Network for Studies on Pensions, Aging and Retirement

NETSPAR ACADEMIC SERIES

Quantifying Ambiguity Bounds Through Hypothetical Statistical Testing

Anne G. Balter and Antoon Pelsser

DP 11/2015-038

Quantifying Ambiguity Bounds Through Hypothetical Statistical Testing

Anne G. Balter
Department of Econometrics and Operations Research, Tilburg University, Warandelaan 2, 5037AB, Tilburg, The Netherlands and Netspar; a.g.balter@uvt.nl

Antoon Pelsser
Department of Quantitative Economics, Maastricht University, Tongersestraat 53, 6211LM, Maastricht, The Netherlands, Netspar and Kleynen Consultants; a.pelsser@maastrichtuniversity.nl

June 20, 2017

Models can be wrong and recognising their limitations is important in financial and economic decision making under uncertainty. Robust strategies, which are least sensitive to perturbations of the underlying model, take uncertainty into account. Finding the explicit set of alternative models surrounding the baseline model has been difficult so far. We specify alternative models by a stochastic change of probability measure and derive a quantitative bound on the uncertainty set. We find an explicit ex ante relation between the choice parameter k, which is the radius of the uncertainty set, and the Type I and II error probabilities on the statistical test that is hypothetically performed to investigate whether the model specification is still appropriate at the future test horizon. The hypothetical test is constructed to obtain all alternative models that cannot be distinguished from the baseline model with enough power. Moreover, we also link the bound k, which is now a function of interpretable variables, to numerical values on several divergence measures. Finally, we illustrate the methodology on a robust investment problem and identify how the robustness multiplier can be numerically interpreted by ascribing meaning to the amount of ambiguity.

Key words: model uncertainty; robustness; stochastic control; optimisation; misspecification; divergence; rectangularity; investment problem

1. Introduction

Models can be wrong and recognising their limitations is important in financial and economic decision making. To capture the uncertainty, an agent is postulated who is concerned about model misspecification. Consequently the agent would like to evaluate the optimal decision rule against plausible alternative models and search for the robust decision. Intuitively, these alternative models come from a set surrounding the baseline model. The literature is mostly silent on how to specify the set of alternative models. Our ambition is to give an explicit interpretable characterisation of such a set. For an econometrician, a natural approach would be to construct a confidence interval based on the estimated parameters of the baseline model. However, the limitation of this approach is that only a very specific class of alternative models is considered: namely only those models that are equivalent to the baseline model with different *constant* parameter values. In this paper we consider a much larger class of alternative models beyond parameter uncertainty of the baseline model. Since we allow for stochastic changes of measure, not only the drift but the distribution itself is also subject to changes. An educated guess of the amount of uncertainty is not straightforward, therefore we use statistical methods to derive the bounds on the uncertainty set. Moreover, the bound that differentiates plausible from implausible models does not depend on historic data since we explicitly take into account that data contains uncertainty which underlies the model misspecification. We do allow the data to be used for estimation or calibration of the baseline model, but we argue that a solution cannot be robust by relying solely on events happened in the past. Therefore we differentiate from data-driven methods which rely mainly on historic confidence sets for a given Type I error. In other words, data can be used to create the baseline model. The baseline can also be justified by other factors such as economic theory. Subsequently the agent acknowledges that the real world is highly unlikely to be described by this baseline model. Therefore he indicates his concern by a robust optimisation problem subject to a set of plausible alternative models. Since past data does not guarantee future robust decisions and since this finite data set that posses uncertainty has already been relied on for calibration of the baseline, the amount of ambiguity we develop is independent of past data to preclude from using the limited data set to create both the baseline models and the robust decision. We disentangle the amount of ambiguity, i.e. how far plausible alternative models are from the baseline model, into variables whose values entail numerical interpretation. And vice versa, given that an agent allows a Type I error of α, that he wants at least a power of $1-\beta$ in order to distinguish between two models and that he assumes that the structural form of the baseline model remains constant for the next $\mathcal{T}$ years, the radius on the ellipsoid of plausible alternative models is determined. The amount of uncertainty can be translated to a bound on the Kullback-Leibler divergence and other ϕ-divergences as well.

We briefly discuss the literature with respect to constructing sets of alternative models. In robust optimisation the worst-case model is picked from an uncertainty set, whereafter the optimal decision rule is determined under the worst-case model. This strategy is accomplished by the maxmin method from Gilboa and Schmeidler (1989). General optimisation problems, examples ranging from finance to operations research, are solved robustly over an uncertainty set.

Ben-Tal et al. (2013) identify confidence sets in terms of a specific ϕ-divergence function. The described divergence functions are Kullback-Leibler divergence, Burg entropy, J-divergence, χ^2-distance, Hellinger distance, Variation distance and Cressie-Read divergence. They derive the asymptotic distribution of the divergence measures from which the quantiles imply the bounds on the uncertainty sets for finite samples. In Section 4 we show how our methodology is able to quantify the bounds on these different divergences explicitly. In the field of operations research the literature that deals with model uncertainty is known as distributionally robust optimisation. Tractability of the optimisation problem is of main focus in for example Bertsimas et al. (2017), Esfahani and Kuhn (2017), Wiesemann et al. (2014), Delage and Ye (2010), Ben-Tal et al. (2010), Bertsimas and Brown (2009), Nilim and El Ghaoui (2005), and for an overview with respect to the recent developments see Gabrel et al. (2014). Note that the non-convex optimisation problem that we solve in this paper is independent from the agent's objective. The optimisation in this paper arises as a by-product of creating the radius on the uncertainty set, since we search for the maximum power among the stochastic changes of measures. This one time optimisation leads to Theorem 1 and serves as a statistically foundation to robustify decisions under a realistic range of uncertainty.

Breuer and Csiszár (2013) base stress tests on plausible sets, which do not include scenarios that are "too implausible". These scenarios are obtained by considering mixed scenarios, known as risk factor distributions. The Kullback-Leibler divergence from these distributions has to fall within an uncertainty ball, where the radius is assumed to coincide with the 1% TVaR or the bound can be calculated based on historical data. However, for all divergences the critical value that distinguishes a plausible from an implausible model is not quantified. Cambou and Filipović (2015) combine model uncertainty based on stress testing with views for regulatory purposes.

Hansen and Sargent (2008) motivate their approach of uncertainty to robust optimisation by choosing models surrounding the baseline model with bounded entropy (i.e. with bounded Kullback-Leibler divergence). However, when implementing their method they make a subtle switch: they replace the endogenous Lagrange multiplier of the entropy bound by a fixed exogenous parameter. Thus they create an entropy penalty for *all* alternative models, whereas we consider a *bounded* set of alternative models. How to choose the fixed entropy penalty parameter is not yet fully determined. Anderson et al. (2003) and Hansen and Sargent (2016) start with an optimisation problem, then they pick an entropy parameter value and calculate the worst-case path which depends on the penalty parameter and the specific optimisation problem. Next they calculate the detection error probability for the specific entropy parameter and the associated worst-case path *ex post*. If the probability of the average of the two incorrect rejections is too high, then the worst-case choice from mother nature is too extreme. Hence this parameter is rejected, a smaller value is considered and the whole procedure is repeated. By this procedure one plausible worst-case based parameter value is selected for each specific optimisation problem. The main difference with our approach is the order of the procedure. Hansen and Sargent (2016) find a penalty parameter for each specific optimisation problem whereas we focus on the characterisation of the set of plausible alternative models that is independent of the particular formulation of the optimisation problem and our set is created *ex ante* since the uncertainty is due to the mismatch between the selected baseline model which is obtained based on theory and data, and the true model that describes the future process. In Hansen and Sargent (2011) they argue that the ability to quantify model ambiguity a priori is important to understand empirical results. They use statistical detection theory on the likelihood of the available data to quantify uncertainty on state transition dynamics and hidden states. In Hansen et al. (2006a) concern is shown about the fact that in applications the the maxmin method of Gilboa and Schmeidler (1989) has been used without any theoretical justification on the specification of the uncertainty set. In Hansen et al. (2006a) four papers on decision making under uncertainty are briefly summarised.

Maccheroni et al. (2006a) show the connection between the multiple priors method from Gilboa and Schmeidler (1989) and the multiplier preferences method from Hansen and Sargent (2001). Both can be combined into a general representation, where the former's utility function equals infinity for models outside the set of plausible alternative models and the latter ranks *all* alternative models based on its divergence, in terms of entropy, relative to the baseline model. Maccheroni et al. (2006a)'s representation is general for the divergence function considered. The mean-variances preferences of Markowitz (1952) and Tobin (1958) are included for the choice of the divergence preference equal to the χ^2-divergence. We relate the set of indistinguishable models to, among others, both these two divergence functions. In Maccheroni et al. (2006b), they provide a condition that ensures time consistency of the variational preferences. From the experimental literature, Eichberger et al. (2016) discuss time consistency of smooth models that is assumed via a preference for ex ante or ex post randomisation. Klibanoff et al. (2005) achieve a separation between ambiguity perception, described by the agent's probability distribution, and ambiguity attitude on top of risk aversion. They obtain smooth indifference curves for decision making under ambiguity.

Hansen et al. (2011) introduce the concept of Model Confidence Set (MCS). This method, or actually algorithm, is a sequential method that starts with explicitly enumerating a finite number of alternative models. The procedure results in a subset that contains the models that cannot be rejected, in terms of a *chosen* test statistic. The MCS method is based on a statistical test and an elimination rule applied to historical data. If the test indicates that the set of models at hand are not indistinguishable, then at least one of the models performs worse than the others and should be removed.

Our contribution is to obtain an explicit characterisation of the set of indistinguishable models based on statistical testing theory. First, the agent selects a baseline model. However, he is worried about the model specification and indicates his uncertainty by plausible alternatives. We use the Neyman and Pearson (1933) Lemma as a device to ascribe meaning to the amount of ambiguity that bounds the set of models that cannot be distinguished statistically from a baseline model. The power of the test is used as selection criterion to determine whether an alternative model belongs to the plausible set or not. We use the Radon-Nikodym derivative which is equivalent to the likelihood ratio test for testing between the two simple hypothesis of the baseline model against a change of measure. By the Neyman-Pearson Lemma this is the most powerful test. Unlike most data-driven methods we do not perform actual tests. We only use the power of the likelihood ratio test, which is computable ex ante, to characterise our set. Therefore the set of indistinguishable models can explicitly be obtained as a function of the Type I error α, the Type II error β and the test horizon $\mathcal{T}$. If the test between the baseline and an alternative model has an insufficient power for a given Type I error, then we define this alternative as *indistinguishable*. The test horizon can be interpreted as the amount of additional information that one would take into consideration or the time span during which one believes that there will not be a structural break in the data generating process. Thus the larger $\mathcal{T}$, the lower the uncertainty. Note that the uncertainty entertained in this paper has to do with potential changes in probability that could occur in the future and are not revealed by history. Thus the decision maker is not actively trying to model testing at the same time as he is doing the portfolio optimisation.

The outline is as follows: in Section 2 we introduce the model and derive the indistinguishable set of models for deterministic alternatives that serves as an intuitive illustration. Then we show an example of a stochastic alternative that fundamentally changes the probability distribution. In Section 3 our main contribution is presented where we create the set of indistinguishable models for stochastic rectangular alternatives. The link to a numerical bound on several divergences is provided in Section 4. Finally, in Section 5 we apply the uncertainty set to the robust investment problem.

