticity and serial correlation. The values are based on annual data.

		Postwar sample			Long sample		
Measure	Horizon (years)	Coefficient	$\sigma(b)$	R^2	Coefficient	$\sigma(b)$	R^2
$\log(REG_t^a)$	1	-0.02	0.03	0.01	0.02	0.03	0.01
"	2	-0.05	0.04	0.02	0.02	0.05	0.00
"	3	-0.09	0.06	0.05	0.00	0.07	0.00
"	5	-0.14	0.09	0.06	-0.06	0.09	0.01
"	7	-0.16	0.10	0.06	-0.11	0.09	0.03
$\log(REG_t^m)$	1	-0.03	0.06	0.01	-0.01	0.04	0.00
"	2	-0.07	0.09	0.02	-0.04	0.08	0.00
"	3	-0.08	0.12	0.02	-0.05	0.12	0.00
"	5	-0.13	0.18	0.03	-0.11	0.17	0.02
"	7	-0.22	0.19	0.05	-0.21	0.17	0.05
$\log(REG_t^{\bar{a}})$	1	-0.04	0.07	0.01	-0.02	0.07	0.00
"	2	-0.12	0.14	0.02	-0.06	0.14	0.00
"	3	-0.17	0.20	0.03	-0.09	0.19	0.00
"	5	-0.30	0.28	0.05	-0.22	0.27	0.02
"	7	-0.43	0.29	0.08	-0.35	0.27	0.05

Overall, this section confirms empirically the three main regret model's predictions: (i) regret is persistent, (ii) prices relative to dividends are a function of regret, and (iii) regret predicts future returns in the time series with a negative with an economically meaningful R^2.

Properties of regret

One may wonder which properties of returns drive regret-averse investors. To provide intuition, this section presents a simple theoretical exercise.

Assume that stock returns are i.i.d. lognormally distributed in the cross section, such that each individual stock return is distributed as $R_i \sim \log N(\mu, \sigma^2)$ in which R_i is the gross return on stock $i = 1, ..., n$. Using the standard properties of a lognormal distribution, the mean and variance of the lognormally distributed returns depend on the parameters μ and σ, while the skewness of the returns depends on the parameter σ only.

Define cross-sectional regret, in line with the earlier measures, as

$$REG = \max_i (R_i)\,. \tag{2.51}$$

Then, the CDF $F_{REG}(x)$ of the random variable REG equals

$$F_{REG}(x) = \Phi\left(\frac{\log(x) - \mu}{\sigma}\right)^n, \tag{2.52}$$

where Φ denotes the CDF of the standard normal distribution. The PDF $f_{REG}(x)$ is

$$f_{REG}(x) = \frac{dF_{REG}}{dx}(x). \tag{2.53}$$

Consequently, we can compute the first moment of the function $f_{REG}(x)$ by

$$m_1 = \int_{-\infty}^{\infty} x f_{REG}(x)dx, \tag{2.54}$$

which is the mean, or expected value, of regret. We can find the variance and skewness of regret by computing the *centralized* moments $n = 2, 3$ respectively:

$$m_n = \int_{-\infty}^{\infty} (x - m_1)^n f_{REG}(x) dx. \tag{2.55}$$

We find that mean regret m_1 is high when the skewness of the underlying returns is positive and high, or when the volatility of the underlying returns is high.[24] Namely, the behavior of regret is dominated by the parameter σ, i.e., the skewness of the returns, while the parameter μ has a small effect on regret. Intuitively, a high cross-sectional dispersion in returns leads to a higher probability of a foregone missed return. Regret itself is substantially less positively skewed than the underlying returns, i.e., the third centralized moment m_3 is small. So, regret-averse investors are mainly concerned about the positive skewness of the returns, which is in line with the findings of Eeckhoudt et al. (2007) and Gollier (2018) who report that regret-averse agents have a preference for positively skewed risks and longshots.

2.4 Conclusion

I explore the idea whether regret can provide a unified explanation for the behavior of asset prices. The central ingredient is regret and the aversion to it, added to an otherwise standard asset pricing model. Regret is based on the intuition that an investor is concerned not only about the outcome she receives, but also about the outcome she could have received, had she invested differently. The central finding is that regret aversion has the potential to explain asset pricing stylized facts in the time series and in the cross section. I provide evidence for the main model's predictions by using an empirical measure of regret, especially that prices relative to dividends are a function of persistent regret.

Three features of the regret model are unique, compared to leading asset pricing models. First, one simple ingredient has the ability to explain several stylized facts in a unifying way. In the current setup, regret cannot yet explain the equity

[24] I used simulations with $n = 1000$ individual stocks and parameter choices for individual stock behaviour with mean $\mu = \{0.01, 0.1\}$ and standard deviation $\sigma = \{0.04, 0.4\}$.

premium puzzle, due to the calibration trade-off between the regret risk premium and excess volatility. Since I am unaware of behavioral studies regarding time-varying volatility of emotions, the regret model does not make any predictions about time-varying risk premiums. Second, regret produces a downward sloping term structure of equity risk premiums, whereas most other models predict the opposite. Third, regret-averse investors produce a value premium and long-term reversal in the cross section. Regret-sorted stocks are persistent and the value premium is persistent, consistent with the empirical findings of Arisoy et al. (2021) and van Binsbergen, Boons, et al. (2021).

A potential avenue for further research is the link between institutional investors, regret, benchmarks and asset prices. Regret, and their counterpart of rejoicing, could potentially be linked to the investment industry as institutional investors typically try to follow an investment benchmark. If they do not achieve their benchmark, then investors could experience regret as they could have made an alternative investment decision. Finally, psychologists interested in finance might be inspired to study the time-varying volatility of regret.

2.5 Appendix

Table 2.10: Properties multiplicative regret-utility function $u(c, x)$. u_1 (u_2) denotes the partial derivative with respect to its first (second) argument. The same rule applies to u_{11}, u_{12} and so on.

Property		
P1a	Utility is increasing (states without regret)	$u_1(c,c) + u_2(c,c) \geq 0$
P1b	Utility is concave (states without regret)	$u_{11}(c,c) + u_{12}(c,c) + u_{21}(c,c) + u_{22}(c,c) \leq 0$
P2a	Utility increases with consumption	$u_1(c,x) \geq 0$
P2b	Utility decreases with regret	$u_2(c,x) \leq 0$
P2c	Utility is globally increasing	$u_1(c,x) + u_2(c,x) \geq 0$
P3	Utility is sensitive to regret	$u_{12}(c,x) = u_{21}(c,x) \geq 0$
P4a	Utility exhibits diminishing marginal consumption utility	$u_{11}(c,x) \leq 0$
P4b	Utility exhibits aversion to the foregone alternative	$u_{22}(c,x) \leq 0$

2.6 Online Appendix

A. Zero-coupon equity

This section derives the prices for zero-coupon equity in the economy along the lines of Lettau and Wachter (2007).

Time t, $n = 1$

Assuming joint log-normality of consumption growth, regret and dividend growth it follows that

$$\begin{aligned}\frac{P_{1,t}}{D_t} &= \mathbb{E}_t\left[M_{t,t+1}\frac{D_{t+1}}{D_t}\right] \\ &= \mathbb{E}_t\left[\delta\left(\frac{C_{t+1}}{C_t}\right)^{-\gamma}\left(\frac{X_{t+1}}{X_t}\right)^{\kappa}\frac{D_{t+1}}{D_t}\right] \\ &= \mathbb{E}_t\left[\delta e^{-\gamma(\mu_c+\varepsilon_{c,t+1})+\kappa(\phi x_t+\mu_x+\varepsilon_{x,t+1})+\mu_d+\varepsilon_{d,t+1}}\right]\end{aligned} \tag{2.56}$$

We can write this as

$$e^{w_{t,1}+x\varepsilon_{c,t+1}+y_1\varepsilon_{x,t+1}+z\varepsilon_{d,t+1}} \tag{2.57}$$

with (please note the difference between the time-independent parameter x with time-varying regret x_t)

$$\begin{aligned}w_{t,1} &= -\gamma\mu_c + \kappa(\phi x_t + \mu_x) + \mu_d \\ &\equiv w_{1,1} + w_{1,2}x_t \\ w_{1,1} &= -\gamma\mu_c + \kappa\mu_x + \mu_d \\ w_{1,2} &= \kappa\phi\end{aligned} \tag{2.58}$$

$$x = -\gamma\sigma_c \tag{2.59}$$

$$y_1 = \kappa\sigma_x + y_0, \quad y_0 = 0 \tag{2.60}$$

$$z = \sigma_d \tag{2.61}$$

Compute conditional expectation

$$\begin{aligned}\frac{P_{1,t}}{D_t} &= \delta e^{-\gamma\mu_c+\kappa(\phi x_t+\mu_x)+\mu_d+\frac{1}{2}\gamma^2\sigma_c^2+\frac{1}{2}\kappa^2\sigma_x^2+\frac{1}{2}\sigma_d^2} \\ &\cdot e^{-\gamma\kappa\sigma_c\sigma_x\rho_{c,x}-\gamma\sigma_c\sigma_d\rho_{c,d}+\kappa\sigma_x\sigma_d\rho_{x,d}}\end{aligned} \tag{2.62}$$

We can write this as

$$\frac{P_{1,t}}{D_t} = F_1(x_t) = \delta e^{a_1+b_1x_t} \tag{2.63}$$

with

$$\begin{aligned}a_1 &= -\gamma\mu_c + \kappa\mu_x + \mu_d + \frac{1}{2}\gamma^2\sigma_c^2 + \frac{1}{2}\kappa^2\sigma_x^2 + \frac{1}{2}\sigma_d^2 \\ &\quad - \gamma\kappa\sigma_c\sigma_x\rho_{c,x} - \gamma\sigma_c\sigma_d\rho_{c,d} + \kappa\sigma_x\sigma_d\rho_{x,d} \\ &\equiv w_{1,1} + \frac{1}{2}(x^2 + y_1^2 + z^2) + xy_1\rho_{c,x} + xz\rho_{c,d} + y_1z\rho_{x,d}\end{aligned} \tag{2.64}$$

$$\begin{aligned}b_1 &= \kappa\phi \\ &\equiv w_{1,2}\end{aligned} \tag{2.65}$$

Time $t+1$, $n=1$

$$\frac{P_{1,t+1}}{D_{t+1}} = F_1(x_{t+1}) = \delta e^{a_1+b_1x_{t+1}} = \delta e^{a_1+b_1(\phi x_t+\mu_x+\varepsilon_{x,t+1})} \tag{2.66}$$

Time t, $n=2$

$$\begin{aligned}\frac{P_{2,t}}{D_t} &= F_2(x_t) = \mathbb{E}_t\left[M_{t,t+1}\frac{P_{1,t+1}}{D_{t+1}}\frac{D_{t+1}}{D_t}\right] \\ &= \mathbb{E}_t\left[\delta^2 e^{-\gamma(\mu_c+\varepsilon_{c,t+1})+\kappa(\phi x_t+\mu_x+\varepsilon_{x,t+1})+\mu_d+\varepsilon_{d,t+1}+a_1+b_1(\phi x_t+\mu_x+\varepsilon_{x,t+1})}\right]\end{aligned} \tag{2.67}$$

We can write this as

$$e^{w_{t,2}+x\varepsilon_{c,t+1}+y_2\varepsilon_{x,t+1}+z\varepsilon_{d,t+1}} \tag{2.68}$$

with

$$\begin{aligned} w_{t,2} &= -\gamma\mu_c + \kappa(\phi x_t + \mu_x) + \mu_d + a_1 + b_1(\phi x_t + \mu_x) \\ &\equiv w_{2,1} + w_{2,2}x_t \\ w_{2,1} &= -\gamma\mu_c + \kappa\mu_x + \mu_d + a_1 + b_1\mu_x \\ w_{2,2} &= \kappa\phi + b_1\phi = (\kappa + b_1)\phi \end{aligned} \tag{2.69}$$

$$x = -\gamma\sigma_c \tag{2.70}$$

$$y_2 = \kappa\sigma_x + b_1\sigma_x = (\kappa + b_1)\sigma_x \tag{2.71}$$

$$z = \sigma_d \tag{2.72}$$

Compute conditional expectation

$$\begin{aligned} \frac{P_{2,t}}{D_t} &= \delta^2 e^{-\gamma\mu_c + \kappa(\phi x_t + \mu_x) + \mu_d + a_1 + b_1(\phi x_t + \mu_x) + \frac{1}{2}\gamma^2\sigma_c^2 + \frac{1}{2}(\kappa+b_1)^2\sigma_x^2 + \frac{1}{2}\sigma_d^2} \\ &\cdot e^{-\gamma(\kappa+b_1)\sigma_c\sigma_x\rho_{c,x} - \gamma\sigma_c\sigma_d\rho_{c,d} + (\kappa+b_1)\sigma_x\sigma_d\rho_{x,d}} \end{aligned} \tag{2.73}$$

We can write this as

$$\frac{P_{2,t}}{D_t} = F_2(x_t) = \delta^2 e^{a_2 + b_2 x_t} \tag{2.74}$$

with

$$\begin{aligned} a_2 &= -\gamma\mu_c + \kappa\mu_x + \mu_d + a_1 + b_1\mu_x + \frac{1}{2}\gamma^2\sigma_c^2 + \frac{1}{2}(\kappa + b_1)^2\sigma_x^2 + \frac{1}{2}\sigma_d^2 \\ &- \gamma(\kappa + b_1)\sigma_c\sigma_x\rho_{c,x} - \gamma\sigma_c\sigma_d\rho_{c,d} + (\kappa + b_1)\sigma_x\sigma_d\rho_{x,d} \\ &\equiv w_{2,1} + \frac{1}{2}(x^2 + y_2^2 + z^2) + xy_2\rho_{c,x} + xz\rho_{c,d} + y_2 z\rho_{x,d} \end{aligned} \tag{2.75}$$

$$\begin{aligned} b_2 &= \kappa\phi + b_1\phi \\ &\equiv w_{2,2} \end{aligned} \tag{2.76}$$

General t, n

$$\frac{P_{n,t}}{D_t} = F_n(x_t) = \delta^n e^{a_n + b_n x_t}, \quad a_0 = b_0 = 0 \tag{2.77}$$

with

$$\begin{aligned} w_{t,n} &= -\gamma\mu_c + \kappa(\phi x_t + \mu_x) + \mu_d + a_{n-1} + b_{n-1}(\phi x_t + \mu_x) \\ &\equiv w_{n,1} + w_{n,2} x_t \\ w_{n,1} &= -\gamma\mu_c + \kappa\mu_x + \mu_d + a_{n-1} + b_{n-1}\mu_x \\ w_{n,2} &= \kappa\phi + b_{n-1}\phi = (\kappa + b_{n-1})\phi \end{aligned} \tag{2.78}$$

$$x = -\gamma\sigma_c \tag{2.79}$$

$$y_n = \kappa\sigma_x + b_{n-1}\sigma_x = (\kappa + b_{n-1})\sigma_x \tag{2.80}$$

$$z = \sigma_d \tag{2.81}$$

$$a_n = w_{n,1} + \frac{1}{2}(x^2 + y_n^2 + z^2) + x y_n \rho_{c,x} + x z \rho_{c,d} + y_n z \rho_{x,d} \tag{2.82}$$

$$b_n = w_{n,2} \tag{2.83}$$

Thus, recursively

$$\begin{aligned} a_n &= -\gamma\mu_c + \kappa\mu_x + \mu_d + a_{n-1} + b_{n-1}\mu_x + \frac{1}{2}(x^2 + y_n^2 + z^2) + x y_n \rho_{c,x} + x z \rho_{c,d} + y_n z \rho_{x,d} \\ &= -\gamma\mu_c + \kappa\mu_x + \mu_d + a_{n-1} + b_{n-1}\mu_x + \frac{1}{2}(\gamma^2\sigma_c^2 + (\kappa + b_{n-1})^2\sigma_x^2 + \sigma_d^2) \\ &- \gamma(\kappa + b_{n-1})\sigma_c\sigma_x\rho_{c,x} - \gamma\sigma_c\sigma_d\rho_{c,d} + (\kappa + b_{n-1})\sigma_x\sigma_d\rho_{x,d} \end{aligned} \tag{2.84}$$

$$
\begin{aligned}
a_{n-1} - a_n &= \gamma\mu_c - \kappa\mu_x - \mu_d - b_{n-1}\mu_x - \frac{1}{2}(x^2 + y_n^2 + z^2) - xy_n\rho_{c,x} - xz\rho_{c,d} - y_n z\rho_{x,d} \\
&= \gamma\mu_c - \kappa\mu_x - \mu_d - b_{n-1}\mu_x - \frac{1}{2}\gamma^2\sigma_c^2 - \frac{1}{2}(\kappa + b_{n-1})^2\sigma_x^2 - \frac{1}{2}\sigma_d^2 \\
&+ \gamma(\kappa + b_{n-1})\sigma_c\sigma_x\rho_{c,x} + \gamma\sigma_c\sigma_d\rho_{c,d} - (\kappa + b_{n-1})\sigma_x\sigma_d\rho_{x,d} \\
&= \gamma\mu_c - \kappa\mu_x - \mu_d - b_{n-1}\mu_x - \frac{1}{2}\gamma^2\sigma_c^2 - \frac{1}{2}\kappa^2\sigma_x^2 - \frac{1}{2}b_{n-1}^2\sigma_x^2 - \kappa b_{n-1}\sigma_x^2 - \frac{1}{2}\sigma_d^2 \\
&+ \gamma\kappa\sigma_c\sigma_x\rho_{c,x} + \gamma b_{n-1}\sigma_c\sigma_x\rho_{c,x} + \gamma\sigma_c\sigma_d\rho_{c,d} - \kappa\sigma_x\sigma_d\rho_{x,d} - b_{n-1}\sigma_x\sigma_d\rho_{x,d}
\end{aligned}
\tag{2.85}
$$

and

$$
b_n = \kappa\phi + b_{n-1}\phi \tag{2.86}
$$

$$
b_{n-1}\phi - b_n = -\kappa\phi \tag{2.87}
$$

□

B. Zero-coupon bonds

The derivation for zero-coupon bond prices is completely analogous to the derivation for zero-coupon equity, but the starting equation is (2.29) rather than (2.11). □

Part II

Household Finance

Chapter 3

Can Estimated Risk and Time Preferences Explain Real-life Financial Choices?*

*We thank Hazel Bateman, Stefano Cassella, Bart Dees, Rob van den Goorbergh, Frank de Jong, Bas Werker, Stefan Zeisberger, and seminar participants at the American Economic Association (2022), American Economic Association (2021, Poster), International Association of Applied Econometrics (2021), European Economic Association (2020), Royal Dutch Economic Association KVS New Paper Sessions (2020), Netspar International Pension Workshop (2020), Netspar Anniversary Meeting (2019), APG Asset Management, and ABP Pension Fund for useful comments.

Risk and time preferences play a role in almost every economic decision. As a consequence, understanding the role of individual risk and time preferences is intimately linked to understanding economic behavior. The economic behavior we study in this paper is actual financial decision making, specifically pension payout choices by retirees. Understanding pension payout choices is important as it converts the pension wealth accrued in retirement plans into long-term income. Typically, these are large amounts for individuals. The average individual pension wealth at the time of retirement is \$302,500 in the United States, \$335,500 in the United Kingdom, and \$1,165,500 in the Netherlands.[1]

In this paper, we analyze the explanatory power of risk and time preferences for actual annuity choices. Directly measuring the empirical link between preferences and actual financial choices is challenging, because individual preferences, such as risk aversion, present bias, and patience, are not readily observable. The present paper provides evidence that risk and time preferences, measured with the Convex Time Budgets (CTB) method of Andreoni and Sprenger (2012a), can explain real-life pension payout choices to a large extent. In particular, we find within an expected utility framework that individually estimated preferences explain for 82% of our sample the choice between a flat annuity and a front-loaded annuity.

This study uses a unique combination of data on individual preferences and administrative data on real-life choices. To elicit individuals' risk and time preferences, we design a purpose-built internet survey that we field at the largest pension fund in the Netherlands. Our survey simultaneously measures risk and time preferences among pension fund participants in the same context as we observe the actual annuitization decisions. Because risk and time preferences are domain specific (Frederick et al., 2002; Schildberg-Hörisch, 2018; Cohen et al., 2020), we measure risk and time preferences in a pension context with large experimental budgets and long decision horizons similar to the observed real-life annuity choices from the administrative data. The administrative data also contain detailed information on the participants personal characteristics and their pension plans.

[1] OECD (2013), Pensions at a Glance 2013: OECD and G20 Indicators, OECD Publishing, Paris, https://doi.org/10.1787/pension_glance-2013-en.

Our population comprises 1062 pension fund participants. The individuals are invited by the pension fund, and in our CTB experiment they allocate €10,000 between an early payment and a late payment ten years in the future. We can expect individuals to spend more effort in thinking about their choice than in a laboratory with small stakes, no pension context, and shorter horizons. The domain-specific preferences that we measure are present bias, long-term patience, and CRRA utility function curvature. Currently, it is mandatory in the Netherlands to convert pension assets to an annuity. The actual pension payout decision concerns a choice between a flat fixed annuity with equal life-long payments throughout the retirement phase, and a front-loaded fixed annuity with higher payments during the first retirement years and actuarially fair lower life-long payments till death.

We study actual annuitization decisions in the context of risk and time preferences, because it appears intuitive that present-biased, impatient, and more risk tolerant individuals (i.e., those with curvature parameters close to unity, who are more tolerant about an unsmoothed consumption path) might prefer a front-loaded annuity. On the other hand, individuals that care more about the future and prefer smooth consumption paths might prefer a flat annuity with equal payments during retirement. A front-loaded annuity bears some resemblances with a lump sum, as it allows the beneficiary to receive pension payments earlier and higher compared to a flat annuity. Thus, in line with Brown (2001) and Inkmann et al. (2011), risk and time preferences are plausible and important channels for annuitization decisions. In addition to the predictive power of preferences for annuity choices, we also quantify the welfare implications that emerge through freedom of choice between a flat and front-loaded annuity.

The Dutch pension fund's data that we use in this paper has several advantages compared to other data sources. First, the dataset includes actual real-life annuity choices rather than incentives, attitudes, or stated preferences and choices on economic decision making. Second, the dataset provides detailed and reliable information on the participants and their pension plans, which is often hard to ask in surveys. Third, the annuity decision involves large stakes with long decision horizons, similar to our CTB experiment on risk and time preferences. Fourth, the annuity decision is less likely to be driven by concerns regarding health sate and medical expenditures due to the universal health care and social security in the Netherlands.[2] Finally, the Dutch annuity decision reflects global pension choices, as near retirees often have to make a choice between an annuity or a lump sum.

We make three main findings. First, our population estimates for risk and time preferences are comparable to previous estimates in the literature (Andreoni and Sprenger, 2012a; Andersen, Harrison, Lau, et al., 2014; Balakrishnan et al., 2020). This is interesting in itself, because previous studies that jointly estimate risk and time preferences typically do not use an explicit real-life pension context. For the quasi-hyperbolic discounting model (also known as the $\beta - \delta$ model, Laibson, 1997), we estimate a median present-bias factor of 0.84, a median annual discount rate of 1.1%, and a median CRRA utility curvature of 0.97. Hence, we find evidence for present bias and a preference for smooth consumption paths, consistent with the general observations in the literature (Frederick et al., 2002; Andreoni and Sprenger, 2012a; Potters et al., 2016).[3] Our estimated annual discount rate is lower than estimates in most previous research, as annual discount rates from 30%-100% are not uncommon (Frederick et al., 2002; Andreoni and Sprenger, 2012a; Cheung, 2020). Potential reasons for our lower, but plausible, estimated discount rate are the magnitude of the experimental budget and the long-term decision horizons, as individuals tend to make more patient choices in these cases (Thaler, 1981).[4]

The second set of results shows that our individually estimated risk and time preferences explain real-life financial decisions to a large extent. Using a simple univariate analysis, we find that patient individuals with a preference for smooth consumption paths choose a flat annuity, while present-biased, impatient, and more risk tolerant individuals choose a front-loaded annuity to withdraw more

[2]See for example de Bresser et al. (2018).

[3]Note that in the literature risk aversion over states of the world tends to deviate more from linear utility (Cheung, 2020).

[4]Other experiments typically have shorter decision horizons that run from several weeks to several months (Andersen, Harrison, M.Lauc, et al., 2010; Tanaka et al., 2010; Augenblick, Niederle, et al., 2015), but do not exceed more than 3 years (Harrison, Lau, et al., 2002; Goda et al., 2015). Moreover, the typical experimental payment equals tens of dollars (Andreoni and Sprenger, 2012a), rather than ten thousand dollars.

pension wealth during the early years of retirement. Using a discounted expected utility framework, we find that risk and time preferences explain actual annuitization decisions for 82% of our population for a utility indifference of at most 2% annual certainty equivalent consumption. This so-called 'indifference bandwidth' resembles a prediction error and indicates the annual consumption loss between the actually chosen and unchosen counterfactual annuity. Because a flat and front-loaded annuity might be observationally equivalent in terms of utility for a retiree, we study the predictive power of preferences while allowing for small consumption losses. Stated differently, individual preferences explain a large fraction of real-life annuity choices conditional on a small annual consumption difference.

The third set of results shows that freedom of choice by means of a front-loaded annuity creates potential welfare gains, but part of the welfare remains unrealized. In particular, given the predictive power of preferences for actual annuity choices, we perform a welfare analysis to investigate the effects of introducing freedom of choice in the annuity decision. We quantify the welfare effects of the front-loaded annuity option from a long-run persistent point of view (Ericson and Laibson, 2019), i.e., setting the present-bias factor to dynamically consistent behavior $\beta = 1$. The estimated mean conditional potential welfare gain of a front-loaded annuity ranges from 1.61% to 2.77% additional annual consumption (i.e., €7659 to €13417 total pension wealth), depending on the indifference bandwidth. The welfare distributions show that realized welfare can be negative and, thus, causes welfare losses. Overall, these findings can have important policy implications, as realized welfare gains can be roughly three times as large — for an indifference bandwidth of 2% — when individuals are guided better during their annuity choices.

Our paper relates to the household finance literature on retirement savings and annuitization (see Gomes et al., 2021 for an overview). Our primary contribution is to combine risk and time preferences with actual annuity choices. To our knowledge, no previous paper has related actual annuity decision making, by means of a utility framework, with simultaneously estimated risk and time preferences as inputs. Inkmann et al. (2011) use microeconomic data from the United Kingdom to study voluntary annuitization, focusing on individual determinants as explanations for annuity decisions but missing out on individually estimated preference parameters. Similar to ours, Bütler and Teppa (2007) study actual annuity and lumpsum decisions at retirement, but their data lacks individual preference parameters as well. Hurwitz and Sade (2020) study the annuity versus lump sum decision through the mechanism of smoking. Dimmock et al. (2021) do elicit individual preferences, but they elicit probability weighting and relate it to household portfolio choice. Andrikogiannopoulou and Papakonstantinou (2020) estimate individual's preferences within a prospect theory paradigm and apply it to explain the disposition effect.[5]

Previous research has been studying the explanatory power of risk and time preferences for economic behavior. However, most previous research relies on stated economic behavior, independently measured risk and time preferences that are context independent, and the explanatory power is studied through correlations (Cohen et al., 2020).[6] Dohmen, Falk, Huffman, and Sunde (2010) relate risk and time preferences separately to cognitive ability, while Dohmen, Falk, Huffman, Sunde, et al. (2011) correlate general and domain specific risk aversion to self-reported field behavior. Chabris et al. (2008) study correlations between laboratory-measured time preferences and self-reported behavior (e.g., BMI, smoking, exercise, saving, and gambling), and Golsteyn et al. (2014) study correlations between children's categorically measured time preferences and observed economic outcomes (e.g., schooling, health, labour, and income) later in life. Sutter et al. (2013) study how independently measured risk and time preferences correlate independently with self-reported behavior (e.g., health, savings, and schooling) amongst children and adolescents. Falk et al. (2018) study how independently measured risk and time preferences correlate separately with economic outcomes

[5]Erner et al. (2013) elicit cumulative prospect theory parameters from German students and find that these have virtually no predictive power for the willingness-to-pay for complex structured financial products, such as one or more derivatives on an underlying, typically a stock or stock index.

[6]Implicit in laboratory elicited general preferences is the assumption that laboratory results are a reliable assessment of general or specific behavior, even though we know that the typical subject pool is different from the population to which they are being applied (Andersen, Harrison, M.Lauc, et al., 2010). We overcome this problem by eliciting preferences and observing behavior directly in the same population and domain.

amongst individuals worldwide.

The rest of the paper is organized as follows. Section 1. describes the measurement of preferences and the data. Section 2. describes the estimation of preferences and shows the estimated results. Section 3. describes the annuity choices and the predictive power of preferences together with the welfare effects. Section 4. concludes the paper.

3.1 Methodology

To measure risk and time preferences, we field a survey at a large pension fund in The Netherlands. The survey implements the experimental CTB method (Andreoni and Sprenger, 2012a; Andreoni and Sprenger, 2012b) and a present-bias task (Frederick, 2005; Rieger et al., 2015). We relate the elicited preferences to actual pension choices of retirees.

3.1.1 Elicitation of risk and time preferences

We use the CTB to elicit patience and utility curvature, and we use an additional present-bias task to elicit present bias. Next, we adopt a simultaneous estimation technique to estimate utility curvature, patience, and present bias together. The advantage of our approach is a simultaneous measurement of risk and time preferences. For this reason, we avoid the assumption of linear utility and, consequently, we avoid upward-biased discount rate estimates if true utility is concave (Andersen, Harrison, Lau, et al., 2008a; Noor, 2009).

The CTB method asks individuals to allocate an initial budget $m =$ €10.000 between payments, available at two points in time: an early payment at time t and a delayed payment at time $t + k$. In line with Potters et al. (2016), the early payment is always one year $t = 1$ from the experimental date, and the late payment is delayed by ten years $k = 10$. The delay length is relatively long and selected such that we can study decision making under uncertainty for long horizons. Subjects receive an interest rate, or investment return, r on delayed payments, which varies between 0% to 8.40% on an annual basis. The allocations must be made such that the budget constraint is satisfied, i.e., the early payment and the present value of the delayed payment must equal the initial budget m. Early payments are certainly paid (i.e., payment probability one), but delayed payments have a payment probability p_{t+k} of 0.5, 0.75, 0.90, or 1.

Individuals make 20 consecutive CTB decisions between early and delayed payments. Our method consists of four different decision sets. Each decision set has a different probability of late payment, and within each set we have five different interest rate scenarios. The difference between the early payment date t and the delayed payment date $t + k$ elicits long-term patience, similar to Andreoni and Sprenger (2012a). We identify risk preferences by sensitivities to variation in the interest rates, similar to Andreoni and Sprenger (2012a), but also by sensitivities to the late payment probability (i.e., states of the world). Thus, we extend the original CTB approach of Andreoni and Sprenger (2012a). Table 3.6 in Appendix A presents an overview of our experimental design.

To identify present bias, we implement a task in our experiment from the INTRA (International Test of Risk Attitudes) study, conducted by the University of Zurich and used by Rieger et al. (2015). This task is inspired by Frederick (2005), and reads as follows:

> *Enter an amount $c_{t+\tau}$ such that option B is as attractive as option A:*
> A. Receive €800 now,
> B. Receive € $c_{t+\tau}$ next year.

Subjects make a trade-off between a direct payment of €800 now or a later certain payment c_1 next year. Due to the implementation of an immediate payment now combined with the long-run decisions from the CTB, we can elicit and estimate the (present-biased) time preferences for every subject while controlling for utility curvature. Table 3.6, Scenario 21, summarizes the present-bias task.[7]

[7]The original question is in US dollars. The monetary payoff of €800 in our scenario is adjusted according to the currency exchange rate in 2018 and the Purchasing Power Parity (PPP) in The Netherlands.

3.1.2 Experimental procedure

The CTB experiment and present-bias task are part of a larger survey. The pension fund wrote a Qualtrics program to implement the survey. In the first part of the survey, we ask subjects for personal information, such as pension attitudes, demographics (age, education), and financial situation (income, housing wealth). The second part of the survey contains the CTB experiment and, then, the present-bias task. Subjects could go through the survey, including the experiment, at their own pace, also going back and forth through the questions. In the email, and at the end of the survey, we announce that subjects are able to receive one out of five vouchers with a value of €50. The voucher will be received via email, implying that subjects need to enter their email address. The survey questions are available from the authors upon request.

Although the questions were not directly incentivized, the pension fund indicated in the instructions that the results would be taken into account to study the desirability of choice options, so participation in the survey was consequential. Our experiment is not incentivized based on the experimental answers of the subjects, which avoids the need for complex equalization of payments, transaction costs and payment confidence. Some researchers argue that answer-based incentives in economic experiments lead to more truthful reveal of preferences, however Cohen et al. (2020) and Hackethal et al. (2022) find little evidence for systematic differences between incentivized and unincentivized risk and time preference experiments. More specifically, Potters et al. (2016) find little differences between financially incentivized and hypothetical decisions in their CTB experiments.[8]

Upon starting the experiment, subjects read through the instructions and a CTB example decision screen. These indicated to the subjects that the budget could be entirely allocated to the early payment (corner), entirely to the later payment (corner) or divided between the two (interior). Figure 3.1 shows an image of a decision screen. The decision screen contains a timeline of the payment

[8]Another review by Camerer and Hogarth (1999) finds that incentives do not reliably change average performance, but tend to decrease the variance of responses. Since our sample is relatively large, this decreases the variance of the preference estimates on an aggregate level.

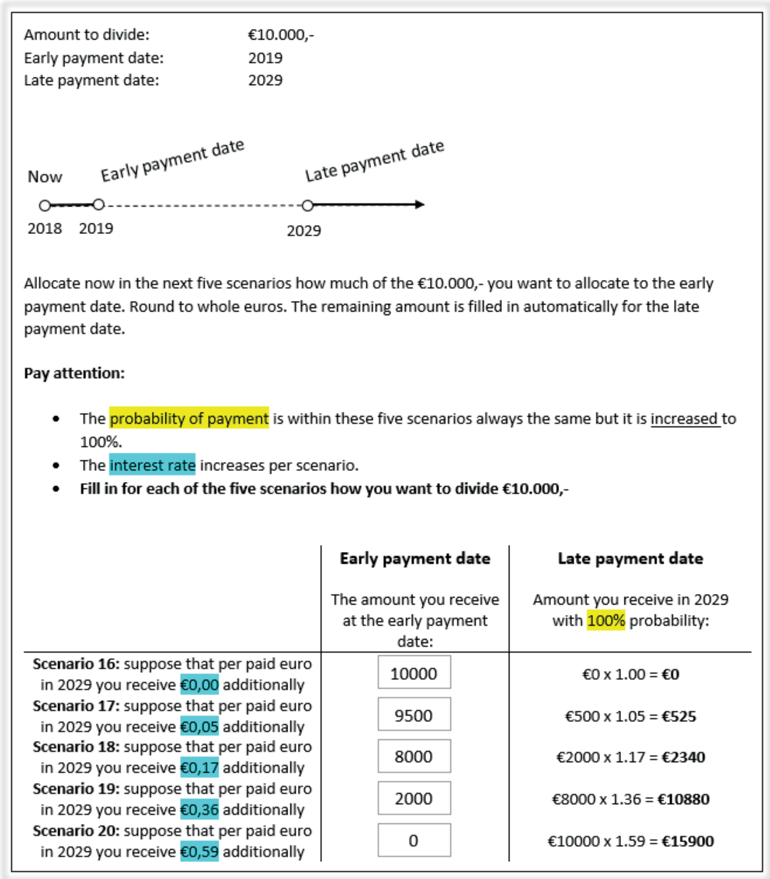

Amount to divide: €10.000,-
Early payment date: 2019
Late payment date: 2029

Allocate now in the next five scenarios how much of the €10.000,- you want to allocate to the early payment date. Round to whole euros. The remaining amount is filled in automatically for the late payment date.