2. Statistically Indistinguishable Models

Let us make our model set-up explicit. We assume that we are considering models that can be described by diffusion processes. This means that we are considering stochastic processes X that are described by stochastic differential equations of the form

$$dX(t,\omega) = \mu(t,\omega)\,dt + \sigma(t,\omega)\,dW(t,\omega). \tag{1}$$

For a control variable π, i.e. the fraction invested in X, the payoff process is $d(\pi X(t,\omega)) = \pi\mu(t,\omega)\,dt + \pi\sigma(t,\omega)\,dW(t,\omega)$. The remainder, $1-\pi$, is often invested in the risk-free asset. Together with the consumption level this typically determines the wealth process. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, where $\mathbb{P}$ defines the baseline model. For the specification of possible alternative models, we consider the Brownian motion with a stochastic drift $dW(t,\omega) + \lambda(t,\omega)\,dt$. The ω indicates that

$\{\lambda(t,\omega)\}_{0\leq t\leq \mathcal{T}}$ will be stochastic and may depend on the whole path. Such an alternative model specification can be captured as a change of probability measure from $\mathbb{P}$ to a new probability measure $\mathbb{L}$. Because of this correspondence we will denote both the alternative model and the alternative probability measure by $\mathbb{L}$.

Robustness with respect to optimisation problems arises by the acknowledgement of model uncertainty, i.e. an agent maximising his expected utility might be concerned about the underlying model specification. Gilboa and Schmeidler (1989) introduced the maxmin method to be safe against perturbations of the underlying processes. A so called "malevolent mother nature" picks the worst-case model among a set of plausible models, whereafter the agent selects the strategy that is optimal given mother nature's choice. Let π be the agent's control variable, let $U(\cdot)$ be an objective function and let the set $\mathcal{L}$ be the set of probability measures $\mathbb{L}$ from which mother nature can choose. The robust optimisation problem is given by

$$\max_{\pi}\min_{\mathbb{L}\in\mathcal{L}}\mathbb{E}^{\mathbb{L}}[U(X)]. \tag{2}$$

The aim of this paper is to characterise the set $\mathcal{L}$.

The likelihood ratio $\left.\frac{d\mathbb{L}}{d\mathbb{P}}\right|_{\mathcal{T}}$, which is based on the information over the interval $[0,\mathcal{T}]$ is given by the random value $L(\mathcal{T},\omega)$ at time $\mathcal{T}$ and can be characterised by the stochastic differential equation

$$dL(t,\omega)=\lambda(t,\omega)L(t,\omega)\,dW^{\mathbb{P}}(t,\omega) \text{ over } [0,\mathcal{T}]. \tag{3}$$

with initial condition $L(0)=1$. The superscript $\mathbb{P}$ denotes the probability measure we are considering. The process $L(t,\omega)$ only represents a valid Radon-Nikodym derivative if $L(t,\omega)$ is a true martingale (and not a local martingale). The random variable $L(\mathcal{T},\omega)$ can be represented as

$$L(\mathcal{T},\omega)=\exp\left\{-\tfrac{1}{2}\int_0^{\mathcal{T}}\lambda(t,\omega)^2\,dt+\int_0^{\mathcal{T}}\lambda(t,\omega)\,dW^{\mathbb{P}}(t,\omega)\right\}. \tag{4}$$

Hence, the value at time $\mathcal{T}$ of the likelihood ratio $L(\mathcal{T},\omega)$ is completely determined by the realisation ω of a path of the Brownian motion $\{W^{\mathbb{P}}(t,\omega)\}_{0\leq t\leq\mathcal{T}}$ and the process $\{\lambda(t,\omega)\}_{0\leq t\leq\mathcal{T}}$ that characterises the alternative model $\mathbb{L}$ along this path.

Based on the realised path of the Brownian motion we could test at time $\mathcal{T}$ if model $\mathbb{P}$ should be rejected in favour of model $\mathbb{L}$ for the specific process $\lambda(t,\omega)$. We are testing two simple hypotheses, $H_0:\mathbb{P}$ versus $H_A:\mathbb{L}$, and the Neyman-Pearson Lemma tells us that the most powerful test between two simple hypotheses is a likelihood ratio test. The form of the optimal test procedure is that we reject model $\mathbb{P}$ if $L(\mathcal{T},\omega)$ is larger than the critical value ξ. It is crucial to realise that before time $\mathcal{T}$ the test statistic is a random variable and we can impose probabilities on the errors. Hence, it is important to emphasis once more that the Radon-Nikodym derivative serves as a *hypothetical* test. In other words, we use the test procedure as an *ex ante* selection criterion to characterise the set of indistinguishable models. In this light, the decision maker is not actively trying to model testing at the same time as he is doing the portfolio optimisation.

To implement our test procedure, we derive the critical value ξ from the equation

$$\mathbb{P}[L(\mathcal{T},\omega) \geq \xi] = \alpha. \tag{5}$$

We set the critical value ξ such that probability of incorrectly rejecting model $\mathbb{P}$ when model $\mathbb{P}$ is the true model is equal to α. This is known as the Type I error. The probability α is the significance level of the test, and is typically set at 5%. [1]

We should also be worried about the Type II error: this is the error of incorrectly rejecting model $\mathbb{L}$ when model $\mathbb{L}$ is the true model. This probability is typically denoted by β and can be computed as

$$\mathbb{L}[L(\mathcal{T},\omega) < \xi] = \beta. \tag{6}$$

The complement of the Type II error is the probability of accepting model $\mathbb{L}$ when model $\mathbb{L}$ is the true model. This is known as the *power* of the statistical test. The power can be computed as

$$\mathbb{L}[L(\mathcal{T},\omega) \geq \xi] = 1 - \beta. \tag{7}$$

A typical value for β is 20%, leading to a statistical power of 80%. It is important to note that the Type I and Type II probabilities can be computed at time 0 for a given alternative model $\mathbb{L}$.

The model selection procedure we propose is based on the *power* of the likelihood ratio test. The intuition is a follows: for small values of $\lambda(t,\omega)$ we will have a model $\mathbb{L}$ that is "close" to the baseline model $\mathbb{P}$. This closeness can be identified by the fact that the likelihood ratio $L(\mathcal{T},\omega)$ will be a random variable with a probability distribution tightly concentrated around the value $L(\mathcal{T},\omega) = 1$. Even though we can define a critical value ξ for any model $\mathbb{L}$, for models that are "close" there will be almost no difference between the $\mathbb{P}$-probability and the $\mathbb{L}$-probability of the event $L(\mathcal{T},\omega) \geq \xi$. Hence, the power of the statistical test will be very low. In the limiting case when $\mathbb{P} = \mathbb{L}$ the power of the (randomised) likelihood ratio test will be as low as α. Thus for models close to $\mathbb{P}$, i.e. low $\lambda(t,\omega)$, it is difficult to distinguish $\mathbb{L}$ from $\mathbb{P}$ and therefore $\mathbb{L}$ is a plausible alternative.

Hence, our model selection criterion will include all models for which the statistical power $\mathbb{L}[L(\mathcal{T},\omega) \geq \xi]$ is below $1-\beta$. We consider these models to be *statistically indistinguishable* from the baseline model $\mathbb{P}$. If the associated power is too high the alternative model $\mathbb{L}$ characterised by $\lambda(t,\omega)$ is excluded from the set of indistinguishable models because with enough statistical power the alternative can be distinguished from the baseline model. Thus the size of the set is determined by α, β and $\mathcal{T}$.

We can express $\mathbb{L}[L(\mathcal{T},\omega)\geq\xi]$ also as $\mathbb{E}^{\mathbb{P}}[L(\mathcal{T},\omega)\mathbb{1}(L(\mathcal{T},\omega)\geq\xi)]$, and we obtain an interpretation for the $\mathbb{P}$-expectation as the Tail-Value-at-Risk (TVaR) or Conditional-Value-at-Risk (CVaR) of the random variable $L(\mathcal{T},\omega)$ with a confidence level of α and a time horizon of $\mathcal{T}$. Hence, if we put an upper bound of $1-\beta$ on the power of the likelihood ratio test, this is equivalent to restricting the TVaR of the random variable $L(\mathcal{T},\omega)$ to $1-\beta$. Tail-Value-at-Risk is a *coherent risk measure* (see, Artzner et al. (1999) and Rockafellar and Uryasev (2000, 2002)), which has attractive properties. For example, the acceptance set (that is the set of all $L(\mathcal{T},\omega)$ for which $\text{TVaR}(L(\mathcal{T},\omega))\leq 1-\beta$) is a closed convex set. Hence, we obtain immediately that our set of indistinguishable models is also a closed convex set.

2.1. Deterministic Drift Term

For deterministic $\lambda(t)$ we can compute the probabilities explicitly, though our main goal is to find the set of indistinguishable models for stochastic $\lambda(t,\omega)$. For alternative models $\mathbb{L}$ with a deterministic drift term $\lambda(t)$, we can compute the probability distribution of the Radon-Nikodym derivative defined in equation (3) explicitly. In this case we obtain

$$L(\mathcal{T},\omega)=\exp\left\{-\tfrac{1}{2}\int_0^{\mathcal{T}}\lambda(t)^2\,dt+\int_0^{\mathcal{T}}\lambda(t)\,dW^{\mathbb{P}}(t,\omega)\right\}. \tag{8}$$

In particular, $\ln L(\mathcal{T},\omega)$ has a normal distribution with mean $-\frac{1}{2}\int_0^{\mathcal{T}}\lambda(t)^2\,dt$ and the variance $\int_0^{\mathcal{T}}\lambda(t)^2\,dt$. The likelihood ratio test procedure $L(\mathcal{T},\omega)>\xi$ is equivalent to performing a test on the statistic

$$l(\mathcal{T},\omega)=\int_0^{\mathcal{T}}\lambda(t)\,dW^{\mathbb{P}}(t,\omega). \tag{9}$$

This test statistic is intuitively appealing: we compute the inner product between the model-drift $\lambda(t)$ and the realised changes $dW(t,\omega)$ in the Brownian motion along the whole path $[0,\mathcal{T}]$. Every time $\lambda(t)$ and $dW(t,\omega)$ have the same sign, this increases the value of $l(\mathcal{T},\omega)$. Hence, if model $\mathbb{L}$ is true, then $l(\mathcal{T},\omega)$ will on average have a positive value. *Ex post*, the realisation of the path of $W(\mathcal{T},\omega)$ is observed and indicates the likelihood whether it was generated by model $\mathbb{P}$ or $\mathbb{L}$. However, *ex ante* the test statistic is a random variable for which we can compute the power.

The test statistic $l(\mathcal{T},\omega)$ has a normal distribution with mean 0 and variance $\int_0^{\mathcal{T}}\lambda(t)^2\,dt$ under the baseline model $\mathbb{P}$. The hypothesis $\mathbb{P}$ is rejected if $l(\mathcal{T},\omega)\geq\zeta$. Under the alternative model $\mathbb{L}$ the test statistic has a normal distribution with mean $\int_0^{\mathcal{T}}\lambda(t)^2\,dt$ and variance $\int_0^{\mathcal{T}}\lambda(t)^2\,dt$. The power of the test can be computed explicitly as

$$\mathbb{L}[l(\mathcal{T},\omega)>\zeta]=\Phi\left(\Phi^{-1}(\alpha)+\left(\int_0^{\mathcal{T}}\lambda(t)^2\,dt\right)^{\frac{1}{2}}\right). \tag{10}$$

For the case $\lambda(t) \equiv 0$ we see that the power is $\Phi\big((\Phi^{-1}(\alpha)\big) = \alpha$. For non-zero values of $\lambda(t)$ the expression $(\int_0^{\mathcal{T}} \lambda(t)^2\,dt)^{\frac{1}{2}}$ is strictly positive and therefore the power will be larger than α. If we consider all models with a power below $1-\beta$ as indistinguishable, then the class of indistinguishable models (with deterministic $\lambda(t)$) is given by all models for which the L_2-norm is below a certain threshold

$$\left(\int_0^{\mathcal{T}} \lambda(t)^2\,dt\right)^{\frac{1}{2}} \leq \Phi^{-1}(1-\beta) - \Phi(^{-1}(\alpha)). \tag{11}$$

If we take for example $\alpha = 5\%$, then $\Phi^{-1}(\alpha) = -1.64$ and if $\beta = 20\%$ then the power is $1-\beta = 80\%$ and we have $\Phi^{-1}(1-\beta) = 0.84$. Hence, the class of all indistinguishable models is then given by all models that satisfy $(\int_0^{\mathcal{T}} \lambda(t)^2\,dt)^{\frac{1}{2}} \leq 0.84 - (-1.64) = 2.48$ for a test horizon of $\mathcal{T}$.

The deterministic example we have formulated can be generalised easily to the *multi-dimensional case*. For a vector-valued Brownian motion all alternative models are specified by the deterministic vector-valued process $\lambda(t)$. The test statistic $l(t,\omega)$ is then given by

$$l(\mathcal{T},\omega) = \int_0^{\mathcal{T}} \lambda(t)\cdot dW^{\mathbb{P}}(t,\omega). \tag{12}$$

This is also a random variable with mean 0 and variance $\int_0^{\mathcal{T}} |\lambda(t)|_2^2\,dt$, where $|\lambda(t)|_2$ denotes the L_2-norm of the vector $\lambda(t)$. Hence, in the multi-dimensional case the set of indistinguishable models is given by all models for which $(\int_0^{\mathcal{T}} |\lambda(t)|_2^2\,dt)^{\frac{1}{2}}$ is below the same threshold as in the one-dimensional case (e.g. 2.48).

Thus the cutoff point that distinguishes plausible alternatives from implausible ones, is a function of the Type I error α, the Type II error β and the test horizon $\mathcal{T}$. The plausible alternatives yield an insufficient power, hence the power is below $1-\beta$. And the test horizon can be interpreted as the amount of additional information one would take into consideration. In other words, if the latter variable goes to infinity this implies that one has an infinite amount of knowledge and is not uncertain at all. Likewise, the uncertainty parameter drops to zero. If $\lambda(t)$ is constant, then the bound that distinguishes plausible from implausible models simplifies further to

$$|\lambda| \leq \frac{\Phi^{-1}(1-\beta) - \Phi(^{-1}(\alpha))}{\sqrt{\mathcal{T}}}. \tag{13}$$

For the right-hand side of equation (13) we introduce the notation k which we call the level of uncertainty or ambiguity bound. Similarly as before, $\mathcal{T}$ represents the future horizon during which no structural break is supposed to occur.

The deterministic $\lambda(t)$ serves as an intuitive illustration of the concepts that we have introduced, but our ambition is to consider a much larger class of alternative models: $\lambda(t,\omega)$. If we allow for stochastic $\lambda(t,\omega)$ then a very large class of alternative models is accessible over an interval $[0,\mathcal{T}]$. By the Martingale Representation Theorem *any* probability distribution (with support on whole $\mathbb{R}$) can be attained over an interval $[t,t+\varepsilon]$ with $\varepsilon > 0$. In other words, a deterministic change of measure yields only other parameters of the same distribution as the baseline model. Actual model uncertainty considers also different distributions that are close to the baseline model to be plausible. This latter set is obtained by stochastic changes of measure. The following example shows that for a change of measure that switches randomly between two bounds, the alternative models have both densities and variances different than the baseline distribution.

2.2. Stochastic Drift Term: Bang-Bang

As an illustrative example, let us consider a model with stochastic $\lambda(t,\omega) = a \cdot \mathrm{sgn}(W^{\mathbb{P}}(t))$, where the sign function is defined as

$$\mathrm{sgn}(x) = \begin{cases} +1 & x > 0 \\ 0 & x = 0 \\ -1 & x < 0. \end{cases} \tag{14}$$

Based on this $\lambda(t,\omega)$ the Radon-Nikodym derivative is

$$\ln L(\mathcal{T}) = \int_0^{\mathcal{T}} a \cdot \mathrm{sgn}(W^{\mathbb{P}}(t)) dW^{\mathbb{P}}(t) - \frac{1}{2}\int_0^{\mathcal{T}} a^2 \mathrm{sgn}^2(W^{\mathbb{P}}(t)) dt. \tag{15}$$

Girsanov Theorem states that the Radon-Nikodym derivative describes how the model changes when the baseline measure $\mathbb{P}$ changes to $\mathbb{L}$. It defines the process $dW^{\mathbb{L}}(t) = dW^{\mathbb{P}} - a \cdot \mathrm{sgn}(W^{\mathbb{P}}(t))dt$ to be a Brownian motion under the measure $\mathbb{L}$. However, under the process $\mathbb{L}$ the process $W^{\mathbb{P}}(t)$ is no longer a Brownian motion. Therefore we introduce $x(t) = W^{\mathbb{P}}(t)$ to facilitate the notation. Hence, the alternative model under $\mathbb{L}$ is

$$dx(t) = a \cdot \mathrm{sgn}(x(t))dt + dW^{\mathbb{L}}(t). \tag{16}$$

The alternative model $\mathbb{L}$ is a mean-repelling process for $a > 0$, which will increase the variance of $x(\mathcal{T})$ under $\mathbb{L}$. And $a < 0$ leads to a mean-reverting model, which will decrease the variance of $x(\mathcal{T})$ under $\mathbb{L}$. The plot of the probability density function of process x under $\mathbb{L}$ is given in Figure 1. This clearly shows that with a stochastic $\lambda(t,\omega)$ we can fundamentally alter the model beyond the normal distribution. This illustrates the rich class of alternative models that can be attained by using stochastic $\lambda(t,\omega)$.

Furthermore, to demonstrate our test procedure once more we need to find the distribution of $L(\mathcal{T})$ to implement the test. Karatzas and Shreve (1984) show that the Radon-Nikodym derivative can be expressed in terms of the absolute Brownian motion $|W(\mathcal{T})|$ and the local time $\ell(\mathcal{T})$

$$\ln L(\mathcal{T}) = a\left(|W(\mathcal{T})| - 2\ell(\mathcal{T})\right) - \frac{1}{2}a^2\mathcal{T}. \tag{17}$$

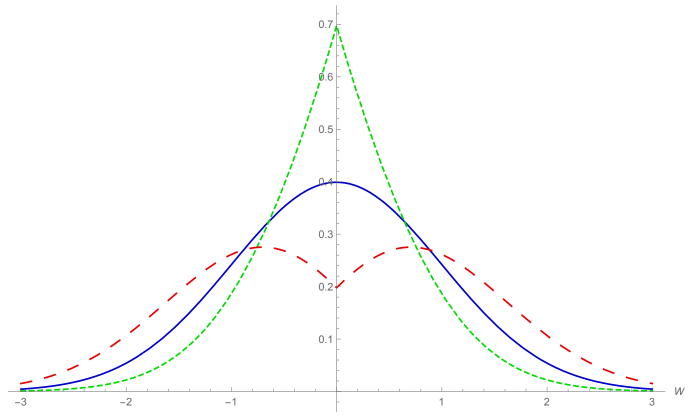

Figure 1 Densities of Bang-Bang drift

The figure compares densities for bang-bang drift processes with $a = 0.5$ (short-dashed green line), $a = -0.5$ (long-dashed red line) and a standard Brownian motion (blue line).

The joint distribution of the absolute Brownian motion and the local time at zero is given by equation (1.5) in Karatzas and Shreve (1984) and on page 45 in Itô and McKean (1974)

$$f_{|W_t|,2\ell_t}(w,l) = \frac{2(w+l)}{\sqrt{2\pi\mathcal{T}^3}} e^{-\frac{(w+l)^2}{2\mathcal{T}}}, \tag{18}$$

for $w, l > 0$. This integrates out to a normal density for $\ln L(\mathcal{T})$ with mean $-\frac{1}{2}a^2\mathcal{T}$ and variance $a^2\mathcal{T}$ under $\mathbb{P}$ and mean $\frac{1}{2}a^2\mathcal{T}$ and variance $a^2\mathcal{T}$ under $\mathbb{L}$. The parameter values for a, which imply indistinguishable models surrounding the baseline model $a \equiv 0$, are

$$|a| \leq \frac{\Phi^{-1}(1-\beta) - \Phi^{-1}(\alpha)}{\sqrt{\mathcal{T}}}. \tag{19}$$

The deterministic case and the stochastic example served as an illustration of our approach. However, we are interested in the generalisation of this method. In other words, we allow for a wide class of alternative models. Since stochastic alternatives can lead to fundamentally different models, we want to characterise the set of stochastic alternative models surrounding the baseline model that are indistinguishable based on an insufficient power and a specific probability on the Type I error.

3. Rectangularity and Coherence

One of the main motivations for studying the set of statistically indistinguishable models is to take model uncertainty explicitly into account, since models can be wrong and recognising their limitations is important in financial and economic decision making. A robust solution is safe against perturbations of the underlying model. This is obtained by the maxmin equation (2) that is introduced by Gilboa and Schmeidler (1989). The original objective is robustified by introducing a so-called malevolent player who minimises the objective over a set of indistinguishable models

$\mathcal{L}$. Based on the worst-case scenario, the agent searches now for the strategy that maximises the penalised objective. Hence, in order to obtain robust solutions to stochastic optimal control problems, a set of alternative models is needed. See Hansen and Marinacci (2016) and Watson and Holmes (2016) for an overview of the literature on decision making under uncertainty.

Time consistency is a basic requirement of rational decision making (Strotz 1955). Epstein and Schneider (2003) impose the rectangularity condition to ensure that dynamic behaviour is determined by preferences only. Riedel (2009) explains that the time consistency axiom requires that the decision maker sticks to the strategy and that the optimal investment at time zero for a sequential problem is not changed when new information arrives or when time elapses. Moreover, Riedel (2009) shows that dynamic consistency is equivalent to the notion of rectangularity in terms of Epstein and Schneider (2003). Note that Hansen and Marinacci (2016) also discuss the recent developments on rectangularity.

We impose time consistency on the set of indistinguishable models by a condition on $\lambda(t,\omega)$ that comes from the backward stochastic differential equation (BSDE) literature of coherent risk measures. Moreover, existence and uniqueness of the value function follows as a direct consequence from the BSDE framework. Hansen et al. (2006b) discuss robustness in combination with rectangularity as well, though they argue that in their setting it would not give rise to interesting restrictions. Riedel (2009) derives the dynamic consistency axiom, which is ensured by backward induction in the context of recursive utility. The existence of the worst-case measure requires rectangularity as well. Riedel (2004) extends the static coherent risk measure to the dynamic counterpart by a coherency axiom.

Thus, when we are solving optimal control problems, we want to consider sets of alternative models that are *rectangular*. The set of indistinguishable models we have defined thus far is *not* rectangular: the set is defined as those models that have sufficiently low power at time $\mathcal{T}$ using the information over the whole path $[0,\mathcal{T}]$. We have established in Section 2 that the set of indistinguishable models is defined by a coherent risk measure, the TVaR on the likelihood ratio $L(\mathcal{T})$. But, this risk measure is not rectangular since it is "static" at time 0 and at time $t>0$ the connection to the TVaR computed at time $t=0$ is lost.

We can however look at a smaller class of risk measures: the class of rectangular risk measures. This class has been extensively studied in recent years, and we know how to characterise this class of risk measures. Delbaen (2006) proves that rectangular (coherent and convex) risk measures are generated by m-stable sets of probability measures. A similar structure (albeit with less mathematical rigour) was already proposed by Epstein and Schneider (2003). An alternative characterisation of rectangular risk measures has been provided by Rosazza Gianin (2006). She proves that every rectangular risk measure is equivalent to a g-expectation $\mathcal{E}^g[]$ which can be represented as the solution of a BSDE with a driver $(g(t,Y,Z))$. A further characterisation has been provided by Barrieu and El Karoui (2007): they prove that rectangular coherent risk measures (such as TVaR) are generated by drivers $g(t,Z)$ that satisfy a Lipschitz growth constraint in Z. Barrieu and El Karoui (2007) also show that the BSDE driver $g(t,Z)$ is connected to $\lambda(t,\omega)$ via the Legendre-Fenchel transform.