Pay attention:

- The probability of payment is within these five scenarios always the same but it is increased to 100%.
- The interest rate increases per scenario.
- **Fill in for each of the five scenarios how you want to divide €10.000,-**

	Early payment date The amount you receive at the early payment date:	**Late payment date** Amount you receive in 2029 with 100% probability:
Scenario 16: suppose that per paid euro in 2029 you receive €0,00 additionally	10000	€0 x 1.00 = **€0**
Scenario 17: suppose that per paid euro in 2029 you receive €0,05 additionally	9500	€500 x 1.05 = **€525**
Scenario 18: suppose that per paid euro in 2029 you receive €0,17 additionally	8000	€2000 x 1.17 = **€2340**
Scenario 19: suppose that per paid euro in 2029 you receive €0,36 additionally	2000	€8000 x 1.36 = **€10880**
Scenario 20: suppose that per paid euro in 2029 you receive €0,59 additionally	0	€10000 x 1.59 = **€15900**

Figure 3.1: Decision screen Convex Time Budgets. In this decision screen, the subject allocates $m = 10.000$ Euro between an early payment with front-end delay $t = 1$ (2019) and a late payment with back-end delay $k = 10$ years (2029). The late payment is with a probability p_{t+k} of 100%. The gross interest rate, or gross investment return, $1 + r$ over k years in the 5 scenarios varies from 1.00 to 1.59. The allocated amounts are for illustration purposes only, the default values were blanks (subjects must actively allocate). The text is translated from Dutch to English.

structure: 2018 is the experimental date, the early payment is received in 2019 and the late payment is received in 2029 after an additional delay of ten years. Subjects are told to divide the amount of €10,000 between the early payment and late payment. Probabilities of late payment and interest rates were highlighted by yellow and blue, respectively. In this particular decision screen, the likelihood that the late payment is paid equals $p_{t+k} = 100\%$ and there are five budget decisions presented in order of increasing gross interest rates from 1.00 to 1.59. Subjects are faced with a total of four such decision screens, corresponding to the four probability decision sets. After the twenty CTB decisions, subjects complete the present-bias task.

We fielded our survey at the pension fund ABP in The Netherlands.[9] The pension fund has a panel for experimental research and communicates via email. The invitations for our experiment and the experiment itself were simultaneously conducted in the period 13 August 2018 till 17 September 2018. Individuals could join the experiment by clicking on a link in the email.

3.1.3 Annuity choices

The Dutch pension system has two main pillars: (i) a publicly financed pay-as-you-go scheme and (ii) a mandatory occupational pension scheme. The first pillar, or General Old-Age Pensions Act, aims at providing a minimum retirement income, and is funded from tax revenues. Individuals receive first-pillar benefits when they reach the statutory retirement age, which is 66 years in 2018. The majority of the active participants with an uninterrupted working career qualify for a state pension benefit close to the maximum yearly amount of €14,000 for single individuals and roughly €18,000 for couples. First pillar benefits are indexed based on price inflation, and always paid out as life-long annuities.

The second pillar is an employer-based occupational pension scheme that features collectivity, mandatory participation, and is not for profit. Pension funds operate on the basis of capital funding: an employee, together with her employer, accrues pension entitlements from the contributions paid in and the return realized by the pension fund over the years through the collective investment of these contributions. The main goal is to maintain the pre-retirement living standards,

[9] ABP is the largest pension fund in The Netherlands. The abbreviation translates to National Civil Pension Fund, and arranges the pensions for mainly civil servants.

together with the benefits from the first pillar. We study the freedom of choice that retirees have in the second pillar through their annuity choices in the occupational pension scheme.

The individual's annuity decision has three key components, and the choice can only be made once. The individual must make a choice regarding (i) the date of retirement, (ii) a bridging pension or not, and (iii) the payment profile.[10] Regarding key decision (i), the individual must decide when to retire, e.g., at the statutory retirement age (i.e., the default) or earlier.[11] Retiring earlier than the statutory retirement age decreases overall monthly life-long benefits at an actuarially fair rate, because the individual starts to withdraw her pension wealth earlier than the statutory retirement age.

Regarding key decision (ii), the pension fund offers the beneficiary the option to receive a so-called bridging pension (i.e., the default option) until the statutory retirement age is reached, i.e., the moment when she receives first-pillar pension benefits. The goal of a bridging pension, only available when retiring early, is to ensure a flat payment stream of benefits before and after the statutory retirement age. When choosing a bridging pension on top of early retirement, the individual depletes her second-pillar pension wealth faster compared to no bridging pension, so that overall monthly life-long benefits are reduced at an actuarially fair rate.

Regarding key decision (iii), the fund offers the possibility to increase benefits for 5 to 10 years at any point during the retirement phase.[12] The idea is that individuals can construct a high-low stream of payments to tailor pension benefits to the individual's needs. A high-low construction frontloads the pension benefits, like a lump sum, and it could be used for paying off a mortgage or travel plans.[13] Of course, a high-low construction depletes second-pillar pension wealth faster than constant annuity payments and, thus, reduces future monthly life-long benefits at

[10] There is also the possibility to exchange partner pension for old-age pension, but we exclude this in our analysis as we study individual decisions.

[11] Individuals can also retire later than the statutory retirement age, but almost no individual does so.

[12] Legally, pension benefits can be increased (or decreased) until the age of 78.

[13] The pension fund also offers the possibility to construct a low-high payment stream that backloads future pension benefits, for example to facilitate later (unexpected) health costs. However, few individuals choose a low-high payment structure in The Netherlands.

an actuarially fair rate. The legal condition states that the lower benefits must at least equal 75% of the higher benefits. The default is no frontloading of pension payments.

So, the retiree constructs her own annuity based on the three choices. At least 6 months before the statutory retirement age the individual receives information from the fund about her annuitization decision (unless she made a choice already). Essentially, the pension fund offers the possibility to withdraw the accumulated capital either as a *flat* life-long annuity or as a *front-loaded* life-long annuity. We label an annuity as front-loaded if the retiree within one year after her pension age has at least 1 year of after-tax pension benefits that are 5% higher than payments in the future years, taking state pension benefits into account.[14] We label an annuity as flat otherwise.

The majority of individuals in our sample that choose a front-loaded annuity construct the annuity such that high payments start within 1 year after retirement with an average duration of 3 years and low payments equalling the legal minimum 75% of the high payments. Examples of a front-loaded annuity include early retirement with bridging pension and high-low payments, or retirement at the statutory retirement age with high-low payments. While examples of a flat annuity include retirement at the statutory retirement age (default), or early retirement with bridging pension and constant payment afterwards (i.e., not front-loading payments). If individuals forego to make an active annuitization decision, then the fund offers by default a flat life-long annuity starting at the statutory retirement age.

3.1.4 Sample

We select pension fund participants between the ages of 50 years and 70 years, as these cohorts are most likely concerned with their pension choices. We have chosen 70 years as an upper cutoff point to minimize potential issues regarding effects of mortality risk. Namely, the CTB decision horizon is 10 years, which

[14] We use a threshold of 5%, because due to administrative reasons flat annuity payments could fluctuate within this bandwidth.

yields a late payment at the age of 80 at most, well within the average Dutch life expectancy of 85 years at the time of the experiment. In total, 3611 pension fund participants clicked on the link in the email to participate in our survey. We exclude 11 retirees with a pension date that is later than the statutory retirement age and we exclude 48 retirees with a back-loaded annuity, as we want to specifically study front-loading behavior, which yields 3552 individuals. We need to drop individuals that did not complete the present-bias task, CTB experiment, or filled the same early payment amount in every CTB scenario. For these individuals we cannot identify the preference parameters without additional assumptions. Finally, we drop individuals that entered an amount next year $c_{t+\tau}$ that is lower than receiving €800 now, or entered an amount next year $c_{t+\tau}$ that is lower than receiving an amount in 10 years, as these answers imply negative interest rates.[15] Overall, this yields a final sample of 1062 individuals.

We are aware that we lose some observations in the sample selection. It is known that response rates for surveys in the pension industry are rather low, since people rarely interact with their pension funds (Bauer et al., 2021; Debets et al., 2021). However, our sample still contains enough heterogeneity. Eventually, it is this heterogeneity that matters for studying the relation between preferences and real-life choices. Moreover, our sample is representative for the pension fund based on several important characteristics.

Table 3.1 compares our sample of subjects with the pension fund's population from 2018, restricted to the ages of 50 and 70.[16] In our sample, 705 respondents are so-called active participants. These active participants actively accrue pension rights at the pension fund through their employer. 357 respondents are retirees, who receive pension benefits from the pension fund. Panel A shows that the male-to-female and active-to-retiree ratios are nearly equal between the fund and our sample. Because we study actual pension choice behavior of retirees later, we present additional summary statistics on the retired population in Panel B. The median age of the retired subjects in our sample is almost similar to the

[15]To check the understanding of participants, subjects also answer how much money they would like to receive in 10 years which makes them indifferent with receiving €800 now.

[16]We focus on old-age pension for the pension fund's retirees.

Table 3.1: Summary statistics. This table presents summary statistics for our sample and the pension fund. Panel A contains all subjects, i.e., active participants and retirees. Panel B contains only retirees. *Male* and *Retired* are dummy variables. *Age* is in years and *Income* is the annual before-tax income in Euros, which includes all employer-related second pillar pension benefits received from the pension fund including state pension benefits. Standard deviation between parentheses.

Panel A: Active participants and retirees				
		Pension fund	Sample mean	N
Male		0.567	0.570 (0.500)	1062
Retired		0.384	0.340 (0.470)	1062
Panel B: Retirees				
		Pension fund	Sample median	N
Age (years)	Male	67.31	67.15 (1.84)	240
	Female	67.14	67.27 (2.36)	117
	Total	67.24	67.20 (2.03)	357
Income (€)	Male	22,670	28,358 (16,587)	217
	Female	16,637	16,739 (12,225)	104
	Total	20,102	23,317 (16,199)	321

pension fund's value of 67. The male respondents in our sample are more likely to have a somewhat higher income, but the female income is nearly identical to the pension fund's value.[17] The median time taken to complete the survey, including the questions on personal and financial information, is 20 minutes. The participants understood the CTB experiment generally well, as the median rating for the difficulty of the CTB experiment is 3 (i.e., "not easy, but also not difficult") on a 5-point Likert scale.[18]

[17] Table 3.7 in the Online Appendix provides additional summary statistics on demographic, financial and pension variables. Tables 3.8 and 3.9 in the Online Appendix describe the definitions of all variables used in our analysis.

[18] The question regarding the difficulty of the CTB experiment follows immediately after the

3.2 Preferences

In this section, we firstly present the aggregate choice behavior in the CTB and present-bias task. Then, we discuss the simultaneous estimation of individual risk and time preferences. Finally, we show the estimation results for the preference parameters.

3.2.1 Descriptive analysis

First, we describe the choice behavior in the CTB and, then, in the present-bias task. Figure 3.2 summarizes aggregate choice behavior in the CTB for the whole population (i.e., actives and retirees combined). We plot the median allocated Euros chosen at the early payment c_t against the gross interest rate $(1 + r)$ for each late payment probability p_{t+k}. The amount of Euros allocated to the early payment declines monotonically with the interest rate, indicating that people wait for the late payment when interest rates are higher. Additionally, as expected, the amount of earlier Euros increases when the late payment probability is lower. So, we observe in Figure 3.2 that individuals respond to changing interest rates and payment probabilities in a predicted way.

Figure 3.3 summarizes aggregate choice behavior in the present-bias task, for actives and retirees separately. The subjects' answers are winsorised at a 5% level from the bottom and the top of the distribution. The dashed red bars depict retirees, while the solid gray bars depict active participants. The upper panel reports the allocated amount $c_{t+\tau}$ in Euros that makes subjects indifferent between receiving €800 now or receiving $c_{t+\tau}$ next year. The fraction of retirees that allocates lower amounts of wealth $c_{t+\tau}$ to next year (e.g., between €800 and €1000) to make them indifferent with €800 directly is larger than for actives. This is preliminary evidence that actives are more impatient than retirees in the short trun.

The bottom panel reports the implied annual interest rates based on the allocated amounts $c_{t+\tau}$. A high interest rate indicates that the subjects discount

CTB experiment. See Table 3.7 in the Online Appendix for additional summary statistics.

Figure 3.2: Choice behavior: Convex Time Budgets. Median allocated Euros at early payment c_t against the gross interest rate $1 + r$ per payout probability p in the Convex Time Budgets.

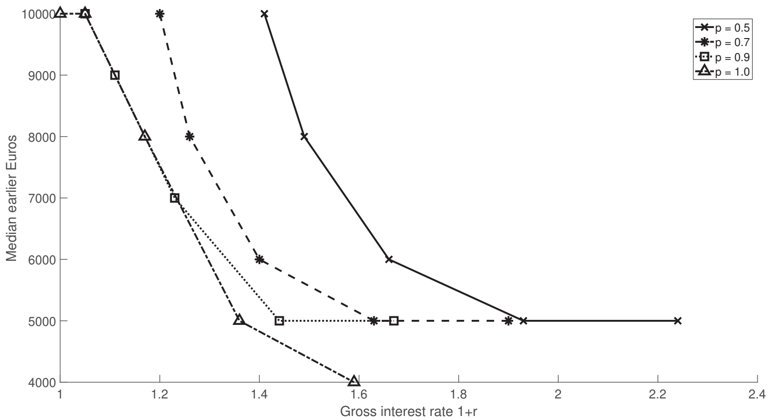

consumption next year heavily. For about 70% (75%) of the actives (retirees), the annual interest rates from the one-year present-bias task are larger than 10%. This is higher than the annual interest rates in the ten-year CTB task, which vary from 0 to 8.40 percent per year.[19] Thus, in line with Thaler (1981), we find that discount rates elicited in the short run are higher than discount rates elicited in the long run. This observation provides evidence for time inconsistency and indicates the possibility of present bias for pension fund participants. More specifically, in line with the upper panel, the bottom panel shows that active participants are more prone to present-biased behavior than retirees as actives discount consumption next year more strongly. The effect is visible between the lower interest rates of 0% to 20%, where the fraction of retirees is higher, while for interest rates larger than 20% the fraction of active participants is higher.

[19]Figure 3.2 shows that the CTB design per se is not an issue, because the median amount allocated to the earlier payment is about €5000 which is not a corner solution.

Figure 3.3: Choice behavior: Present-bias task. Distribution of allocated Euros $c_{t+\tau}$ in the present-bias task, together with the implied annual interest rate. Implied annual interest rate calculated as $(c_{t+\tau}/800 - 1) \times 100$.

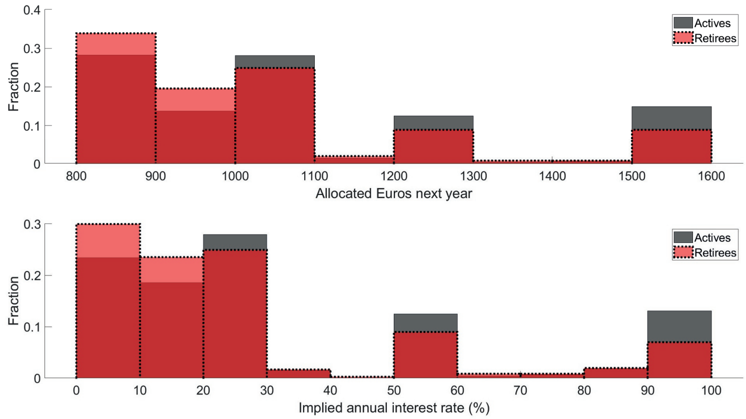

3.2.2 Simultaneous estimation of risk and time preferences

To estimate risk and time preferences, we identify the experimental allocated payments as solutions to standard intertemporal optimization problems. These solutions are supposed to be functions of our parameters of interest (present bias, discounting, and utility curvature), and experimentally varied parameters (interest rates, delay lengths and payment probabilities). Given assumptions on the functional form of utility and the nature of discounting, our experimental tasks provide a natural context to jointly estimate individual preferences.

In the CTB, subjects choose an amount c_t, available at time t, and an amount c_{t+k}, available after a delay of k periods, continuously along a convex budget set

$$c_t + \frac{c_{t+k}}{1+r} = m, \tag{3.1}$$

where $(1 + r)$ is the experimental gross interest rate and m is the experimental budget. Money allocated to the early payment has a value of c_t, while money allocated to the late payment has a present value of $c_{t+k}/(1 + r)$. c_{t+k}/c_t defines the gross interest rate $1+r$ over k years, so $(1+r)^{1/k}-1$ gives the standardized annual interest rate r. Multiplication by the payment probability p_{t+k} defines the risk-adjusted interest rates.

Using the quasi-hyperbolic $\beta - \delta$ model of intertemporal decision making (Phelps and Pollak, 1968; Laibson, 1997), the subject maximizes discounted expected utility over the early payment c_t and late payment c_{t+k}

$$\max_{c_t, c_{t+k}} \delta^t U(c_t + w_t) + \beta\delta^{t+k} \left[p_{t+k} U(c_{t+k} + w_{t+k}) + (1 - p_{t+k}) U(w_{t+k})\right], \qquad (3.2)$$

where δ is the one period (i.e., annual) discount factor and β is the present-bias factor. The quasi-hyperbolic form captures the notion of time-inconsistent behavior, since $\beta < 1$ indicates present bias. Moreover, it nests exponential discounting (i.e. standard time-consistent behavior, Samuelson, 1937) when $\beta = 1$. Early payments are certain, while late payments can be uncertain such that with probability $1 - p_{t+k}$ no delayed payment is received. The terms w_t and w_{t+k} could be interpreted as background consumption or income (see, e.g., Andersen, Harrison, Lau, et al., 2008a).

In line with former CTB experiments (Andreoni and Sprenger, 2012a; Balakrishnan et al., 2020), we posit the agent has a time separable Constant Relative Risk Aversion (CRRA) utility function of the form

$$U(x) = \frac{1}{\alpha} x^{\alpha}, \qquad (3.3)$$

where $\alpha < 1$ is the curvature of the CRRA utility function, giving rise to a preference to smooth payoffs. Because we assume a CRRA utility function, we assume that utility for risk (i.e., risk aversion) also represents instantaneous utility for time (i.e., intertemporal substitution), in line with Andersen, Harrison, Lau, et al. (2008a). Although these are conceptually distinct preferences, in our setting were both risk and time are present it is common to assume that utility for risk is one and the same as instantaneous utility for time (Cheung, 2020).[20]

[20] It is not the goal of the current paper to separate parameter estimates for utility for time and risk, see Cheung (2020) for this.

Solving the subject's standard intertemporal maximization problem (3.2) subject to the budget constraint (3.1) yields the first-order condition

$$\left(\frac{c_t + w_t}{c_{t+k} + w_{t+k}}\right)^{\alpha-1} = \begin{cases} \beta\delta^k p_{t+k}(1+r) & \text{if } t \in [0,1] \\ \delta^k p_{t+k}(1+r) & \text{if } t \geq 1 \end{cases} \tag{3.4}$$

Clearly, the experimental allocations depend on the parameters of interest (present bias, discounting, and curvature), and the experimentally varied parameters (interest rates, delay length, and payment probabilities). The present, i.e., $t \in [0,1]$, runs from the experimental date at the end of 2018 (i.e., $t = 0$) till next year at the end of 2019 (i.e., $t = 1$), and afterwards the future starts.

Please note that for some CTB scenarios with uncertain late payment probabilities, decision sets 1 till 3 in Table 3.6 in the Appendix, the risk-adjusted interest rates are negative and the expected payment values are not always constant between the decision sets. Additionally, the present-bias task involves only payments with certainty.

Thus, taking the natural logarithm of (3.4), we find

$$\begin{aligned} \ln\left(\frac{c_t + w_t}{c_{t+k} + w_{t+k}}\right) &= \frac{1}{\alpha - 1}\left(\log(\beta) \cdot \mathbb{1}_{t\in[0,1]} + \log(\delta) \cdot k\right) \\ &+ \frac{1}{\alpha - 1}\left(\log(p_{t+k}) + \log(1+r)\right). \end{aligned} \tag{3.5}$$

The variation in payment probabilities and interest rates identifies the utility curvature parameter α. Stated differently, the utility curvature parameter is identified by the smoothness of payoffs over time (i.e., instantaneous utility for time) and for different states of the world (i.e., utility for risk). Because the front-end delay t and back-end delay k are fixed in our CTB design, we cannot separate present bias from long-term patience using only CTB scenarios. Therefore, we use the present-bias task to separate the present-bias factor β from the discount factor δ, while simultaneously correcting for potential utility curvature α. To identify present bias, we assume that the payment c_t is received during the present, the payment $c_{t+\tau}$ marks the end of the present, and c_{t+k} is received during the future.

The subject during the present-bias task solves

$$U(800 + w_0) = \beta\delta U(c_{t+\tau} + w_{t+\tau}). \tag{3.6}$$

In words, the subject considers a trade-off between a direct early payment of €800 at the experimental date of end 2018 (i.e., $t = 0$), or a discounted payment $c_{t+\tau}$ one year later at the end of 2019 (i.e., $t = 1$). Solving explicitly for the present-bias factor yields

$$\beta = \frac{1}{\delta}\left(\frac{800 + w_0}{c_{t+\tau} + w_{t+\tau}}\right)^{\alpha}. \tag{3.7}$$

Clearly, the present-bias factor β is identified by the payment $c_{t+\tau}$, is corrected for the utility curvature α, and is separated from the long-term discount factor δ. Note that a high discount factor induces a lower present-bias factor.

Substituting the expression for β in (3.5), we find the following equation

$$\begin{aligned} \ln\left(\frac{c_t + w_t}{c_{t+k} + w_{t+k}}\right) = \frac{\alpha}{\alpha - 1}\log\left(\frac{800 + w_0}{c_{t+\tau} + w_{t+\tau}}\right) + \frac{1}{\alpha - 1}\log(\delta)\cdot(k-1) \\ + \frac{1}{\alpha - 1}\left(\log(p_{t+k} + \log(1+r)\right). \end{aligned} \tag{3.8}$$

Given an additive error structure and assumptions on background consumption, such a linear equation is easily estimated with parameter estimates for β, δ, α obtained via nonlinear combinations of coefficient estimates. We estimate the parameters $\hat{\beta}, \hat{\delta}, \hat{\alpha}$ by two-limit tobit and, as robustness check, by OLS.

To limit the number of estimated parameters and facilitate comparison with previous literature, our results use a predetermined background income level. In line with Andreoni and Sprenger (2012a) and Potters et al. (2016), we set $w_0 = w_t = w_{t+\tau} = w_{t+k} = 0.01$. That is, the budget offered in the experiment is the only source of income that participants consider when making their early and late payment allocations. Essentially, we assume that the experimental budget compromises the state and second pillar pensions. This is consistent with the real-life annuity choice, because the actual annuitization possibilities on the pension fund's personal web page that are shown visually to the individuals treat the state and second pillar pensions together. We winsorize the estimated preference parameters at the bottom and top of the distribution for a 5% level.[21]

3.2.3 Estimated preference parameters

Table 3.2 presents our estimation results for the present-bias factor β, the discount factor δ, and the CRRA curvature parameter α. For each individual, we estimate the preference parameters according to equations (3.7) and (3.8) and, then, we compute summary statistics for the population. We show estimation result for our complete sample, and for actives and retirees separately. We make three observations.

First, echoing the results from our descriptive analysis, we find evidence for present bias since $\beta < 1$. We estimate the median and mean present-bias factor β respectively at 0.836 and 0.819. Active pension fund participants have a lower present-bias factor than retirees, such that actives are more subject to present bias. The difference between the median present-bias factors of retirees and actives is about 0.08. Roughly 14% of our sample is future biased (i.e., $\beta > 1$). This is in line with the observation of future-biased participants in the sample of Andersen, Harrison, Lau, et al. (2014), and similar to 19% of the subjects being future biased in Bleichrodt, Gao, et al. (2016).

A common finding in the literature is a (substantial) present bias, see for example Frederick et al. (2002), Tanaka et al. (2010), and Laibson, Maxted, et al. (2020). Our estimated present-bias value is similar to those estimated by other researchers. Balakrishnan et al. (2020) also use the CTB design, with also a monetary experiment, and they estimate present-bias factors between 0.902 to 0.924. Other papers have used nonmonetary experiments such as job search for estimating discounting behavior. For example, Paserman (2008) estimates a present-bias factor of 0.894 for high income workers. DellaVigna and Paserman (2005) often find a present-bias factor near 0.9. Using experiments on real effort tasks, Augen-

[21] Our results below are robust to a different winsorization level, for example a level of 1%. See Table 3.13 in the Online Appendix.

Table 3.2: Present bias, annual discounting, and curvature parameter estimates. Two-limit tobit maximum likelihood and Ordinary Least Squares (OLS) estimates for present-bias factor β, discount factor δ, and CRRA utility curvature α.

	Median	Mean	Standard Deviation	25^{th} Percentile	75^{th} Percentile	N
Tobit: All						
Present-bias factor $\hat{\beta}$	0.836	0.819	0.184	0.695	0.953	1062
Discount factor $\hat{\delta}$	0.989	1.004	0.092	0.962	1.039	1062
Annual discount rate	0.011	0.004	0.089	-0.037	0.040	1062
CRRA curvature $\hat{\alpha}$	0.965	0.915	0.252	0.905	0.987	1062
Tobit: Actives						
Present-bias factor $\hat{\beta}$	0.820	0.802	0.185	0.672	0.938	705
Discount factor $\hat{\delta}$	0.991	1.006	0.091	0.963	1.039	705
Annual discount rate	0.009	0.001	0.086	-0.038	0.039	705
CRRA curvature $\hat{\alpha}$	0.965	0.923	0.248	0.911	0.988	705
Tobit: Retirees						
Present-bias factor $\hat{\beta}$	0.902	0.853	0.178	0.767	0.971	357
Discount factor $\hat{\delta}$	0.986	0.998	0.095	0.961	1.039	357
Annual discount rate	0.014	0.010	0.094	-0.037	0.041	357
CRRA curvature $\hat{\alpha}$	0.964	0.899	0.260	0.876	0.984	357
OLS: All						
Present-bias factor $\hat{\beta}$	0.860	0.835	0.194	0.714	0.963	1062
Discount factor $\hat{\delta}$	0.989	0.994	0.105	0.949	1.039	1062
Annual discount rate	0.011	0.018	0.111	-0.037	0.054	1062
CRRA curvature $\hat{\alpha}$	0.936	0.887	0.296	0.863	0.963	1062

blick, Niederle, et al. (2015) and Augenblick and Rabin (2019) find a present-bias factor ranging from 0.83 to 0.89.

Second, the estimated annual discount factor δ has a median value of 0.989 and 50% of the sample has a discount factor between 0.962 and 1.039. The annual discount factor translates to an annual discount rate of 1.1%.[22] About 25% of our sample has long-term negative annual discount rates, such that these participants are extremely patient as they are willing to pay, rather than generate interest, to receive a payment in the future. Differences between active participants and retirees are negligible. Our median estimated annual discount rate is in line with

[22] The annual discount rate follows from $(1/\delta) - 1$, since the discount factor is measured in years.

(long-term) market interest rates and lower than most previous studies. Estimates of annual discount rates over hundred percent are not uncommon, as shown by the overview article of Frederick et al. (2002). Cheung (2020) estimates an annual discount rate of 62.6%, when controlling for CRRA curvature. The CTB design of Andreoni and Sprenger (2012a) corrects for CRRA curvature and present bias, but they still estimate an annual discount rate of 27.5%. A close estimate is that of Andersen, Harrison, Lau, et al. (2014), who report an annual discount rate of 7.3% in the quasi-hyperbolic model, while controlling for classical risk aversion over states of the world.

A potential reason for our lower annual discount rate is the magnitude of the experimental budget and the long-term decision horizon. Thaler (1981) already shows that discount rates drop sharply as the size of wealth increases, which is known as the magnitude effect. Additionally, he reports that discount rates drop sharply as the length of time increase. We confirm both findings in our large non-student sample while controlling for risk preferences. The experimental budget of €10,000 and a decision horizon of 10 years are both (much) larger than many of the previous studies. Horizons are frequently used up to several weeks (Augenblick, Niederle, et al., 2015), 3 months (Tanaka et al., 2010), 6 months (Andersen, Harrison, M.Lauc, et al., 2010), 1 year (Dohmen, Falk, Huffman, and Sunde, 2010; Andersen, Harrison, Lau, et al., 2014), 2 years (Goda et al., 2015) and 3 years (Harrison, Lau, et al., 2002). A paper that comes close to ours in terms of large stakes and long decision horizons is Potters et al. (2016). They use an experimental budget of €1,000 with a decision horizon up to retirement age and report an annual discount rate of 1%. [23]

Our third finding is that the median CRRA utility curvature α is 0.965, implying that subjects have concave utility because $\alpha < 1$. Individuals have a preference to smooth payoffs over time. A minority has a convex utility function, which implies that these individuals prefer less smoothed payoffs over time.[24]. Curvature

[23]Another reason might be that not all previous studies correct for utility curvature when estimating time preferences, such that discount rates might be upward biased (Andreoni and Sprenger, 2012a). However, based on high income workers, Paserman (2008) estimates a yearly discount factor of 0.9989 not corrected for curvature.

[24]CRRA curvature comes much closer to linear utility than estimates of classical risk aversion,

estimates for active participants and retirees are identical at the median. Our estimated utility curvature is in line with previous CRRA curvature estimates (Andreoni and Sprenger, 2012a; Potters et al., 2016).

Notice that OLS and tobit parameter estimates are very similar for all preferences parameters. This indicates that censored corner solutions do not seem to be a major issue. Indeed, the percentage of responses that are at corners equals 46% and the number of subjects that made zero interior allocations is only 8%. Compared to the literature, these percentages are low. Andreoni and Sprenger (2012a) find that "roughly 70 percent of responses are at corners, but only 36 of 97 subjects [37%] made zero interior allocations."

Figure 3.5 in the Appendix visualizes the distributions of the present-bias factor, the annual discount factor and CRRA curvature. Clearly, there is individual heterogeneity in risk and time preferences. Due to the winsorization we observe a higher fraction of subjects at the boundaries of the distributions.

Regarding concerns on mortality risk, Table 3.11 in the Online Appendix shows the results of regressing the individually estimated preference parameters on personal characteristics. The main takeaway is that risk and time preferences are not associated with self-reported life expectancy in a linear way, similar to the findings of Chao et al. (2009). Thus, beliefs regarding one's life expectancy are not driving our preference estimates linearly, which corroborates our earlier point to only invite participants of at most 70 years to avoid mortality effects on preferences. Note that the number of observations is rather low for this variable, as participants are not very willing to report this information to the pension fund. In addition, we find that the present-bias factor correlates positively with male, age, and savings. The discount factor correlates negatively with male and lower savings. The curvature parameter does not correlate with any observed individual characteristics.

Table 3.12 in the Online Appendix shows the preference parameters when individual annual after-tax income is used as background consumption w, assuming that income w remains constant from the experimental date to future date $t + k$.

---

as employed by Holt and Laury (2002) and Eckel and Grossman (2008)

The discount factor δ remains similar to the estimations without background income, but the present-bias factor β is somewhat higher and individuals act more risk averse.

3.3 Real-life choices

This section uses administrative micro data from the pension fund to study actual annuitization decisions of retirees ($N = 357$) in relation to their individually estimated preferences. The combination of the administrative data on actual decision making with the experimental survey is a unique feature of our research. We first study how predictive preferences are for financial decision making by using a discounted expected utility model.[25] Secondly, we quantify the welfare effects of freedom of choice in annuitization decisions by studying flexibility in the payout phase of pension schemes.

3.3.1 Predictivity of annuity choices by preferences

This section studies how well risk and time preferences explain individual annuitization decisions. We use a simple discounted expected utility model in which we include the individually estimated preferences.

Utility of annuity choices

To determine the utility of annuity choices, we follow 3 steps. First, we compute the utility value of the actual observed real-life annuitization decision at retirement. Secondly, we compute the utility value of the annuity that has not been chosen. This is the foregone alternative or the counterfactual. For example, if a retiree chooses a front-loaded annuity, then the foregone alternative is a flat annuity. Finally, we determine the expected annuity choice by comparing whether the actual or alternative annuity yields the highest utility.

[25]Since annuity payments are no guarantee, we use the term *expected* as well.

If the actual chosen annuity yields higher total utility during the retirement phase than the foregone alternative annuity, then the individual made a choice in line with the model and the measured risk and time preferences. The discounted expected utility model, using the individual preferences as inputs, is able to explain actual choice behavior since the observed annuity choice coincides with the expected annuity choice. If the actual chosen annuity yields lower total utility during the retirement phase than the foregone alternative annuity, then the choice of the individual deviates from the model and the measured preferences. In this case, the discounted expected utility model suffers from a prediction error since the actual annuity choice differs from the expected annuity choice. If the difference in utility levels between the actual and expected annuity choices is large, then the prediction error is larger, and individually estimated preferences and the model have more difficulty with explaining actual choice behavior. If the difference in utility levels between the actual and expected annuity choices is small (i.e., small prediction error), then measured preferences are not much in favor of one of the annuities.

To determine the total utility of the actual chosen annuity during the retirement phase, we compute the discounted expected utility of the annuity payments at retirement $t = 0$ by

$$U = \sum_{t=0}^{T} p(t)\phi(t; \hat{\beta}, \hat{\delta})u(x_t; \hat{\alpha}). \tag{3.9}$$

Thus, the annuity's utility value depends on the individually estimated risk and time preferences. $u(x_t; \hat{\alpha})$ is the CRRA utility function, with estimated curvature parameter $\hat{\alpha}$, from the after-tax annuity payment x_t for $t = 0, \ldots, T$ with T the maximum time of death. $\phi(t; \hat{\beta}, \hat{\delta})$ is the individually estimated quasi-hyperbolic discount structure. Quasi-hyperbolic discounting requires a distinction between the present and the future. In line with our experimental approach and the observed front-loaded annuity characteristics, we set the present-bias interval equal to one year. So, one year after retirement consumption is valued less by an amount equal to the present-bias factor β. Since the moment of death is unknown, we include fund-specific survival probabilities $p(t)$ at each time t, which are cohort and gender specific. As the number of observations regarding self-reported subjective survival probabilities is very low, we favor to use the fund-specific survival probabilities instead. Note that in these utility calculations we correct for life expectancy, whereas the CTB experiment on purpose does not (as argued before), as the actual cash flows on average have a longer horizon than the ten years in the experiment.

To compute the utility value of the annuity that has not been chosen, we need to construct the payment scheme of the unobserved foregone alternative. To construct the payment scheme of the counterfactual, we need the individual's pension wealth at retirement. We find the individual's pension wealth by computing the present value of all future payments of the actual chosen annuity. In line with the actual fund's present value calculations, we use (i) the fund specific survival probabilities $p(t)$ for every date, cohort and gender, and (ii) an actuarial interest rate of 1.39% to discount future payments, as set by the Dutch Central Bank in 2018 based on the yield curve.[26] Ultimately, we convert the individual's pension wealth into the unchosen foregone annuity.

If the retiree actually chooses a front-loaded annuity, then for the counterfactual we convert pension wealth into the default flat annuity. We assume that the flat annuity starts at the observed date of retirement. For example, in case of early retirement, the retiree still retires early, but receives a flat life-long annuity rather than the chosen front-loaded life-long annuity. If the retiree actually chooses a flat annuity, then for the counterfactual we convert pension wealth into a front-loaded annuity. Again, we assume that the front-loaded payments start at the observed date of retirement. In line with our earlier observations regarding pension choices, we assume that front-loaded annuities start with high payments at retirement for a duration of 3 years, and low payments are equal to the legal minimum of 75% of the high payments.

[26] We do not use the self-reported life-expectancies as the number of observations would become too low.

Table 3.3: Annuity choices, preferences, and welfare effects This table presents the median present-bias factor $\hat{\beta}$, the median long-run discount factor $\hat{\delta}$, and the median curvature parameter $\hat{\alpha}$ for the actual and expected annuity choices according to the observed preferences. Between parentheses the mean time preference parameter values. Potential and realized welfare effects, associated with actual and expected annuity choices, are shown in the last two columns.

	Observed preferences			Welfare effects	
	$\hat{\beta}$	$\hat{\delta}$	$\hat{\alpha}$	Potential	Realized
1. Actual **flat**, expected **front-loaded**	0.94	0.96	0.97	+	0
	(0.91)	(0.94)	(0.89)		
2. Actual **front-loaded**, expected **front-loaded**	0.93	0.96	0.95	+	+
	(0.91)	(0.95)	(0.88)		
3. Actual **front-loaded**, expected **flat**	0.87	1.01	0.96		-
	(0.84)	(1.02)	(0.88)		
4. Actual **flat**, expected **flat**	0.82	1.03	0.96		
	(0.79)	(1.05)	(0.92)		

Observed and expected annuity choices

We now study the relation between preferences, observed annuity choices and expected annuity choices. We distinguish between four groups, because the observed annuity choice can be in line with the expected annuity choice, or not: "actual flat, expected front-loaded", "actual front-loaded, expected front-loaded", "actual front-loaded, expected flat", and "actual flat, expected flat".

Table 3.3 shows the relation between annuity choices and median measured preferences. First, we discuss the relation between preferences and expected annuity choices, then we include the actual annuity choices. We observe that an expected front-loaded annuity is accompanied by a lower median discount factor δ and a higher median curvature parameter α relative to an expected flat annuity. Namely, a lower discount factor δ implies stronger long-run impatience, while a higher curvature parameter α implies a preference for a less smoothed consumption path. A front-loaded annuity therefore fits individuals with stronger impatience, and a front-loaded annuity is also less smooth than a flat annuity. On the other hand, an expected flat annuity is accompanied by a median discount factor close to one. Individuals that are expected to choose a flat annuity are individuals with a preference for smooth consumption paths and they are more patient than individuals that prefer a front-loaded annuity.

From the perspective of actual annuity choices, we see that the group "actual front-loaded, expected flat" has a relatively low median present-bias factor. These individuals are relatively present biased and, thus, tempted to actually choose front-loaded annuity payments. The group "actual front-loaded, expected front-loaded" has similar preferences to the group "actual flat, expected front-loaded". Note that the group "actual flat, expected flat" has a low median present-bias factor as well, however the relatively high median long-term discount factor and the relatively low median curvature parameter are dominating for this group, which pulls individuals towards a flat annuity. The discount factor and curvature parameter matter for each year in the utility calculations, while the present-bias factor only matters the first few years. Because the median discount factor is relatively high, the median present-bias factor is lower due to the experimental answers and the estimation methodology, as shown in the expression for β in equation (3.7). The last two columns of Table 3.3 are discussed in the section on welfare effects.

Table 3.4, Panel A, shows the number of individuals for the actual observed annuity choices and the expected annuity choices. The total sample of retirees is $N = 357$. 248 retirees actually choose a flat annuity, while 109 retirees actually choose a front-loaded annuity. Based on the individually estimated preferences, the discounted expected utility model expects that 195 retirees choose a flat annuity and 162 retirees choose a front-loaded annuity. Thus, we observe that too many individuals actually choose a flat annuity compared to their expected choice based on individually estimated preferences. A potential reason for this difference is that the pension fund offers the flat annuity as a default.

Using the numbers on the diagonal, we observe that for 52% of the retirees (i.e., 184 retirees out of 357, of which 135 choosing flat and 49 choosing front-loaded) the individually estimated preferences explain actual annuity choices according to the expected annuity choices based on the discounted expected utility model. However, asking our simplified model to explain annuity choices with perfect utility indifference might be too strict. So, we study the utility differences between actual and expected annuity choices in the case annuities might be perceived as observationally equivalent.

Table 3.4: Explanatory power of preferences for annuity choices. In Panel A, the observed preferences show the actual annuity choices against the expected utility choices, according to the individually estimated risk and time preferences. The long-run preferences show the actual annuity choices against the expected-utility choices according to the individually estimated risk and time preferences under the persistent long-run view, i.e., present-bias factor $\hat{\beta}_i = 1$ for each retiree i. Panel B shows the actual annuity choices against the expected utility choices for several utility indifference bandwidths (i.e., the "observed preferences" from Panel A).

Panel A: No utility indifference					
	Observed preferences				
			Actual		
			Flat	Front-loaded	Total
	Expected	Flat	135	60	195
		Front-loaded	113	49	162
		Total	248	109	357
	Long-run preferences				
			Actual		
			Flat	Front-loaded	Total
	Expected	Flat	142	64	206
		Front-loaded	106	45	151
		Total	248	109	357
Panel B: Utility indifference bandwidths					
			Actual		
			Flat	Front-loaded	Total
Annual cons. difference ≤ 0.01	Expected	Flat	180	35	215
		Front-loaded	68	74	142
Percentage explained: 71%		Total	248	109	357
			Actual		
			Flat	Front-loaded	Total
Annual cons. difference ≤ 0.02	Expected	Flat	206	21	227
		Front-loaded	42	88	130
Percentage explained: 82%		Total	248	109	357
			Actual		
			Flat	Front-loaded	Total
Annual cons. difference ≤ 0.03	Expected	Flat	222	14	236
		Front-loaded	26	95	121
Percentage explained: 89%		Total	248	109	357
			Actual		
			Flat	Front-loaded	Total
Annual cons. difference ≤ 0.04	Expected	Flat	232	11	243
		Front-loaded	16	98	114
Percentage explained: 92%		Total	248	109	357
			Actual		
			Flat	Front-loaded	Total
Annual cons. difference ≤ 0.05	Expected	Flat	236	6	242
		Front-loaded	12	103	115
Percentage explained: 95%		Total	248	109	357

Indifference bandwidths

We now allow for the possibility that expected annuity choices lie within an indifference bandwidth of the actual annuity choices.[27] Within this bandwidth, we argue that actual and expected annuity choices are observationally equivalent for the individual in terms of utility. The analysis with the indifference bandwidths can be interpreted as conditional predictions.

We compute the bounds of the indifference bands by the difference between the actual (x_t^{act}) and expected (x_t^{exp}) annuity payments. The bound of the bandwidth determines the maximum allowed utility difference between the actual and expected annuity choice to be observationally equivalent. The bound of the bandwidth ε is defined as the annual percentage consumption loss and determined by

$$\sum_{t=0}^{T} p(t)\phi(t;\hat{\beta},\hat{\delta})u\left(x_t^{act}\cdot(1+\varepsilon);\hat{\alpha}\right) = \sum_{t=0}^{T} p(t)\phi(t;\hat{\beta},\hat{\delta})u\left(x_t^{exp};\hat{\alpha}\right). \tag{3.10}$$

If the bound of the indifference bandwidth ε is zero or negative, then the discounted expected utility model — with individual preferences as inputs — explains the actual choice of the retiree entirely successful. Namely, the actual annuity choice yields equal utility or higher utility than the expected utility choice. Thus, the individual makes an actual annuity choice that maximizes her utility given her preferences. In Table 3.4, this holds true for the 52% of our sample, namely 184 retirees out of the 357 (i.e., the retirees on the diagonal).

If the bound of the indifference bandwidth is not too large and positive, i.e., $\varepsilon > 0$, then actual and expected annuity choices are observationally equivalent for the individual as utility differences are small. That is, individual preferences explain the actual choice of the retiree with some prediction error. The severity of misprediction is given by the magnitude ε in terms of annual certainty equivalent consumption. In case the indifference interval is not too wide, then individually estimated preferences explain actual choices. Stated differently, individual preferences explain real-life annuity choices conditional on a small annual consumption

[27] One interpretation is that preferences might be measured with some error or the discounted expected utility model might be misspecified.

difference. In Table 3.4, the groups "actual flat, expected front-loaded" and "actual front-loaded, expected flat" suffer from some prediction error. We now study the severity of these predictions error by means of indifference bands.

Table 3.4, Panel B, presents the number of retirees for each indifference bandwidth ε. We create indifference bandwidths from 0% to 5% annual consumption loss. Using the numbers on the diagonals we can assess the explanatory power of preferences for actual choices. Panel B shows that individually estimated preferences explain actual annuity choices for 82% of the retirees if the indifference bandwidth is at most 2% annual consumption loss (i.e., 206+88 retirees out of 357). Thus, the explanatory power of individually estimated preferences is 82% when the indifference bandwidth equals at most 2% annual consumption loss. Or, the other way around, the percentage of cases with a severe prediction error, e.g., larger than 2%, is only 18%. Of course, if the indifference bandwidth becomes larger (smaller), then individually estimated preferences explain actual annuity choices to a larger (lower) extent.

Figure 3.4 shows the distribution of explained annuity choices for each indifference interval, excluding the correct predictions $\varepsilon = 0$. Stated differently, we display the distribution of explained annuity choices for the groups "actual flat, expected front-loaded" and "actual front-loaded, expected flat". The fraction of explained annuity choices clusters mainly around zero or close to zero, which supports the idea that risk and time preferences explain financial decision making. The severity of prediction errors is distributed similarly amongst both groups.

The actual annuity choice can be made only once, so the participant's real-life choice in general could have been made earlier than the date of our experiment. Since we find that risk and time preferences help explain intertemporal financial choices, our findings speak to a stability of preferences during the retirement phase. This is in line with the finding of Schildberg-Hörisch (2018) that risk preferences are stable on average over the life-cycle from age 60 onwards. Meier and Sprenger (2015) find that estimates of discount parameters are unchanged over a 2 year period, while Kureishi et al. (2021) find that estimates of discount parameters decrease linearly with age over the life-cycle.

Figure 3.4: Distribution of explained annuity choices. The figure displays the fraction of explained annuity choices by individually estimated preferences for each indifference bandwidths. The distribution excludes the group with a prediction error of zero, i.e., $\varepsilon = 0$.

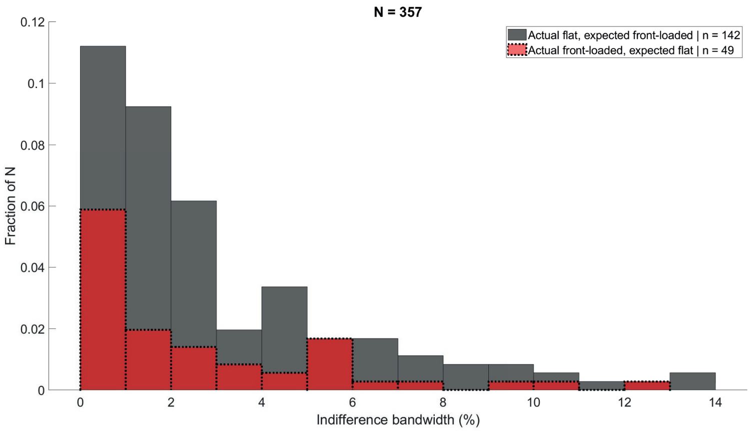

3.3.2 Welfare effects

Given that individually estimated preferences are to a certain extent able to explain annuity choices, we study in this final section the welfare effects of the possibility to choose a front-loaded annuity. The option to take a front-loaded annuity, besides the default flat annuity, creates freedom of choice. Freedom of choice may generate welfare gains, but also welfare losses. Specifically, a front-loaded annuity may cause potential and realized welfare gains or losses. This section computes the potential and realized welfare effects of adding the option to take a front-loaded annuity, rather than a default flat annuity, to the annuity choice menu.

To evaluate the policy of freedom of choice in annuitization decisions, a welfare criterion is needed. A common choice to evaluate welfare is from a long-run perspective, on the grounds that these are the preferences that are persistent (Ericson and Laibson, 2019). In line with Ericson and Laibson (2019), this implies that we study choice behavior if individuals are dynamically consistent, i.e., retirees do not suffer from present bias. So, for each individual we set the estimated present-bias factor $\hat{\beta} = 1$. Using the discounted expected utility model in equation (3.9) we compute the actual and expected annuities' utility values.

Table 3.4, Panel A, shows the actual observed annuity choice and the expected annuity choice from a long-run welfare perspective, i.e., setting $\hat{\beta} = 1$. Of course, the actual observed annuity choices are identical to the "observed preferences": 248 retirees actually choose flat, while 109 retirees choose a front-loaded annuity. According to the long-run preferences (i.e., $\hat{\beta} = 1$), it is expected that 206 retirees choose a flat annuity and 151 retirees choose a front-loaded annuity. Compared to the "observed preferences", the expected utility model using "long-run preferences" predicts that a higher number of retirees prefers a flat annuity, while a lower number of retirees prefers a front-loaded annuity. This is intuitive, because not being subject to present bias pulls individuals away from the possibly tempting choice of a front-loaded annuity. Still, we observe that relatively too many individuals choose actually a flat annuity compared to their expected annuity choices, which may be due to the flat annuity being the default.

Potential and realized welfare effects

To analyse the welfare effects of the option to choose a front-loaded annuity, we distinguish between potential and realized welfare effects. Furthermore, we split these welfare effects in gains (+), losses (-) and no effect (0). Table 3.3 summarizes our explanations below.

The potential welfare gains (+) from a front-loaded annuity come from individuals that are expected to choose a front-loaded annuity, based on their long-run preferences. Namely, from a long-run welfare perspective (i.e., $\hat{\beta} = 1$), it increases the total utility of these retirees to choose a front-loaded annuity. Thus, potential welfare gains of a front-loaded annuity come from the groups "actual flat, expected front-loaded" and "actual front-loaded, expected front-loaded". As Table 3.3 confirms, the long-run discount factors are the lowest among these 2 groups, indicating that these retirees are the most impatient in the long run and prefer front-loaded annuities. Potential welfare effects are defined as the sum of these two groups.

Realized welfare effects are defined as the sum of realized gains (+), losses (-), and no effects (0). Realized welfare gains (+) from a front-loaded annuity come from individuals that are expected to choose a front-loaded annuity and actually do so. Stated differently, the group "actual front-loaded, expected front-loaded" chooses their actual annuity in line with their long-run preferences. Realized welfare losses (-) from the option to take a front-loaded annuity stem from individuals that actually choose a front-loaded annuity, but are expected to choose a flat annuity given their long-run preferences. The main mechanism here is that the group "actual front-loaded, expected flat" chooses a front-loaded annuity because the present-bias factor is relatively low (i.e., 0.87 at the median), while from a long-run perspective this group is patient as the long-run discount factor is close to one (i.e., 1.01 at the median). The group "actual flat, expected front-loaded" has no effect (0) on the realized welfare of a front-loaded annuity, as these individuals realized a flat annuity and, therefore, the potential welfare is unrealized. Finally, we leave the group "actual flat, expected flat" outside the welfare analysis because it has no implications for the welfare effects of a front-loaded annuity.

To determine the magnitude of the welfare effects, we calculate the annual percentage consumption effect ε using equation (3.10) with $\hat{\beta} = 1$. For the group "actual flat, expected front-loaded", the potential welfare gain follows naturally from equation (3.10) as ε is positive since the expected choice always yields equal or higher utility than the actual choice. For the group "actual front-loaded, expected flat", the realized welfare loss follows from equation (3.10) by converting the positive ε to its negative counterpart. Namely, we want to know the consumption loss that occurs by foregoing to choose a flat annuity. For the group "actual front-loaded, expected front-loaded", we compute the potential and realized welfare gains as follows. We counterfactually assume as if this group actually chooses a flat annuity. Then, we quantify the potential and realized consumption gains by directly computing ε in equation (3.10). Again, the potential and realized welfare gains follow naturally from equation (3.10) as ε is positive since the expected choice always yields equal or higher utility than the actual choice, such that these positive values indicate the annual percentage consumption gains of a front-loaded annuity.

Table 3.5: Potential and realized welfare gains and losses due to freedom of choice. This table presents the annual consumption effects (CE) in percentage (%) and the monetary welfare effects in Euro (€) for each indifference interval.

	Indifference interval (%)													
	0		[0,1]		[0,2]		[0,3]		[0,4]		[0,5]		[0,∞)	
	CE	€	CE	€	CE	€	CE	€	CE	€	CE	€	CE	€
Panel A: Potential welfare														
Mean	2.77	13417	1.66	7969	1.61	7659	1.74	8160	1.87	8993	1.94	9355	2.29	11599
Median	1.89	7624	0.63	2884	0.95	4195	1.18	5100	1.32	5790	1.37	5895	1.63	6974
Std. Dev.	3.12	14919	2.59	12487	2.27	11012	2.15	10454	2.12	10607	2.12	10683	2.38	13392
5% perc.	0.13	324	0.03	93	0.04	112	0.04	117	0.04	120	0.04	121	0.04	124
95% perc.	10.70	41982	6.31	38092	6.17	33674	5.99	31877	5.79	30549	5.68	29911	6.18	37613
Percentage population	24	24	33	33	37	37	39	39	41	41	41	41	42	42
Observations pot. welfare	45	45	83	83	109	109	125	125	135	135	139	139	151	151
Observations interval	191	191	254	254	294	294	317	317	330	330	339	339	357	357
Panel B: Realized welfare														
Mean	2.50	12141	0.98	4715	0.56	2586	0.39	1771	0.31	1358	0.17	663	-0.07	-262
Median	1.81	7114	0.00	0	0.00	0	0.00	0	0.00	0	0.00	0	0.00	0
Std. Dev.	3.12	14928	2.47	11892	2.27	11110	2.20	10796	2.18	10736	2.28	11334	2.63	12136
5% perc.	-0.55	-2301	-0.78	-4588	-1.63	-10395	-2.06	-11718	-2.40	-11994	-2.67	-14306	-4.33	-18655
95% perc.	10.35	41379	6.15	33251	5.31	28019	4.82	26254	4.72	24807	4.56	23998	4.17	22852
Percentage population	26	26	44	44	52	52	55	55	57	57	58	58	60	60
Observations real. welfare	49	49	112	112	152	152	175	175	188	188	197	197	215	215
Observations interval	191	191	254	254	294	294	317	317	330	330	339	339	357	357
Panel C: Potential welfare for population														
Mean	0.65	3161	0.54	2604	0.60	2840	0.69	3218	0.77	3679	0.80	3836	0.97	4906
Median	0.45	1796	0.21	943	0.35	1555	0.46	2011	0.54	2369	0.56	2417	0.69	2950
Std. Dev.	0.73	3515	0.85	4080	0.84	4083	0.85	4122	0.87	4339	0.87	4381	1.01	5664
5% perc.	0.03	76	0.01	31	0.01	41	0.02	46	0.02	49	0.02	49	0.02	53
95% perc.	2.52	9891	2.06	12447	2.29	12484	2.36	12570	2.37	12498	2.33	12264	2.62	15909
Panel D: Realized welfare for population														
Mean	0.64	3115	0.43	2079	0.29	1337	0.22	978	0.18	774	0.10	385	-0.04	-158
Median	0.46	1825	0.00	0	0.00	0	0.00	0	0.00	0	0.00	0	0.00	0
Std. Dev.	0.80	3830	1.09	5244	1.17	5744	1.21	5960	1.24	6116	1.32	6587	1.58	7309
5% perc.	-0.14	-590	-0.34	-2023	-0.85	-5374	-1.14	-6469	-1.37	-6833	-1.55	-8314	-2.61	-11235
95% perc.	2.66	10615	2.71	14662	2.75	14486	2.66	14493	2.69	14133	2.65	13946	2.51	13763

Table 3.5 presents an overview of the potential and realized welfare effects. We compute the welfare effects for different subsamples. Namely, we want to assure that preferences have predictive power for annuity choices when doing a welfare analysis. The smaller the interval, the more strict the individual determines indifference between her actual and expected annuity choice. Panels A and B display the conditional potential and realized welfare effects, i.e., only for the affecting the potential and realized welfare. Panels C and D display the unconditional potential and realized welfare effects, i.e., including the group of retirees "actual flat, expected flat". We do the latter, because welfare effects can be substantial for a small number of affected individuals, but smoothed over the other individuals in society, welfare effects might appear differently. Besides the effects on annual percentage consumption, we also compute the monetary welfare effects in terms of additional present value pension wealth at retirement.

For an indifference bandwidth of 2%, the mean potential welfare effect of a front-loaded annuity is a gain of 1.61%.[28] Unconditional potential welfare in Panel C yields an average potential welfare gain of 0.60%. Realized welfare gains are lower, but still positive on average: 0.56% for the conditional sample (Panel B), and on average 0.29% for the unconditional sample (Panel C). In terms of money, Panels A (C) and B (D) show that the average individual has a potential gain of €7659 (€2840), but only realizes €2586 (€1337). The 5%-percentile shows that realized welfare is negative. Welfare effects are similar for other indifference intervals, where the highest average potential welfare of 2.77% is attained in the indifference interval with zero prediction error. Hence, a takeaway is that policy making can be improved to guide individuals in annuitization decisions, because there is still unrealized welfare in the economy.

[28] The number is based on 83 observations from the groups "actual flat, expected front-loaded" and "actual front-loaded" as a fraction of the total number of 294 retirees in the indifference interval.

3.4 Conclusion

To the best of our knowledge, our paper is the first to relate domain specific and simultaneously measured risk and time preferences to real-life annuitization decisions through a utility framework. We simultaneously measure risk and time preferences in a real-life pension context, with long horizons, for a large group of pension fund participants. We base our method on the Convex Time Budgets of Andreoni and Sprenger (2012a) with an additional present-bias task (Rieger et al., 2015). We use the individually estimated preferences as inputs in a discounted expected utility framework to predict actual observed annuity decisions. Given the predictive power of preferences for actual annuity choices, we quantify the welfare effects of freedom of choice in the annuitization decision between a flat and front-loaded annuity.

We find that pension fund participants are present biased, but retirees are less present biased than active participants. In the context of pension decision making, involving long horizons and large stakes, we find annual discount rates close to 1% and utility curvature close to unity (i.e., risk neutral). The front-loaded annuity from the Dutch pension fund, bearing similarities with a lump sum, is actually chosen over a flat annuity by present biased, impatient, and more risk tolerant individuals. The discounted expected utility model, with measured preferences as inputs, explains for 82% of the retirees actual annuitization decisions when allowing for an indifference bandwidth of at most 2%. Within this bandwidth, we argue that the actual and utility-expected annuity choices are observationally equivalent to the individual in terms of utility. For an indifference interval of 2%, conditional individual potential welfare gains are on average 1.61% (€7659) but only 0.56% (€2586) is realized. Welfare losses realize at the lower end of the distribution, because individuals make suboptimal choices due to present bias compared to a time-consistent expected utility maximizer.

For policy making, we can conclude that offering freedom of choice in the annuity decision, by means of an option to take a front-loaded annuity in addition to a flat annuity, yields on average welfare gains. But, the realized welfare gains can be roughly three times as large — for an indifference bandwidth of 2% — when individuals are guided better during their annuity choices.

Our study is based on a unique dataset of individuals from a large Dutch pension fund compromising detailed data on retirement plans. We augment the dataset by survey data, which includes our experiments on risk and time preferences. The measurement of preferences in the same domain, context (i.e., amounts at stake), and population as the actual decision making is a novel contribution of our research. The data deals with real annuity choices rather than self-reported intentions. Our analysis is based on administrative records of the Dutch pension fund, including fund specific survival probabilities.

3.5 Appendix

Table 3.6: Overview experimental design: Convex Time Budgets and present bias. Choice sets in the Convex Time Budgets and present-bias task. t and k are front and end delays in years, and c_t and c_{t+k} are allocated amounts in Euros. $1+r$ is the implied gross interest rates. Annual r is the yearly interest rate in percent and calculated as $((1+r)^{1/k}-1)\times 100$. For the present bias task subjects enter an amount (in €) for $c_{t+\tau}$.

Task	Scenario	Set	t	k	p_{t+k}	c_t	c_{t+k}	$1+r$	Annual r
	1	1	1	10	0.5	10,000	14,100	1.41	3.50
	2	1	1	10	0.5	10,000	14,900	1.49	4.07
	3	1	1	10	0.5	10,000	16,600	1.66	5.20
	4	1	1	10	0.5	10,000	19,300	1.93	6.80
	5	1	1	10	0.5	10,000	22,400	2.24	8.40
	6	2	1	10	0.7	10,000	12,000	1.20	1.84
	7	2	1	10	0.7	10,000	12,600	1.26	2.34
	8	2	1	10	0.7	10,000	14,000	1.40	3.42
Convex	9	2	1	10	0.7	10,000	16,300	1.63	5.01
Time	10	2	1	10	0.7	10,000	19,000	1.90	6.63
Budgets	11	3	1	10	0.9	10,000	10,500	1.05	0.49
	12	3	1	10	0.9	10,000	11,100	1.11	1.05
	13	3	1	10	0.9	10,000	12,300	1.23	2.09
	14	3	1	10	0.9	10,000	14,400	1.44	3.71
	15	3	1	10	0.9	10,000	16,700	1.67	5.26
	16	4	1	10	1.0	10,000	10,000	1.00	0.00
	17	4	1	10	1.0	10,000	10,500	1.05	0.49
	18	4	1	10	1.0	10,000	11,700	1.17	1.58
	19	4	1	10	1.0	10,000	13,600	1.36	3.12
	20	4	1	10	1.0	10,000	15,900	1.59	4.75
Present bias	21	5	0	1	1.0	800	$c_{t+\tau}$	$800/c_{t+\tau}$	$800/c_{t+\tau}-1$

Figure 3.5: Distributions of estimated individual present-bias factor, annual discount factor and CRRA curvature. These distributions are based on the estimated parameters from Table 3.2.

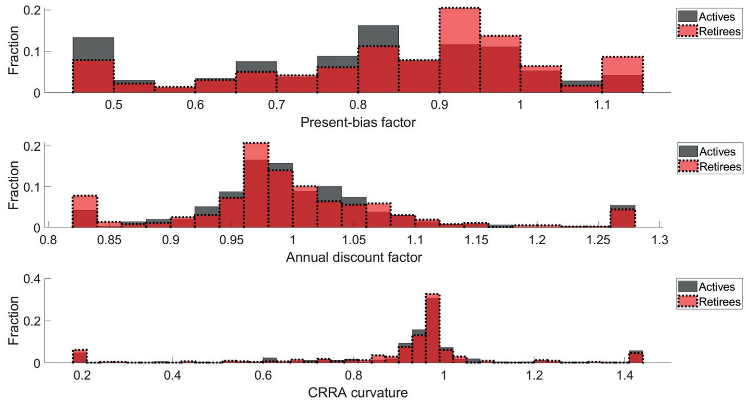

3.6 Online appendix

A. Estimation

There are N experimental subjects and P convex budget decisions, where we combine the present-bias task with the CTB decisions. We assume that each subject j makes her allocation decision $c_{t_{i,j}}$, $i = 1, ..., P$, according to the relationship in (3.5), but that each decision is made with some additive mean-zero (potentially correlated) error. That is,

$$\begin{aligned} \ln\left(\frac{c_t + w_t}{c_{t+k} + w_{t+k}}\right)_{i,j} &= \frac{1}{\alpha - 1}\left(\log(\beta) \cdot \mathbb{1}_{t\in[0,1]} + \log(\delta) \cdot k\right) \\ &+ \frac{1}{\alpha - 1}\left(\log(p_{t+k_i}) + \log(1 + r_i)\right) + \varepsilon_{i,j}. \end{aligned} \tag{3.11}$$

Stacking the P observations per individual j, we have

$$\begin{aligned} \ln\left(\frac{\boldsymbol{c_t} + \boldsymbol{w_t}}{\boldsymbol{c_{t+k}} + \boldsymbol{w_{t+k}}}\right)_{\boldsymbol{j}} &= \frac{1}{\alpha - 1}\left(\log(\beta) \cdot \mathbb{1}_{t\in[0,1]} + \log(\delta) \cdot k\right) \\ &+ \frac{1}{\alpha - 1}\left(\log(\boldsymbol{p_{t+k}}) + \log(\boldsymbol{1} + \boldsymbol{r})\right) + \boldsymbol{\varepsilon_j}. \end{aligned} \tag{3.12}$$

The vector $\boldsymbol{\varepsilon_j}$ is zero in expectation with variance-covariance matrix $\boldsymbol{\Sigma_j}$, a $P \times P$ matrix, allowing for arbitrary correlation in the errors $\varepsilon_{i,j}$. For each subject j, we assume that all decisions i are subject to an error with mean zero and variance σ_i^2. So, $\boldsymbol{\Sigma_j}$ is a (homogeneous) diagonal variance-covariance matrix with entries σ_i^2 on the diagonal and zeros off diagonal. In other words, the error term is the same within subject j for each decision i, but the error term may vary across individuals.

Equation (3.12) is easily estimated with ordinary least squares. However, the log-consumption ratio is censored by the corner responses on the budget constraint

$$\ln\left(\frac{\boldsymbol{c_t} + \boldsymbol{w_t}}{\boldsymbol{c_{t+k}} + \boldsymbol{w_{t+k}}}\right)_{\boldsymbol{j}} \in \left(\ln\left(\frac{\boldsymbol{0} + \boldsymbol{w_t}}{(\boldsymbol{m} \cdot (\boldsymbol{1} + \boldsymbol{r})) + \boldsymbol{w_{t+k}}}\right)_{\boldsymbol{j}}, \ln\left(\frac{\boldsymbol{m} + \boldsymbol{w_t}}{\boldsymbol{0} + \boldsymbol{w_{t+k}}}\right)_{\boldsymbol{j}}\right). \tag{3.13}$$

Namely, either the subject allocates the complete budget m to the late payment at the vector of gross interest rates $\boldsymbol{1+r}$ (and allocates nothing to the early payment), or the subject allocates the complete budget m to the early payment (and allocates nothing to the late payment). These corner solutions motivate the use of censored regression techniques such as the two-limit tobit model.

Finally, the risk- and time-preference parameters for each individual j can be estimated via the linear equation

$$\ln\left(\frac{\boldsymbol{c_t}+\boldsymbol{w_t}}{\boldsymbol{c_{t+k}}+\boldsymbol{w_{t+k}}}\right)_{\boldsymbol{j}} = \eta_{j,0}+\eta_{j,1}\cdot(\ln(\boldsymbol{1+r})+\ln(\boldsymbol{p_{t+k}}))+\boldsymbol{\varepsilon_j}, \tag{3.14}$$

where $\eta_{j,0}$ and $\eta_{j,1}$ are the individual specific intercept and regression coefficient, respectively. For each individual j, the preference estimates for utility curvature, long-term discounting and present bias are found via the non-linear combinations

$$\begin{aligned}
\hat{\alpha} &= \frac{1}{\hat{\eta_{j,1}}}+1,\\
\hat{\delta} &= \exp\left[\frac{\hat{\alpha}-1}{k-1}\left(\hat{\eta}_{j,0}-\frac{\hat{\alpha}}{\hat{\alpha}-1}\log\left(\frac{800+w_0}{c_{t+\tau}+w_{t+\tau}}\right)\right)\right],\\
\hat{\beta} &= \frac{1}{\hat{\delta}}\left(\frac{800+w_0}{c_{t+\tau}+w_{t+\tau}}\right)^{\hat{\alpha}}.
\end{aligned} \tag{3.15}$$

A point of attention is that the background consumption parameters are known or fixed and, secondly, that the consumption ratio $(c_t+w_t)/(c_{t+k}+w_{t+k})_{i,j}$ is strictly positive, such that the log transform is well-defined. The strength is that corner solutions are easily addressed by censoring models such as two-limit tobit maximum likelihood regression.

B. Additional summary statistics and tax levels

Table 3.7: Additional summary statistics of the sample.