To summarize these results: our set of indistinguishable models is characterised by a test on $\lambda(t,\omega)$. The power of the test procedure can be interpreted as the TVaR, which is a coherent risk measure. The upper bound of $1-\beta$ on the power of the Radon-Nikodym derivative is equivalent to a restriction on the TVaR of the random variable $L(\mathcal{T},\omega)$. The transition from the static test to the dynamic equivalent is accomplished by the set $\mathcal{L}$ that is coherent and rectangular. Using the Legendre-Fenchel duality we find that the class of rectangular coherent risk measures are generated by BSDE's with drivers that satisfy the Lipschitz growth condition $g(t,Y,Z) \le k|Z|$. These BSDE-drivers correspond to Radon-Nikodym derivatives with kernels $|\lambda(t,\omega)| \le k$ (Barrieu and El Karoui 2007, Chen and Epstein 2002, Epstein and Schneider 2003).

3.1. Maximum Power Calculation

We obtain an explicit characterisation on the rectangular set of alternative models in the following way. We consider the class of alternative models with[2] $|\lambda(t,\omega)| \le k$ and we link the value of k explicitly to the Type I and II error probabilities α and β respectively. We want to investigate the maximum power that can be achieved within the class of Radon-Nikodym derivatives with $|\lambda(t,\omega)| \le k$, such that the Type I error is equal to α. We can formulate this as a stochastic optimisation problem of the form

$$\begin{aligned}
\max_{\xi,\lambda(t,\omega)} \; & \mathbb{E}\left[L(\mathcal{T},\omega)\mathbb{1}(L(\mathcal{T},\omega) \ge \xi)\right] \\
\text{s.t. } & \mathbb{E}\left[\mathbb{1}(L(\mathcal{T},\omega) \ge \xi)\right] = \alpha \\
& dL(t,\omega) = \lambda(t,\omega)L(t,\omega)dW(t,\omega), L(0,\omega) = 1 \\
& |\lambda(t,\omega)| \le k.
\end{aligned} \tag{MP}$$

The objective function is the power of the test $L(\mathcal{T},\omega) \ge \xi$ formulated as a $\mathbb{P}$-expectation. The second line gives the Type I error (also formulated as a $\mathbb{P}$-expectation), and the third line describes the stochastic process for the Radon-Nikodym derivative given the control variable $\lambda(t,\omega)$. Note that in this general setting the Radon-Nikodym derivative is no longer lognormal distributed like in the deterministic case. Based on the derivation for the maximum power, all $\lambda(t,\omega)$'s with $|\lambda(t,\omega)| \le k$ represent alternative models that are indistinguishable. We summarise this in Theorem 1.

THEOREM 1. *(Rectangular and Coherent Sets of Indistinguishable Models)*
Consider a baseline model $dX(t,\omega) = \mu(t,\omega)dt + \sigma(t,\omega)dW(t,\omega)$. The set $\mathcal{L}$ of all models with $dW(t,\omega) + \lambda(t,\omega)dt$ and $|\lambda(t,\omega)| \leq k$ is rectangular and coherent. For

$$k = \frac{\Phi^{-1}(1-\beta) - \Phi^{-1}(\alpha)}{\sqrt{\mathcal{T}}} \tag{20}$$

every alternative in the set $\mathcal{L}$ is indistinguishable from the baseline model for a Type I error of α, a Type II error of β and a test horizon $\mathcal{T}$.

Proof of Theorem 1. The objective function is non-convex due to the indicator functions. Since it is easier to work with convex functions, we introduce the auxiliary function[3]

$$\begin{aligned} F_\alpha(L,\xi) &= \alpha\xi + \mathbb{E}\left[(L(\mathcal{T}) - \xi)^+\right] \\ &= \mathbb{E}\left[L(\mathcal{T})\mathbb{1}(L(\mathcal{T}) \geq \xi)\right] + (\alpha - \mathbb{E}\left[\mathbb{1}(L(\mathcal{T}) \geq \xi)\right])\xi, \end{aligned} \tag{21}$$

where $(L-\xi)^+$ denotes $\max(L-\xi,0)$. The functional $F_\alpha(L,\xi)$ is convex and continuous as a function of $\xi \in \mathbb{R}$ and $L(\mathcal{T}) \in L_2(\mathcal{T}) : \mathbb{E}[L(\mathcal{T})^2] < \infty$. See Rockafellar and Uryasev (2000, 2002) who introduce a similar auxiliary function to minimise the Conditional-Value-at-Risk (CVaR).
In order to solve the constrained optimisation problem (MP) we solve

$$\begin{aligned} &\max_{|\lambda(t,\omega)|\leq k} \min_{\xi} \; F_\alpha(L,\xi) \\ &\text{s.t. } dL = \lambda(t,\omega)LdW. \end{aligned} \tag{MaMi}$$

The optimisation (MaMi) is equivalent to (MP), which is shown by Rockafellar and Uryasev (2000, 2002). The Type I error constraint is satisfied by minimisation of the auxiliary function over ξ

$$\frac{\partial F_\alpha(L,\xi)}{\partial \xi} = \alpha - \mathbb{E}\left[\mathbb{1}(L(\mathcal{T}) \geq \xi)\right] = 0. \tag{22}$$

Since $F_\alpha(L,\xi)$ is convex in ξ, the extreme value is a minimum. In particular, after minimisation we have

$$\min_\xi F_\alpha(L,\xi) = \mathbb{E}\left[L(\mathcal{T})\mathbb{1}(L(\mathcal{T}) \geq \xi^*)\right].$$

This is the power that we would like to maximise. Henceforth we continue by maximising the minimised $F_\alpha(L,\xi^*)$. However, because this problem is difficult to solve we prefer to change the order of the optimisation. Solving the initial non-convex optimisation problem is identical with solving the max min. The optimal value from the (MaMi) is always lower than or equal to the reversed order optimisation by the max min inequality. Let (MiMa) be

$$\begin{aligned} &\min_{\xi} \max_{|\lambda(t,\omega)|\leq k} F_\alpha(L,\xi) \\ &\text{s.t. } dL = \lambda(t,\omega)LdW. \end{aligned} \tag{MiMa}$$

We can summarise the relations between the different optimisation formulations by (MP) = (MaMi) $\leq$ (MiMa). Hence, if we solve the (MiMa) we find an upper bound on (MP).

Let us start with the inner maximisation problem of (MiMa)

$$\begin{aligned} &\max_{|\lambda(t,\omega)|\leq k} F_\alpha(L,\xi) \\ &\text{s.t. } dL = \lambda(t,\omega)LdW. \end{aligned} \tag{23}$$

The inner maximisation is solved by formulating it as a Hamilton-Jacobi-Bellman (HJB) problem. Note that for every fixed $\xi \in \mathbb{R}$ the function $F_\alpha(L,\xi)$ is convex in L.

Before we give the mathematical proof, we first describe the intuitive outline. First, the expectation of the Radon-Nikodym derivative over all alternative models $\lambda(t,\omega)$ is equal to one by definition. Second, $F_\alpha(L,\xi)$ is convex in L for all ξ. Hence by Jensen's inequality, the maximum power is obtained by increasing the volatility of the Radon-Nikodym process as much as possible and set $|\lambda(t,\omega)| \equiv k$. We now proceed with the formal proof.

We introduce the value function $V(t,l,\xi) = \mathbb{E}[F_\alpha(L(\mathcal{T}),\xi)|L(t)=l]$ with boundary condition $V(\mathcal{T},l,\xi) = F_\alpha(L(\mathcal{T}),\xi)$. To maximise $V(t,l,\xi)$, for any value of α and ξ the optimised V function with respect to λ for $t \leq \mathcal{T}$ solves the HJB-equation

$$V_t + \max_{|\lambda(t,l)|\leq k} \tfrac{1}{2}\lambda(t,l)^2 l^2 V_{ll} = 0. \tag{24}$$

Optimal control problems of this sort have been studied in the literature in the context of uncertain volatility models, see Avellaneda et al. (1995) and Vanden (2006). We propose $|\lambda(t,\omega)| \equiv k$ as candidate solution based on the heuristic argument given above. The analytical expression for the value function is

$$V(t,l,\xi) = \alpha\xi + lN(d_1) - \xi N(d_2), \tag{25}$$

where $d_1 = \frac{1}{k\sqrt{\mathcal{T}-t}}\left(\ln\left(\frac{l}{\xi}\right) + \frac{1}{2}k^2(\mathcal{T}-t)\right)$ and $d_2 = d_1 - k\sqrt{\mathcal{T}-t}$ and $N(\cdot)$ is the cumulative standard normal distribution. This function solves the HJB-equation for the boundary condition. The value function is convex since

$$V_{ll}(t,l,\xi) = \frac{n(d_1)}{lk\sqrt{\mathcal{T}-t}} > 0 \ \forall\ l, t < \mathcal{T}, \tag{26}$$

where $n(\cdot)$ is the standard normal density function.

The Verification Theorem 11.2.2 of Øksendal (2003) states that if $V(t,l,\xi)$ is uniformly integrable then a solution of the HJB is an optimal control. The value function $V(t,l,\xi)$ is uniformly integrable since $L(\mathcal{T})$ itself is uniformly integrable for all $\lambda(t,\omega) < \infty$, consequently also the partial moment $V(t,l,\xi) = \mathbb{E}[F_\alpha(L,\xi)|L(t)=l]$ is uniformly integrable. Thus $|\lambda(t,\omega)| = k$ is indeed an optimal control that leads to the maximum power of (23). The optimal value function is $V(0,1,\xi) = \mathbb{E}\left[F_\alpha(L(\mathcal{T}),\xi)|\mathcal{F}_0\right] = \alpha\xi + \mathbb{E}\left[(L(\mathcal{T})-\xi)^+|t=0, L_0=1\right]$.

The remaining part is the outer minimisation. In particular, the calculation $\min_\xi V(0,1,\xi)$ has as optimal ξ^* the $(1-\alpha)$-quantile of the lognormal $L(\mathcal{T})$. This proves that the Radon-Nikodym process with $|\lambda(t,\omega)| \equiv k$ for all $0 \le t \le \mathcal{T}$ achieves the highest possible power while fulfilling the constraint on the Type I error.

We have the result that an upper bound on (MP) is the lognormal power of (MiMa). Since $|\lambda^*(t,\omega)| \equiv k$ is a feasible solution of the original problem (MP), it is also the lower bound on (MP). Thus we can conclude that $|\lambda^*(t,\omega)| \equiv k$ gives the optimal power for (MP).

Given that the optimal $L^*(\mathcal{T},\omega)$ is a lognormal martingale with volatility k, then the optimal value for ξ is equal to the $(1-\alpha)$-quantile of $L^*(\mathcal{T},\omega)$. The optimised power at time $t=0$ and $L(0,\omega)=1$ is therefore equal to $\Phi\left(\Phi^{-1}(\alpha) + k\sqrt{\mathcal{T}}\right)$. If we want to construct a set of rectangular coherent risk measures with stochastic $\lambda(t,\omega)$ such that every element in the set is statistically indistinguishable, then a sufficient condition is to choose a k such that the "worst-case" power does not exceed $(1-\beta)$. $\square$

If we choose $\alpha = 5\%$ and $\beta = 20\%$, then Theorem 1 implies that all $\lambda(t,\omega)$'s are considered as indistinguishable alternatives that are in between $-2.48/\sqrt{\mathcal{T}}$ and $2.48/\sqrt{\mathcal{T}}$. All previous derivations also hold for vector Brownian motions and vector $\lambda(t,\omega)$. The multi-dimensional equivalent is a bound on the L_2-norm; $|\lambda(t,\omega)|_2 \le \frac{2.48}{\sqrt{\mathcal{T}}}$. One can recognise that in the classical confidence interval approach without model uncertainty the bound reduces to $-\frac{\Phi^{-1}(\alpha)}{\sqrt{N}}$. This implies that the power is set to $1-\beta = 50\%$, and the forward interpretation that is attached to the hypothetical test to quantify uncertainty is taken over by a backward reliance on historical data (N represents the number of data points) to quantify risk.