	Mean	Median	Standard Deviation	N
Panel A: Demographics				
Male	0.57	1.00	0.50	1062
Age (years)	60.45	60.60	5.89	1062
Education (classes)	3.90	4.00	0.97	1058
Retired	0.34	0.00	0.47	1062
Partner	0.82	1.00	0.39	1055
Children	1.83	2.00	1.23	1058
Panel B: Financial				
Income	53119	51329	22451	1000
Savings	49198	20000	80970	945
Homeowner	0.90	1.00	0.29	1059
Plan to buy	0.13	0.00	0.34	85
Rent price	744	693	282	74
House price	300750	269000	244620	893
Mortgage	0.82	1.00	0.38	947
Expect inheritance	0.26	0.00	0.44	979
Inheritance amount	88393	56250	113280	224
Leave bequest	0.58	1.00	0.49	755
Bequest amount	151020	150000	145730	343
Panel C: Pension				
Pension income	22546	20916	13735	1000
Pension income (max)	29357	28463	15243	1000
Other pension income	7466	500	16548	280
Individual pension income	50000	30000	41569	22
Part-time pension	0.00	0.00	0.00	357
AOW bridge	0.40	0.00	0.49	357
Flexible pension	0.35	0.00	0.48	357
Transfer partner pension	0.23	0.00	0.42	357
Intended retirement year	-2.58	-3.00	2.21	627
Attitude pension choice	4.43	5.00	0.75	1060
Attitude premium stop	2.64	3.00	1.04	1036
Attitude flexible pension age	4.50	5.00	0.70	1056
Attitude flexible pension benefits	3.81	4.00	1.07	1050
Panel D: Other				
Life expectancy	84.01	85.00	4.29	395
Duration (min.)	311.71	19.78	1728.70	1062
Complexity	2.86	3.00	1.01	1028

Table 3.8: Definition of variables. [a] Participants could easily access this information via a provided link directing to house price administration

Variable	Definition
	Panel A: Demographics
Male	Dummy; 1 = male; 0 = female
Age	Age in years (pension fund administration)
Education	Classes; 0 = primary school; 1 = secondary school; 2 = pre-vocational education and training (LBO); 3 = vocational education and training (MBO); 4 = university of applied sciences (HBO); 5 = university
Retired	Dummy; 1 = retired participant (retiree); 0 = active participant (worker)
Partner	Dummy; 1 = married, registered partnership or cohabitation; 0 = no partner
Children	Number of children
	Panel B: Financial
Income	Individual annual before tax income. For retirees, all employer-related second pillar pension benefits received from the pension fund including state pension benefits. For workers, salary corrected for part-time work.
Private savings	Self-reported total individual amount of voluntary liquid savings (e.g. a bank account and/or investments) in one of the classes: (0-5,000), (5,001-10,001),(10,001-30,000), (30,001-50,000), (50,001-100,000), (100,001-200,000), (200,001-400,000), (> 400,000). Excluding house and pension savings.
Homeowner	Dummy; 1 = House owner, 0 = rent a house
Plan to buy	Dummy; 1 = Rent a house, but planning to buy a house, 0 = rent a house, but not planning to buy a house
Rent price	Self-reported current rent price of house (on household level) for tenants (including service fees, excluding gas, water and electricity costs)
House price	Self-reported current house price (on household level) for homeowners; 0 = renting a house[a]
Mortgage	Dummy; 1 = currently one or more mortgage loans; 0 = currently no mortgage loans
Expect inheritance	Dummy; 1 = expect to receive an inheritance (money, real estate or other possessions) during remaining life cycle; 0 = no
Inheritance amount	Individual expected inherited amount in one of the classes: (< 25,000), (25,001-50,000), (50,001-100,000), (100,001-300,000), (300,001-500,000), (> 500,000)
Leave bequest	Dummy; 1 = wish to leave a bequest (savings, house or other possessions) when passing away; 0 = no
Bequest amount	Individual expected bequest amount in one the classes: (< 25,000), (25,001-50,000), (50,001-100,000), (100,001-300,000), (300,001-500,000), (> 500,000)

Table 3.9: Definition of variables (continued). Att. is abbreviation for attitude, and AOW is abbreviation for state pension. [a] Participants could easily access this information via a provided link directing to the pension government administration. [b] Participants were provided that per year of early retirement pension benefits decrease by 6%, and per year of later retirement pension benefits increase by 8%. [c] Participants were shown that the average life expectancy in The Netherlands equals approximately 85 years.

Variable	Definition
	Panel C: Pension
Pension income	Individual annual before tax second pillar accrued pension rights.
Pension income max	Projected individual annual before tax second pillar accrued pension rights.
Other pension income	Self-reported individual annual before tax second pillar pension benefits received from other pension funds (e.g. accrued in the past) in one of the classes: 0 = none, (< 1,000), (1,002-5,000), (5,001-10,000), (10,001-20,000), (20,001-30,000), (30,001-50,000), (50,001-100,000), (> 100,000)[a]
Individual pension income	Self-reported individual annual before tax pension benefits received from insurance companies or banks in one of the classes: 0 = none, (< 5000), (5,001-10,000), (10,001-30,000), (30,001-50,000), (50,001-100,000), (100,001-200,000), (> 200,000)
Part-time pension	Dummy; 1 = administrated part-time pesion; 0 = no part-time pension
AOW bridge	Dummy; 1 = administrated AOW bridge (second pillar financial compensation in case of early retirement); 0 = no AOW bridge
Flexible pension	Dummy; 1 = administrated flexible pension in the form of a high-low or low-high annuity; 0 = no flexible pension
Transfer partner pension	Dummy; 1 = administrated transfer of partner pension to old age pension; 0 = no transfer of partner pension
Intended retirement age	Intended retirement year with respect to the statutory retirement age in one of the classes (negative values indicate early retirement, positive values indicate later retirement): (< -5), (-5), (-4), (-3), (-2), (-1), (1), (2), (3), (> 3).
Att. pension choices	Classes; 1 = strongly disagree with more freedom of pension choices, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree with more freedom of pension choices
Att. premium stop	Classes; 1 = strongly disagree with the choice premium stop, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree with the choice premium stop
Att. flexible pension age	Classes; 1 = strongly disagree with the choice of a flexible pension age, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree with the choice of a flexible pension age
Att. flexible pension benefits	Classes; 1 = strongly disagree with the choice a flexible pension benefits, 2 = disagree, 3 = neutral, 4 = agree, 5 = strongly agree with the choice a flexible pension benefits
	Panel D: Other
Life expectancy	Expected life expectancy in years reported in one of the classes: (< 75), (75-84), (85), (86-90), (> 90)[c]
Duration	Minutes between starting and ending the survey
Complexity	Classes; 1 = very easy survey, 2 = easy, 3 = neutral, 4 = difficult, 5 = very difficult survey

Table 3.10: Median individual tax levels in The Netherlands for active participants and retirees. Tax levels are based on individual annual before tax income. We constructed the tax levels for actives by adding 10 percentage points to the tax level of the corresponding income level for retirees.

	Tax (fraction income)	
Income (€)	Active participants	Retirees
<17,802	0.19676	0.09676
<20,018	0.18671	0.08671
<21,849	0.20572	0.10572
<23,731	0.23090	0.13090
<26,327	0.24774	0.14774
<29,729	0.26721	0.16721
<34,250	0.28571	0.18571
<40,542	0.29940	0.19940
<51,792	0.35565	0.25565
<65,000	0.39650	0.29650
<80,000	0.42698	0.32698
≥80,000	0.48345	0.38345

C. Robustness checks

Table 3.11: Preferences and personal characteristics. The table presents correlations of three Ordinary Least Squares (OLS) regressions with the individually estimated preferences as dependent variables: present-bias factor $\hat{\beta}$, discount factor $\hat{\delta}$, and curvature $\hat{\alpha}$. Controls include the duration and reported complexity of the survey. ** and * indicate statistical significance at the 5% and 10% level, respectively. Robust standard errors (MacKinnon and White, 1985) between parentheses.

	Present-bias factor		Discount factor		Curvature	
Male	0.029**		-0.016**		0.014	
	(0.015)		(0.007)		(0.02)	
Age	0.003**		0		-0.002	
	(0.001)		(0.001)		(0.002)	
Edu. medium	-0.034		-0.023		0.02	
	(0.042)		(0.031)		(0.09)	
Edu high	0.009		-0.038		0.042	
	(0.04)		(0.031)		(0.088)	
Partner	-0.009		0.007		0.003	
	(0.016)		(0.008)		(0.023)	
Income ($\times$1000)	-0.001		0.001		0.004	
	(0.004)		(0.002)		(0.006)	
Savings 5k-10k	0.053**		-0.022*		0.014	
	(0.023)		(0.012)		(0.033)	
Savings 10k-30k	0.056**		-0.018*		-0.028	
	(0.02)		(0.01)		(0.027)	
Savings 30k-50k	0.062**		-0.011		-0.028	
	(0.022)		(0.011)		(0.03)	
Savings $\geq$ 50k	0.091**		-0.014		-0.011	
	(0.02)		(0.01)		(0.028)	
Life expectancy		0.004*		0		-0.001
		(0.002)		(0.001)		(0.003)
Constant	0.55**	0.457**	1.061**	0.995**	0.953**	0.991**
	(0.085)	(0.199)	(0.051)	(0.086)	(0.14)	(0.22)
Observations	862	395	862	395	862	395
Controls	Yes	No	Yes	No	Yes	No

Table 3.12: Individual present bias, annual discounting, and risk aversion parameter estimates under different background income assumptions. Two-limit tobit maximum likelihood estimates for CRRA risk aversion γ, present-bias factor β, and discount factor δ. The CRRA utility function is $x^{(1-\gamma)}/(1-\gamma)$ for $\gamma \neq 1$: $\gamma = 0$ denotes risk neutral behavior, $\gamma > 0$ denotes risk aversion and $\gamma < 0$ denotes risk seeking behavior. Panel A presents estimation results without background income, i.e., $w = 0$. Panel B presents the preference parameters when individual annual after-tax income is used in the estimation.

	Median	Mean	Standard Deviation	25th Percentile	75th Percentile	N
Panel A: Without background income						
Present-bias factor $\hat{\beta}$	0.836	0.819	0.184	0.695	0.953	1062
Discount factor $\hat{\delta}$	0.989	1.004	0.092	0.962	1.039	1062
Annual discount rate	0.011	0.004	0.089	-0.037	0.040	1062
CRRA risk aversion $\hat{\gamma}$	0.035	0.085	0.252	0.013	0.095	1062
Panel B: With background income						
Present-bias factor $\hat{\beta}$	1.025	1.021	0.103	0.978	1.058	1000
Discount factor $\hat{\delta}$	0.975	0.998	0.114	0.950	1.029	1000
Annual discount rate	0.025	0.014	0.105	-0.028	0.052	1000
CRRA risk aversion $\hat{\gamma}$	1.430	1.911	4.724	0.587	3.028	1000

Table 3.13: Robustness winsorization level of 1%. Panel A shows the two-limit tobit maximum likelihood estimates for all participants for the present-bias factor, discount factor, and CRRA utility, similar to Table 3.2 but now preferences are winsorized at a 1% level. Panel B reports the number of retirees for which individually estimated preferences (winsorized at a 1% level) explain actual annuity choices using an indifference bandwidth of at most $\varepsilon = 2\%$, similar to Panel B in Table 3.4.

Panel A: Present bias, annual discounting, and curvature parameter estimates

	Median	Mean	Standard Deviation	25^{th} Percentile	75^{th} Percentile	N
Present-bias factor	0.836	0.832	0.318	0.694	0.953	1062
Discount factor	0.989	1.050	0.394	0.962	1.043	1062
Annual discount rate	0.011	0.006	0.193	-0.041	0.039	1062
CRRA curvature	0.965	0.919	0.577	0.905	0.987	1062

Panel B: Explanatory power of preferences for annuity choices

			Actual		
			Flat	Front-loaded	Total
Annual cons. difference ≤ 0.02	Expected	Flat	205	23	228
		Front-loaded	43	86	129
Percentage explained: 82%		Total	248	109	357

Chapter 4

Time-varying Risk and Time Preferences: Relation with Trading Behavior*

*We thank Hazel Bateman, Stefano Cassella, Rob van den Goorbergh, Frank de Jong, Jona Linde, Eduard Ponds, Jan Potters, Nikolaus Schweizer, Bas Werker, Stefan Zeisberger, and seminar participants at the Netherlands Economists Day (2022), Research in Behavioral Finance Conference (2022), Royal Dutch Economic Association KVS New Paper Sessions (2022), Experimental Finance Conference (2021), Netspar International Pension Workshop (2021), Tilburg University (2022, 2021), International Pension Research Association (2020), Netspar Pension Day (2020), and APG Asset Management for useful comments.

Are individuals' risk and time preferences affected during extreme events? And if so, do economic outcomes, such as trading behavior, relate to changes in preferences? Preferences form the foundation in almost any intertemporal-choice model. Standard models in economics and finance typically assume that individuals have stable and persistent preferences over time (Stigler and Becker, 1977). Research on the stability of preferences is conceptually at the heart of microeconomics. Changes in the stability of preferences have vital real-world consequences for economic outcomes, financial decision making, policy options and welfare analyses.

The literature on the relation between preferences, economic behavior, and exogenous shocks such as natural disasters, violent conflicts, economic crises, and pandemics is relatively new, but growing rapidly (Chuang and Schechter, 2015; Schildberg-Hörisch, 2018; Drichoutis and Nayga, 2021). We contribute to this literature by examining empirically how individuals' preferences in financial-economic decision making develop *during* an exogenous shock and, at the same time, whether economic behavior is related to changes in preferences. The exogenous shock we study is the COVID-19 crisis, i.e., a pandemic, and the economic behavior we study is the disposition effect. The disposition effect is the stylized fact that paper losses are realized less than paper gains (Odean, 1998). Besides, we contribute to the literature by studying whether the disposition effect is stable *during* an exogenous shock. Overall, we make three contributions to the literature (i) by studying how preferences and economic behavior develop over time during a crisis rather than by means of a 'before-after' analysis based on waves, (ii) by using multiple elicitation methods to measure preferences, and (iii) by analyzing economic behavior.

We use experimental approaches in two surveys to elicit preferences and trading behaviour on a daily basis during the COVID-19 crisis. We fielded our surveys during March 2020 (i.e., emergence COVID-19 and first lockdown) and December 2020 (i.e., second lockdown) using the Longitudinal Internet Studies for the Social Sciences (LISS). LISS yields a representative sample of the Dutch population. The first part of each survey elicits risk and time preferences, and the second part of each survey elicits the disposition effect. Our main measure of risk and time preferences is based on the Convex Time Budgets (CTB) approach of

Andreoni and Sprenger (2012a) to simultaneously elicit and estimate preferences. We estimate risk aversion, present bias, and patience, and we control for probability weighting. Using cognitive simpler but coarser methods (Dave et al., 2010) we also independently measure (i) risk aversion by the lottery task of Eckel and Grossman (2008), (ii) present bias and patience by the matching task of Rieger et al. (2015) and Wang (2017), and (iii) we use qualitative statements to measure these preferences. The second part of the survey is a trading experiment building upon the design of Ploner (2017) to measure the disposition effect. By combining preferences with trading behavior, we are able to study how elicited preferences are related to observed economic outcomes during an exogenous shock.[1]

Even if individuals do not experience the shock directly themselves (i.e., get infected or hospitalized by COVID-19), fear and uncertainty can be activated by watching and reading news about the exogenous shock. Our experiments took place during the emergence of COVID-19 and the following lockdowns, times in which fear and uncertainty could be born naturally as there were growing concerns about the future state of the world. National COVID-19 hospitalization numbers were communicated daily on the news and via push notifications on mobile phones. Hospitalizations were especially salient, because there was a general concern that hospitals might reach full capacity with amongst others a potential health crisis consequently. So, we measure the severity of COVID-19 by the relative change in hospitalizations and we argue that a potential mechanism for time-varying risk and time preferences is through fear (Meier, 2022; Guiso et al., 2018; Cohn et al., 2015) and uncertainty (Carroll et al., 2019; Giavazzi and McMahon, 2012; Mody et al., 2012).

Our results show that risk and time preferences strongly correlate with the severity of the COVID-19 crisis, indicating instability of preferences during an exogenous shock. Specifically, if daily national COVID-19 hospitalizations in The Netherlands increase, then risk aversion, time consistency, and long-term patience increase as well. Thus, individuals are less willing to take financial risks and prefer

[1]We study trading behavior in our paper, but it could actually have been any other economic decision. We merely use the observed trading behavior to demonstrate how preferences relate to financial-economic decision making.

to save more for the future after an increase in the daily COVID-19 hospitalizations. The effect on time preferences is economically large: a two standard deviation increase (approximately 75%) in *daily* COVID-19 hospitalizations, decreases the *annual* discount rate from 4.3% to 2.6%, increases the *annual* present-bias factor by 0.05, and increases risk aversion by 0.11.

The increase in risk aversion is consistent with the fear-based mechanism and the findings of Cohn et al. (2015) and Guiso et al. (2018). Additionally, Meier (2022) identifies fear as a significant correlate with within-person increases in risk aversion, using a large panel dataset, and Haushofer and Fehr (2014) finds that fear decreases the amount invested in a risky asset. The increase in time consistency and patience is consistent with the uncertainty-based mechanism (Carroll et al., 2019; Giavazzi and McMahon, 2012; Mody et al., 2012) and the findings of Parker et al. (2022). The former studies find that precautionary savings increase during uncertain times, and the latter finds that households lack spending during COVID-19. The increase in patience in our experiment is identified by an increase in savings for the future and delaying current consumption. Individuals want to insure themselves against bad states of the economy in the future and try to retain a smooth consumption path by precautionary savings and lower consumption in the current period.

Contemporaneously with changes in preferences, the disposition effect decreases if there is an increase in daily COVID-19 hospitalizations. During the disposition effect experiment, investors are asked to sell an asset immediately or to hold it for one more year after they have observed the gains and losses on their chosen assets. We find a decrease of -10.58% in immediately selling assets that experienced a gain if COVID-19 hospitalizations increase. Consistent with intertemporal realization utility (Ingersoll and Jin, 2013), the disposition effect decreases as a result of the investor's higher risk aversion and higher patience during rising COVID-19 hospitalizations. Specifically, investors hold on to winning assets relatively more. Present bias plays no significant role in the selling behavior. So, this suggests that the changes in the disposition effect during the crisis are at least partially driven by time-varying risk aversion and patience, through the domain of selling gains.

Our results are robust to alternative specifications. Decomposing within and between effects allows us to take into account unobserved individual specific heterogeneity and confirms our results. Moreover, our results are not driven by weekly seasonality effects in hospitalizations, nor by changing beliefs in life expectancy, nor by changes in income, and they are robust to changes in COVID-19 *ICU* hospitalizations as dependent variable. We also study the behavior of preferences with changes in COVID-19 hospitalizations on a province level rather than the national level. These province-level estimations yield insignificant results, in line with the fact that the reported COVID-19 hospitalizations on the news were based on national numbers. This suggests that the dynamics in preferences are triggered by the salient national news concerning COVID-19, as province-level COVID-19 hospitalizations were not reported directly in the usual media unless you actively searched for these numbers. We also provide some insight into how preferences develop during the crisis, that is, into differences between the start of the COVID-19 crisis in March 2020 and 10 months later when the COVID-19 crisis manifested itself tightly in December 2020.[2] We find that the level of risk and time preferences is statistically similar in both months, and the sensitivity of preferences to COVID-19 hospitalizations is similar in both months.

We find that cognitive simpler experimental measures, both quantitative and qualitative, are unable to capture the instability of preferences compared to the CTB method. Simpler, but coarser, experimental measures might give too broad ranges for capturing the daily effect of the crisis on preferences, or these methods make too restrictive simplifying assumptions such that these measures have overall lower (predictive) accuracy (Dave et al., 2010). The cognitive more complex, but finer, CTB method might perform better in capturing changing preferences, especially in a developed country such as The Netherlands because the population on average exhibits relatively high numerical skills (Dave et al., 2010; Chuang and Schechter, 2015).

Our contribution to the literature is threefold. First, we study how prefer-

[2]The crisis in The Netherlands started roughly 6 March 2020 with the first COVID-19 death, 1 April 2020 1173 deaths, and 31 December 2020 the total number of COVID-19 deaths was 11843.

ences develop *during* an exogenous shock. The review papers of Chuang and Schechter (2015) and Schildberg-Hörisch (2018) classify shocks in natural catastrophes (i.e., earthquakes, tsunami's, famines, floods, droughts, hurricanes and epidemics), violent conflict (i.e., wars and political violence) and economic shocks (i.e., macroeconomic conditions, financial crises, job search, income and education). The COVID-19 crisis is a combination of a pandemic and an economic shock. Research on the effects of the COVID-19 crisis on preferences finds divergent and inconclusive empirical results. Studies suggest that the COVID-19 crisis increases risk aversion (Bu et al., 2020) for men only (Lohnmann et al., 2020), decreases risk aversion (Ikeda et al., 2020), has no effect on risk preferences (Angrisani et al., 2020; Drichoutis and Nayga, 2021; Bokern et al., 2021), or has no consistent effect on risk preferences (Shachat et al., 2020); decreases impatience for men (Lohnmann et al., 2020), or has no effect on time preferences (Drichoutis and Nayga, 2021; Bokern et al., 2021). [3] Additional research is needed to understand theoretical predictions about the circumstances when we should expect an increase or decrease in preferences.

However, all these studies analyze the effect of the COVID-19 crisis by means of surveys before and after the outbreak, or by means of short waves during the crisis (i.e., a few days to respond to the questionnaire within a specific time interval).[4] Typically, researchers study whether preferences are different from one wave to another. We differ here from the previous studies, because we study how preferences develop over time within and during the COVID-19 crisis itself by measuring preferences over a relatively long period. We find that preferences change with respect to the severity of the COVID-19 crisis, but by comparing aggregate preference levels between both our two waves (i.e., March 2020 and December 2020) we observe no differences in risk and time preferences. Our findings relate to short-term systematic but temporary variations in preferences (Schildberg-Hörisch, 2018).

Second, our article uses multiple experimental methods to measure preferences

[3] Drichoutis and Nayga (2021) give an overview of papers studying the stability of preferences during the COVID-19 pandemic.

[4] Another possibility is a long-term view such as life-cycle or cohort effects on preferences as studied by Malmendier and Nagel (2011), but we abstain from this in our current research.

during the COVID-19 crisis. While we see that there are many other papers studying the stability of preferences, most of these focus only on one type of preference (with a majority studying risk preferences) or one type of elicitation method. We contribute to the literature here by eliciting and estimating risk and time preferences simultaneously including life-expectancy beliefs, and by studying the methodological stability of preferences over multiple elicitation methods such as Eckel and Grossman (2002), Andreoni and Sprenger (2012a), and Rieger et al. (2015). Despite its economic importance, especially relatively little is known about the stability of time preferences (Meier and Sprenger, 2015; Meier, 2019). An exception is the paper of Harrison, Hofmeyr, et al. (2022). They find stability of risk preferences during the COVID-19 crisis in the United States under expected utility theory but not under rank-dependent utility, and stability of time preferences under exponential discounting but more instability under quasi-hyperbolic discounting, although they estimate negative long-term discount rates on average.

The paper closest to our main measure is Lohnmann et al. (2020), as they also use the CTB method. However, the authors only use it to measure time preferences (i.e., present bias and long-term patience), since they measure risk preferences independently by the methods of Eckel and Grossman (2002) and Gneezy and Potters (1997). We simultaneously elicit and estimate risk and time preferences with CTB including a correction for probability weighting (Potters et al., 2016) and risk capacity through the channel of background income in our preference estimations.[5] Compared to the previous experimental literature (e.g., see the review of Andersen, Harrison, Lau, et al., 2014), our CTB setup uses long horizons up to 6 years with large stakes of €10,000 and a large non-student sample representative for the Dutch population. Lohnmann et al. (2020) study a smaller sample ($N = 793$) of Chinese students in Beijing using waves before and after the outbreak of COVID-19. Our approach arguably yields more individual heterogeneity in the exposure to the crisis than studying students only.

Third, our research contributes to the literature regarding (in)stability of the

[5] Rather than assuming two separate mental accounts for experimental and real-life earnings, we assume that participants integrate experimental earnings with real-life earnings in their survey decisions.

disposition effect, and the relation between preferences and the disposition effect. Existing studies implicitly assume the disposition effect to be unaffected by market cycles, while Bernard et al. (2022) find that the disposition effect is time varying. We add to this literature by confirming that the disposition effect varies over time and potentially driven by fluctuations in risk and time preferences. Specifically, we show that time-varying risk aversion and time-varying patience form at least a part of the explanation for the variation in the disposition effect during the COVID-19 crisis such that we provide additional evidence for asset pricing models with time-varying risk aversion (Campbell and Cochrane, 1999) and intertemporal realization utility (Barberis and Xiong, 2012; Ingersoll and Jin, 2013).

The rest of the paper is organized as follows. Section I. describes the experimental designs, including potential mechanisms. Section II. presents the estimated preferences and how they develop during the COVID-19 crisis. Section III. confirms the existence of an aggregate disposition effect and shows the time variation during the crisis. It also presents the relation between preferences and trading behavior. Section IV. concludes the paper.

4.1 Methodology

We adopt an experimental approach in a survey to elicit preferences and trading behavior. The first part of the survey elicits risk and time preferences. Our main measure of risk and time preferences is based on the CTB approach (Andreoni and Sprenger, 2012a). We additionally use simpler quantitative and qualitative methods to elicit risk and time preferences. The second part of the survey uses the experimental setup of Ploner (2017) to investigate the disposition effect. The survey questions are available from the authors upon request.

This section describes the elicitation of preferences and trading behavior, including the experimental procedure. We then describe our sample, and the COVID-19 data. We conclude with the preference parameter estimation based on the CTB.

4.1.1 Measuring preferences

In the first part of the survey, we use the CTB (Andreoni and Sprenger, 2012a) to measure risk aversion, present bias and patience simultaneously. We also measure these preferences independently. To measure risk aversion separately, we use the Eckel and Grossman (2008) lottery task and a qualitative question from the Dutch Central Bank Household Survey (DHS). To measure present bias and patience separately, we use the matching task from the INTRA study of Rieger et al. (2015) and qualitative questions from Gathergood (2012). We randomize the order of presentation of the elicitation methods at the individual level to control potential order effects.

Simultaneous elicitation

An important advantage of the CTB is that it allows to measure risk and time preferences simultaneously. For this reason, we avoid the assumption of linear utility and we avoid upward biased discount rate estimates if true utility is concave (Andersen, Harrison, Lau, et al., 2008b) We simultaneously measure risk aversion, present bias and patience, and we also correct for probability weighting. Therefore, our approach is a combination of the original approach of Andreoni and Sprenger (2012a) and Potters et al. (2016). Andreoni and Sprenger (2012a) measure risk aversion, present bias and patience, while Potters et al. (2016) measure risk aversion, patience and probability weighting.

The simultaneous elicitation method asks individuals to allocate an initial budget $m =$ €$10,000$ between payments, available at two points in time: an early payment at time t and a delayed payment at time $t + k$. The early payment is either today $t = 0$ or next year $t = 1$, and the late payment is delayed by either one year $k = 1$ or five years $k = 5$. Subjects receive an interest rate r on delayed payments, which varies between 0% to 800% interest on an annual basis. The allocations must be made such that their budget constraint is satisfied, i.e., the early payment and the present value of the delayed payment must equal the initial budget m. Early payments are certainly paid (i.e., payment probability 1), but delayed payments have a payment probability p_{t+k} of 0.5, 0.75 or 1.

Individuals make 20 consecutive CTB decisions between early and delayed payments. Our method consists of five different decision sets, and within each set we have four different interest rate scenarios. The first three choice sets use $t = 0$ and $k = 1$, and the three choice sets differ in the delayed payment probability $p_{t+k} = \{0.5, 0.75, 1\}$. The fourth choice set uses $t = 0$ and $k = 5$, and the fifth choice set uses $t = 1$ and $k = 5$, both sets having certain early and late payment probabilities. Table 4.10 in Appendix A presents an overview of our experimental design. In addition to randomizing the order of simultaneous and independent elicitation methods, we also randomize the order of presentation of the five CTB decision sets.

Differences between the early payment dates t (i.e., front-end delay) elicit present bias, while differences between the delayed payment dates $t+k$ (i.e., back-end delay) elicit long-term patience. Sensitivity to variation in the interest rates identifies curvature of the utility function, while variation in the payment probabilities enables the measurement of probability weighting. Figure 4.3 in Appendix A summarizes aggregate choice behavior in the CTB. We plot the mean and median allocated Euros at the early payment, c_t, against the gross interest rate, $(1 + r)$, for each of the five decision sets.

The amount of Euros allocated to the early payment declines monotonically with the interest, showing that our participants are willing to wait for a late payment when interest rates are higher. Additionally, the amount of Euros allocated to the early payment increases when the late payment probability is lower. Both observations reveal that choices respond to changing interest rates and payment probabilities in an intuitive predicted way. Evidence for strong present bias would be observed if the earlier allocated Euros are higher when $(t = 0, k = 5, p_{t+k} = 1)$ compared to $(t = 1, k = 5, p_{t+k} = 1)$. We observe that the early allocated budgets for these decision sets are roughly constant at each interest rate, indicating not too much evidence for strong present bias.

To estimate risk and time preferences, we identify the experimental allocated payments as solutions to standard intertemporal optimization problems. These solutions are supposed to be functions of our parameters of interest (present bias, discounting, risk aversion and probability weighting), and experimentally varied parameters (interest rates, delay lengths and payment probabilities). Given assumptions on the functional form of utility, the nature of discounting, and the nature of probability weighting, our experimental tasks provide a natural context to jointly estimate individual.

We assume that the agent has a standard CRRA utility function with curvature parameter γ, that the agent is a quasi-hyperbolic discounter with $\beta - \delta$ preferences, and that the agent distorts probabilities according to a simple Prelec weighting function with parameter η. We estimate time preferences, i.e., present-bias factor β and long-term discount factor δ, and we estimate risk preferences, i.e., risk aversion γ and probability weighting η. As such, we combine the CTB approaches of Andreoni and Sprenger (2012a) and Potters et al. (2016). Andreoni and Sprenger (2012a) estimate present bias, patience, and CRRA curvature, while Potters et al. (2016) estimate patience, CRRA curvature, and probability weighting. Additionally, we also control for background consumption through annual income in our estimations of preferences. Appendix D provides more details on the estimation.

Independent elicitation

In addition to the CTB method described above, we measure risk aversion, present bias and patience independently by using arguably simpler methods (Dave et al., 2010). We measure risk aversion with a lottery task as developed by Eckel and Grossman (2002) and Eckel and Grossman (2008). Table 4.11 in Appendix A shows the tasks.

The task involves a single choice among six gambles, all with probability 0.5 of winning a higher prize. The range of gambles includes a safe choice involving a sure payoff of €5600 with zero risk. Then, moving from Gamble 1 to 5, the gambles increase in both expected return and risk (standard deviation). Gamble 6 involves only an increase in risk, with an expected return equal to Gamble 5. More risk averse subjects choose low risk, low return gambles; risk-neutral subjects choose Gamble 5 or 6; risk-seeking subjects choose Gamble 6. This simple but coarser method only allows categorization of individuals into six risk categories, whereas the more complex but finer CTB allow categorization of individuals' risk aversion on a continuous scale.

We measure present bias and patience together, using a matching task from Rieger et al. (2015) and Wang (2017), which was originally inspired by Frederick (2005). Table 4.12 in Appendix A shows the task. The task asks participants to give an amount for a delayed payment which makes them indifferent with an immediate payment of €10,000. Participants give an amount €X_1 for a delayed payment of 1 year and an amount €X_5 for a delayed payment of 5 years. Assuming risk neutrality, the present-bias factor and discount factor can together be inferred from these two responses.[6]

Finally, we use three qualitative statements to measure financial risk-taking behavior, financial impulsiveness and financial patience. These statements proxy respectively for risk aversion, present bias and patience. Subjects answer the questions on a 7-point Likert scale from strongly disagree to strongly agree. The risk taking questions comes from the Dutch Central Bank Household Survey, and the impulsiveness and patience questions are taken from Gathergood (2012). We also ask the subjects a question about their trust in insurers. Table 4.12 in Appendix A shows the qualitative questions.

4.1.2 Measuring the disposition effect

In the second part of the survey, we select a random fraction of our total sample to participate in the trading experiment. Only those individuals that are head of the household and make the main financial household's decisions are eligible. The trading experiment is specifically designed to measure the disposition effect. The disposition effect is the well-known stylized fact that paper losses are realized less than paper gains (Odean, 1998). We implement the experimental method developed by Ploner (2017) to assess the existence of a disposition effect. For

[6]Assuming risk neutrality, the two matching questions can be represented as two equations with two unknowns: $10000 = \beta\delta X_1$ and $10000 = \beta\delta^5 X_5$, which can simply be solved for β and δ.

clarity, we explain Ploner's methodology now.

Subjects in the experiment make four investment choices regarding risky investment products. Per investment choice, the investors chooses between two investment products.[7] All investment choices are simple win/loss gambles with the same probability assigned to the win and the loss outcomes. The win and loss outcomes are both 50%, and determined by a coin toss (heads or tails). For investment choices 1. to 3., the magnitudes of the win and loss outcomes are symmetric between investment products within, but differ across investment choices. The loss outcome is always equal to -€4000, while the win outcome assumes the following values: +€3000, +€4000, and +€5000. It follows that choice 1 has a negative expected value, choice 2 is a fair prospect, and choice 3 has a positive expected value. For investment choice 4., the magnitudes of the win and loss outcomes are asymmetric between products. Product X has a win outcome (heads) of +€6000 and a loss outcome (tails) of -€5000, while product Y has a win outcome of +€4000 (tails) and a loss outcome (heads) of -€2000. Thus, product Y has a higher expected value (+€1000) than product X (+€500), while product X has also a higher standard deviation. See Table 4.13 in Appendix A for an overview.

The investors in our experiment are given an endowment of €10,000 for each of the four investment choices. In each choice, the investor must choose to invest the complete endowment in one product. In choice 1. the investor chooses to invest in product A or in product B, in choice 2. the investor chooses to invest in product C or product D, in choice 3. the investor chooses to invest in product E or product F, in choice 4. the investor chooses to invest in product X or product Y. The investors are aware that products A, C, and E warrant a win if the outcome of a coin toss is heads, and a loss otherwise. The opposite holds for products B, D, and F. As such, the assets are ex-ante identical and perfectly anti-correlated. The investors are ex-ante unaware of the sizes of the gains or losses for the products in investment choices 1. to 3. We do this because the investment products are identical and symmetric within an investment choice, and we are only interested in the selling behavior of investors after a gain or loss. In choice 4., the investor

[7]We deliberately use the neutral wording of 'product' in the experiment to avoid framing.

chooses to invest in product X or product Y. In this case, since the payoffs are asymmetric, the investors are aware of the sizes of the gains and losses when choosing between products X and Y.

After the investors have invested in their four chosen products, we toss a coin (virtually). Investors become aware of the outcome of the coin toss, and each investor must choose whether she wants to hold the product for one more year or whether she wants to sell the product immediately. If the investor sells her chosen product immediately, the earnings are immediately paid to the investor. If the investor holds her chosen product for one more year, we perform a second coin toss next year (hypothetically) and earnings are computed as in the first coin toss. As an example, consider investment choice 1. yielding a gain of +€3000 when the outcome is favorable: the investor can choose to receive €13,000 immediately, or she can choose to toss a second coin next year with potential equally likely earnings of €10,000 or €16,000 (known to the investor). The four investment choices are implemented as independent investments, and an overview of total earnings is shown after all sell and hold decisions are made.

Thus, in summary, the investors go through the following sequence of events. First, investors choose between two investment products per investment choice. Second, we toss a coin to define the investment returns, which are shown to the investors. Third, investors choose to hold the product for one more year or to sell the investment immediately. Fourth, a second coin toss is performed to define the investment returns for those who chose to hold on to their investment. Finally, all investors obtain a summary about the returns of their investments.

This setting yields a straightforward measure of the disposition. Simply compare the fraction of sell choices among those winning and those losing after the (first) coin toss. Evidence for a disposition effect would be found if investors sell their products more after gains than after losses (i.e., the sell rate after gains is higher). A loss, or gain, is by construction of the experimental design defined with respect to the initial purchase price of the asset, i.e., the complete invested endowment of €10,000 per asset.

4.1.3 Experimental procedure

Upon starting the online experiment, subjects read through the instructions for the CTB experiment. The instructions indicate that the budget should be allocated to an early payment or a later payment. The instructions state that there is no inflation. We also avoid arbitrage opportunities by stating that the allocated budget could be consumed or saved in a deposit account without interest, but could not be used to invest or to payoff a mortgage.[8]

Figure 4.4 in Appendix A shows a screenshot of a decision set in the experiment. Subjects are told to divide an amount of €10,000 between the early payment today (i.e., no front-end delay) and a late payment next year. The likelihood that the late payment is paid equals 100% in this particular decision screen. The subjects have to make four budget decisions presented in order of increasing interest rates from 1.00 to 4.50. In subsequent decision screens, the varying early and later payment dates are emphasized by underlining the dates, and probabilities of uncertain late payment were underlined as well. Subjects face a total of five such decision screen sets, such that they complete 20 decisions.

After the 20 CTB decisions, subjects answer the questions regarding the independent elicitation methods. They additionally answer questions about their estimated life expectancy and financial literacy.

Finally, the randomly selected household heads start with the disposition effect experiment by reading through the instructions. The investors read that they themselves are going to invest €10,000 each time in several products. These products yield a gain or loss, which depends on the outcome of toin cosses. The instructions state that they will invest three times: once in product A or B, once in product C or D, and once in product E or F. The instructions end by stating that for each investment a coin will be tossed. After these three choices and realizations of outcomes, the investors are told to invest once more €10,000 in product X or Y and that their chosen product can make a gain or loss, depending on a coin toss.

[8] See Andreoni and Sprenger (2012b) and Augenblick, Niederle, et al. (2015) for a detailed discussion on arbitrage opportunities in discounting studies.

4.1.4 Data

Here we first describe our sample. Then, we sketch the COVID-19 timeline and we introduce our main measure for studying time variation in preferences and trading behavior.

Sample

For the online experiment, we use the LISS (Longitudinal Internet Study in the Social Sciences) panel gathered by CentERdata in The Netherlands. The panel is recruited through address based sampling (no self-selection), and households without a computer and/or internet connection receive an internet connection and computer free of charge. This household panel, representative for the Dutch population, receives online questionnaires each month on different topics. When respondents complete a questionnaire, they receive a monthly incentive. Our experiment is not incentivized based according to the experimental answers of the subjects, which avoids the need for complex equalization of payments, transaction costs and payment confidence. Some researchers argue that answer-based incentives in economic experiments lead to more truthful reveal of preferences, however Cohen et al. (2020) in their overview study find little evidence for systematic differences between incentivized and unincentivized time preference experiments. More specifically, Potters et al. (2016) find little differences between financially incentivized and hypothetical decisions in their CTB experiments.

We invite a total of 2998 LISS panel members between the ages of 40 and 70 during March 2020 and December 2020. A total of 2631 panel members responded during both months, so we have a response rate of about 88% across both months. We drop 385 individuals because they have no reported income, they did not fully complete the CTB, or they made exactly the same allocations — being strictly corner solutions also — in each CTB decision. We require to have individual's income, because we use it as a proxy for background consumption in the estimation of the CTB preference parameters. Incomplete CTB experiments and individuals that did not alter their decision from specific corner solutions provide insufficient

Table 4.1: Summary statistics. *Partner* equals 1 if the participant lives together with a partner (married or unmarried). *Education low*, *Education medium* and *Education high* are education dummies, and the ordering is based on the categories of Statistics Netherlands. *Income* is individual monthly after-tax income. $1 - year\ life\ expectancy$ and $5 - year\ life\ expectancy$ are self-reported probabilities for reaching at least your current *age* plus 1 year and your current *age* plus 5 years, respectively. $\Delta Hosp$ is the daily percentage change in national hospitalizations.

	Mean	St. dev.	Min	Max	N
Male	0.48	0.50	0.00	1.00	2240
Age (years)	56.61	8.56	40.00	70.00	2240
Partner	0.71	0.45	0.00	1.00	2240
Education low	0.25	0.43	0.00	1.00	2240
Education medium	0.37	0.48	0.00	1.00	2240
Education high	0.38	0.48	0.00	1.00	2240
Income (€)	1884	1143	0.00	10000	2240
$\Delta Hosp$	16.02	36.49	-52.63	144.44	2240
1-year life expectancy	93.02	15.61	0.00	100.00	2239
5-year life expectancy	85.36	18.36	0.00	100.00	2239

variation for the calculation of preference parameters. Thus, similar to Andreoni, Kuhn, et al. (2015) we drop these observations. We additionally drop 6 individuals because we require to have data on gender, age, partner status, and education as these function as main controls in our analysis. Overall, we drop a total of 391 individuals such that our final sample contains 2240 subjects.

Table 4.1 presents descriptive statistics for the main characteristics of our sample. The male to female ratio is nearly 50%, and the average age is 56.61 years. The sample is roughly uniformly distributed across education levels, 38% has a degree from a higher vocational education or a university. The average individual monthly after-tax income is €1884. Participants on average estimate that they have 93% chance of surviving one more year and 85% chance of surviving the next five years. Respondents took on average 15 minutes to complete the survey. Using a 5-point Likert scale ('1 = definitely not' to '5 = definitely yes'), participants at the end of the survey answer the question "Did you find the questions clear?" at the median with a 4.0. This indicates that it was clear to the participants what was expected from them and that they understood their tasks.

COVID-19

The experiment took place in the period between 2 March 2020 and 31 March 2020, and between 7 December 2020 and 29 December 2020. We have 1997 observations during March, and 243 observations during December. The observations in December are lower, because we only invited the household heads that participated in the trading experiment in March as well. The average amount of respondents per day during March is 67 and during December 11. Figure 4.5 in Appendix B shows the number of daily observations throughout March and December.

Our first survey (i.e., in March 2020) took place during the emergence of the COVID-19 crisis in The Netherlands, the peak of global stock market crashes, and severe lockdown measures in The Netherlands. On March 1 The Netherlands had 0 deaths, 10 confirmed cases and no lockdown measures, while April 1 The Netherlands had 1.173 deaths, 13.614 contaminations, and a so-called intelligent lockdown. Our second survey (i.e., in December 2020) took place during the second lockdown in The Netherlands, and the total number of COVID-19 deaths was 11.843 at 31 December 2020.

To study how extreme events, in our case COVID-19, affect preferences and trading behavior, we use the daily percentage change in national COVID-19 hospitalizations in The Netherlands, $\Delta Hosp$, to proxy for the severity of the exogenous shock.[9] During March and December 2020, national hospitalizations were communicated daily on the news and via push notifications on phones. We download the national hospitalizations from the website of the National Institute for Public Health and the Environment, based on the Osiris database which uses the data reported by the Public Health Services (GGD). Until 16th December 2020 the official COVID-19 hospitalization numbers on the governmental corona-dashboard were based on the Osiris database. The hospitalization numbers include COVID-19 patients on the ICU and the nursing wards, and only those individuals that are in

[9]We do not use COVID-19 infected cases, because test capacity was absent during the first few months of the crisis. We do not use COVID-19 deaths nor the reproduction number, because these measure lag behind the actual COVID-19 situation; the reproduction number always has a lag of two weeks by definition and was based on the national COVID-19 hospitalizations till June 2020.

the hospital because of COVID-19 (i.e., excluding patients with additional reasons, besides COVID-19, for being in the hospital). To download the hospitalizations on a province level, we need the National Intensive Care Evaluation (NICE) reported numbers.[10]

If $\Delta Hosp > 0$, then the number of COVID-19 hospitalizations from day $t-1$ to day t is increasing. Table 4.1 shows that, on average, daily COVID-19 hospitalizations are increasing during March 2020 and December 2020. The minimum and maximum show that the daily percentage change in hospitalizations ranges between roughly -53% and +144%. Table 4.14 in Appendix B shows that on average the increase in COVID-19 hospitalizations is about 3 percentage points higher in March 2020 (i.e., Panel B) than December 2020 (i.e., Panel C).

We use hospitalizations rather than the number of infected individuals, because test capacity during especially the emergence of COVID-19 in March 2020 was absent or too constrained and, therefore, forms an imperfect measure of the severity of the crisis. Another possibility would be to use COVID-19 death rates, however deaths severe lag behind the actual situation and, therefore, is imperfect as well. We use hospitalizations because that was the most salient number during the COVID-19 crisis because citizens, the government, and institutions had the fear that hospitals would become overcrowded. We use the percentage change in hospitalizations rather than levels, since COVID expanded exponentially. In the beginning of March 2020 the absolute level of hospital admissions was low, but the impact of the large relative hospital changes was big. Besides, the daily percentage change in hospitalizations is temporarily uncorrelated, while levels violate the temporal uncorrelatedness.

4.1.5 Potential mechanisms

Before turning to the results, this part describes the potential mechanisms for the observed time variation in preferences and, then, their relation with trading

[10]These COVID-19 numbers include patients on the ICU and nursing wards for individuals being in the hospital for at least having COVID-19, including possibly other reasons for being in the hospital.

behavior.

Risk preferences

In the event of extreme negative shocks, such as COVID-19, people may experience fear concerning the uncertain future. Guiso et al. (2018) find that individuals become more risk averse through fear during an extreme negative exogenous shock (i.e., the financial crisis of 2007-2008 and a horror movie), and Cohn et al. (2015) find that investors primed with a financial bust are more fearful and, therefore, more risk averse. Additionally, Meier (2022) identifies fear as a significant correlate with within-person increases in risk aversion, and Haushofer and Fehr (2014) find that fear decreases the amount invested in a risky asset. Thus, individuals risk aversion can be altered through a fearful negative shock.

Even if individuals do not experience the shock directly themselves (i.e., get infected or hospitalized by COVID-19), fear can be activated by watching and reading news about the shock (Guiso et al., 2018). Our experiments took place during the emergence of COVID-19 and lockdowns, times in which fear could be born naturally as there were growing concerns about the uncertain future. National COVID-19 hospitalization numbers were communicated daily on the news and via push notifications on mobile phones. Hospitalizations were especially salient, because there was a general fear that hospitals might reach full capacity with amongst others a potential health crisis consequently.

Thus, as potential mechanism we hypothesize that risk aversion increases when the severity of the COVID-19 crisis increases, i.e., when national COVID-19 hospitalizations increase, as individuals become more fearful when seeing higher reported COVID-19 hospitalizations. Figure 4.1, left panel, shows that fear during the COVID-19 crisis in 2020 is higher than the years before and after the crisis in 2020. This provides suggestive evidence that fear could play a role in the potential mechanism.[11]

[11] Unfortunately, data on fear and uncertainty is only available on an annual basis, and unavailable on a daily basis.

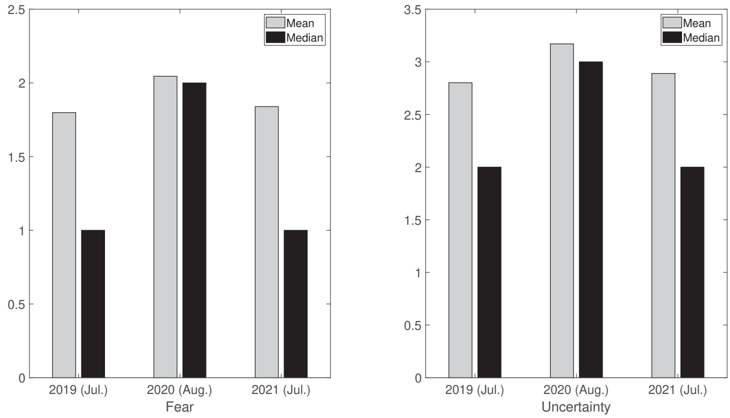

Figure 4.1: Fear and uncertainty during the crisis. The left panel shows mean and median self-reported fear on a 7-point Likert scale (0 = 'not at all' and 6 = 'extremely') the year before, during, and after the COVID-19 crisis. Similarly, the right panel shows uncertainty, proxied by self-reported distress.

Time preferences

The literature on the connection between time preferences and emotions is limited (Meier, 2019). Experimental settings, especially in an investment context, regarding fear and time preferences are to the best of our knowledge absent. However, there are studies that research the connection between precautionary savings and uncertain times (Carroll et al., 2019; Giavazzi and McMahon, 2012; Mody et al., 2012). Albeit on a micro- or macroeconomic level, these studies find that precautionary savings increase during uncertain times, such as the Great Depression of 2007-2008. Individuals want to insure themselves against bad states of the economy in the future and try to retain a smooth consumption path by precautionary savings and lower consumption in the current period. Thus, individuals saving behavior can be altered by uncertainty.

Our experiments took place during uncertain times, namely the emergence of COVID-19 and lockdowns. Our CTB experiment identifies time preferences directly through consumption and savings decisions. Individuals are elicited to be more patient when they delay consumption and increase savings in the current period. Thus, as a potential mechanism we hypothesize that patience increases when national COVID-19 hospitalizations increase, i.e., when uncertainty increases. Individuals increase their precautionary savings for the uncertain future and defer current consumption to retain a smooth consumption path. Figure 4.1, right panel, shows that uncertainty during the COVID-19 crisis in 2020 is higher than the years before and after the crisis in 2020. This provides suggestive evidence that uncertainty could play a role in the potential mechanism.

Trading behavior

Based on the hypotheses that risk aversion and patience increase as COVID-19 hospitalizations increase, we are also able to hypothesize about the effects on trading behavior. Because we elicit risk and time preferences only in the gain domain, we assume any effects in the loss domain to be constant (i.e., ceteris paribus).

The intertemporal realization utility model of Ingersoll and Jin (2013) makes two testable predictions.[12] First, a more patient investor is subject to a smaller disposition effect as she wants to realize gains later. Second, a more risk averse investor is subject to a smaller disposition effect as she wants to realize gains at a higher value than a less risk averse investor. Namely, a more risk averse investor is also willing to sell her assets at a larger loss, so her optimal selling point for gains is higher to outweigh the disutility of a loss compared to a less risk averse investor. Thus, contrary to static realization utility (i.e., the reflection effect), intertemporal realization utility with a concave utility function does not create a disposition effect, but actually reduces it.

We can test these two predictions in our trading experiment directly. Regarding the first prediction, we expect that a more impatient investor immediately sells her gains while more a more patient investor holds on to her gains for one more year. For a more patient investor, the marginal utility of a future gain is higher than a current gain. Regarding the second prediction, we expect that a more risk averse investor holds on to her gains at while a less risk averse investor sells her gains.

[12]The model of Barberis and Xiong (2012) is a special case with piecewise linear utility.

For a more risk averse investor, the marginal utility of a larger gain is higher than a smaller gain, compared to a more risk tolerant investor. Overall, for a more risk averse and more patient investor the marginal utility of a future gain is higher and the expected larger gain offsets the extra risk.

Thus, the hypotheses and their relations are as follows. When COVID-19 hospitalizations increase, we expect that risk aversion and patience increase, as a result of fear and uncertainty. Consequently, we expect that a more risk averse and more patient investor has a smaller disposition effect, as she holds on to her gains on an asset.

4.2 Results preferences

This section presents evidence that the estimated risk and time preferences are related to the severity of COVID-19 during the crisis by means of daily changes in national hospitalizations. Our results are robust to taking individual specific unobserved heterogeneity into account. Furthermore, we study whether preferences behave differently in the months during the crisis, and whether preferences react to COVID-19 hospitalizations on a province level. Finally, we present evidence that the less cognitive demanding, but coarser, independent elicited preferences do not show time variation.

4.2.1 Aggregated preferences

Table 4.2, Panel A, shows the estimated risk and time preferences for the population. The preferences are aggregated, meaning that for each subject we estimate the preferences by (4.10) and, then, we compute the 25$^{\text{th}}$, 50$^{\text{th}}$ and 75$^{\text{th}}$ percentiles of the population's distribution. Echoing the results in Figure 4.3, we do not find evidence for strong present bias at the median. We estimate a median present-bias factor β of 0.968. The median annual discount factor δ equals 0.958, which yields an annual discount rate of 4.3%.[13] The median CRRA risk aversion γ equals 1.520,

[13] In line with Potters et al. (2016), the discount factors are already yearly, so the annual discount rate follows from $1/\hat{\delta} - 1$.

implying that respondents behave risk averse. Finally, we estimate a probability weighting parameter η strictly larger than 1, which implies that respondents underweight probabilities larger than 50%. Overall, the 25^{th} and 75^{th} percentiles reveal heterogeneity in risk and time preferences.

Table 4.2: Aggregate risk and time preferences from Convex Time Budgets. Two-limit Tobit maximum likelihood estimates for quasi-hyperbolic β, δ discounting, CRRA utility $U(x) = \frac{x^{1-\gamma}}{1-\gamma}$ and Prelec-weighting function $\pi(p) = p^{\eta}$. Background consumption equals annual after-tax income, which varies across subjects.

	Median	25th Percentile	75th Percentile	N
Panel A: Aggregate				
Risk aversion $\hat{\gamma}$	1.520	1.117	2.061	2240
Present-bias factor $\hat{\beta}$	0.968	0.854	1.222	2240
Discount factor $\hat{\delta}$	0.958	0.876	1.038	2240
Annual discount rate	0.043	-0.037	0.141	2240
Probability weighting $\hat{\eta}$	1.063	-0.134	2.139	2240
Panel B: March				
Risk aversion $\hat{\gamma}$	1.515	1.112	2.065	1997
Present-bias factor $\hat{\beta}$	0.970	0.853	1.227	1997
Discount factor $\hat{\delta}$	0.959	0.875	1.039	1997
Annual discount rate	0.043	-0.038	0.143	1997
Probability weighting $\hat{\eta}$	1.030	-0.155	2.124	1997
Panel C: December				
Risk aversion $\hat{\gamma}$	1.538	1.159	2.026	243
Present-bias factor $\hat{\beta}$	0.961	0.869	1.204	243
Discount factor $\hat{\delta}$	0.956	0.888	1.022	243
Annual discount rate	0.046	-0.021	0.126	243
Probability weighting $\hat{\eta}$	1.321	0.036	2.275	243

Our estimated present-bias factor is similar to the value of Augenblick, Niederle, et al. (2015), who estimate $\beta = 0.97$ in the financial domain using a CTB design. Our estimated annual discount rate is similar to prevailing market interest rates and lower than most previous studies. Frederick (2005) show in their overview article that estimates of annual discount rates in the literature over hundred percent are not uncommon. Cheung (2020) estimates an annual discount rate of 62.2% in the discounted utility model, when controlling for CRRA curvature. Andreoni and Sprenger (2012a) estimate an annual discount rate of 27.5% in their CTB design, when controlling for CRRA curvature and present bias. Our estimated CRRA risk aversion parameter γ is similar to the value of Balakrishnan et al. (2020), they estimate a CRRA risk aversion parameter $\gamma = 1.4$ using also a CTB design with individually varying background income and two-limit Tobit regressions. Finally, our finding that subjects underweight higher probabilities is consistent with prospect theory (Tversky and Kahneman, 1992) and our estimate is in line with the estimated value of Potters et al. (2016) who also use a CTB design, late payment probabilities larger than 50% and a simple Prelec weighting function.

A potential reason for our lower, but plausible, annual discount rate is the magnitude of the experimental budget and the long-term decision horizon. Thaler (1981) already shows that discount rates drop sharply as the size of wealth increases, which is known as the magnitude effect, and he reports that discount rates drop sharply as the length of time increase. The experimental budget of €10,000 for our subjects, combined with the time delays of 5 years, are both larger than in most previous studies. Horizons are frequently used up to several weeks (Augenblick, Niederle, et al., 2015), 3 months (Tanaka et al., 2010), 6 months (Andersen, Harrison, M.Lauc, et al., 2010), 1 year (Dohmen, Falk, Huffman, and Sunde, 2010; Andersen, Harrison, Lau, et al., 2014), 2 years (Goda et al., 2015) and 3 years (Harrison, Lau, et al., 2002). A paper that comes close to ours in terms of large stakes and long decision horizons is Potters et al. (2016), who use a lower, but still relatively high, experimental budget of €1,000 with a decision horizon up to retirement age. They report an annual discount rate of 1%.

Panel B and Panel C in Table 4.2 show the estimated population's preferences in March 2020 (i.e., emergence of COVID-19 and first lockdown) and December 2020 (i.e., second lockdown), respectively. The population's preferences do not differ much between these two periods of the crisis. Table 4.15 in Appendix B displays the risk and time preferences for only the sample that participated in the trading experiment. Again the distribution of preferences is similar. Overall, by studying the preferences at an aggregated level during these two periods of the crisis, we find that the distribution of preferences remains similar based on a partial before-after analysis, which is in line with Bokern et al. (2021).[14]

Table 4.14, Panel A, in Appendix B shows the outcomes for the independent preference elicitation methods. The qualitative measures show moderate risk taking behavior and low impulsiveness, in line with our estimated risk aversion parameter and present-bias factor. The quantitative risk aversion measure (Eckel and Grossman, 2008) confirms that our respondents are risk averse. Finally, the quantitative measures for time preferences (Rieger et al., 2015; Wang, 2017), assuming linear utility, yield a somewhat higher present bias and a somewhat higher discount rate than in the CTB experiment. These findings indeed corroborate the claim that time preferences are upward biased if true utility is concave (Andersen, Harrison, Lau, et al., 2008b), which is true since our subjects are risk averse and, therefore, simultaneous estimation of risk and time preferences is preferred over independent measurements. Panels B and C show that on a population level the preferences are stable throughout March 2020 and December 2020.

4.2.2 Preferences during the crisis

Our main is result is shown in Figure 4.2. The graphs show the estimated risk and time preferences from the CTB experiment on a daily frequency throughout March 2020, i.e., the emergence of COVID-19 in The Netherlands and the first lockdown. From top to bottom, the panels displays the present-bias factor β, the discount factor δ, and the the risk aversion γ. Clearly, the preferences (solid lines) vary over time. Specifically, the preference parameters show a strong positive correlation with the daily percentage change in COVID-19 hospitalizations (dotted lines).[15] Individuals become more risk averse, less present biased, and more patient when COVID-19 hospitalizations rise. The correlation of the daily percentage change in COVID-19 hospitalizations (i) with risk aversion equals 0.36 (p-value < 0.05),

[14]We use the term partial before-after analysis, because the COVID measures and, thereby, the pandemic lasted actually to at least March 2022 in The Netherlands.

[15]To fit both the change in hospitalizations and the preference parameters in the graphs, we normalize both variables. Let $X = (X_1, X_2, ..., X_n)$ be a vector of n observations, then the normalized value equals $\frac{X_i - \min(X)}{\max(X) - \min(X)}$ such that both series in each graph lie in the interval of 0 and 1.

(ii) with the present-bias factor equals 0.53 (p-value < 0.01), and (iii) with the discount factor equals 0.51 (p-value < 0.01).

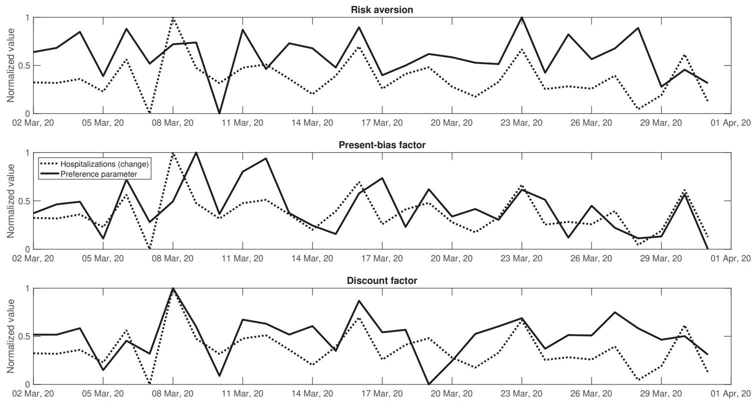

Figure 4.2: Risk and time preferences with COVID-19 hospitalizations during the crisis. Preferences (solid lines) are simultaneously estimated from Convex Time Budgets. Hospitalizations (dotted lines) are daily percentage changes in COVID-19 hospitalizations on a national level.

Figure 4.6 in Appendix B shows the estimated risk and time preferences from the CTB experiment on a daily frequency throughout December 2020, i.e., the second lockdown. There appears to be some time variation in preferences, but the correlation with the daily percentage change in hospitalizations is insignificant at all reasonable significance levels.[16] Changes in preferences during March 2020 might indeed have been stronger than during December 2020. The pandemic was completely new during March 2020 and most implications were unknown, but during December 2020 people got more used to the situation. Thus, implying stronger correlations between preferences and the COVID-19 crisis during March 2020 and less so during December 2020.

To formalize this suggestive evidence for time-varying preferences, we regress the estimated preferences on hospitalizations while controlling for multiple vari-

[16]The correlation of the daily percentage in COVID-19 hospitalizations (i) with the risk aversion equals 0.09 (p-value $= 0.68$), (ii) with the present-bias factor equals -0.16 (p-value $= 0.46$), and (iii) with the discount factor equals 0.00 (p-value $= 0.99$).

ables. Specifically, we analyze the effect of the percentage change in COVID-19 hospitalizations on individuals' preferences by estimating the following regressions equation:

$$y_{i,t} = a_0 + a_1 \Delta Hosp_{i,t} + bX_{i,t} + \varepsilon_{i,t}, \tag{4.1}$$

where observations occur at the individual (i) and day (t) level. $y_{i,t}$ is the estimated preference parameter (i.e., present-bias factor, discount factor, or risk aversion) for individual i measured at day t. $\Delta Hosp_{i,t}$ is the daily percentage change in national COVID-19 hospitalizations from day $t-1$ to day t. $X_{i,t}$ is a vector of control variables. The control variables compromise gender, age, having a partner, education, income, a dummy for answering the survey in December, and day fixed effects. We include day fixed effects to address potential concerns regarding daily seasonality in hospitalizations. We use quantile regressions to estimate the conditional median of the preference parameters, because it is more robust to extreme observations.

Our coefficient of main interest is a_1. Based on the suggestive evidence from Figure 4.2, we expect a_1 to be positive: if the number of COVID-19 hospitalizations is increasing, then the risk aversion parameter, the present-bias factor, and the discount factor increase. Individuals become less willing to take risks and want to save more. Because individuals' decisions are likely to be cross-sectionally correlated, we cluster standard errors at the individual level in all regressions.

Table 4.3 shows that our coefficient of interest, a_1, is positive and statistically significant for risk and time preferences. Hence, when national COVID-19 hospitalizations rise, individuals are more risk averse, less present biased and more patient. Specifically the change in time preferences is economically sizeable. A two-standard deviation increase in the change in COVID-19 hospitalizations (about 73%) leads to an increase of 0.11 in the risk-aversion parameter γ, 0.047 in the present-bias factor β, and 0.017 in the annual discount factor δ. Regarding the latter, this yields a decrease in the annual discount rate of 1.7%. Thus, individuals discount the future less and behave more patient, namely by 1.7 percentage points compared to a median annual discount rate of 4.3%.[17]

[17] A potential concern might be that the increased patience is not driven by precautionary

Table 4.3: Risk and time preferences during the crisis. This table reports all coefficients of the pooled median regressions $y_{i,t} = a_0 + a_1 \Delta Hosp_{i,t} + bX_{i,t} + \varepsilon_{i,t}$. $y_{i,t}$ represents the preference parameter for individual i at day t (per column): risk aversion γ, present-bias factor β, annual discount factor δ, 1-year and 5-year self-reported life expectancy. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Risk aversion	Present-bias factor	Discount factor	1-year life exp.	5-year life exp.
$\Delta Hosp$ ($\times$100)	0.152***	0.064***	0.023**	-0.932	-0.477
	(0.049)	(0.023)	(0.010)	(1.197)	(1.432)
December	-0.060	-0.018	-0.003	-2.346**	-2.743**
	(0.052)	(0.016)	(0.007)	(1.160)	(1.252)
Male	-0.072**	-0.007	-0.001	-0.609	-1.350
	(0.034)	(0.012)	(0.006)	(0.755)	(0.891)
Age (years)	0.010***	0.002***	0.000	0.050	-0.237***
	(0.002)	(0.001)	(0.000)	(0.040)	(0.046)
Partner	-0.116***	-0.009	-0.009	-0.511	0.548
	(0.033)	(0.012)	(0.006)	(0.760)	(0.890)
Edu. medium	-0.103**	-0.010	-0.017**	2.425**	2.677**
	(0.044)	(0.019)	(0.008)	(1.062)	(1.181)
Edu. high	-0.242***	-0.039**	-0.030***	4.116***	4.840***
	(0.043)	(0.018)	(0.008)	(1.053)	(1.201)
Income ($\times$1000)	0.396***	-0.014***	0.008***	0.618*	0.894**
	(0.019)	(0.004)	(0.002)	(0.327)	(0.386)
Constant	0.365***	0.897***	0.963***	88.031***	95.784***
	(0.116)	(0.041)	(0.018)	(2.680)	(3.067)
Observations	2240	2240	2240	2239	2239
Day FE	Yes	Yes	Yes	Yes	Yes

Table 4.16 in Appendix C corroborates this finding: we regress the experimentally allocated amounts to the late payments on changes in COVID-19 hospitalizations, and we find that the experimentally allocated amounts to the late payments increase the median (mean) by €1759 (€971) when COVID-19 hospitalizations rise.

savings, but rather by the inability of households to spend income during the COVID crisis as a result of shop colsures and lockdowns. However, we find that risk and time preferences react similarly to $\Delta Hosp$ during the beginning of March 2020, where everything was still 'normal', and the end of March 2020, where shops were closed as a result of an intelligent lockdown. Hence, this alleviates such concerns.

Our results are not driven by simultaneous changes in beliefs regarding individual's life expectancy, as shown in the last two columns of Table 4.3. There is no statistically significant correlation between self-reported life expectancy probabilities and changes in COVID-19 hospitalizations. Table 4.17 in Appendix C shows that our results are robust to using OLS rather than median regressions, using monthly background income w in the estimated preference parameters rather than annual background income, different sets of controls (i.e., only demographic variables, life expectancy, and financial literacy), and the unbalanced panel test of Verbeek and Nijman (1992).

Additionally, Table 4.18 in Appendix C shows that our results remain similar for the group of subjects with an equal salary during February, March, and December 2020. This indicates that the changes in preferences are not driven by changes in income. Finally, Table 4.19 in Appendix C shows that our results are similar when using national COVID-19 *ICU* hospitalizations as main independent variable.[18] Although COVID-19 ICU hospitalizations were less salient than COVID-19 hospitalizations during the emergence the crisis, individuals' preferences react similarly.

Though not the main focus of the paper, Table 4.3 shows that education and income consistently affect preferences and beliefs. Higher educated individuals are less risk averse, more present biased and less patient. Income is correlated with preferences, which might not be too surprising since we use income to proxy for background consumption in the estimation of the preference parameters. Additionally, higher education and higher income yield higher beliefs regarding the individuals' 1-year and 5-year life expectancy probabilities. Finally, we find that risk aversion increases with age — consistent with Schildberg-Hörisch (2018) — the present-bias factor increases with age, and females are more risk averse than males — consistent with Eckel and Grossman (2002) and Eckel and Grossman (2008).

[18]The national COVID-19 ICU hospitalizations are downloaded from the website of the National Institute for Public Health and the Environment, and they are based on the NICE reported numbers.

Alternative specifications

Model (4.1) assumes the error term to be independent of the explanatory variables. For example, it is assumed that individuals who fill in the questionnaire on a day in which hospital admissions increase are comparable to individuals who fill in the questionnaire on a day in which hospital admissions decline. To test for this, we separate within and between effects below. The specification in equation (4.2) relaxes the assumption. Analyzing the within effect allows individual specific unobserved heterogeneity to be correlated with the explanatory variables. We are particularly interested in our main explanatory variable of interest $\Delta Hosp_{i,t}$. We estimate the between and within effects for each preference parameter using the approach of Allison (2009). Specifically, we estimate the following regression equation

$$y_{i,t} = \tilde{a}_0 + \tilde{a}_1(\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}) + \tilde{a}_2\overline{\Delta Hosp_i} + \tilde{b}X_{i,t} + \tilde{\varepsilon}_{i,t}. \tag{4.2}$$

The between effect is given by $\overline{\Delta Hosp_i} = n_i^{-1}\sum_{t=1}^{n_i} Hosp_{i,t}$ and the within effect is given by $\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}$. In this model, $\tilde{a}_1$ is the within effect (comparable to a fixed-effects estimator) and $\tilde{a}_2$ is the between effect (Mundlak, 1978). An advantage of this approach is that it allows us to test for the equivalence of within and between estimates using a Wald test. If between and within effects are the same, then it should hold under the null hypothesis that $\tilde{a}_1 = \tilde{a}_2$. Again, we use quantile regressions to estimate the conditional medians of the preference parameters as dependent variables.

Table 4.4, Panel A, shows the results of the estimated hybrid model. If we compare the between and within estimates for each preference, then we firstly observe that the between and within effects are very similar, although standard errors are larger for within effects. Secondly the estimated coefficients are almost identical to the estimates from model (4.1), as shown in Table 4.3. The Wald test suggests strongly that we can not reject the null hypothesis of equality for between and within estimates. This suggests that individual specific unobserved variables are not correlated with hospitalizations. This is in favor of model (4.1). Furthermore,

Table 4.4: Alternative specifications. This table reports the coefficients of median regressions from within and between analyses among individuals, between months, and at a province level. Panel A reports the estimated coefficients $\tilde{a}_1$ and $\tilde{a}_2$ from $y_{i,t} = \tilde{a}_0 + \tilde{a}_1(\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}) + \tilde{a}_2\overline{\Delta Hosp_i} + \tilde{b}X_{i,t} + \tilde{\varepsilon}_{i,t}$. Panel B reports the estimated coefficients $\tilde{a}_1$, $\tilde{a}_2$ and $\tilde{a}_3$ from $y_{i,t} = \tilde{a}_0 + \tilde{a}_1(\Delta Hosp_{i,t} \times Dec) + \tilde{a}_2\Delta Hosp_{i,t} + \tilde{a}_3 Dec + \tilde{b}X_{i,t} + \tilde{\varepsilon}_{i,t}$. Panel C reports the estimated coefficients $\tilde{a}_1$ from $y_{i,t,p} = \tilde{a}_0 + \tilde{a}_1\Delta Hosp_{t,p} + \tilde{b}X_{i,t} + \tilde{\varepsilon}_{i,t,p}$, in which $\Delta Hosp_{t,p}$ is measured on a province level. Controls $X_{i,t}$ include *December*, *Male*, *Age*, *Partner*, *Edu. medium*, *Edu. high*, and *Income*. The Wald tests show the null hypotheses between parentheses. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Risk aversion	Present-bias factor	Discount factor	1-year life exp.	5-year life exp.
Panel A: Within and between					
$\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}$ ($\times$100)	0.162*	0.058	0.022	-0.498	0.055
	(0.098)	(0.041)	(0.016)	(2.420)	(2.685)
$\overline{\Delta Hosp_i}$ ($\times$100)	0.152***	0.063***	0.023**	-0.990	-0.548
	(0.055)	(0.023)	(0.010)	(1.268)	(1.514)
Observations	2240	2240	2240	2239	2239
Controls	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes
Wald test ($\tilde{a}_1 = \tilde{a}_2$)	0.9286	0.8887	0.9506	0.8468	0.8299
Panel B: March and December					
$\Delta Hosp \times Dec$ ($\times$100)	0.029	-0.071	-0.021	-3.893	1.120
	(0.163)	(0.050)	(0.018)	(3.526)	(4.087)
$\Delta Hosp$ ($\times$100)	0.152***	0.079***	0.024**	-0.640	-0.561
	(0.049)	(0.026)	(0.010)	(1.188)	(1.441)
December	-0.077	-0.001	0.002	-1.829	-2.892**
	(0.056)	(0.021)	(0.007)	(1.187)	(1.384)
Observations	2240	2240	2240	2239	2239
Controls	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes
Wald test ($\tilde{a}_1 = \tilde{a}_3 = 0$)	0.3789	0.1934	0.4970	0.0914	0.0895
Panel C: Province					
$\Delta Hosp_p$ ($\times$100)	0.049*	0.003	0.004	-0.694	-0.353
	(0.025)	(0.008)	(0.003)	(0.520)	(0.641)
Observations	2230	2230	2230	2229	2229
Controls	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes	Yes

self-reported life expectancies are not particularly affected by between or within effects.