4. Bounds on Divergences

Ben-Tal et al. (2013) discuss several possible divergences to generate robust results in optimisation problems. A ϕ-divergence (or f-divergence) function measures the distance between two probability distributions weighted by the specific function. The choice which measure should be picked remains an issue in their and many other papers, plus the question when the distance is "too far" is rarely investigated. In this paper we explicitly focus on the size of the set of alternatives, which answers the latter question. We can directly link this to a critical value for each measure.

The continuous ϕ-divergence is defined as

$$D_\phi(L(\mathcal{T},\omega)) = \mathbb{E}^{\mathbb{L}}\left[\phi\left(\frac{1}{L(\mathcal{T})}\right)\right] = \mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right], \tag{27}$$

where the functions $\phi(\cdot)$ are given for each divergence, which are convex by definition. We focus on the size of the uncertainty which is quantified by c

$$D_\phi(L(\mathcal{T},\omega)) \leq c. \tag{28}$$

The definitions of the divergences in Table 1 are expressed in terms of $x = \frac{1}{L(\mathcal{T},\omega)}$ under the measure $\mathbb{L}$. These can easily be converted to $\mathbb{P}$ by multiplication with $L(\mathcal{T},\omega)$.

Ben-Tal et al. (2013) list a number of divergence functions for which they derive the asymptotic distributions in a discrete setting. The second column of Table 1 consists of the definitions of the different divergences. So far, these divergences are not rectangular. For each of these functions we show how to obtain the rectangular and coherent set of indistinguishable models.

The lay-out of the proof for each divergence function is similar as the proof of the maximum power calculation. First we postulate the lognormal Radon-Nikodym derivative as candidate solution. Then the analytical expression for the value function is used to prove its convexity, which implies that the proposed solution $|\lambda(t,\omega)| = k$ solves the HJB. By the Verification Theorem 11.2.2 of Øksendal (2003) and the uniform integrability of the value function the optimal control is obtained. See Appendix A for the proofs.

The numerical value for the set of rectangular models that cannot be distinguished between with a Type I error of 5% and a power less than 80% is displayed in the last column of Table 1. The third column shows the bound on the divergence measure, $D_\phi(L(\mathcal{T},\omega)) < c$ which is implied by $|\lambda(t,\omega)| \equiv k$.

Table 1 ϕ-Divergences

Divergence	$\phi(x)$	$c \Leftrightarrow \lvert\lambda(t,\omega)\rvert \equiv k$	$k\sqrt{\mathcal{T}} = 2.48$
Kullback-Leibler	$x\ln x - x + 1$	$\frac{1}{2}k^2\mathcal{T}$	3.08
Burg entropy	$-\ln x + x - 1$	$\frac{1}{2}k^2\mathcal{T}$	3.08
J-divergence	$(x-1)\ln x$	$k^2\mathcal{T}$	6.15
χ^2-divergence	$\frac{1}{x}(x-1)^2$	$e^{k^2\mathcal{T}} - 1$	467.90
Modified χ^2-divergence	$(x-1)^2$	$e^{k^2\mathcal{T}} - 1$	467.90
Hellinger distance	$(\sqrt{x}-1)^2$	$2 - 2e^{-\frac{1}{8}k^2\mathcal{T}}$	1.07
Variation distance	$\lvert x-1\rvert$	$4N(\frac{1}{2}k\sqrt{\mathcal{T}}) - 2$	1.57
χ-divergence of order $\vartheta > 1$	$\lvert x-1\rvert^\vartheta$	*numerical*	*Table 2*
Cressie-Read $\vartheta \neq 0, 1$	$\frac{1-\vartheta+\vartheta x - x^\vartheta}{\vartheta(1-\vartheta)}$	$\frac{1}{\vartheta(1-\vartheta)}\left(1 - e^{-\frac{1}{2}k^2\vartheta(1-\vartheta)\mathcal{T}}\right)$	*Table 2*

The second column shows the divergence function, the third column the optimal maximum divergence and the last column the numerical bound on the divergence for a Type I error of 5% and a power of 80%.

Both the χ-divergence of order $\vartheta > 1$ and the Cressie-Read divergence depend on the additional parameter ϑ. For $k\sqrt{\mathcal{T}} = 2.48$ we get the bounds on both measures displayed in Table 2.

Table 2 Numerical bounds

Divergence $\diagdown\vartheta$	1.5	2.0	2.5	3.0
χ-divergence of order ϑ	10.40	467.90	1.02×10^6	1.03×10^8
Cressie-Read	12.05	233.95	2.72×10^4	1.72×10^7

For a given ϑ we calculated the maximum divergence possible by imposing a Type I error of 5% and Type II error of 20%.

Possible applications of these measures are optimal control problems, e.g. the portfolio investment problem by Maenhout (2004) and Biagini and Pınar (2017) or optimisation problems in a macroeconomic setting by Anderson et al. (2003) and Hansen and Sargent (2008). There, the agent is worried about possible misspecification of the baseline model. The uncertainty could be described by an additional constraint such as that the Kullback-Leibler divergence between the baseline model and all alternatives has to be bounded. Maccheroni et al. (2006a) show the connection of the divergence preferences with the robust maxmin method of Gilboa and Schmeidler (1989). The smooth ambiguity that is initially founded by Klibanoff et al. (2005) is for example applied by Gollier (2014) to design insurance contracts under the acknowledgment of uncertainty of the loss distribution. And Cambou and Filipović (2015) construct alternative models for solvency regulation, by a combination of expert views on a given set of scenarios with a baseline model that is within a specified minimum divergence.

The cutoff point that differentiates plausible alternatives from implausible ones is what we derived. If one would impose a probability of incorrectly rejecting the baseline model of 5% and if one would need at least a power of 80% to correctly accept the alternative, then for all models that yield a lower power, the information up to $\mathcal{T}$ would be insufficient to distinguish the alternative from the baseline. Hence the equivalent robustness is obtained when the objective is optimised subject to bounded divergence measure. Thus by the linkage between $\lambda(t,\omega)$ and $\{\alpha,\beta,\mathcal{T}\}$, an intuitive interpretation is attached to the numerical bound on the Kullback-Leibler divergence (also known as entropy), the Burg entropy, the J-divergence, the (modified) χ^2-divergence, the Hellinger distance, the variation distance, χ-divergence of higher orders, and the Cressie-Read divergence. For example, the set of alternative models that yield a χ^2-divergence of less than 467.90 is identical with the set described by a maximum entropy of 3.08. This holds for any choice of divergence function, since all numerical bounds shown in Table 1 and 2 imply a statistical power lower than 80%.

5. Application

The contribution of this paper is that we derive quantitative bounds on uncertainty sets based on statistical testing theory. The explicit characterisation of the set of indistinguishable models distinguishes plausible from implausible models relative to the baseline model. The testing procedure is utilised to find the relation between the quantitative bound and the variables $\alpha, \beta, \mathcal{T}$. In many robust optimisation problems either the bound is unspecified, the variable remains a symbol instead of a quantity, or it should be picked by intuition or "common sense" rather than rigorous statistical considerations. Examples of such situations are for instance portfolio problems in which the agent acknowledges his concern about the underlying diffusion process. A direct application of our theory is the Bang-Bang control problem in Section 2.2. In Balter and Pelsser (2016)[4] we provide an application of a robust optimisation problem in which the robust price and hedge position in an incomplete market setting is obtained for a given uncertainty parameter k. Here we illustrate the quantitative effect on the extended Merton investment problem, explaining the equity premium puzzle.

5.1. Investment Problem with Model Uncertainty

The Merton (1969) problem is an optimisation problem where an agent's objective is to maximise expected utility from terminal wealth by allocating some of his wealth to the risky asset and the rest to the risk-free bank account. Note that "terminal" refers to a certain time point in the future. Hence, the terminal time T is different than the test horizon $\mathcal{T}$. We introduced the test horizon to give the quantification of uncertainty an intuitive interpretation. The bound that distinguishes plausible from implausible models depends on the Type I error, the Type II error and the test horizon. These three parameters determine the level of uncertainty k. If the test horizon becomes large, the uncertainty goes down as the agent believes that there will not be a structural break in the near future. On the other hand, if the terminal time at which the wealth is evaluated becomes larger, one could argue that the agent becomes more uncertain since he has to decide which investment strategy will be beneficial for a long future. He can indicate this additional concern by adjusting $\mathcal{T}$.

First, we solve the classical Merton problem without model uncertainty. A risk aversion parameter γ that fits the realistic (relatively low) interest rate, leads to an unrealistic low equity premium which is known as the equity premium puzzle. While a γ that fits the realistic (relatively high) equity premium, leads to an unrealistic high interest rate which is known as the risk-free rate puzzle. A realistic equity premium $\mu - r$ is in the range between 4% and 6%, and a realistic risk aversion γ is between 0 and 10. Only an extreme risk aversion could explain the empirical much lower investment in risky assets than the optimal strategy implied by the Merton setting. Several hypotheses see the acknowledgment of model uncertainty as a possible explanation, i.e. additional to risk aversion agents might be concerned about the underlying model specification. Since first moments are considerably hard to estimate, we apply the set of indistinguishable models to the

Merton investment problem to investigate the effect of model uncertainty besides risk aversion. We compare the classical investment problem with the robust version based on a constraint and a penalty (multiplier) approach.

We derive the robust Merton problem by the dynamic and time-consistent optimisation, similarly as Biagini and Pınar (2017). The result coincides with optimal investment strategy of the mean-variance portfolio model including uncertainty, which is also solved by Garlappi et al. (2007) in a static rather than a dynamic way, and Stein (1956) considered the shrinkage on the drifts at first.

Hansen and Sargent and co-authors (Hansen et al. 2006b, Anderson et al. 2003, Hansen et al. 2002, 1999) have several papers on robustness in which they motivate the constraint approach based on the maxmin framework of Gilboa and Schmeidler (1989). However, they transform the constraint into a penalty term. In the growth optimal portfolio setting, the bound on the constraint, k, is directly linked to the constant θ which is the Lagrange multiplier. However, for utility functions other than the natural logarithm the one-to-one relation is lost.

Maenhout (2004) starts with the entropy penalty and transforms the multiplier into a function that depends on the value function itself. The motivation behind this step is that the homothetic condition increases the analytical tractability. However, due to this step the connection with the initial motivation of ellipsoid uncertainty is lost. Pathak (2002) has an interesting discussion on the implications of these changes. By solving the adjusted objective, one is optimising a fundamentally different problem. The solution of an unknown objective is derived. Sargent (2007) comments on this issue from a macroeconomic side. Namely, the optimisation is either related to the nonrobust model that uses recursive utilities or that the problem is solved robustly based on different preferences.

Nonetheless, we compare the (i) classical Merton solution, with the (ii) robust constraint solution and the (iii) robust penalty solution.