Additionally, one might wonder whether the time variation in preferences is different between March 2020 (i.e., emergence COVID-19 and first lockdown) and December 2020 (i.e., second lockdown). Panel B, Table 4.4, shows that the time variation in preferences during December 2020 is not different from March 2020, since the coefficient for the interaction between hospitalizations and December is statistically insignificant.

Finally, we test whether preferences react to news regarding COVID-19 hospitalizations on a province level, rather than using news regarding COVID-19 hospitalizations on a national level.[19] We hypothesize that province level COVID-19 hospitalizations matter less for time-varying preferences as the reported hospitalizations on TV, internet, and smartphones were based on national COVID-19 hospitalizations. To explore this, we regress preferences on province level COVID-19 hospitalizations $\Delta Hosp_p$ and we control for province fixed effects. Indeed, COVID-19 hospitalizations on a province level are unrelated to the time variation in preferences (and neither for life expectancy). Thus, this provides additional suggestive evidence that the saliently reported national news induces time variation in risk and time preferences.

4.2.3 Independent measures

In addition to the simultaneously estimated preference parameters from the more cognitive demanding, but finer, convex time budgets we measured risk and time preferences independently using less cognitively demanding but coarser elicitation methods. Panels A and B in Table 4.5 show that the simpler qualitative and quantitative preference measures are unable to capture the time variation in preferences. National COVID-19 hospitalizations do not correlate with these preference measures. A potential reason can found in the study of Dave et al. (2010): the cognitive simpler tasks are too coarse to capture time variation, while the cognitive more complex measure has overall superior accuracy as the measure is finer (potentially at the cost of more noisy behavior).

Table 4.20 in Appendix C reveals that within and between effects are similar, but remain insignificant. We also observe no differences between March and

[19] A province in The Netherlands is very similar to a county in the U.S.

December 2020, and province level hospitalizations have no consistent effect on preferences as well. Only trust in insurers seems to vary over time. Specifically, individuals indicate to have more trust in insurers during December 2020 compared to March 2020. Correlations between the CTB measure, qualitative measures, and quantitative measures are small and weak, as shown by Table 4.21 in Appendix C.

4.3 Results trading behavior

First, we confirm the existence of an aggregate disposition effect in our sample. Then, we present evidence that the disposition effect is time varying and correlates negatively with COVID-19 hospitalizations. Investors are less likely to realize gains during rising COVID-19 hospitalizations. Finally, we test the potential mechanisms for the relation between preferences and the disposition effect.

4.3.1 Aggregate findings

Investors make a sell and hold decision for four investment products after they have observed their losses or gains on each product. If investors sell a product, then they realize the trading outcome immediately and if investors hold a product, then they hold the product for another year and realize the trading outcome next year. The outcome of a loss or gain is determined by a fair coin toss and, thus, should be around 50% in our sample. 537 individuals participated in the trading experiment aggregated over March 2020 (i.e., emergence COVID-19 and first lockdown) and December 2020 (i.e., second lockdown). We lose five individuals because they did not fully complete the disposition effect experiment, so we have a total of 2128 trading observations for 532 unique individuals.

Table 4.6 shows that 1040 gains (i.e., 49% of the total trading decisions) and 1088 losses (i.e., 51% of the total trading decisions) occurred in our sample. The table provides suggestive evidence for a disposition effect. We want to test whether individuals exhibit a higher tendency to sell assets that are at a gain rather than those that are at a loss, or vice versa holding losses more than holding gains. Indeed, the conditional fraction of products sold for a gain is 76%, while losses

Table 4.5: Preferences according to simpler measures during the crisis. This table reports the estimated coefficients $\bar{a}_1$ and $\bar{a}_2$ of the pooled OLS regressions $\bar{y}_{i,t} = \bar{a}_0 + \bar{a}_1 \Delta Hosp_{i,t} + \bar{a}_2 Dec + \bar{b}X_{i,t} + \bar{\varepsilon}_{i,t}$. Panel A shows the qualitatively-measured binary dependent variable $\bar{y}_{i,t}$ (per column): risk taking, impulsiveness, impatience and trust. Panel B shows the quantitatively-measured dependent variable $\bar{y}_{i,t}$ (per column): risk aversion (Eckel and Grossman, 2008), present-bias factor $\hat{\beta}$ and annual discount factor $\hat{\delta}$ assuming risk neutrality (Rieger et al., 2015; Wang, 2017). Controls $X_{i,t}$ include $Male$, Age, $Partner$, $Edu.\ medium$, $Edu.\ high$, and $Income$. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Panel A: Qualitative measures			
	Risk taking	Impulsiveness	Impatience	Trust
$\Delta Hosp$ ($\times 100$)	-0.035	-0.024	-0.061*	-0.027
	(0.033)	(0.018)	(0.032)	(0.030)
December	0.034	-0.000	-0.008	0.074**
	(0.031)	(0.016)	(0.028)	(0.030)
Constant	0.222***	0.275***	0.272***	0.119
	(0.083)	(0.047)	(0.074)	(0.075)
Observations	2154	2210	2203	2172
Controls	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes
	Panel B: Quantitative measures			
	Risk aversion (EG)	Present-bias factor (RN)	Discount factor (RN)	
$\Delta Hosp$ ($\times 100$)	-0.084	-0.006	-0.002	
	(0.107)	(0.018)	(0.007)	
December	-0.069	0.015	0.002	
	(0.094)	(0.013)	(0.006)	
Constant	2.294***	0.790***	0.846***	
	(0.252)	(0.041)	(0.017)	
Observations	2240	1764	1764	
Controls	Yes	Yes	Yes	
Day FE	Yes	Yes	Yes	

Table 4.6: Sell and hold rates after gains and losses.

	Sell	Hold	N
Unconditional			
Gain	37%	12%	2128
Loss	26%	25%	2128
Conditional			
Gain	76%	24%	1040
Loss	51%	49%	1088

are only realized for 51%. Likewise, investors hold on to losses more often than holding on to gains. On average, investors are more reluctant to hold products (i.e., 37%) than to sell products (i.e., 63%).

Odean (1998) proposes a proportion-based measure to calculate the disposition, but thereby possibly neglects other variables affecting investor's trading behavior, as is also the case for our suggestive evidence. Thus, we follow the regression technique approach of Chang et al. (2016) and we estimate the regressions equation

$$Sell_{i,t} = d_0 + d_1 Gain_{i,t} + f X_{i,t} + e_{i,t}, \tag{4.3}$$

where observations occur at individual level i and day t. The dependent variable *Sell* is a dummy variable that equals one if the investor sells the product, and zero if the investor holds the product for one more year. *Gain* is a dummy variable that is equal to one if the investor experienced a gain after the first coin toss, and zero if the investor suffered a loss. X is a vector of controls, similar to the set of controls in regression equation (4.1) for time-varying preferences. Additionally, we now consider asset fixed effects. These capture the expected value and volatility (i.e., standard deviation) of each asset.

Table 4.7 confirms that those experiencing a gain are more likely to sell the investment than those experiencing a loss. Specifically, column (1) shows that investors are 25% more likely to sell a gain compared to a loss. Conditional on suffering a loss, investors sell directly 51% of their investments, while conditional on experiencing a gain, investors sell directly 76% of their investments. Columns (1) to (4) confirm a clean randomisation as the coefficient for individuals experiencing a gain remains similar throughout the different specifications. As shown by column (4), the selling behavior of assets is heterogeneous across levels of the population, because the decision to sell is uncorrelated to socio-demographic variables.

However, socio-demographic variables may have an asymmetric effect on assets trading at a gain or loss, i.e., the disposition effect. In Table 4.22 in Appendix C, we redo the regressions for the sample of assets trading at a gain and those trading at a loss. Interestingly, we observe that age correlates positively with selling gains and income correlates negatively with selling gains, however socio-demographic variables are heterogeneous in the domain of losses. Stated differently, older investors are more likely to sell assets after a gain (i.e., larger disposition effect), and wealthy investors are less likely to sell assets after a gain (i.e., smaller disposition effect).

4.3.2 Disposition effect during the crisis

We want to study whether COVID-19 hospitalizations influence the trading behavior of investors. This provides insights in the time variation of the disposition effect. Similar to our analyses above, we regress the decision to sell an investment immediately on national COVID-19 hospitalizations while controlling for multiple variables:

$$Sell_{i,t} = \tilde{d}_0 + \tilde{d}_1 \Delta Hosp_{i,t} + \tilde{f} X_{i,t} + \tilde{e}_{i,t}. \tag{4.4}$$

Since the effects on trading might be asymmetric, we perform the analysis for the assets trading at a gain and for those trading at a loss. Our coefficient of interest is $\tilde{d}_1$.

Table 4.8 shows that the change in COVID-19 hospitalizations affects the selling of gains (Panel A), but not the selling of losses (Panel B). From Panel A we conclude that an increase in COVID-19 hospitalizations decreases the assets sold after a gain. In other words, if COVID-19 hospitalizations increase, then the

Table 4.7: Aggregate disposition effect. This table reports all coefficients of the panel OLS regressions $Sell_{i,t} = d_0 + d_1 Gain_{i,t} + fX_{i,t} + e_{i,t}$. The dependent variable $Sell_{i,t}$ is a dummy variable equal to one if the investor sells an asset. $Gain_{i,t}$ is a dummy variable equal to one when the investor experienced a gain after the toin coss, and equal to zero otherwise. Columns (1) - (4) use different sets of control variables. Demographics include Dec, $Male$, Age, and Controls additionally include $Partner$, $Edu.\ medium$, $Edu.\ high$, and $Income$. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Dependent variable: Sell	(1)	(2)	(3)	(4)
Gain	0.251***	0.248***	0.247***	0.246***
	(0.028)	(0.028)	(0.028)	(0.027)
December		-0.019	-0.020	-0.019
		(0.025)	(0.025)	(0.026)
Male		-0.047	-0.036	-0.035
		(0.029)	(0.033)	(0.034)
Age (years)		0.000	0.000	0.000
		(0.002)	(0.002)	(0.002)
Partner			0.029	0.029
			(0.031)	(0.031)
Edu. medium			0.018	0.017
			(0.038)	(0.038)
Edu. high			0.002	0.002
			(0.041)	(0.041)
Income ($\times$1000)			-0.011	-0.012
			(0.017)	(0.017)
Constant	0.513***	10.083	9.266	9.260
	(0.022)	(41.259)	(41.186)	(41.095)
Observations	2128	2128	2128	2128
Demographics	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
Asset FE	No	No	No	Yes
Day FE	No	Yes	Yes	Yes

disposition effect decreases. A two-standard deviation increase in the change in COVID-19 hospitalizations (about 75%) yields a decrease of -10.58% in selling assets after a gain. The mechanism is through a higher reluctance for selling gains rather than through the domain of losses, as investors prefer to hold their winning assets for one more year.

4.3.3 Preferences and disposition effect

This section studies the relation between the elicited CTB preferences and the disposition effect. We are interested in the relations of risk and time preferences with the selling behavior of investors as predicted by the intertemporal realization utility model.

In Table 4.9 we report the median risk aversion, present-bias factor, and discount factor for investors that sell their assets after a gain and a loss. The sample in Panel A includes all observations from March 2020 and December 2020, and Panel B includes the observations from March 2020 only. Investors that hold their asset after experiencing a gain are more risk averse and more patient. The difference with selling after a gain is statistically significant at any reasonable significance level, as indicated by the Wilcoxon rank-sum test p-values, for the aggregated sample in Panel A and, to a somewhat less extent, for the March sample in Panel B. Consistent with the absence of time variation in selling losses, and the fact that only preferences are elicited in the gain domain, we do not observe any relationship between preferences and selling losses.

Our findings are in line with the potential mechanisms regarding intertemporal realization utility (Ingersoll and Jin, 2013). We have documented that an increase in the change of COVID-19 hospitalizations induces more risk aversion and more patience, and an increase in the change of COVID-19 hospitalizations yields a smaller disposition effect as investors are less likely to sell their gains. Now, we consistently find that indeed more risk averse, $\gamma = 1.67$, and more patient, $\delta = 0.97$, investors are subject to a smaller disposition effect as they are more likely to hold their gains.

A potential concern might be that our disposition effect results are driven by another explanation than preferences. For example, the gambler's fallacy might alter the beliefs of an investor about the probability of a certain draw from a fair coin toss. To address this concern, we show participants in our survey the history of two independent fair coins that have been tossed five times. One series showing

Table 4.8: Disposition effect during the crisis. This table reports the results of the panel OLS regressions $Sell_{i,t} = \tilde{d}_0 + \tilde{d}_1 \Delta Hosp_{i,t} + \tilde{f} X_{i,t} + \tilde{e}_{i,t}$. Panel A contains all assets trading at a gain (i.e., $Gain_{i,t} = 1$), and Panel B contains all assets trading at a loss. The dependent variable $Sell_{i,t}$ is a dummy variable equal to one if the investor sells an asset. Columns (1) - (4) use different sets of control variables. Demographics include *Dec*, *Male*, *Age*, and Controls additionally include *Partner*, *Edu. medium*, *Edu. high*, and *Income*. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Dependent variable: Sell	(1)	(2)	(3)	(4)
Panel A: Gains				
$\Delta Hosp$ ($\times 100$)		-0.132**	-0.148**	-0.145**
		(0.061)	(0.062)	(0.062)
December		-0.057*	-0.059*	-0.058*
		(0.032)	(0.032)	(0.032)
Constant	.763***	110.154**	92.021*	89.601*
	(.018)	(49.739)	(50.620)	(50.578)
Observations	1040	1040	1040	1040
Demographics	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
Asset FE	No	No	No	Yes
Day FE	No	Yes	Yes	Yes
Panel B: Losses				
$\Delta Hosp$ ($\times 100$)		0.083	0.082	0.082
		(0.078)	(0.078)	(0.078)
December		0.008	0.006	0.006
		(0.036)	(0.036)	(0.036)
Constant	.513***	-88.393	-79.390	-74.513
	(.022)	(59.206)	(59.133)	(58.873)
Observations	1088	1088	1088	1088
Demographics	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
Asset FE	No	No	No	Yes
Day FE	No	Yes	Yes	Yes

Table 4.9: Preferences and disposition effect. This table reports the medians of the preference parameters if an investor holds an asset after a gain and sells an asset after a gain. The column 'Difference' reports the difference in preference parameters for holding and selling a gain. Wilcoxon rank-sum test p-values are reported between parentheses. Likewise, the table reports the values of the preferences parameters for holding or selling after a loss. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Aggregated

	Gain			Loss		
	Hold	Sell	Difference (p-value)	Hold	Sell	Difference (p-value)
Risk aversion	1.67	1.54	0.13	1.61	1.56	0.05
			(0.00***)			(0.35)
Present-bias factor	0.96	0.96	0.00	0.95	0.96	-0.02
			(0.24)			(0.50)
Discount factor	0.97	0.95	0.02	0.96	0.95	0.01
			(0.01***)			(0.16)
Observations	246	794	1040	530	558	1088

Panel B: March

	Gain			Loss		
	Hold	Sell	Difference (p-value)	Hold	Sell	Difference (p-value)
Risk aversion	1.68	1.55	0.12	1.62	1.63	0.00
			(0.01**)			(0.82)
Present-bias factor	0.98	0.95	0.02	0.94	0.97	-0.03
			(0.69)			(0.90)
Discount factor	0.97	0.96	0.01	0.96	0.94	0.02
			(0.07*)			(0.22)
Observations	131	452	583	279	298	577

five times the consecutive outcome heads, and the other series shows heads and tails in an alternating way with the fifth outcome being heads. We consequently ask the probability of tails in the sixth coin toss. We construct a new dummy variable *Correct Probability* that equals one when the participant answered 50% to each question.

Tables 4.23 and 4.24 in Appendix C show that our results are robust when considering this variable in our analysis. First, Panel A in Table 4.23 shows that the aggregate disposition effect remains substantial and Panel B in Table 4.23 shows that the disposition effect remains time varying. The variable *Correct Probability* and its interaction with the main explanatory variables of interest are statistically insignificant in the analyses. Table 4.24 displays the relation between preferences and trading behavior *only* for the group of investors not subject to the gambler's fallacy (i.e., *Correct Probability* equals one). The magnitudes and statistical significance remain similar to our above findings. Overall, this indicates that our results are not driven by the gambler's fallacy, as our results remain similar for the group of investors correctly estimating the objective probabilities of a coin toss.

4.4 Conclusion

Typically, individuals' preferences are assumed to be stable and persistent over time (Stigler and Becker, 1977). However, we show that preferences and economic outcomes are related to the severity of COVID-19 during the crisis. Furthermore, changes in economic outcomes are related to changes in preferences.

We elicit and estimate risk and time preferences during the COVID-19 crisis, and we find a strong correlation between preferences and daily changes in the national COVID-19 hospitalizations. In particular, we show that risk aversion, the present-bias factor, and the discount factor correlate positively with daily changes in national COVID-19 hospitalizations. Individuals become more risk averse, more time consistent, and more patient when hospitalizations increase. Individuals decrease their long-term annual discount rate from 4.3% to 2.6% when COVID-19 hospitalizations increase by two standard deviations. Extensive robustness checks confirm our findings. Preferences measured by the CTB are only weakly correlated with cognitive simpler measures, which show no association with COVID-19 hospitalizations.

A potential mechanism for the time variation in risk preferences is fear, while a potential mechanism for the time variation in time preferences is precautionary savings borne by uncertainty. At the same time, the disposition effect declines when COVID-19 hospitalizations increase. We observe that investors are hold assets that experienced a gain, while we find no effects in the loss domain. We find that investors who are subject to a smaller disposition effect are more risk averse and more patient. This finding is in line with intertemporal realization utility (Ingersoll and Jin, 2013), which provides evidence that at least part of the time variation in the disposition effect is driven by risk and time preferences. Present bias plays no significant role in the trading behavior of investors.

Overall, our findings cast doubt on the perfect stability of preferences and the disposition effect during a crisis. Our results support studies arguing that individuals' preferences and selling behavior vary with negative exogenous shocks. By linking preferences to trading behavior, we highlight the importance of preferences for financial-economic decision making. Changes in the stability of preferences and, consequently, economic outcomes have vital real-world consequences for policy making and welfare analyses.

4.5 Appendix

A. Experimental design

Table 4.10: Overview experimental design: Convex Time Budgets. Choice sets in the Convex Time Budgets. t and k are front and end delays in years, and c_t and c_{t+k} are allocated amounts in Euros. $1+r$ is the implied gross interest rate. Annual r is the yearly interest rate in percent and calculated as $((1+r)^{1/k}-1)\times 100$. r' is the interest rate adjusted for the late payment probability p_{t+k}.

						Interest			Risk adjusted interest	
Decision	Set	t	k	c_t	c_{t+k}	$1+r$	Annual r	p_{t+k}	$1+r'$	Annual r'
1	1	0	1	10,000	10,000	1	0	1	1	0
2	1	0	1	10,000	15,000	1.5	50	1	1.5	50
3	1	0	1	10,000	25,000	2.5	150	1	2.5	150
4	1	0	1	10,000	45,000	4.5	350	1	4.5	350
5	2	0	1	10,000	20,000	2	100	0.5	1	0
6	2	0	1	10,000	30,000	3	200	0.5	1.5	50
7	2	0	1	10,000	50,000	5	400	0.5	2.5	150
8	2	0	1	10,000	90,000	9	800	0.5	4.5	350
9	3	0	1	10,000	13,300	1.33	33.33	0.75	1	0
10	3	0	1	10,000	20,000	2	100	0.75	1.5	50
11	3	0	1	10,000	33,300	3.33	233.33	0.75	2.5	150
12	3	0	1	10,000	60,000	6	500	0.75	4.5	350
13	4	0	5	10,000	10,000	1	0	1	1	0
14	4	0	5	10,000	15,000	1.5	8.45	1	1.5	8.45
15	4	0	5	10,000	45,000	4.5	35.1	1	4.5	35.1
16	4	0	5	10,000	85,000	8.5	53.42	1	8.5	53.42
17	5	1	5	10,000	10,000	1	0	1	1	0
18	5	1	5	10,000	15,000	1.5	8.45	1	1.5	8.45
19	5	1	5	10,000	45,000	4.5	35.1	1	4.5	35.1
20	5	1	5	10,000	80,000	8	51.57	1	8	51.57

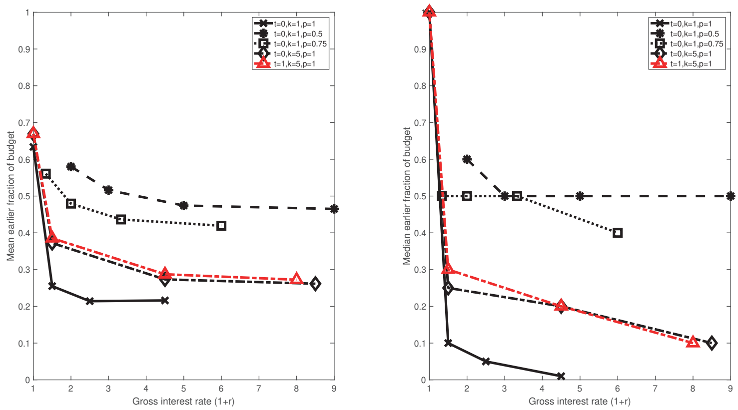

Figure 4.3: Mean and median experimental allocated payments in the Convex Time Budgets.

Table 4.11: Eckel-Grossman risk aversion task. Subjects choose which gamble to play, all of which involve a 50/50 chance of a low or high payoff. The implied Coefficient of Relative Risk Aversion (CRRA) range is based on the power utility function $U(x) = \frac{x^{1-\gamma}}{1-\gamma}$. Each range is calculated by equalizing the gamble to its neighbors, and computing the value of γ that makes the individual indifferent in utility between each adjacent gamble.

Choice	Low payoff	High payoff	Exp. return	St. Dev.	Implied CRRA range
Gamble 1	5600	5600	5600	0	$\gamma > 3.46$
Gamble 2	4800	7200	6000	1200	$1.16 < \gamma < 3.46$
Gamble 3	4000	8800	6400	2400	$0.71 < \gamma < 1.16$
Gamble 4	3200	10400	6800	3600	$0.50 < \gamma < 0.71$
Gamble 5	2400	12000	7200	4800	$0 < \gamma < 0.50$
Gamble 6	400	14000	7200	7000	$\gamma < 0$

Table 4.12: Time preferences (risk neutrality) and qualitative statements. Panel A shows the matching task of Rieger et al. (2015) and Wang (2017), in which subjects fill in the amount X. Panel B displays the qualitative statements, and subjects answer the questions on a 7-point Likert scale.

Panel A: Time preferences (under the assumption of risk neutrality)
Assume for this question that prices in the future remain equal to the prices today (no inflation).
Fill in an amount X_1 such that option B is as attractive as option A.
A. Receive €10,000 now
B. Receive $\mathbf{X_1}$ over 1 year
Assume for this question that prices in the future remain equal to the prices today (no inflation).
Fill in an amount X_5 such that option B is as attractive as option A.
A. Receive €10,000 now
B. Receive $\mathbf{X_5}$ over 5 years
Panel B: Qualitative questions (7-point Likert scale: strongly disagree to strongly agree)
Risk taking
I am prepared to take the risk to lose money, when there is also a chance to gain money
Impulsiveness
I am impulsive and tend to buy things even when I can't
Patience
I am prepared to spend now and let the future take care of itself
Trust
I have trust in insurers

Table 4.13: Products in disposition effect experiment. This table shows for each product the win and loss outcomes, which follow from a coin toss (heads or tails).

Choice	Product	Heads	Tails	Exp. return	St. Dev.
1.	A	+3000	-4000	-500	4950
	B	-4000	+3000	-500	4950
2.	C	+4000	-4000	0	5657
	D	-4000	+4000	0	5657
3.	E	+5000	-4000	+500	6364
	F	-4000	+5000	+500	6364
4.	X	+6000	-5000	+500	7778
	Y	-2000	+4000	+1000	4243

Divide €10,000 below each time between today and 1 year later.

	Euros today (with certainty)	**Euros that you receive in 1 year with certainty**
Suppose you receive an extra €0.00 per euro that you have paid out in 1 year	0 … 10000	
Suppose you receive an extra €0.50 per euro that you have paid out in 1 year	0 … 10000	
Suppose you receive an extra €1.50 per euro that you have paid out in 1 year	0 … 10000	
Suppose you receive an extra €3.50 per euro that you have paid out in 1 year	0 … 10000	

Figure 4.4: Decision set Convex Time Budgets. In this decision screen, the subject allocates $m = 10,000$ Euro between an early payment today and a late payment with delay $k = 1$ year. The late payment is with a probability p_{t+k} of 100%. The gross interest rate $1 + r$ over k years in the 4 scenarios varies from 1.00 to 4.50. The allocated amount of €0 today is for illustration purposes only, the default values were blanks (subjects must actively allocate). The text is translated from Dutch to English.

B. Additional data

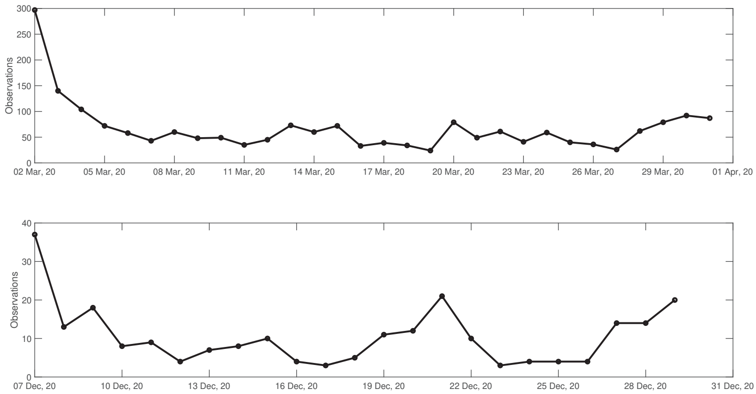

Figure 4.5: Daily survey observations during March 2020 (top) and December 2020 (bottom).

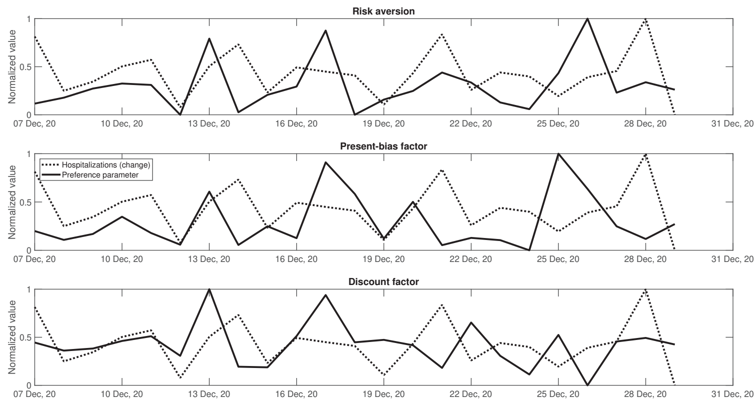

Figure 4.6: Time-varying risk and time preferences with COVID-19 hospitalizations during December 2020.

Table 4.14: Aggregate preferences from independent measures. *Risk taking*, *Impulsiveness* and *Impatience* measure qualitatively self-stated risk and time preferences on a 7-points Likert scale. *Trust* measures the trust in insurance companies with a 7-points Likert scale. *Risk aversion* (*EG*) measures quantitatively risk aversion following Eckel and Grossman (2008), while $Present-bias\ factor\ (Wang)$ and $Discount\ factor\ (Wang)$ measure quantitatively time preferences (assuming risk neutrality) following Rieger et al. (2015) and Wang (2017).

	Mean	St. Dev.	Min.	Max.	N
Panel A: Aggregate					
Qualitative measures					
Risk taking	3.52	1.76	1.00	7.00	2154
Impulsiveness	1.94	1.30	1.00	7.00	2210
Impatience	3.03	1.75	1.00	7.00	2203
Trust	3.35	1.49	1.00	7.00	2172
Quantitative measures					
Risk aversion (EG)	2.01	1.44	1.00	6.00	2240
Present-bias factor (RN)	0.93	0.21	0.10	1.78	1764
Discount factor (RN)	0.92	0.09	0.56	1.00	1764
Panel B: March					
$\Delta Hosp$	16.37	36.70	-52.63	144.44	1997
1-year life expectancy	93.25	15.26	0.00	100.00	1996
5-year life expectancy	85.70	18.14	0.00	100.00	1996
Qualitative measures					
Risk taking	3.50	1.76	1.00	7.00	1917
Impulsiveness	1.93	1.31	1.00	7.00	1969
Impatience	3.03	1.76	1.00	7.00	1963
Trust	3.33	1.49	1.00	7.00	1937
Quantitative measures					
Risk aversion (EG)	2.02	1.44	1.00	6.00	1997
Present-bias factor (RN)	0.93	0.21	0.10	1.78	1556
Discount factor (RN)	0.92	0.09	0.56	1.00	1556
Panel C: December					
$\Delta Hosp$	13.15	34.64	-44.73	74.16	243
1-year life expectancy	91.13	18.18	0.00	100.00	243
5-year life expectancy	82.63	19.92	0.00	100.00	243
Qualitative measures					
Risk taking	3.62	1.80	1.00	7.00	237
Impulsiveness	1.98	1.24	1.00	7.00	241
Impatience	3.00	1.66	1.00	7.00	24
Trust	3.54	1.48	1.00	7.00	235.00
Quantitative measures					
Risk aversion (EG)	1.95	1.45	1.00	6.00	243
Present-bias factor (RN)	0.95	0.17	0.38	1.50	208
Discount factor (RN)	0.92	0.08	0.67	1.00	208

Table 4.15: Aggregate risk and time preferences from Convex Time Budgets: Trading sample. This table shows the preferences for the only sample that participated in the disposition effect experiment. Two-limit Tobit maximum likelihood estimates for quasi-hyperbolic β, δ discounting, CRRA utility $U(x) = \frac{x^{1-\gamma}}{1-\gamma}$ and Prelec-weighting function $\pi(p) = p^{\eta}$. Background consumption equals annual after-tax income, which varies across subjects.

	Median	25th Percentile	75th Percentile	N
Panel A: All, trading experiment				
Risk aversion $\hat{\gamma}$	1.590	1.215	2.111	537
Present-bias factor $\hat{\beta}$	0.961	0.835	1.203	537
Discount factor $\hat{\delta}$	0.958	0.875	1.043	537
Annual discount rate	0.044	-0.041	0.143	537
Probability weighting $\hat{\eta}$	1.192	-0.007	2.309	537
Panel B: March, trading experiment				
Risk aversion $\hat{\gamma}$	1.599	1.232	2.316	294
Present-bias factor $\hat{\beta}$	0.962	0.784	1.203	294
Discount factor $\hat{\delta}$	0.961	0.869	1.055	294
Annual discount rate	0.041	-0.052	0.150	294
Probability weighting $\hat{\eta}$	1.085	-0.083	2.331	294

C. Robustness checks

Table 4.16: Time-varying Convex Time Budget allocations. This table reports the coefficients of the regressions $y_{i,t} = a_0 + a_1 \Delta Hosp_{i,t} + bX_{i,t} + \varepsilon_{i,t}$. $y_{i,t}$ is the average per individual i over the allocated amounts to the late payment at time t. Controls $X_{i,t}$ include *Dec*, *Male*, *Age*, *Partner*, *Edu. medium*, *Edu. high*, *Income* and day fixed effects. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Median regression	OLS
$\Delta Hosp$ ($\times$100)	1759.414***	971.356*
	(659.977)	(556.944)
December	-33.117	85.591
	(507.393)	(508.178)
Constant	26887.068***	26342.921***
	(1422.588)	(1290.979)
Observations	2240	2240
Controls	Yes	Yes
Day FE	Yes	Yes

Table 4.17: Robustness time-varying preferences. This table reports the coefficients of the regressions $y_{i,t} = a_0 + a_1 \Delta Hosp_{i,t} + bX_{i,t} + \varepsilon_{i,t}$. $y_{i,t}$ represents the preference parameter for individual i at day t (per column): risk aversion γ, present-bias factor β and annual discount factor δ. Controls $X_{i,t}$ include Dec, $Male$, Age, $Partner$, $Edu.\ medium$, $Edu.\ high$, $Income$ and day fixed effects. Panel A uses OLS rather than median regressions to estimate our main regression equation, while Panels B-F use median regressions. Panel B uses the preference parameter estimates based on monthly background income rather than annual background income. Panel C only uses the set of controls $X_{i,t}$: Dec, $Male$, Age, and day fixed effects. Panel D controls for life expectancy, added to the standard set of controls. Panel E controls for financial literacy, added to the standard set of controls. Panel F tests for unbalancedness of the panel dataset using the test of Verbeek and Nijman (1992). Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Risk aversion	Present-bias factor	Discount factor
Panel A: OLS			
$\Delta Hosp$ ($\times 100$)	0.390**	1.362**	0.052***
	(0.160)	(0.554)	(0.019)
Observations	2240	2240	2240
Panel B: Monthly background income			
$\Delta Hosp$ ($\times 100$)	0.046*	0.081***	0.025**
	(0.027)	(0.030)	(0.013)
Observations	2240	2240	2240
Panel C: Controlling for demographics only			
$\Delta Hosp$ ($\times 100$)	0.098*	0.056***	0.025***
	(0.055)	(0.020)	(0.010)
Observations	2246	2246	2246
Panel D: Controlling for life expectancy			
$\Delta Hosp$ ($\times 100$)	0.147***	0.066***	0.025**
	(0.047)	(0.021)	(0.010)
1-year life expectancy	-0.001	-0.000	-0.000
	(0.002)	(0.001)	(0.000)
5-year life expectancy	-0.002	-0.001	-0.000
	(0.002)	(0.001)	(0.000)
Observations	2239	2239	2239
Panel E: Controlling for financial literacy			
$\Delta Hosp$ ($\times 100$)	0.083**	0.069***	0.022**
	(0.038)	(0.024)	(0.009)
Financial literacy	-0.319***	-0.058**	-0.041***
	(0.032)	(0.024)	(0.010)
Observations	2240	2240	2240
Panel F: Unbalanced panel test			
$\Delta Hosp$ ($\times 100$)	0.156***	0.065**	0.023**
	(0.048)	(0.025)	(0.010)
Times observed	-0.030	-0.002	-0.008
	(0.041)	(0.015)	(0.009)
Observations	2240	2240	2240

Table 4.18: Robustness income. This table reports the coefficients of the regressions $y_{i,t} = a_0 + a_1 \Delta Hosp_{i,t} + bX_{i,t} + \varepsilon_{i,t}$. $y_{i,t}$ represents the preference parameter for individual i at day t (per column): risk aversion γ, present-bias factor β and annual discount factor δ. Controls $X_{i,t}$ include *Dec*, *Male*, *Age*, *Partner*, *Edu. medium*, *Edu. high*, *Income* and day fixed effects. The analysis contains only the subjects that have an equal salary during February, March, and December 2020. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Risk aversion	Present-bias factor	Discount factor
$\Delta Hosp$ ($\times 100$)	0.204***	0.102***	0.025**
	(0.075)	(0.030)	(0.012)
Constant	0.399**	0.905***	0.989***
	(0.155)	(0.060)	(0.027)
Observations	1386	1386	1386
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes

Table 4.19: Robustness ICU. This table reports the coefficients of the regressions $y_{i,t} = a_0 + a_1 \Delta ICU_{i,t} + bX_{i,t} + \varepsilon_{i,t}$. $y_{i,t}$ represents the preference parameter for individual i at day t (per column): risk aversion γ, present-bias factor β, annual discount factor δ, and annual discount rate. $\Delta ICU_{i,t}$ is the daily percentage change in national COVID-19 ICU hospitalizations from day $t-1$ to day t. Controls $X_{i,t}$ include *Dec*, *Male*, *Age*, *Partner*, *Edu. medium*, *Edu. high*, *Income* and day fixed effects. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Risk aversion	Present-bias factor	Discount factor	Discount rate
ΔICU ($\times 100$)	0.445***	0.209***	0.042	-0.050*
	(0.143)	(0.056)	(0.027)	(0.029)
Constant	0.344***	0.885***	0.958***	0.044**
	(0.124)	(0.040)	(0.019)	(0.021)
Observations	2240	2240	2240	2240
Controls	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes

Table 4.20: Within and between analyses for simpler measures. This table reports the coefficients of OLS regressions from within and between analyses for the qualitative and quantitative preference measures. Panel A shows the qualitatively-measured binary dependent variable $y_{i,t}$. Panel B shows the quantitatively-measured dependent variable $y_{i,t}$. Subpanel 1 estimates $y_{i,t} = a_0 + a_1(\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}) + a_2\overline{\Delta Hosp_i} + bX_{i,t} + \varepsilon_{i,t}$. Subpanel 2 estimates $y_{i,t} = a_0 + a_1(\Delta Hosp_{i,t} \times Dec) + a_2\Delta Hosp_{i,t} + a_3 Dec + bX_{i,t} + \varepsilon_{i,t}$. Subpanel 3 estimates $y_{i,t,p} = a_0 + a_1\Delta Hosp_{t,p} + bX_{i,t} + \varepsilon_{i,t,p}$, in which $\Delta Hosp_{t,p}$ is measured on a province level. Controls $X_{i,t}$ include *Dec*, *Male*, *Age*, *Partner*, *Edu. medium*, *Edu. high*, and *Income*. The Hausman and Wald tests show the null hypothesis between parentheses. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Qualitative				
	Risk taking	Impulsiveness	Impatience	Trust
Panel A.1: Individuals				
$\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}$ ($\times$100)	-0.093	-0.033	-0.123**	-0.180***
	(0.058)	(0.034)	(0.060)	(0.062)
$\overline{\Delta Hosp_i}$ ($\times$100)	-0.027	-0.022	-0.052	-0.005
	(0.035)	(0.019)	(0.034)	(0.032)
Observations	2154	2210	2203	2172
Controls	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes
Wald test ($a_1 = a_2$)	0.2903	0.7549	0.2684	0.0083
Panel A.2: March and December				
$\Delta Hosp \times Dec$ ($\times$100)	-0.122	0.027	-0.003	-0.251***
	(0.097)	(0.048)	(0.076)	(0.090)
$\Delta Hosp$ ($\times$100)	-0.026	-0.026	-0.060*	-0.008
	(0.034)	(0.019)	(0.032)	(0.030)
December	0.050	-0.004	-0.008	0.106***
	(0.033)	(0.017)	(0.030)	(0.033)
Observations	2154	2210	2203	2172
Controls	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes
Wald test ($a_1 = a_3 = 0$)	0.2361	0.8497	0.9540	0.0018
Panel A.3: Province				
$\Delta Hosp_{t,p}$ ($\times$100)	0.002	-0.006	-0.028**	-0.029**
	(0.014)	(0.006)	(0.013)	(0.013)
Observations	2144	2200	2193	2162
Controls	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes
Province FE	Yes	Yes	Yes	Yes

Table continued.

Panel B: Quantitative			
	Risk aversion (EG)	Present-bias factor (RN)	Discount factor (RN)
Panel B.1: Individuals			
$\Delta Hosp_{i,t} - \overline{\Delta Hosp_i}$ ($\times$100)	0.031	0.017	-0.001
	(0.204)	(0.034)	(0.012)
$\overline{\Delta Hosp_i}$ ($\times$100)	-0.099	-0.009	-0.002
	(0.112)	(0.019)	(0.008)
Observations	2240	1764	1764
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Wald test ($a_1 = a_2$)	0.5311	0.4534	0.9524
Panel B.2: March and December			
$\Delta Hosp \times Dec$ ($\times$100)	0.089	0.061	0.009
	(0.285)	(0.039)	(0.017)
$\Delta Hosp$ ($\times$100)	-0.091	-0.011	-0.002
	(0.111)	(0.019)	(0.007)
December	-0.081	0.007	0.000
	(0.100)	(0.013)	(0.006)
Observations	2240	1764	1764
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Wald test ($a_1 = a_3 = 0$)	0.7211	0.1532	0.8588
Panel B.3: Province			
$\Delta Hosp_{t,p}$ ($\times$100)	-0.024	-0.011	0.002
	(0.046)	(0.008)	(0.003)
Observations	2230	1756	1756
Controls	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Province FE	Yes	Yes	Yes

Table 4.21: Correlations between preference measures. This table reports the Spearman rank correlations between preference measures. Symbol * indicates significance at the 5% level.

Panel A: Convex Time Budgets and qualitative measures						
	Risk taking	Impulsiveness	Impatience	Risk aversion $\hat{\gamma}$	Present-bias factor $\hat{\beta}$	Discount factor $\hat{\beta}$
Risk taking	1					
Impulsiveness	0.0918*	1				
Impatience	0.1825*	0.2692*	1			
Risk aversion $\hat{\gamma}$	0.038	-0.0422	0.0387	1		
Present-bias factor $\hat{\beta}$	-0.0618*	0.008	-0.0351	0.2611*	1	
Discount factor $\hat{\delta}$	-0.0136	-0.0027	-0.0281	0.4211*	0.3157*	1
Panel B: Convex Time Budgets and quantitative measures						
	Risk aversion (EG)	Present-bias factor (Wang)	Discount factor (Wang)	Risk aversion $\hat{\gamma}$	Present-bias factor $\hat{\beta}$	Discount factor $\hat{\beta}$
Risk aversion (EG)	1					
Present-bias factor (Wang)	-0.0618*	1				
Discount factor (Wang)	-0.0363	0.1282*	1			
Risk aversion $\hat{\gamma}$	0.0168	0.0119	0.0367	1		
Present-bias factor $\hat{\beta}$	-0.0013	0.0433	0.0107	0.2656*	1	
Discount factor $\hat{\delta}$	0.0272	0.0708*	0.0816*	0.4168*	0.3159*	1

Table 4.22: Aggregate disposition effect, conditional on gains and losses. This table reports all coefficients of the panel OLS regressions $Sell_{i,t} = d_0 + fX_{i,t} + e_{i,t}$. The dependent variable $Sell_{i,t}$ is a dummy variable equal to one if the investor sells an asset. Columns (1) - (4) contain all assets trading at a gain (i.e., $Gain_{i,t} = 1$), and Columns (5) - (8) contain all assets trading at a loss. The columns use different sets of control variables. Demographics include *Dec*, *Male*, *Age*, and Controls additionally include *Partner*, *Edu. medium*, *Edu. high*, and *Income*. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Gains				Losses			
Dependent variable: Sell	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
December		-0.045	-0.045	-0.044		0.000	-0.002	-0.002
		(0.033)	(0.033)	(0.033)		(0.036)	(0.036)	(0.036)
Male		-0.053	-0.028	-0.026		-0.046	-0.045	-0.043
		(0.036)	(0.040)	(0.040)		(0.043)	(0.047)	(0.046)
Age (years)		0.005**	0.004*	0.004*		-0.004	-0.003	-0.003
		(0.002)	(0.002)	(0.002)		(0.003)	(0.003)	(0.003)
Partner			0.046	0.048			0.014	0.013
			(0.037)	(0.037)			(0.045)	(0.045)
Edu. medium			-0.023	-0.022			0.051	0.050
			(0.046)	(0.045)			(0.062)	(0.061)
Edu. high			-0.024	-0.023			0.029	0.027
			(0.050)	(0.050)			(0.064)	(0.064)
Income (× 1000)			-0.034*	-0.034*			0.001	0.000
			(0.019)	(0.019)			(0.024)	(0.024)
Constant	0.763***	106.553**	89.025*	86.433*	0.513***	-84.857	-75.778	-70.898
	(0.018)	(50.686)	(51.593)	(51.561)	(0.022)	(59.018)	(58.964)	(58.744)
Observations	1040	1040	1040	1040	1088	1088	1088	1088
Demographics	No	Yes	Yes	Yes	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes	No	No	Yes	Yes
Asset FE	No	No	No	Yes	No	No	No	Yes
Day FE	No	Yes	Yes	Yes	No	Yes	Yes	Yes

Table 4.23: Gambler's fallacy: Aggregate and time-varying disposition effect. This table reports the results for the additional control variable *Correct Probability*, which is a dummy variable equal to one when the investor correctly estimated (i.e., 50%) the two probabilities of tails in two independent sequences of fair coin tosses. Panel A shows the aggregate (stable) disposition effect, and Panel B the disposition effect during the crisis. Panel A reports coefficients of the panel OLS regressions $Sell_{i,t} = d_0 + d_1 Gain_{i,t} + fX_{i,t} + e_{i,t}$, and Panel B reports the results of the panel OLS regressions $Sell_{i,t} = \tilde{d}_0 + \tilde{d}_1 \Delta Hosp_{i,t} + \tilde{f}X_{i,t} + \tilde{e}_{i,t}$ for all assets trading at a gain (i.e., $Gain_{i,t} = 1$). The dependent variable $Sell_{i,t}$ is a dummy variable equal to one if the investor sells an asset. $Gain_{i,t}$ is a dummy variable equal to one when the investor experienced a gain after the toin coss, and equal to zero otherwise. Columns (1) - (4) use different sets of control variables. Demographics include *Dec*, *Male*, *Age*, and Controls additionally include *Partner*, *Edu. medium*, *Edu. high*, and *Income*. Robust standard errors, corrected for clustering of observations at the individual level, are in parentheses. Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Aggregate disposition effect				
Dependent variable: Sell	(1)	(2)	(3)	(4)
Gain× *Correct probability*	-0.032	-0.033	-0.031	-0.029
	(0.060)	(0.060)	(0.060)	(0.060)
Gain	0.274***	0.271***	0.269***	0.267***
	(0.051)	(0.051)	(0.051)	(0.051)
Correct probability	0.066	0.069	0.072	0.070
	(0.047)	(0.047)	(0.047)	(0.047)
Constant	0.465***	15.528	14.040	13.929
	(0.040)	(40.738)	(40.592)	(40.546)
Observations	2128	2128	2128	2128
Demographics	No	Yes	Yes	Yes
Controls	No	No	Yes	Yes
Asset FE	No	No	No	Yes
Day FE	No	Yes	Yes	Yes
Panel B: Gains				
$\Delta Hosp$ (×100)× *Correct probability*		0.105	0.100	0.105
		(0.103)	(0.103)	(0.102)
$\Delta Hosp$ (×100)		-0.191**	-0.203**	-0.204**
		(0.095)	(0.096)	(0.095)
Correct probability		0.022	0.026	0.026
		(0.047)	(0.047)	(0.047)
Constant		116.303**	97.679*	95.394*
		(49.685)	(50.667)	(50.615)
Observations		1040	1040	1040
Demographics		Yes	Yes	Yes
Controls		No	Yes	Yes
Asset FE		No	No	Yes
Day FE		Yes	Yes	Yes

Table 4.24: Gambler's fallacy: Preferences and disposition effect. This table reports the medians of the preference parameters if an investor holds (i.e., Sell Gain is 'No') an asset after a gain and sells (i.e., Sell Gain is 'Yes') an asset after a gain, *only* for investors not subject to the gambler's fallacy (i.e., *Correct Probability* equal to one). The column 'Difference' reports the difference in preference parameters for holding and selling a gain. Wilcoxon rank-sum test p-values are reported between parentheses, which tests the null hypothesis H_0 : preference(sell gain = No) = preference(sell gain = Yes). Symbols ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	Sell Gain		
	No	Yes	Difference (p-value)
Risk aversion	1.67	1.51	0.16 (0.01**)
Present-bias factor	.96	.96	0.00 (0.23)
Discount factor	.97	.95	0.02 (0.06*)
Observations	172	585	757

D. Estimation

This section shows how we estimate risk and time preferences from the CTB. In the CTB experiment, subjects choose a payment c_t, available at time t, and a payment c_{t+k} available after delay k, continuously along a convex budget constraint

$$c_t + \frac{c_{t+k}}{1+r} = m, \tag{4.5}$$

where $(1 + r)$ is the experimental gross interest rate and m is the experimental budget.

Using the quasi-hyperbolic model of intertemporal decision making (Phelps and Pollak, 1968; Laibson, 1997), the subject maximizes discounted expected utility over the early payment c_t and late payment c_{t+k} (including interest)

$$\begin{aligned}\max_{c_t, c_{t+k}} \delta^t \left[\pi(p_t)U(c_t + w_t) + (1 - \pi(p_t))U(w_t)\right] \\ + \beta\delta^{t+k} \left[\pi(p_{t+k})U(c_{t+k} + w_{t+k}) + (1 - \pi(p_{t+k}))U(w_{t+k})\right],\end{aligned} \tag{4.6}$$

where δ is the one period discount factor and β is the present-bias factor. The quasi-hyperbolic form is able to capture time-inconsistent behavior. $\beta < 1$ indicates present bias, and if $\beta = 1$ the model equals exponential discounting (i.e., standard time-consistent behavior). p_t and p_{t+k} are the corresponding probabilities of payment. The terms w_t and w_{t+k} are additional utility parameters which could be interpreted as background consumption or income (Andersen, Harrison, Lau, et al., 2008b).

Hence, the CTB method asks subjects to maximize a utility function $U(c_t, c_{t+k})$. We assume that subjects have a time-separable Constant Relative Risk Aversion (CRRA) utility function of the form

$$U(x) = \frac{x^{1-\gamma}}{1-\gamma}, \tag{4.7}$$

with γ the coefficient of relative risk aversion.[20] Money allocated to the early payment has a value of c_t, while money allocated to the late payment has a present value of $c_{t+k}/(1+r)$.[21] Since early payments are always certain, it holds that $p_t = 1$. In some decision sets, the late payment is uncertain with probability p_{t+k}. For instance, when p_{t+k} is 0.7, the late payment is paid with a chance of 70%, and nothing is paid with a chance of 30%.

Given the evidence regarding probability distortions (Kahneman and A.Tversky, 1979; Tversky and Kahneman, 1992), we use $\pi(p_{t+k})$ as the subjective probabilities of a late payment. We use a simple Prelec probability weighting

[20] With this functional form, $\gamma = 0$ denotes risk neutral behavior, $\gamma > 0$ denotes risk aversion and $\gamma < 0$ denotes risk seeking behavior.

[21] c_{t+k}/c_t defines the gross interest rate $(1+r)$ over k years, so $(1+r)^{1/k} - 1$ gives the standardized annual interest rate r. Multiplication by the payment probability p_{t+k} denotes the risk-adjusted interest rates.

function

$$\pi(p) = p^{\eta}, \tag{4.8}$$

where p is the objective probability and $\pi(p)$ is the subjective (distorted) probability. Since our experimental setup contains late payment probabilities of 50% or larger, it is sufficient to identify the underweighting of high probabilities as there is no need to capture both the over- and underweighting of low and high probabilities, respectively.

Solving the subject's maximization problem (4.6) subject to her budget constraint (4.5) yields the first-order condition

$$\frac{c_t + w_t}{c_{t+k} + w_{t+k}} = \begin{cases} \left(\beta\delta^k(1+r)\pi(p_{t+k})\right)^{-\frac{1}{\gamma}}, & \text{if } t = 0 \\ \left(\delta^k(1+r)\pi(p_{t+k})\right)^{-\frac{1}{\gamma}}, & \text{if } t > 0 \end{cases} \tag{4.9}$$

which shows that the experimentally allocated payments depend on the preference parameters and the experimentally varied parameters.

Taking the logarithm and using the Prelec weighting function (4.8), we find

$$\begin{aligned} \ln\left(\frac{c_t + w_t}{c_{t+k} + w_{t+k}}\right) = &\left(\frac{\ln\beta}{-\gamma}\right) \cdot \mathbb{1}_{t=0,\ p_{t+k}=1} + \left(\frac{\ln\delta}{-\gamma}\right) \cdot k \\ &+ \left(\frac{1}{-\gamma}\right) \cdot \ln(1+r) + \left(\frac{\eta}{-\gamma}\right) \cdot \ln(p_{t+k}), \end{aligned} \tag{4.10}$$

where $\mathbb{1}_{t=0,\ p_{t+k}=1}$ is an indicator function for the time period $t = 0$ and a sure probability of late payment. In other words, we presume that present bias enters after today and to avoid interference between present bias and probability weighting, we estimate present bias only in the scenarios where late payment is guaranteed.

Given an additive error structure and assumptions on background income w_t, w_{t+k}, such an equation is easily estimated with parameter estimates for β, δ, γ and η obtained via non-linear combinations of coefficient estimates. The equation shows indeed that the present-bias factor β is identified through the front-end delay in t, the long-term annual discount factor is identified through the back-end delay via k, the CRRA risk aversion follows from changes in the gross interest rate $(1+r)$ and probability weighting follows from changes in the late payment probability p_{t+k}. We estimate the preference parameters per individual using two-limit Tobit maximum likelihood regressions to account for corner solutions. To limit the number of estimated parameters, we set background income for each individual to her annual after-tax income and we assume that $w_t = w_{t+k}$. After the estimation, we winsorize the individually estimated parameters at the 1% level of the bottom and top of the overall distribution.

Bibliography

Abel, A. (1999). “Risk Premia and Term Premia in General Equilibrium”. In: *Journal of Monetary Economics* 43.1, pp. 3–33.

Ainslie, G. (1992). *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person.* Cambridge, UK: Cambridge University Press.

Allison, P. (2009). *Fixed Effects Regression Models*. Thousand Oaks, CA: SAGE Publications, Inc.

Andersen, S., F. Harrison, M. Lau, and E. Rutström (2008a). “Eliciting risk and time preferences”. In: *Econometrica* 76.3, pp. 583–618.

— (2008b). “Eliciting risk and time preferences”. In: *Econometrica* 76.3, pp. 583–618.

Andersen, S., G. Harrison, M.Lauc, and E.Rutström (2010). “Preference heterogeneity in experiments: Comparing the field and laboratory”. In: *Journal of Economic Behavior & Organization* 73, pp. 209–224.

Andersen, S., G.W. Harrison, M.I. Lau, and E.E. Rutström (2014). “Discounting behavior: A reconsideration”. In: *European Economic Review* 71, pp. 15–33.

Andreoni, J., M. Kuhn, and C. Sprenger (2015). “Measuring time preferences: A comparison of experimental methods”. In: *Journal of Economic Behavior* 116, pp. 451–464.

Andreoni, J. and C. Sprenger (2012a). “Estimating Time Preferences from Convex Budgets”. In: *American Economic Review* 102.7, pp. 3333–3356.

— (2012b). “Risk Preferences Are Not Time Preferences”. In: *American Economic Review* 102.7, pp. 3357–3376.

Andrikogiannopoulou, A. and F. Papakonstantinou (2020). "History-Dependent Risk Preferences: Evidence from Individual Choices and Implications for the Disposition Effect". In: *Review of Financial Studies* 33.8, pp. 3674–3718.

Angeletos, G.M., D. Laibson, A. Repetto, J. Tobacman, and S. Weinberg (2001). "The hyperbolic consumption model: Calibration, simulation and empirical evaluation". In: *Journal of Economic Perspectives* 15.3, pp. 47–68.

Angrisani, M., M. Cipriani, A. Guarino, R. Kendall, and J. Ortiz de Zarate (2020). "Risk Preferences at the Time of COVID-19: An Experiment with Professional Traders and Students". In: *FRB of New York Staff Report* 927.

Arisoy, Y., T. Bali, and Y. Tang (2021). "Investor regret and stock returns". In: *Working paper*.

Atmaz, A. (2021). "Stock Return Extrapolation, Option Prices, and Variance Risk Premium". In: *Review of Financial Studies*, pp. 1–46.

Augenblick, N., M. Niederle, and C. Sprenger (2015). "Working over Time: Dynamic Inconsistency in Real Effort Tasks". In: *Quarterly Journal of Economics*, pp. 1067–1115.

Augenblick, N. and M. Rabin (2019). "An Experiment on Time Preference and Misprediction in Unpleasant Tasks". In: *Review of Economic Studies* 86, pp. 941–975.

Augustin, P. and R. Tédongap (2020). "Disappointment Aversion, Term Structure, and Predictability Puzzles in Bond Markets". In: *Management Science* 67.10, pp. 6266–6293.

Backus, D., A. Gregory, and S. Zin (1989). "Risk Premiums in the Term Structure: Evidence from Artificial Economies". In: *Journal of Monetary Economics* 24, pp. 371–399.

Balakrishnan, U., J. Haushofer, and P. Jakiela (2020). "How Soon Is Now? Evidence of Present Bias from Convex Time Budget Experiments". In: *Experimental Economics* 23, pp. 294–321.

Bali, T., N. Cakici, and R. Whitelaw (2011). "Maxing out: Stocks as lotteries and the cross-section of expected returns". In: *Journal of Financial Economics* 99.2, pp. 427–446.

Balter, A. and B. Werker (2021). “The Effect of the Assumed Interest Rate and Smoothing on Variable Annuities”. In: *ASTIN Bulletin* 50.1, pp. 131–154.

Bansal, R., R. Dittmar, and C. Lundblad (2005). “Consumption, Dividends, and the Cross Section of Equity Returns”. In: *Journal of Finance* 60.4, pp. 1639–1672.

Bansal, R., S. Miller, D. Song, and A. Yaron (2021). “The term structure of equity risk premia”. In: *Journal of Financial Economics (forthcoming).*

Bansal, R. and I. Shaliastovich (2013). “A long-run risks explanation of predictability puzzles in bond and currency markets”. In: *Review of Financial Studies* 26.1, pp. 1–33.

Bansal, R. and A. Yaron (2004). “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles”. In: *Journal of Finance* 59.4, pp. 1481–1509.

Barberis, N. (2018). “Psychology-Based Models of Asset Prices and Trading Volume”. In: *Handbook of Behavioral Economics: Applications and Foundations 1.* Ed. by B. Bernheim, S. DellaVigna, and D. Laibson. Vol. 1. North-Holland. Chap. 2, pp. 79–175.

Barberis, N. and M. Huang (2001). “Mental Accounting, Loss Aversion and Individual Stock Returns”. In: *Journal of Finance* 56.4, pp. 1247–1292.

Barberis, N., M. Huang, and T. Santos (2001). “Prospect Theory and Asset Prices”. In: *Quarterly Journal of Economics* 116.1, pp. 1–53.

Barberis, N., M. Huang, and R. Thaler (2006). “Individual Preferences, Monetary Gambles, and Stock MarketParticipation: A Case for Narrow Framing”. In: *American Economic Review* 96.4, pp. 1069–1090.

Barberis, N., L. Jin, and B. Wang (2021). “Prospect Theory and Stock Market Anomalies”. In: *Journal of Finance* 76.5, pp. 2639–2687.

Barberis, N. and W. Xiong (2012). “Realization utility”. In: *Journal of Financial Economics* 104.2, pp. 251–271.

Barro, R. (2006). “Rare Disasters and Asset Markets in the Twentieth Century”. In: *Quarterly Journal of Economics* 121.3, pp. 823–866.

Basu, S. (1983). "The relationship between earnings' yield, market value and return for NYSE common stocks: Further evidence". In: *Journal of Financial Economics* 12.1, pp. 129–156.

Bauer, R., T. Ruof, and P. Smeets (2021). "Get Real! Individuals Prefer More Sustainable Investments". In: *Review of Financial Studies* 34.8, pp. 3976–4043.

Baule, R., O. Korn, and L.C. Kuntz (2019). "Markowitz with regret". In: *Journal of Economic Dynamics & Control* 103, pp. 1–24.

Beattie, J., J. Baron, J. Hershey, and M. Spranca (1994). "Psychological determinants of decision attitude". In: *Journal of Behavioral Decision Making* 7, pp. 129–144.

Beeler, J. and J. Campbell (2012). "The Long-Run Risks Model and Aggregate Asset Prices: An Empirical Assessment". In: *Critical Finance Review* 1.1, pp. 141–182.

Bell, D. (1982). "Regret in Decision Making under Uncertainty". In: *Operations Research* 30.5, pp. 961–981.

Benartzi, S. and R. Thaler (1995). "Myopic Loss Aversion and the Equity Premium Puzzle". In: *Quarterly Journal of Economics* 110.1, pp. 73–92.

Bernard, S., B. Loos, and M. Weber (2022). "The Disposition Effect in Boom and Bust Markets". In: *Working paper*.

Bleichrodt, H., A. Cillo, and E. Diecidue (2010). "A Quantitative Measurement of Regret Theory". In: *Management Science* 56.1, pp. 161–175.

Bleichrodt, H., Y. Gao, and K. Rohde (2016). "A measurement of decreasing impatience for health and money". In: *Journal of Risk and Uncertainty* 52, pp. 213–231.

Bokern, P., J. Linde, A. Riedl, and P. Werner (2021). "The effect of the COVID-19 crisis on economic and social preferences". In: *Netspar discussion paper*.

Boudoukh, J., M. Richardson, T. Smith, and R. Whitelaw (1999). "Ex Ante Bond Returns and the Liquidity Preference Hypothesis". In: *Journal of Finance* 104.54, pp. 1153–1167.

Bourgeois-Gironde, S. (2010). "Regret and the rationality of choices". In: *Phil. Trans. R. Soc. B* 365, pp. 249–257.

Braun, M. and A. Muermann (2004). "The Impact of Regret on the Demand for Insurance". In: *Journal of Risk and Insurance* 71.4, pp. 737–767.

Brennan, M.J. and Y. Xia (2002). "Dynamic Asset Allocation under Inflation". In: *Journal of Finance* 57.3, pp. 1201–1238.

Brown, J. (2001). "Private pensions, mortality risk, and the decision to annuitize". In: *Journal of Public Economics* 82.1, pp. 29–62.

Brunnermeier, M., C. Gollier, and J. Parker (2007). "Decision making : a psychological analysis of conflict, choice, and commitment". In: *American Economic Review* 97.2, pp. 159–165.

Bu, D., T. Hanspal, Y. Liao, and Y. Liu (2020). "Risk taking during a global crisis: Evidence from Wuhan". In: *Working Paper*.

Bushee, B. (2001). "Do Institutional Investors Prefer Near-Term Earnings over Long-Run Value?" In: *Contemporary Accounting Research* 18, pp. 207–246.

— (1998). "The Influence of Institutional Investors on Myopic R&D Investment Behavior". In: *Accounting Review* 73, pp. 305–333.

Bütler, M. and F. Teppa (2007). "The choice between an annuity and a lump sum: Results from Swiss pension funds". In: *Journal of Public Economics* 91.10, pp. 1944–1966.

Camerer, C. and R. Hogarth (1999). "The Effects of Financial Incentives in Experiments: A Review and Capital-Labor-Production Framework". In: *Journal of Risk and Uncertainty* 19, pp. 7–42.

Camille, N., G. Coricelli, J. Sallet, P. Pradat-Diehl, JR. Duhamel, and A. Sirigu (2004). "The involvement of the orbitofrontal cortex in the experience of regret". In: *Science* 304.5674, pp. 1167–1170.

Campbell, J. and J. Cochrane (1995). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior". In: *NBER Working paper*.

— (1999). "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior". In: *Journal of Political Economy* 107.2, pp. 205–251.

Campbell, J. and R. Shiller (1988). "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors". In: *Review of Financial Studies* 1.3, pp. 195–228.

— (1991). “Yield spreads and interest rate movements: a bird’s eye view”. In: *Review of Economic Studies* 58.3, pp. 495–514.

Carroll, C., J. Slacalek, and M. Sommer (2019). “Dissecting saving dynamics: Measuring wealth, precautionary, and credit effects”. In: *NBER Working Paper*.

Chabris, C., D. Laibson, C. Morris, J. Schuldt, and D. Taubinsky (2008). “Individual laboratory-measured discount rates predict field behavior”. In: *Journal of Risk and Uncertainty* 37, pp. 237–269.

Chang, T., D. Solomon, and M. Westerfield (2015). “Looking for Someone to Blame: Delegation, Cognitive Dissonance, and the Disposition Effect”. In: *Journal of Finance* 71.1, pp. 267–302.

— (2016). “Looking for Someone to Blame: Delegation, Cognitive Dissonance, and the Disposition Effect”. In: *Journal of Finance* 71.1, pp. 267–302.

Chao, L., H. Szrek, N. Pereira, and M. Pauly (2009). “Time preference and its relationship with age, health, and survival probability.” In: *Judgement and Decision Making* 4.1, pp. 1–19.

Cheung, S. L. (2020). “Eliciting utility curvature in time preference”. In: *Experimental Economics* 23, pp. 493–525.

Chin, E., D. Nel, and S. Ólafsson (2014). *Problems and Solutions in Mathematical Finance. Volume 1: Stochastic Calculus*. Chichester: John Wiley & Sons, Ltd.

Chuang, Y. and L. Schechter (2015). “Stability of experimental and survey measures of risk, time, and social preferences: A review and some new results”. In: *Journal of Development Economics* 117, pp. 151–170.

Chung, T. and A. Low (2017). “The impact of investor impatience and environmental turbulence on myopic marketing management and stock performance”. In: *International Journal of Research in Marketing* 34.3, pp. 660–677.

Cochrane, J. (2017). “Macro-Finance”. In: *Review of Finance* 21.3, pp. 945–985.

— (2008). “The Dog That Did Not Bark: A Defense of Return Predictability”. In: *Review of Financial Studies* 21.4, pp. 1533–1575.

Cohen, J., K. Ericson, D. Laibson, and J. White (2020). “Measuring Time Preferences”. In: *Journal of Economic Literature* 58.2, pp. 299–347.

Cohn, A., J. Engelmann, E. Fehr, and M. Maréchal (2015). "Evidence for Countercyclical Risk Aversion: An Experiment with Financial Professionals". In: *American Economic Review* 105.2, pp. 860–885.

Connolly, T. and D. Butler (2006). "Regret in economic and psychological theories of choice". In: *Journal of Behavioral Decision Making* 19.2, pp. 139–154.

Connolly, T. and M. Zeelenberg (2002). "Regret in Decision Making". In: *Current Directions in Psychological Science* 11.6, pp. 212–216.

Coricelli, G., H. Critchley, M. Joffily, J. O'Doherty, A. Sirigu, and R. Dolan (2005). "Regret and its avoidance: a neuroimaging study of choice behavior." In: *Nature Neuroscience* 8, pp. 1255–1262.

Cox, J.C. and C.F. Huang (1989). "Optimal Consumption and Portfolio Policies when Asset Prices Follow a Diffusion Process". In: *Journal of Economic Theory* 49, pp. 33–83.

Creal, D. and J. Wu (2020). "Bond risk premia in consumption-based models". In: *Quantitative Economics* 11.4, pp. 1461–1481.

Dave, C., C. Eckel, and C. Johnson (2010). "Eliciting risk preferences: When is simple better?" In: *Journal of Risk and Uncertainty* 41, pp. 219–243.

De Bondt, W. and R. Thaler (1985). "Does the Stock Market Overreact?" In: *Journal of Finance* 40.3, pp. 793–805.

de Bresser, J., M. Knoef, and R. van Ooijen (2018). "Health and household expenditures". In: *Netspar design paper.*

De Jong, F. (2000). "Time Series and Cross-section Information in Affine Term-Structure Models". In: *Journal of Business & Economic Statistics* 18.3, pp. 300–314.

de Vries, M. (2021). "Time-Varying Risk Preferences and the Equity Term Structure". In: *Working paper.*

Debets, S., H. Prast, M. Rossi, and A. van Soest (2021). "Pension Communication in the Netherlands and Other Countries". In: *Working paper, Tilburg University.*

DellaVigna, S. and M. Paserman (2005). "Job Search and Impatience". In: *Journal of Labour Economics* 23.3, pp. 527–588.

Diecidue, E. and J. Somasundaram (2017). "Regret theory: A new foundation". In: *Journal of Economic Theory* 172, pp. 88–119.

Dikolli, S., S. Kulp, and K. Sedatole (2009). "Transient Institutional Ownership and CEO Contracting". In: *The Accounting Review* 84.3, pp. 737–770.

Dimmock, S., R. Kouwenberg, O. Mitchell, and K. Peijnenburg (2021). "Household Portfolio Underdiversification and Probability Weighting: Evidence from the Field". In: *Review of Financial Studies* 34.9, pp. 4524–4563.

Dodonova, A. and Y. Khoroshilov (2005). "Applications of Regret Theory to Asset Pricing". In: *Working paper*.

Dohmen, T., A. Falk, D. Huffman, and U. Sunde (2010). "Are Risk Aversion and Impatience Related to Cognitive Ability?" In: *American Economic Review* 100.3, pp. 1238–1260.