Let $A(t)$ be the wealth at time t and let there be one[5] risky asset which can be described by the stochastic differential equation

$$dS(t) = \mu S(t)dt + \sigma S(t)dW(t). \tag{29}$$

Let $\pi(t)$ denote the fraction of wealth invested in the risky asset $S(t)$. Then by the self-financing condition the total wealth evolves as

$$dA(t) = (\pi(t)(\mu - r) + r)\, A(t)dt + A(t)\pi(t)\sigma dW(t). \tag{30}$$

The robust version of the investment problem is

$$dA(t) = (\pi(t)(\mu - r) + r + \pi(t)\sigma\lambda(t,\omega))\, A(t)dt + A(t)\pi(t)\sigma dW(t). \tag{31}$$

We consider the wealth process under alternative measures to indicate the possibility that the baseline model might be misspecified. The robust solution is obtained as the investment strategy that is least sensitive to perturbations of the model. Strategically this is ensured by giving a so called mother nature the role to minimise the objective that the agent tries to maximise. She can do this by picking the worst case probability measure among the set of plausible alternatives. In this section we assume that the agent wants to maximise the utility from his terminal wealth at the finite horizon T. The uncertainty set or equivalently the set of indistinguishable models are those models surrounding the baseline model that one cannot distinguish between with enough power for a given probability on the Type I error. This takes the form of a bounded interval on the change of measure for which we quantified a plausible cutoff point in the previous sections. The robust objective is

$$\max_{\pi(t)} \min_{\lambda(t,\omega)\in\mathcal{L}} \quad \mathbb{E}[U(A(T))], \tag{32}$$

and we assume power utility

$$U(A(t)) = \frac{A(t)^{1-\gamma}}{1-\gamma}. \tag{33}$$

The stochastic differential equation of the wealth process is given by (31) and the uncertainty region is given by

$$\mathcal{L} = \left\{ |\lambda(t,\omega)| \leq k \,\middle|\, k = \frac{\Phi^{-1}(1-\beta) - \Phi^{-1}(\alpha)}{\sqrt{\mathcal{T}}} \right\}, \tag{34}$$

where $\lambda(t,\omega)$ is allowed to be both deterministic and stochastic (for the stochasticity allowance we use the ω-notation). The optimisation problem summarises as

$$\begin{aligned} \max_{\pi(t)} \min_{\lambda(t,\omega)} \quad & \mathbb{E}\left[\frac{A(T)^{1-\gamma}}{1-\gamma}\right] \\ \text{s. t.} \quad & dA(t) = (\pi(t)\tilde{\mu} + r)\, A(t)dt + \pi(t)\sigma A(t)dW(t) \\ & \lambda(t,\omega)^2 \leq k^2, \end{aligned} \tag{35}$$

where $\tilde{\mu} = \mu - r + \sigma\lambda(t,\omega)$. The optimal investment strategy is

$$\pi^C(t) = \max\left(\frac{\mu - r - \sigma k}{\gamma\sigma^2}, 0\right). \tag{36}$$

If there is no uncertainty, $k = 0$ then we are back in the classical Merton setting with

$$\pi^M(t) = \frac{\mu - r}{\gamma\sigma^2}. \tag{37}$$

If the agent is uncertain he will invest less in the risky asset, consequently he invests more on the bank account. This behaviour goes on until he is so uncertain that he shall not invest in the risky assets anymore.

The ellipsoid constraint is equivalent to having a constraint on the Kullback-Leibler divergence. In the previous section we derived the connection between the bounds. Moreover, since $\mathbb{E}^{\mathbb{P}}\left[\ln L(\mathcal{T})^{-1}\right] = \mathbb{E}^{\mathbb{P}}\left[\frac{1}{2}\int_0^{\mathcal{T}} \lambda(t)^2\,dt - \int_0^{T} \lambda(t,)\,dW^{\mathbb{P}}(t)\right] = \frac{1}{2}\int_0^{\mathcal{T}} \lambda(t)^2\,dt$ it follows that $(\int_0^{\mathcal{T}} \lambda(t)^2 dt)^{\frac{1}{2}} \leq 2.48$ corresponds to $D_{KL}(L(\mathcal{T},\omega)) \leq 3.08$.

The addition of a penalty in form of the relative entropy to the Merton problem is what Maenhout (2004) and Anderson et al. (2003) investigate. On top of that, Maenhout assumes homotheticity, i.e. that the optimal portfolio weight is independent from the initial wealth position. In Maenhout (2006) he extends this by a mean-reverting risk premium. He incorporates the entropy term by a penalty on the utility premultiplied by the strength of preference for robustness. Similarly as for the constraint approach, we express the dependence of the objective and the penalty on the Radon-Nikodym derivative in terms of an endogenous drift $\lambda(t)$[6]. The penalty formulation of the optimisation problem becomes

$$\begin{aligned} \max_{\pi(t)} \min_{\lambda(t)} \quad & \mathbb{E}[U(A(T))] + \frac{1}{2\theta}\lambda^2(T) \\ \text{s. t.} \quad & dA(t) = (r + \pi(t)(\mu - r) + \pi(t)\sigma\lambda(t))\, A(t)dt + \pi(t)\sigma A(t)dW(t). \end{aligned} \tag{38}$$

The strength of preference for robustness is measured by θ, $\theta = 0$ corresponds to expected utility maximisation. The larger θ, the more uncertain the agent is. Hence, there is a one-to-one correspondence between the bound on the uncertainty set k and the multiplier θ. Maenhout (2004) fixes the functional form of θ as a function of the value function specified in the HJB formulation[7]. The transformation assures homotheticity. Stated differently, without this functional form $\pi(t)$ is not analytically solvable. The optimal investment strategy according to the penalty approach with the additional homotheticity condition is

$$\pi^P(t) = \frac{1}{\gamma + \theta}\frac{\mu - r}{\sigma^2}. \tag{39}$$

Comparing the three optimal strategies shows that the constraint approach leads to an adjustment of the risk premium while the penalty approach leads to an adjustment of the volatility. In both approaches uncertainty has exactly the same influence on the classical Merton solution. Uncertainty decreases the risk premium or increases the volatility, both lowering the investment strategy. The relation between k and θ is

$$k = \frac{\mu - r}{\sigma}\frac{\theta}{\gamma + \theta}. \tag{40}$$

The higher the preference for robustness, the closer the bound on the constraint converges towards the Sharpe ratio which coincides with refraining from risky investment whose model specification is uncertain.

Since the equity premium puzzle yields a risk aversion which is rejected empirically, the incorporation of uncertainty transfers the question from whether γ is too large to whether θ is too large. Whether the robustness preference is unrealistically high, boils down to the same question as to what a realistic bound on the uncertainty set is. Therefore we show the optimal investment fractions for $\mu = 8\%$ and $\sigma = 16\%$ in Table 3, where the results in the left panel of the table are based on $r = 4\%$ and the results in the right panel are based on $r = 2\%$ yielding an equity premium of 4% and 6% respectively. The standard Merton investment strategy implies that the agent has to borrow money in order to invest more than 100% of his wealth into the risky asset for a risk aversion below 3. This extreme optimal strategy is not in line with empirical findings. A constraint with a numerical bound equal to $k = \frac{\Phi^{-1}(1-\beta)-\Phi^{-1}(\alpha)}{\sqrt{\mathcal{T}}}$ for $\alpha = 5\%, \beta = 20\%$ and $\mathcal{T} = 100$ or $\mathcal{T} = 200$ leads to lower investment strategies. For each case we also depicted the implied preference towards robustness. All values in the fifth and tenth column and the sixth and eleventh column represent the same level of uncertainty. Both the values 124 and 5.86 for $\theta_{\mathcal{T}=100}$ imply the agent to allow for a Type I error on the hypothetical test of 5%, to require at least a power of 80% and to assume that the model specification holds for the next 100 years. Moreover, the robust investment strategy is to invest 1.25% in risky assets if the equity premium is 4% compared to 156% if the agent believes that his model specification is correct. And when the equity premium equals 6% then the ignorance of uncertainty leads to an exposure to risk of 78.1% while acceptance of misspecification leads to an investment strategy of 26.5%. As can be seen the parameter k is independent from the baseline assumption, whereas there is no general interpretation for the numerical value of θ. In other words, if we express the bounds as a function of the dependent variables then $k(\alpha, \beta, \mathcal{T})$ and $\theta(\mu, r, \sigma, \gamma, x)$ where x is an unknown parameter that incorporates the amount of ambiguity. Hence, if we set $x = k$ then the robustness parameter can be numerically interpreted to ascribe meaning to the amount of ambiguity.

The empirical application in Maenhout (2004) is based on the data from Campbell (1999). Calibration of the equilibrium model gives the elasticity of intertemporal substitution, the risk aversion and uncertainty aversion. Together these values can explain the equity premium puzzle and risk-free rate puzzle. The initial risk aversion γ of 21 that Maenhout shows in his Table 2, is split into a risk aversion γ of 7 plus an uncertainty aversion θ of 14. We consider the long annual sample from 1891 to 1994, where $\mu - r = 6.258\%, r = 1.955\%$, $\sigma = 0.18534$ and $(\gamma, \theta) = (21, 0)$ or $(\gamma, \theta) = (7, 14)$ with and without robustness respectively. Maenhout tests the likelihood of the uncertainty aversion via the detection error probability (Anderson et al. 2003). Based on the amount of historical data the alternative risk premium cannot be rejected from the baseline estimation. The implied Type I and II errors belonging to the calibration in Maenhout imply that the investment mix is robust for an agent who characterises the uncertainty set by those models that are indistinguishable if

Table 3 Robust Optimal Investments

γ	π^M	$\pi^C_{\mathcal{T}=100}$	$\pi^C_{\mathcal{T}=200}$	$\theta_{\mathcal{T}=100}$	$\theta_{\mathcal{T}=200}$	π^M	$\pi^C_{\mathcal{T}=100}$	$\pi^C_{\mathcal{T}=200}$	$\theta_{\mathcal{T}=100}$	$\theta_{\mathcal{T}=200}$
0.5	313%	2.5%	93.3%	62	1.17	469%	159%	250%	0.98	0.44
1	156%	1.25%	46.7%	124	2.35	234%	79.4%	125%	1.95	0.88
3	52.1%	0.42%	15.6%	372	7.05	78.1%	26.5%	41.6%	5.86	2.64
5	31.3%	0.25%	9.33%	620	11.8	46.9%	15.9%	25.0%	9.76	4.39
10	15.6%	0.13%	4.66%	1240	23.5	23.4%	7.94%	12.5%	19.53	8.78

The optimal robust investment strategy for the constraint approach (π^C) is shown for $\alpha = 5\%, 1-\beta = 80\%$ and $\mathcal{T} = \{100, 200\}$ with $k = \frac{2.48}{\sqrt{\mathcal{T}}}$. The classical Merton solution without model uncertainty is denoted by π^M. The multipliers θ imply the same robust strategy as the constraint approach.

the power is less than 73% for a Type I error of 5% and a test horizon of 100. Stated differently for a power of 80%, the test horizon that corresponds to the calibrated θ is 122 years. Hence the calibrated preference towards robustness is separated from the baseline model (μ, r, σ, γ) and the ambiguity part is further disentangled into the implied Type I error α, the Type II error β and the test horizon $\mathcal{T}$.

Appendix A: Divergences

The proofs for the bounds on the different divergence measures follow the same steps as the initial proof of the rectangular and coherent set of indistinguishable models that can be characterised by $|\lambda(t,\omega)| \leq k$.

A.1. Kullback-Leibler

The Kullback-Leibler divergence (also known as *entropy*) is defined as

$$D(\mathbb{P}\|\mathbb{L}) = \mathbb{E}^{\mathbb{P}}[-\ln L(\mathcal{T}) - 1 + L(\mathcal{T})] = \mathbb{E}^{\mathbb{L}}\left[-\frac{\ln L(\mathcal{T}) + 1 - L(\mathcal{T})}{L(\mathcal{T})}\right]. \tag{41}$$

We want to investigate the maximum entropy that can be achieved within the class of Radon-Nikodym derivatives with $|\lambda(t,\omega)| \leq k$. We can formulate this as a stochastic optimisation problem of the form

$$\begin{aligned} &\max_{|\lambda(t,\omega)|\leq k} \mathbb{E}[-\ln L(\mathcal{T}) - 1 + L(\mathcal{T})] \\ &\text{s.t. } dL = \lambda(t,\omega) L dW. \end{aligned} \tag{42}$$

This optimisation problem admits the following HJB representation. If we set $V(t,l) := \mathbb{E}[-\ln L(\mathcal{T}) - 1 + L(\mathcal{T}) \mid L(t) = l]$, then the optimised value function $V(t,l)$ for $t \leq \mathcal{T}$ is given by the HJB-equation

$$V_t + \max_{|\lambda(t,l)|\leq k} \tfrac{1}{2}\lambda(t,l)^2 l^2 V_{ll} = 0. \tag{43}$$

The terminal condition $V(\mathcal{T},l) = -\ln l - 1 + l$ is a convex payoff in l. Hence we propose $|\lambda(t,\omega)| = k$ to solve the HJB. The implied value function for the candidate solution is

$$V(t,l) = -\ln l + \tfrac{1}{2}k^2(T-t) - 1 + l. \tag{44}$$

This function solves the HJB-equation for the boundary condition and maximises the objective in (43) as V is convex in l. Similarly as in the maximum power calculation in Section 3.1, the value function $V(t,l)$ is uniformly integrable and thus the optimal control is $|\lambda(t,\omega)| = k$. Hence, the lognormal Radon-Nikodym derivative with $|\lambda(t,\omega)| \equiv k$ for all $0 \le t \le \mathcal{T}$ achieves the maximal entropy of $V(0,1) = \frac{1}{2}k^2\mathcal{T}$. Thus we can characterise the class of rectangular indistinguishable models with $|\lambda(t,\omega)| \le 2.48/\sqrt{\mathcal{T}}$ by the implied maximum attainable entropy of $\frac{1}{2}(2.48)^2 = 3.08$.

A.2. Burg Entropy

Burg entropy or also called minimum discrimination information is defined as

$$D(\mathbb{L}\|\mathbb{P}) = \mathbb{E}^{\mathbb{L}}[\ln L(\mathcal{T}) + \frac{1}{L(\mathcal{T})} - 1] = \mathbb{E}^{\mathbb{P}}[L(\mathcal{T})\ln L(\mathcal{T}) + 1 - L(\mathcal{T})]. \tag{45}$$

We can formulate this as a stochastic optimisation problem of the form

$$\begin{aligned} &\max_{|\lambda(t,\omega)|\le k} \mathbb{E}\left[L(\mathcal{T})\ln L(\mathcal{T}) + 1 - L(\mathcal{T})\right] \\ &\text{s.t. } dL = \lambda(t,\omega) L dW. \end{aligned} \tag{46}$$

This optimisation problem admits the following HJB representation. If we set $V(t,l) := \mathbb{E}\left[L(\mathcal{T})\ln L(\mathcal{T}) + 1 - L(\mathcal{T}) \mid L(t) = l\right]$, then the optimised value function $V(t,l)$ for $t \le \mathcal{T}$ is given by the HJB-equation

$$V_t + \max_{|\lambda(t,l)|\le k} \tfrac{1}{2}\lambda(t,l)^2 l^2 V_{ll} = 0. \tag{47}$$

As the terminal condition $V(\mathcal{T},l) = l\ln l + 1 - l$ is a strictly convex function in l, we can follow the same procedure as the proof of the maximum power and entropy calculation. The lognormal Radon-Nikodym derivative with $|\lambda(t,\omega)| \equiv k$ for all $0 \le t \le \mathcal{T}$ achieves the maximal Burg entropy of $V(0,1) = \frac{1}{2}k^2\mathcal{T}$. Since for $q \in \mathbb{R}$

$$\begin{aligned} \mathbb{E}\left[L(\mathcal{T})^q \mid \mathcal{F}_t\right] &= \mathbb{E}\left[l^q e^{-\frac{1}{2}qk^2(\mathcal{T}-t) + qk(W^{\mathbb{P}}(\mathcal{T}) - W^{\mathbb{P}}(t))} \,\middle|\, \mathcal{F}_t\right] \\ &= l^q e^{\frac{1}{2}k^2(\mathcal{T}-t)(q^2-q)}. \end{aligned} \tag{48}$$

The derivative with respect to q on both sides leads to

$$\begin{aligned} \frac{\partial \mathbb{E}\left[L(\mathcal{T})^q \mid \mathcal{F}_t\right]}{\partial q} &= \mathbb{E}\left[L(\mathcal{T})^q \ln L(\mathcal{T}) \mid \mathcal{F}_t\right] \\ \frac{\partial l^q e^{\frac{1}{2}k^2(\mathcal{T}-t)(q^2-q)}}{\partial q} &= l^q \ln l e^{\frac{1}{2}k^2(\mathcal{T}-t)(q^2-q)} + \\ &\quad l^q e^{\frac{1}{2}k^2(\mathcal{T}-t)(q^2-q)} \tfrac{1}{2}k^2(\mathcal{T}-t)(2q-1). \end{aligned} \tag{49}$$

For $q=1$ we get $V(t,l)=l\left(\ln l+\frac{1}{2}k^2(\mathcal{T}-t)\right)+1-l$. Hence, we can characterise the class of rectangular indistinguishable models with $|\lambda(t,\omega)|\leq 2.48/\sqrt{\mathcal{T}}$ by the implied maximum attainable Burg entropy of $\frac{1}{2}(2.48)^2=3.08$.

A.3. J-Divergence

Jeffreys'(1946) J-divergence is $D(\mathbb{P}\|\mathbb{L})+D(\mathbb{L}\|\mathbb{P})$ which equals $\mathbb{E}^{\mathbb{P}}[-\ln L(\mathcal{T})]+\mathbb{E}^{\mathbb{P}}[L(\mathcal{T})\ln L(\mathcal{T})]$. We can formulate this as a stochastic optimisation problem of the form

$$\begin{aligned}&\max_{|\lambda(t,\omega)|\leq k}\mathbb{E}\left[-\ln L(\mathcal{T})\right]+\mathbb{E}\left[L(\mathcal{T})\ln L(\mathcal{T})\right]\\&\quad\text{s.t. } dL=\lambda(t,\omega)LdW.\end{aligned}\tag{50}$$

Again the associated value function is convex in l at terminal time $\mathcal{T}$, $V(\mathcal{T},l)=(l-1)\ln l$. The lognormal candidate leads to $V(t,l)=(l-1)\ln l+\frac{1}{2}k^2(\mathcal{T}-t)(l+1)$, and hence the maximum J-divergence is obtained for $V(0,1)=k^2\mathcal{T}=(2.48)^2=6.15$.

A.4. χ^2-Distance

Actually for all $\frac{\partial^2\phi(l)}{\partial l^2}\geq 0$ the associated value function is convex for $l>0$ and hence $|\lambda(t,\omega)|=2.48/\sqrt{\mathcal{T}}$ leads to the maximum attainable distance. The $\mathbb{P}$-expectation of the χ^2-distance is

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right]&=\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})^2-2L(\mathcal{T})+1\right]\\&=\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})^2\right]-2\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\right]+1.\end{aligned}\tag{51}$$

For the associated optimisation problem the HJB representation is similar as the ones before. If we set $V(t,l):=\mathbb{E}\left[L(\mathcal{T})^2-2L(\mathcal{T})+1\mid L(t)=l\right]$, then the terminal condition $V(\mathcal{T},l)=l^2-2l+1$ is convex. For $|\lambda(t,\omega)|=k$ the value function is

$$\begin{aligned}V(t,l)&=l^2e^{k^2(\mathcal{T}-t)}-2l+1\\V(0,1)&=e^{2.48^2}-1=467.90.\end{aligned}\tag{52}$$

A.5. Modified χ^2-Distance

The objective function of the modified χ^2-distance under $\mathbb{P}$ is

$$\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right]=\mathbb{E}^{\mathbb{P}}\left[\frac{1}{L(\mathcal{T})}-2+L(\mathcal{T})\right].\tag{53}$$

The value function is $V(t,l):=\mathbb{E}^{\mathbb{P}}\left[\frac{1}{L(\mathcal{T})}-2+L(\mathcal{T})\,\middle|\,L(t)=l\right]$ which has a convex terminal condition. For the lognormal Radon-Nikodym derivative the modified χ^2-distance's value function is

$$V(t,l)=\frac{1}{l}e^{k^2(\mathcal{T}-t)}-2+l.\tag{54}$$

Henceforth, $V(0,1)=e^{k^2\mathcal{T}}-1=467.90$ is the bound that includes the indistinguishable models for a 5% probability that a Type I error occurs and have a power less than 80%. In order to derive (54) we applied Itô's Lemma, $d\left(\frac{1}{L(t)}\right)=\frac{1}{L(t)}\left(k^2dt-kW^{\mathbb{P}}(t)\right)$.

A.6. Hellinger Distance

The Hellinger distance equals

$$\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right] = \mathbb{E}^{\mathbb{P}}\left[1-2\sqrt{L(\mathcal{T})}+L(\mathcal{T})\right]. \tag{55}$$

The associated value function is $V(t,l) := \mathbb{E}^{\mathbb{P}}\left[1-2\sqrt{L(\mathcal{T})}+L(\mathcal{T}) \,\middle|\, L(t)=l\right]$. Since

$$\begin{aligned}\sqrt{L(\mathcal{T})} &= \sqrt{L(t)}e^{-\frac{1}{4}k^2(\mathcal{T}-t)+\frac{1}{2}k(W^{\mathbb{P}}(\mathcal{T})-W^{\mathbb{P}}(t))} \\ d\sqrt{L(\mathcal{T})} &= -\tfrac{1}{8}k^2\sqrt{L(t)}dt+\tfrac{1}{2}k\sqrt{L(t)}dW^{\mathbb{P}}(t).\end{aligned} \tag{56}$$

The optimal control solves for $|\lambda(t,\omega)| = k$ is

$$\begin{aligned}V(t,l) &= 1-2\sqrt{l}e^{-\frac{1}{8}k^2\mathcal{T}}+l \\ V(0,1) &= 1.07.\end{aligned} \tag{57}$$

A.7. Variation Distance

The variation distance can be decomposed into a call and put option and is defined as follows

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right] &= \mathbb{E}^{\mathbb{P}}\left[\left|\frac{1}{L(\mathcal{T})}-1\right|L(\mathcal{T})\right] \\ &= \mathbb{E}^{\mathbb{P}}\left[\max\left(\frac{1}{L(\mathcal{T})}-1,0\right)L(\mathcal{T})\right]+ \\ &\quad\ \mathbb{E}^{\mathbb{P}}\left[\max\left(1-\frac{1}{L(\mathcal{T})},0\right)L(\mathcal{T})\right] \\ &= \mathbb{E}^{\mathbb{P}}\left[\max\left(1-L(\mathcal{T}),0\right)\right]+\mathbb{E}^{\mathbb{P}}\left[\max\left(L(\mathcal{T})-1,0\right)\right].\end{aligned} \tag{58}$$

The objective is convex by convexity of the max-operator. The associated value function is $V(t,l) := \mathbb{E}^{\mathbb{P}}\left[\left|\frac{1}{L(\mathcal{T})}-1\right|L(\mathcal{T}) \,\middle|\, L(t)=l\right]$. Therefore we propose $dL(t) = kL(t)dW^{\mathbb{P}}(t)$ for the value function

$$\begin{aligned}V(t,l) &= N(d_1)l-N(d_2)+N(-d_2)-N(-d_1)l \\ d_1 &= \frac{1}{k\sqrt{\mathcal{T}-t}}\left(\ln l+\tfrac{1}{2}k^2(\mathcal{T}-t)\right) \\ d_2 &= d_1-k\sqrt{\mathcal{T}-t}.\end{aligned} \tag{59}$$

The second derivative is positive thus the optimal value is obtained for

$$\begin{aligned}V(0,1) &= N(d_1)-N(d_2)+N(-d_2)-N(-d_1) \\ d_1 &= \tfrac{1}{2}k\sqrt{\mathcal{T}} \\ d_2 &= -d_1 \\ N(-d_1) &= 1-N(d_1) \\ V(0,1) &= 4N(d_1)-2.\end{aligned} \tag{60}$$

For $d_1 = \frac{1}{2}k\sqrt{\mathcal{T}} = 1.24$, the bound on the variation distance is $V(0,1) = 1.57$.