Dohmen, T., A. Falk, D. Huffman, U. Sunde, J. Schupp, and G. Wagner (2011). "Individual Risk Attitudes: Measurement, Determinants, and Behavioral Consequences". In: *Journal of the European Economic Association* 9.3, pp. 522–550.

Drerup, T., M. Wibral, and C. Zimpelmann (2022). "Skewness Expectations and Portfolio Choice". In: *IZA Discussion Paper*.

Drichoutis, A. and R. Nayga (2021). "On the stability of risk and time preferences amid the COVID-19 pandemic". In: *Experimental Economics*, pp. 1–36.

Eckel, C. and P. Grossman (2002). "Sex differences and statistical stereotyping in attitudes toward financial risk". In: *Evolution and Human Behavior* 23.4, pp. 281–295.

Eckel, C. C. and P. J. Grossman (2008). "Forecasting risk attitudes: An experimental study using actual and forecast gamble choices". In: *Journal of Economic Behavior & Organization* 68, pp. 1–17.

Eeckhoudt, L., B. Rey, and H. Schlesinger (2007). "A Good Sign for Multivariate Risk Taking". In: *Management Science* 53.1, pp. 117–124.

Ericson, K. and D. Laibson (2019). "Intertemporal Choice". In: *Handbook of Behavioral Economics - Foundations and Applications 2*. Ed. by D. Bernheim, D. Laibson, and S. DellaVigna. Elsevier.

Erner, C., A. Klos, and T. Langer (2013). "Can prospect theory be used to predict an investor's willingness to pay?" In: *Journal of Banking and Finance* 37.6, pp. 1960–1973.

Falk, A., A. Becker, T. Dohmen, B. Enke, D. Huffman, and U. Sunde (2018). "Global Evidence on Economic Preferences". In: *Quarterly Journal of Economics* 133.4, pp. 1645–1692.

Fama, E. and K. French (1996). "Multifactor explanations of asset pricing anomalies". In: *Journal of Finance* 51.1, pp. 55–84.

— (1988). "Permanent and Temporary Components of Stock Prices". In: *Journal of Political Economy* 96.2, pp. 246–273.

— (1992). "The Cross-Section of Expected Stock Returns". In: *Journal of Finance* 47.2, pp. 427–465.

Fioretti, M., A. Vostroknutov, and G. Coricelli (2021). "Dynamic Regret Avoidance". In: *American Economic Journal: Microeconomics (forthcoming).*

Fogel, S. and T. Berry (2006). "The Disposition Effect and Individual Investor Decisions: The Roles of Regret and Counterfactual Alternatives". In: *Journal of Behavioral Finance* 7.2, pp. 107–116.

Frederick, S. (2005). "Cognitive Reflection and Decision Making". In: *Journal of Economic Perspectives* 19.4, pp. 25–42.

Frederick, S., G. Loewenstein, and T. O'Donoghue (2002). "Time Discounting and Time Preference: A Critical Review". In: *Journal of Economic Literature* 40, pp. 351–401.

Frehen, R., R. Hoevenaars, F. Palm, and P. Schotman (2008). "Regret aversion and annuity risk in defined contribution pension plans". In: *Insurance: Mathematics and Economics* 42.3, pp. 1050–1061.

Frijda, N. (1988). "The laws of emotion". In: *American Psychologist* 43.5, pp. 349–3587.

— (2007). *The laws of emotion.* Lawrence Erlbaum Associates Publishers.

Frydman, C. and C. Camerer (2016). "Neural Evidence of Regret and Its Implications for Investor Behavior". In: *Review of Financial Studies* 29.11, pp. 3108–3139.

Gabaix, X. (2012). "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance". In: *Quarterly Journal of Economics* 127.2, pp. 645–700.

Gabillon, E. (2011). "A Theory of Regret and Information". In: *Cahiers du GREThA* 15, pp. 1–28.

Gallmeyer, M., B. Hollifield, F. Palomino, and S. Zin (2017). "Term Premium Dynamics and the Taylor Rule". In: *Quarterly Journal of Finance* 7.4.

Gao, F., A. Xi He, and P. He (2018). "A theory of intermediated investment with hyperbolic discounting investors". In: *Journal of Economic Theory* 177, pp. 70–100.

Gathergood, J. (2012). "Self-control, financial literacy and consumer over-indebtedness". In: *Journal of Economic Psychology* 33.3, pp. 590–602.

Gerakos, J. and J. Linnainmaa (2018). "Decomposing value". In: *Review of Financial Studies* 31.5, pp. 1825–1854.

Geyer, A. and S. Pichler (1999). "A State-Space Approach to Estimate and Test Multifactor Cox-Ingersoll-Ross models of the Term Structure". In: *Journal of Financial Research* 22.1, pp. 107–130.

Giavazzi, F. and M. McMahon (2012). "Political uncertainty and household savings". In: *Review of Economics and Statistics* 94.2, pp. 517–531.

Gilovich, T. and V. Medvec (1995). "The experience of regret: What, when, and why". In: *Psychological Review* 102.2, pp. 379–395.

Gneezy, U. and J. Potters (1997). "An Experiment on Risk Taking and Evaluation Periods". In: *Quarterly Journal of Economics* 112.2, pp. 631–645.

Goda, G.S., M.R. Levy, C.F. Manchester, A. Sojourner, and J. Tasoff (2015). "The Role of Time Preferences and Exponential-Growth Bias in Retirement Savings". In: *NBER Working Paper*.

Gollier, C. (2018). "Aversion to risk of regret and preference for positively skewed risks". In: *Economic Theory*.

Gollier, C. and B. Salanié (2012). "Individual decisions under risk, risk sharing and asset prices with regret". In: *Working paper*.

Golsteyn, B., H. Grönqvist, and L. Lindahl (2014). “Adolescent Time Preferences Predict Lifetime Outcomes”. In: *Economic Journal* 124, F739–F761.

Gomes, F., M. Haliassos, and T. Ramadorai (2021). “Household Finance”. In: *Journal of Economic Literature* 59.3, pp. 919–1000.

Gomez-Cram, R. and A. Yaron (2021). “How Important Are Inflation Expectations for the Nominal Yield Curve?” In: *Review of Financial Studies* 34.2, pp. 985–1045.

Guiso, L., P. Sapienza, and L. Zingales (2018). “Time varying risk aversion”. In: *Journal of Financial Economics* 128.3, pp. 403–421.

Gul, F. (1991). “A Theory of Disappointment Aversion”. In: *Econometrica* 59.3, pp. 667–686.

Guo, H. and J. Wachter (2021). ““Superstituous” Investors”. In: *NBER Working paper*.

Hackethal, A., M. Kirchler, C. Laudenbach, M. Razen, and A. Weber (2022). “On the role of monetary incentives in risk preference elicitation experiments”. In: *Journal of Risk and Uncertainty*, pp. 637–648.

Hajcak, G. and D. Olvet (2008). “The persistence of attention to emotion: Brain potentials during and after picture presentation”. In: *Emotion* 8.2, pp. 250–255.

Haldane, A. (2010). “Patience and Finance”. In: *Speech Bank of England*.

Harris, C. and D. Laibson (2013). “Instantaneous Gratification”. In: *The Quarterly Journal of Economics* 128.1, pp. 205–248.

Harrison, G., A. Hofmeyr, H. Kincaid, B. Monroe, D. Ross, M. Schneider, and J. Swarthout (2022). “Subjective beliefs and economic preferences during the COVID-19 pandemic”. In: *Experimental Economics* 25, pp. 795–823.

Harrison, G. W., M. I. Lau, and M.B. Williams (2002). “Estimating Individual Discount Rates in Denmark: A Field Experiment”. In: *American Economic Review* 92.5, pp. 1606–1617.

Haushofer, J. and E. Fehr (2014). “On the psychology of poverty”. In: *Science* 344.6186, pp. 862–867.

Holt, C. A. and S. K. Laury (2002). “Risk Aversion and Incentive Effects”. In: *American Economic Review* 92.5, pp. 1644–1655.

Horvath, F., F. de Jong, and B. Werker (2017). "Robust Pricing of Fixed Income Securities". In: *Netspar Working Paper*.

Humphrey, S. (2004). "Feedback-conditional regret theory and testing regret-aversion in risky choice". In: *Journal of Economic Psychology* 25, pp. 839–857.

Hurwitz, A. and O. Sade (2020). "An investigation of time preferences, life expectancy, and annuity versus lump sum choices: Can smoking harm long-term saving decisions?" In: *Journal of Economic Behavior & Organization* 180, pp. 812–825.

Ikeda, S., E. Yamamura, and Y. Tsutsui (2020). "COVID-19 Enhanced Diminishing Sensitivity in Prospect-Theory Risk Preferences: A Panel Analysis". In: *ISER Discussion Paper* 1106.

Ingersoll, J. and L. Jin (2013). "Realization Utility with Reference-Dependent Preferences". In: *Review of Financial Studies* 26.3, pp. 723–767.

Inkmann, J., P. Lopes, and A. Michaelides (2011). "How Deep Is the Annuity Market Participation Puzzle?" In: *Review of Financial Studies* 24.1, pp. 279–319.

Janis, I. and L. Mann (1997). *Decision making : a psychological analysis of conflict, choice, and commitment*. New York : Free Press.

Kahneman, D. and A.Tversky (1979). "Prospect Theory: An Analysis of Decision under Risk". In: *Econometrica* 47.2, pp. 263–291.

Köszegi, B. and M. Rabin (2006). "A Model of Reference-Dependent Preferences". In: *Quarterly Journal of Economics* 121.4, pp. 1133–1165.

Kureishi, W., H. Paule-Paludkiewicz, H. Tsujiyama, and M. Wakabayashi (2021). "Time preferences over the life cycle and household saving puzzles". In: *Journal of Monetary Economics* 124, pp. 123–139.

Laibson, D. (1997). "Golden Eggs and Hyperbolic Discounting". In: *Quarterly Journal of Economics* 112.2, pp. 443–477.

Laibson, D., P. Maxted, and A. Repetto (Oct. 2015). "Estimating Discount Functions with Consumption Choices over the Lifecycle". Mimeo, Harvard University.

Laibson, D., P. Maxted, and A. Repetto (2020). “Estimating Discount Functions with Consumption Choices over the Lifecycle”. Mimeo, Harvard University.

Larrick, R. (1993). “Motivational factors in decision theories: The role of self-protection”. In: *Psychological Bulletin* 113.3, pp. 440-450–1170.

Larrick, R. and T. Bowles (1995). “Avoiding regret in decisions with feedback: A negotiation example”. In: *Organizational Behavior and Human Decision Processes* 63.1, pp. 87–97.

Laverty, K. (1996). “Economic "Short-Termism": The Debate, the Unresolved Issues, and the Implications for Management Practice and Research”. In: *The Academy of Management Review* 21.3, pp. 825–860.

Lehmann, D. (2004). “Metrics for Making Marketing Matter”. In: *Journal of Marketing* 68.4, pp. 73–75.

Lettau, M., M. Maggiori, and M. Weber (2014). “Conditional risk premia in currency markets and other asset classes”. In: *Journal of Financial Economics* 114.2, pp. 197–225.

Lettau, M. and J. Wachter (2007). “Why Is Long-Horizon Equity Less Risky? A Duration-Based Explanation of the Value Premium”. In: *Journal of Finance* 62.1, pp. 55–92.

Lin, CH., WH. Huang, and M. Zeelenberg (2006). “Multiple reference points in investor regret”. In: *Journal of Economic Psychology* 27, pp. 781–792.

Lin, T. and X. Liu (2018). “Individual Investor Preference, and the Cross-section of Stock Returns”. In: *Review of Finance* 22.5, pp. 1841–1876.

Liu, Y. and J. Wu (2021). “Reconstructing the yield curve”. In: *Journal of Financial Economics* 142.3, pp. 1395–1425.

Loewenstein, G. and D. Prelec (1992). “Anomalies in Intertemporal Choice: Evidence and an Interpretation”. In: *Quarterly Journal of Economics* 107.2, pp. 573–597.

Lohnmann, P., E. Gsottbauer, J. You, and A. Kontoleon (2020). “Social Preferences and Economic Decision-Making in the Wake of COVID-19: Experimental Evidence from China”. In: *Working Paper*.

Lohrenz, T., K. McCabe, C. Camerer, and P. Montague (2007). "Neural signature of fictive learning signals in a sequential investment task". In: *Proceedings of the National Academy of Sciences* 104, pp. 9493–9498.

Loomes, G. and R. Sugden (1982). "Regret Theory: An Alternative Theory of Rational Choice Under Uncertainty". In: *The Economic Journal* 92.368, pp. 805–824.

Lowenstein, L. (1988). *What's wrong with Wall Street: Short-term gain and the absentee shareholder*. Addison-Wesley.

MacKinnon, J. and H. White (1985). "Some Heteroskedasticity-Consistent Covariance Matrix Estimators with Improved Finite Sample Properties". In: *Journal of Econometrics* 29, pp. 305–325.

Magron, C. and M. Merli (2015). "Repurchase behavior of individual investors, sophistication and regret". In: *Journal of Banking & Finance* 61, pp. 15–26.

Malmendier, U. (2021). "Exposure, Experience, and Expertise: Why Personal Histories Matter in Economics". In: *NBER Working Paper*.

Malmendier, U. and S. Nagel (2011). "Depression Babies: Do Macroeconomic Experiences Affect Risk Taking?" In: *Quarterly Journal of Economics* 126.1, pp. 373–416.

Marín-Solano, J. and J. Navas (2010). "Consumption and portfolio rules for time-inconsistent investors". In: *European Journal of Operational Research* 201, pp. 860–872.

McClure, S., K. Ericson, D. Laibson, G. Loewenstein, and J. Cohen (2007). "Time Discounting for Primary Reward". In: *Journal of Neuroscience* 27.21, pp. 5796–5804.

Mehra, R. (2012). "Consumption-Based Asset Pricing Models". In: *Annual Review of Financial Economics* 4.1, pp. 385–409.

Mehra, R. and E. Prescott (1985). "The equity premium: A puzzle". In: *Journal of Monetary Economics* 15.2, pp. 145–161.

Meier, A. (2022). "Emotions and risk attitudes". In: *American Economic Journal: Applied Economics (forthcoming)*.

— (2019). "Emotions, risk attitudes, and patience". In: *Working paper*.

Meier, S. and C. Sprenger (2015). "Temporal Stability of Time Preferences". In: *Review of Economics and Statistics* 97.2, pp. 273–286.

Merton, R.C. (1969). "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case". In: *Review of Economics and Statistics* 51.3, pp. 247–257.

Michenaud, S. and B. Solnik (2008). "Applying regret theory to investment choices: Currency hedging decisions". In: *Journal of International Money and Finance* 27, pp. 677–694.

Mody, A., F. Ohnsorge, and D. Sandri (2012). "Precautionary Savings in the Great Recession". In: *IMF Economic Review* 60.1, pp. 114–138.

Muermann, A., O. Mitchell, and J. Volkman (2006). "Regret, portfolio choice, and guarantees in defined contribution schemes". In: *Insurance: Mathematics and Economics* 39, pp. 219–229.

Mundlak, Y. (1978). "On the Pooling of Time Series and Cross Section Data". In: *Econometrica* 46.1, pp. 69–85.

Neumann, J. von and O. Morgenstern (1947). *Theory of Games and Economic Behaviour*. 2nd ed. Princeton, N.J.: Princeton University Press.

Noor, J. (2009). "Hyperbolic discounting and the standard model: Eliciting discount functions". In: *Journal of Economic Theory* 144, pp. 2077–2083.

O'Donoghue, T. and M. Rabin (2015). "Present Bias: Lessons Learned and To Be Learned". In: *American Economic Review* 105.5, pp. 273–279.

Odean, T. (1998). "Are Investors Reluctant to Realize Their Losses?" In: *Journal of Finance* 40.3, pp. 1775–1798.

Pagel, M. (2016). "Expectations-Based Reference-Dependent Preferences and Asset Pricing". In: *Journal of the European Economic Association* 14.2, pp. 468–514.

Parker, J., J. Schild, L. Erhard, and D. Johnson (2022). "Economic Impact Payments and Household Spending During the Pandemic". In: *NBER working paper*.

Paserman, M. (2008). "Job Search and Hyperbolic Discounting: Structural Estimation and Policy Evaluation". In: *Economic Journal* 118, pp. 1418–1452.

Phelps, E. and R. Pollak (1968). "On Second-Best National Saving and Game-Equilibrium Growth". In: *The Review of Economic Studies* 35.2, pp. 185–199.

Piazzesi, M. and M. Schneider (2006a). "Equilibrium yield curves". In: *NBER Macroeconomics Annual*. Ed. by D. Acemoglu, K. Rogoff, and M. Woodford. Vol. 21. Cambridge: MIT Press, pp. 389–472.

— (2006b). "Equilibrium yield curves". In: *NBER Macroeconomics Annual*. Ed. by D. Acemoglu, K. Rogoff, and M. Woodford. Vol. 21. Cambridge: MIT Press, pp. 389–472.

Ploner, M. (2017). "Hold on to it? An experimental analysis of the disposition effect". In: *Judgment and Decision Making* 12.2, pp. 118–127.

Porter, M. (1992). *Capital choices: Changing the way America Invests in Industry*. Council on Competitiveness/Harvard Business School.

Poterba, J. and L. Summers (1988). "Mean reversion in stock prices: Evidence and implications". In: *Journal of Financial Economics* 22.1, pp. 27–59.

Potters, J., A. Riedl, and P. Smeets (2016). "Towards a Practical and Scientifically Sound Tool for Measuring Time and Risk Preferences in Pension Savings Decisions". In: *Netspar Industry Paper* 59.

Qin, J. (2015). "A model of regret, investor behavior, and market turbulence". In: *Journal of Economic Theory*, pp. 150–174.

— (2020). "Regret-based capital asset pricing model". In: *Journal of Banking and Finance* 114, pp. 1–8.

Quiggin, J. (1994). "Regret Theory with General Choice Sets". In: *Journal of Risk and Uncertainty* 8, pp. 153–165.

Riedel, F. (2004). "Heterogeneous time preferences and interest rates – the preferred habitat theory revisited". In: *European Journal of Finance* 10.1, pp. 3–22.

Rieger, M., M. Wang, and T. Hens (2015). "Risk Preferences Around the World". In: *Management Science* 61.3, pp. 637–648.

Roese, N. and J. Olson (1995). "Counterfactual thinking: A critical overview". In: *What might have been: The social psychology of counterfactual thinking*. Ed. by N. Roese and J. Olson. Lawrence Erlbaum Associates, Inc, pp. 1–55.

Roese, N. and J. Olson (1997). "Counterfactual thinking: The intersection of affect and function". In: *Advances in experimental social psychology*. Ed. by M. Zanna. Vol. 29. Academic Press, pp. 1–59.

Rudebusch, G. and E. Swanson (2008). "Examining the bond premium puzzle with a DSGE mode". In: *Journal of Monetary Economics* 55, S111–S126.

Saffrey, C., N. Roese, and A. Summerville (2008). "Praise for regret: People value regret above other negative emotions". In: *Motivation and Emotion* 32, pp. 46–54.

Samuelson, P. (1937). "A Note on Measurement of Utility". In: *Review of Economic Studies* 4, pp. 155–161.

Santos, T. and P. Veronesi (2010). "Habit formation, the cross section of stock returns and the cash-flow risk puzzle". In: *Journal of Financial Economics* 98.2, pp. 385–413.

Schildberg-Hörisch, H. (2018). "Are Risk Preferences Stable?" In: *Journal of Economic Perspectives* 32.2, pp. 135–154.

Shachat, J., M. Walker, and L. Wei (2020). "The impact of the Covid-19 pandemic on economic behaviours and preferences: Experimental evidence from Wuhan". In: *ESI Working Paper*, pp. 20–33.

Shafir, E., P. Diamond, and A. Tversky (1997). "Money Illusion". In: *Quarterly Journal of Economics* 112.2, pp. 341–374.

Shefrin, H. and M. Statman (1984). "Explaining Investor Preference for Cash Dividends". In: *Journal of Financial Economics* 13, pp. 253–282.

— (1985). "The Disposition to Sell Winners Too Early and Ride Losers Too Long: Theory and Evidence". In: *Journal of Finance* 40.3, pp. 777–790.

Shiller, R. (1981). "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" In: *American Economic Review* 71.3, pp. 421–436.

Solnik, B. and L. Zuo (2012). "A Global Equilibrium Asset Pricing Model with Home Preference". In: *Management Science* 58.2, pp. 273–292.

Stigler, G. and G. Becker (1977). "De Gustibus Non Est Disputandum". In: *American Economic Review* 67.2, pp. 76–90.

Stout, L. (2012). *The shareholder value myth: How putting shareholders first harms investors, corporations, and the public*. Berrett-Koehler Publishers.

Strahilevitz, M., B. Barber, and T. Odean (2011). "Once burned, twice shy: How naive learning, counterfactuals, and regret affect the repurchase of stocks previously sold". In: *Journal of Marketing Research* 48, pp. 102–120.

Strotz, R. (1956). "Myopia and Inconsistency in Dynamic Utility Maximisation". In: *Review of Economic Studies* 23.3, pp. 165–180.

Sutter, M., M. Kocher, D. Glätzle-Rützler, and S. Trautmann (2013). "Impatience and Uncertainty: Experimental Decisions Predict Adolescents' Field Behavior". In: *American Economic Review* 103.1, pp. 510–531.

Tanaka, T., C. Camerer, and Q. Nguyen (2010). "Risk and Time Preferences: Linking Experimental and Household Survey Data from Vietnam". In: *American Economic Review* 100.1, pp. 557–571.

Thaler, R. (1985). "Mental Accounting and Consumer Choice". In: *Marketing Science* 4.3, pp. 199–214.

— (1981). "Some empirical evidence on dynamic inconsistency". In: *Economics Letters* 8.3, pp. 201–207.

Tversky, A. and D. Kahneman (1992). "Advances in Prospect Theory: Cumulative Representation of Uncertainty". In: *Journal of Risk and Uncertainty* 5, pp. 297–323.

van Binsbergen, J., M. Boons, C. Opp, and A. Tamoni (2021). "Dynamic asset (mis)pricing: Build-up versus resolution anomalies". In: *Working paper.*

Van Binsbergen, J., M. Brandt, and R. Koijen (2008). "Optimal Decentralized Investment Management". In: *Journal of Finance* 63.4, pp. 1849–1895.

van Binsbergen, J., M. Brandt, and R. Koijen (2012). "On the Timing and Pricing of Dividends". In: *American Economic Review* 102.4, pp. 1596–1618.

Van Binsbergen, J., J. Fernández-Villaverde, R. Koijen, and J. Rubio-Ramírez (2012). "The termstructure of interest rates in a DSGE model with recursive preferences". In: *Journal of Monetary Economics* 59, pp. 634–648.

van Binsbergen, J. and R. Koijen (2017). "The term structure of returns: Facts and theory". In: *Review of Financial Studies* 124.1, pp. 1–21.

Vasicek, O. (1977). “An equilibrium characterization of the term structure”. In: *Journal of Financial Economics* 5, pp. 177–188.

Vayanos, D. and J. Villa (2021). “A Preferred-Habitat Model of the Term Structure of Interest Rates”. In: *Econometrica* 89.1, pp. 77–112.

Verbeek, M. and T. Nijman (1992). “Testing for selectivity bias in panel data models”. In: *International Economic Review* 33.3, pp. 681–703.

Wachter, J. (2006). “A consumption-based model of the term structure of interest rates”. In: *Journal of Financial Economics* 79.2, pp. 365–399.

Wang, H. (2017). “Robust asset pricing with stochastic hyperbolic discounting”. In: *Finance Research Letters* 21, pp. 178–185.

Weil, P. (1989). “The equity premium puzzle and the risk-free rate puzzle”. In: *Journal of Monetary Economics* 24.3, pp. 401–421.

Wilson, T. and D. Gilbert (2005). “Affective Forecasting: Knowing What to Want”. In: *Current Directions in Psychological Science* 14.3, pp. 131–134.

Yan, X. and Z. Zhang (2009). “Institutional Investors and Equity Returns: Are Short-term Institutions Better Informed?” In: *Review of Financial Studies* 22.2, pp. 893–924.

Zeelenberg, M. (1999). “Anticipated regret, expected feedback and behavioral decision making”. In: *Journal of Behavioral Decision Making* 12.2, pp. 93–106.

— (2018). “Anticipated regret: A prospective emotion about the future past”. In: *The Psychology of Thinking about the Future*. Ed. by G. Oettingen, A. Sevincer, and P. Gollwitzer. New York: The Guilford Press, pp. 276–295.

— (2020). “Regret is a Basic Emotion”. In: *Working paper*.

Zeelenberg, M., J. Beattie, J. van der Pligt, and N. de Vries (1996). “Consequences of regret aversion: Effects of expected feedback on risky decision making”. In: *Organizational Behavior and Human Decision Processes* 65, pp. 148–158.

Zeelenberg, M. and R. Pieters (2007). “A Theory of Regret Regulation 1.0”. In: *Journal of Consumer Psychology* 17.1, pp. 3–18.

— (2004). “Consequences of regret aversion in real life: The case of the Dutch postcode lottery”. In: *Organizational Behavior and Human Decision Processes* 93.2, pp. 155–168.

Zhao, Q., Y. Shen, and J. Wei (2014). “Consumption–investment strategies with non-exponential discounting and logarithmic utility”. In: *European Journal of Operational Research* 238.3, pp. 824–835.

Zou, Z., S. Chen, and L. Wedge (2014). “Finite horizon consumption and portfolio decisions with stochastic hyperbolic discounting”. In: *Journal of Mathematical Economics* 52, pp. 70–80.

CENTER DISSERTATION SERIES

CentER for Economic Research, Tilburg University, the Netherlands

No.	Author	Title	ISBN	Published
638	Pranav Desai	Essays in Corporate Finance and Innovation	978 90 5668 639 0	January 2021
639	Kristy Jansen	Essays on Institutional Investors, Asset Allocation Decisions, and Asset Prices	978 90 5668 640 6	January 2021
640	Riley Badenbroek	Interior Point Methods and Simulated Annealing for Nonsymmetric Conic Optimization	978 90 5668 641 3	February 2021
641	Stephanie Koornneef	It's about time: Essays on temporal anchoring devices	978 90 5668 642 0	February 2021
642	Vilma Chila	Knowledge Dynamics in Employee Entrepreneurship: Implications for parents and offspring	978 90 5668 643 7	March 2021
643	Minke Remmerswaal	Essays on Financial Incentives in the Dutch Healthcare System	978 90 5668 644 4	July 2021
644	Tse-Min Wang	Voluntary Contributions to Public Goods: A multi-disciplinary examination of prosocial behavior and its antecedents	978 90 5668 645 1	March 2021
645	Manwei Liu	Interdependent individuals: how aggregation, observation, and persuasion affect economic behavior and judgment	978 90 5668 646 8	March 2021
646	Nick Bombaij	Effectiveness of Loyalty Programs	978 90 5668 647 5	April 2021
647	Xiaoyu Wang	Essays in Microeconomics Theory	978 90 5668 648 2	April 2021
648	Thijs Brouwer	Essays on Behavioral Responses to Dishonest and Anti-Social Decision-Making	978 90 5668 649 9	May 2021
649	Yadi Yang	Experiments on hold-up problem and delegation	978 90 5668 650 5	May 2021
650	Tao Han	Imperfect information in firm growth strategy: Three essays on M&A and FDI activities	978 90 5668 651 2	June 2021

No.	Author	Title	ISBN	Published
651	Johan Bonekamp	Studies on labour supply, spending and saving before and after retirement	978 90 5668 652 9	June 2021
652	Hugo van Buggenum	Banks and Financial Markets in Microfounded Models of Money	978 90 5668 653 6	August 2021
653	Arthur Beddock	Asset Pricing with Heterogeneous Agents and Non-normal Return Distributions	978 90 5668 654 3	September 2021
654	Mirron Adriana Boomsma	On the transition to a sustainable economy: Field experimental evidence on behavioral interventions	978 90 5668 655 0	September 2021
655	Roweno Heijmans	On Environmental Externalities and Global Games	978 90 5668 656 7	August 2021
656	Lenka Fiala	Essays in the economics of education	978 90 5668 657 4	September 2021
657	Yuexin Li	Pricing Art: Returns, Trust, and Crises	978 90 5668 658 1	September 2021
658	Ernst Roos	Robust Approaches for Optimization Problems with Convex Uncertainty	978 90 5668 659 8	September 2021
659	Joren Koëter	Essays on asset pricing, investor preferences and derivative markets	978 90 5668 660 4	September 2021
660	Ricardo Barahona	Investor Behavior and Financial Markets	978 90 5668 661 1	October 2021
660	Stefan ten Eikelder	Biologically-based radiation therapy planning and adjustable robust optimization	978 90 5668 662 8	October 2021
661	Maciej Husiatyński	Three essays on Individual Behavior and New Technologies	978 90 5668 663 5	October 2021
662	Hasan Apakan	Essays on Two-Dimensional Signaling Games	978 90 5668 664 2	October 2021
663	Ana Moura	Essays in Health Economics	978 90 5668 665 9	November 2021
664	Frederik Verplancke	Essays on Corporate Finance: Insights on Aspects of the General Business Environment	978 90 5668 666 6	October 2021

No.	Author	Title	ISBN	Published
665	Zhaneta Tancheva	Essays on Macro-Finance and Market Anomalies	978 90 5668 667 3	November 2021
666	Claudio Baccianti	Essays in Economic Growth and Climate Policy	978 90 5668 668 0	November 2021
667	Hongwei Zhang	Empirical Asset Pricing and Ensemble Machine Learning	978 90 5668 669 7	November 2021
668	Bart van der Burgt	Splitsing in de Wet op de vennootschapsbelasting 1969 Een evaluatie van de Nederlandse winstbelastingregels voor splitsingen ten aanzien van lichamen	978 90 5668 670 3	December 2021
669	Martin Kapons	Essays on Capital Markets Research in Accounting	978 90 5668 671 0	December 2021
670	Xolani Nghona	From one dominant growth mode to another: Switching between strategic expansion modes	978 90 5668 672 7	December 2021
671	Yang Ding	Antecedents and Implications of Legacy Divestitures	978 90 5668 673 4	December 2021
672	Joobin Ordoobody	The Interplay of Structural and Individual Characteristics	978 90 5668 674 1	February 2022
673	Lucas Avezum	Essays on Bank Regulation and Supervision	978 90 5668 675 8	March 2022
674	Oliver Wichert	Unit-Root Tests in High-Dimensional Panels	978 90 5668 676 5	April 2022
675	Martijn de Vries	Theoretical Asset Pricing under Behavioral Decision Making	978 90 5668 677 2	June 2022
676	Hanan Ahmed	Extreme Value Statistics using Related Variables	978 90 5668 678 9	June 2022
677	Jan Paulick	Financial Market Information Infrastructures: Essays on Liquidity, Participant Behavior, and Information Extraction	978 90 5668 679 6	June 2022
678	Freek van Gils	Essays on Social Media and Democracy	978 90 5668 680 2	June 2022

No.	Author	Title	ISBN	Published
679	Suzanne Bies	Examining the Effectiveness of Activation Techniques on Consumer Behavior in Temporary Loyalty Programs	978 90 5668 681 9	July 2022
680	Qinnan Ruan	Management Control Systems and Ethical Decision Making	978 90 5668 682 6	June 2022
681	Lingbo Shen	Essays on Behavioral Finance and Corporate Finance	978 90 5668 683 3	August 2022
682	Joshua Eckblad	Mind the Gales: An Attention-Based View of Startup Investment Arms	978 90 5668 684 0	August 2022
683	Rafael Greminger	Essays on Consumer Search	978 90 5668 685 7	August 2022
684	Suraj Upadhyay	Essay on policies to curb rising healthcare expenditures	978 90 5668 686 4	September 2022
685	Bert-Jan Butijn	From Legal Contracts to Smart Contracts and Back Again: An Automated Approach	978 90 5668 687 1	September 2022
686	Sytse Duiverman	Four essays on the quality of auditing: Causes and consequences	978 90 5668 688 8	October 2022
687	Lucas Slot	Asymptotic Analysis of Semidefinite Bounds for Polynomial Optimization and Independent Sets in Geometric Hypergraphs	978 90 5668 689 5	September 2022
688	Daniel Brosch	Symmetry reduction in convex optimization with applications in combinatorics	978 90 5668 690 1	October 2022
689	Emil Uduwalage	Essays on Corporate Governance in Sri Lanka	978 90 5668 691 8	October 2022
690	Mingjia Xie	Essays on Education and Health Economics	978 90 5668 692 5	October 2022
691	Peerawat Samranchit	Competition in Digital Markets	978 90 5668 693 2	October 2022
692	Jop Schouten	Cooperation, allocation and strategy in interactive decision-making	978 90 5668 694 9	December 2022

No.	Author	Title	ISBN	Published
693	Pepijn Wissing	Spectral Characterizations of Complex Unit Gain Graphs	978 90 5668 695 6	November 2022
694	Joris Berns	CEO attention, emotion, and communication in corporate financial distress	978 90 5668 696 3	November 2022
695	Tom Aben	The (long) road towards smart management and maintenance: Organising the digital transformation of critical infrastructures	978 90 5668 697 0	December 2022
696	Gülbike Mirzaoğlu	Essays in Economics of Crime Prevention and Behavior Under Uncertainty	978 90 5668 698 7	February 2023
697	Suwei An	Essays on incentive contracts, M&As, and firm risk	978 90 5668 699 4	February 2023
698	Jorgo Goossens	Non-standard Preferences in Asset Pricing and Household Finance	978 90 5668 700 7	February 2023

This dissertation is a collection of four papers using non-standard preferences to better understand the behavior of asset prices and households. The first paper shows that present bias in an equilibrium model has the ability to explain multiple features of bond behavior. The second paper investigates the consequences of regret aversion for asset prices in an otherwise standard model of financial markets. The third paper combines experimentally elicited preferences with administrative microdata and explains actual annuitization decisions. The fourth paper demonstrates that risk and time preferences are time varying, and that they are related to trading behavior.

Jorgo Theodorus Godefridus Goossens (Breda, Netherlands, 1993) received his bachelor's degree in Econometrics and Operations Research at Tilburg University in 2015. He obtained his first master's degree cum laude in Quantitative Finance and Actuarial Science in 2016 at Tilburg University. He obtained his second master's degree in Science, Education, and Communication in 2017 at Eindhoven University of Technology, Faculty of Applied Mathematics. In 2018 he obtained his research master's degree with distinction in Finance at Tilburg University. His four years of Ph.D. were funded by a NWO Industrial Doctorate with cooperation between Tilburg University and APG. In the fall of 2021, he visited the University of Chicago, Booth School of Business. For several years, he has been affiliated with APG and the Network for Studies on Pensions, Aging and Retirement (Netspar). Since September 2022, he is an Assistant Professor in Finance at Radboud University Nijmegen.

ISBN: 978 90 5668 700 7
DOI: 10.26116/7mnr-hz54

TILBURG SCHOOL OF ECONOMICS
AND MANAGEMENT

CentER Dissertation Series