A.8. χ-Divergence of Order $\theta > 1$

For general θ strictly above 1 the χ-divergence can only be solved numerically, though for integers analytically. The divergence functions is defined by

$$\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right] = \mathbb{E}^{\mathbb{P}}\left[\left|\frac{1}{L(\mathcal{T})} - 1\right|^{\theta} L(\mathcal{T})\right]. \tag{61}$$

For different values of $\theta = \{1.5, 2, 2.5, 3\}$ and $k^2\mathcal{T} = 2.48^2$, we integrate

$$\mathbb{E}^{\mathbb{P}}\left[\left|\frac{1}{L(\mathcal{T})} - 1\right|^{\theta} L(\mathcal{T})\right] = \int_{-\infty}^{\infty} \left|e^{\frac{1}{2}k^2\mathcal{T} - k\sqrt{\mathcal{T}}z} - 1\right|^{\theta} e^{-\frac{1}{2}\mathcal{T} + k\sqrt{\mathcal{T}}z} n(z)dz.$$

If $\theta = 2$ this coincides with modified χ^2-divergence.

A.9. Cressie-Read

The Cressie-Read divergence is

$$\begin{aligned}\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T})\phi\left(\frac{1}{L(\mathcal{T})}\right)\right] &= \mathbb{E}^{\mathbb{P}}\left[\frac{1-\theta+\theta\frac{1}{L(\mathcal{T})} - \left(\frac{1}{L(\mathcal{T})}\right)^{\theta}}{\theta(1-\theta)} L(\mathcal{T})\right] \\ &= \frac{1}{\theta(1-\theta)}\mathbb{E}^{\mathbb{P}}\left[L(\mathcal{T}) - \theta L(\mathcal{T}) + \theta - L(\mathcal{T})^{1-\theta}\right].\end{aligned} \tag{62}$$

This is convex in l therefore if $L(\mathcal{T})$ is lognormal distributed with volatility k,

$$V(t,l) = \frac{1}{\theta(1-\theta)}\left(l - \theta \cdot l + \theta - l^{1-\theta} e^{-\frac{1}{2}k^2\theta(1-\theta)(\mathcal{T}-t)}\right). \tag{63}$$

At time 0 this equals

$$V(0,1) = \frac{1}{\theta(1-\theta)}\left(1 - e^{-\frac{1}{2}k^2\theta(1-\theta)\mathcal{T}}\right), \tag{64}$$

for general $k^2\mathcal{T}$ and

$$V(0,1) = \frac{1}{\theta(1-\theta)}\left(1 - 0.05e^{\theta(1-\theta)}\right),$$

for $\alpha = 5\%$ and $1 - \beta = 80\%$.

Endnotes

1. Inclusion/exclusion of the equality sign in the expression $L(\mathcal{T},\omega) \geq \xi$ is potentially relevant when there are point-masses in the probability distribution of $L(\mathcal{T},\omega)$. For ease of exposition, we assume this is not the case. When we do have point-masses we can still handle this mathematically, by considering randomised tests.
2. This ensures that the Novikov condition is satisfied for all alternative models in the rectangular and coherent set and thus stochastic differential equation (3) represents a true martingale.

3. To lighten the notation, we drop the dependence on $t, \mathcal{T}$ and ω.

4. An early version was known under the name *Robustness, model uncertainty and pricing.*

5. The multidimensional case is a straightforward extensions, see e.g. Biagini and Pınar (2017).

6. In Maenhout (2004) he uses the notation $u(A(t))$. To prevent misunderstandings with the utility function we use $\lambda(t)$. Moreover, note the difference with the extra multiplication by the term in front of the Brownian motion. Hence, $\lambda(t) = \pi(t)\sigma A(t)u(A(t))$.

7. For $\gamma = 1$, the Maenhout result holds without transformation. In Appendix B of Maenhout (2004) he explicitly derives the log case. Here θ is redefined to be $\frac{\theta}{(1-\gamma)v(A,t)}$, where $v(A,t)$ is the value function.

Acknowledgments

We would like to thank Franz Palm, Bas Werker, Eric Renault, Ludger Rüschendorf, Thorsten Schmidt, Damir Filipovic, Bernt Øksendal, Ernst Eberlein, David Taylor and Erik Schlögl for their valuable comments. And for presenting at the 13th International Paris Finance Meeting, the 9th Conference on Computational and Financial Econometrics in London, the Netspar Pension Day in Utrecht, the 7th General Advanced Mathematical Methods in Finance and Swissquote Conference in Lausanne, the 8th Nordic Econometric Meeting in Helsinki and the seminars at Maastricht University, Tilburg University, the Albert Ludwigs University of Freiburg, the University of Cape Town, the University of Southern Denmark and the University of Glasgow. We gratefully acknowledge the financial support from The Netherlands Organisation for Scientific Research (NWO) TOP grant.

References

Anderson EW, Hansen LP, Sargent TJ (2003) A quartet of semigroups for model specification, robustness, prices of risk, and model detection. *Journal of the European Economic Association* 1(1):68–123.

Artzner P, Delbaen F, Eber JM, Heath D (1999) Coherent measures of risk. *Mathematical Finance* 9(3):203–228.

Avellaneda M, Levy A, Parás A (1995) Pricing and hedging derivative securities in markets with uncertain volatilities. *Applied Mathematical Finance* 2(2):73–88.

Balter AG, Pelsser A (2016) Pricing and hedging in incomplete markets with model uncertainty. *Tilburg University Working Papers, SSRN: 2564521* .

Barrieu PM, El Karoui N (2007) Pricing, hedging and optimally designing derivatives via minimization of risk measures. *Indifference Pricing: Theory and Applications*, 77–146 (Princeton University Press).

Ben-Tal A, Bertsimas D, Brown DB (2010) A soft robust model for optimization under ambiguity. *Operations Research* 58(4-part-2):1220–1234.

Ben-Tal A, Den Hertog D, De Waegenaere A, Melenberg B, Rennen G (2013) Robust solutions of optimization problems affected by uncertain probabilities. *Management Science* 59(2):341–357.

Bertsimas D, Brown DB (2009) Constructing uncertainty sets for robust linear optimization. *Operations Research* 57(6):1483–1495.

Bertsimas D, Gupta V, Kallus N (2017) Data-driven robust optimization. *Mathematical Programming* 1–58.

Biagini S, Pınar MÇ (2017) The robust merton problem of an ambiguity averse investor. *Mathematics and Financial Economics* 11(1):1–24.

Breuer T, Csiszár I (2013) Systematic stress tests with entropic plausibility constraints. *Journal of Banking & Finance* 37(5):1552–1559.

Cambou M, Filipović D (2015) Model uncertainty and scenario aggregation. *Mathematical Finance* .

Campbell JY (1999) Asset prices, consumption, and the business cycle. *Handbook of Macroeconomics* 1:1231–1303.

Chen Z, Epstein L (2002) Ambiguity, risk, and asset returns in continuous time. *Econometrica* 1403–1443.

Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* 58(3):595–612.

Delbaen F (2006) The structure of m–stable sets and in particular of the set of risk neutral measures. *In Memoriam Paul-André Meyer*, 215–258 (Springer).

Eichberger J, Grant S, Kelsey D (2016) Randomization and dynamic consistency. *Economic Theory* 62(3):547–566.

Epstein LG, Schneider M (2003) Recursive multiple-priors. *Journal of Economic Theory* 113(1):1–31.

Esfahani PM, Kuhn D (2017) Data-driven distributionally robust optimization using the wasserstein metric: Performance guarantees and tractable reformulations. *arXiv preprint arXiv:1505.05116* .

Gabrel V, Murat C, Thiele A (2014) Recent advances in robust optimization: An overview. *European Journal of Operational Research* 235(3):471–483.

Garlappi L, Uppal R, Wang T (2007) Portfolio selection with parameter and model uncertainty: A multi-prior approach. *Review of Financial Studies* 20(1):41–81.

Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2):141–153.

Gollier C (2014) Optimal insurance design of ambiguous risks. *Economic Theory* 57(3):555–576.

Hansen LP, Maenhout P, Rustichini A, Sargent TJ, Siniscalchi MM (2006a) Introduction to model uncertainty and robustness. *Journal of Economic Theory* 128(1):1–3.

Hansen LP, Marinacci M (2016) Ambiguity aversion and model misspecification: An economic perspective. *Statistical Science* 31(4):511–515.

Hansen LP, Sargent TJ (2001) Robust control and model uncertainty. *The American Economic Review* 91(2):60–66.

Hansen LP, Sargent TJ (2008) *Robustness* (Princeton University Press).

Hansen LP, Sargent TJ (2011) Robustness and ambiguity in continuous time. *Journal of Economic Theory* 146(3):1195–1223.

Hansen LP, Sargent TJ (2016) Sets of models and prices of uncertainty. *Manuscript, University of Chicago* .

Hansen LP, Sargent TJ, Tallarini TD, et al. (1999) Robust permanent income and pricing. *Review of Economic Studies* 66(4):873–907.

Hansen LP, Sargent TJ, Turmuhambetova G, Williams N (2006b) Robust control and model misspecification. *Journal of Economic Theory* 128(1):45–90.

Hansen LP, Sargent TJ, Wang NE (2002) Robust permanent income and pricing with filtering. *Macroeconomic Dynamics* 6(01):40–84.

Hansen PR, Lunde A, Nason JM (2011) The model confidence set. *Econometrica* 79(2):453–497.

Itô K, McKean HP (1974) *Diffusion Processes and their Sample Paths* (Springer).

Jeffreys H (1946) An invariant form for the prior probability in estimation problems. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 186, 453–461 (The Royal Society).

Karatzas I, Shreve SE (1984) Trivariate density of brownian motion, its local and occupation times, with application to stochastic control. *The Annals of Probability* 819–828.

Klibanoff P, Marinacci M, Mukerji S (2005) A smooth model of decision making under ambiguity. *Econometrica* 73(6):1849–1892.

Maccheroni F, Marinacci M, Rustichini A (2006a) Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica* 74(6):1447–1498.

Maccheroni F, Marinacci M, Rustichini A (2006b) Dynamic variational preferences. *Journal of Economic Theory* 128(1):4–44.

Maenhout PJ (2004) Robust portfolio rules and asset pricing. *Review of Financial Studies* 17(4):951–983.

Maenhout PJ (2006) Robust portfolio rules and detection-error probabilities for a mean-reverting risk premium. *Journal of Economic Theory* 128(1):136–163.

Markowitz H (1952) Portfolio selection. *The Journal of Finance* 7(1):77–91.

Merton RC (1969) Lifetime portfolio selection under uncertainty: The continuous-time case. *The Review of Economics and Statistics* 247–257.

Neyman J, Pearson ES (1933) The testing of statistical hypotheses in relation to probabilities a priori. *Mathematical Proceedings of the Cambridge Philosophical Society*, volume 29, 492–510 (Cambridge Univ Press).

Nilim A, El Ghaoui L (2005) Robust control of markov decision processes with uncertain transition matrices. *Operations Research* 53(5):780–798.

Øksendal B (2003) *Stochastic differential equations* (Springer).

Pathak P (2002) Notes on robust portfolio choice. *Manuscript, Harvard University* 3.

Riedel F (2004) Dynamic coherent risk measures. *Stochastic Processes and Their Applications* 112(2):185–200.

Riedel F (2009) Optimal stopping with multiple priors. *Econometrica* 77(3):857–908.

Rockafellar RT, Uryasev S (2000) Optimization of conditional value-at-risk. *Journal of Risk* 2:21–42.

Rockafellar RT, Uryasev S (2002) Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance* 26(7):1443–1471.

Rosazza Gianin E (2006) Risk measures via g-expectations. *Insurance: Mathematics and Economics* 39(1):19–34.

Sargent TJ (2007) Commentary. *Federal Reserve Bank of St. Louis Review* 89(4):301–303.

Stein C (1956) Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, 197–206.

Strotz RH (1955) Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies* 23(3):165–180.

Tobin J (1958) Liquidity preference as behavior towards risk. *The Review of Economic Studies* 25(2):65–86.

Vanden JM (2006) Exact superreplication strategies for a class of derivative assets. *Applied Mathematical Finance* 13(01):61–87.

Watson J, Holmes C (2016) Approximate models and robust decisions. *Statistical Science* 31(4):465–489.

Wiesemann W, Kuhn D, Sim M (2014) Distributionally robust convex optimization. *Operations Research* 62(6):1358–1376.