

School of Finance



University of St.Gallen

AMBIGUITY AND REALITY

FABIO TROJANI

CHRISTIAN WIEHENKAMP

JAN WRAMPELMEYER

WORKING PAPERS ON FINANCE NO. 2014/18

SWISS INSTITUTE OF BANKING AND FINANCE (S/BF – HSG)

DECEMBER 2014



Ambiguity and Reality*

Fabio Trojani

Christian Wiehenkamp

Jan Wrampelmeyer

Swiss Finance Institute

risklab[‡]

University of St. Gallen[§]

University of Lugano[†]

December 4, 2014

*We thank Jérôme Detemple, Itamar Drechsler, Robert Hodrick, Kerstin Kehrlé, Lorian Mancini, Massimo Marinacci, Franz Palm, Angelo Ranaldo, Paulo Rodrigues, Matthias Rupprecht, Andras Sali, Mirela Sandulescu, Christian Schlag, Peter Schotman, Andreas Schrimpf, Paul Söderlind, Monika Trapp, and Maxim Ulrich, as well as seminar participants at the 60th Midwest Finance Association Annual Meeting, the 17th Annual Meeting of the German Finance Association, the 2010 CEPR/Study Center Gerzensee European Summer Symposium in Financial Markets, the Ninth Swiss Doctoral Workshop in Finance, Maastricht University, and University of Geneva for helpful comments. The views expressed herein are those of the authors and not necessarily those of risklab GmbH, which does not accept any responsibility for the contents and opinions expressed in this paper. Financial support by the National Centre of Competence in Research “Financial Valuation and Risk Management” (NCCR FINRISK) is gratefully acknowledged.

[†]Fabio Trojani, University of Lugano, Via Buffi 13, 6000 Lugano, Switzerland. Email: Fabio.Trojani@usi.ch

[‡]Christian Wiehenkamp, risklab GmbH, Seidlstr. 24/24a, 80335 Munich, Germany. Email: christian.wiehenkamp@risklab.com.

[§]Jan Wrampelmeyer, University of St. Gallen, Swiss Institute of Banking and Finance, Rosenbergstrasse 52, 9000 St. Gallen, Switzerland. Email: jan.wrampelmeyer@unisg.ch.

Ambiguity and Reality

First version: July 2010

This version: December 4, 2014

Abstract

Model builders face ambiguity about the true data generating process. Consequently, they need to deal with ambiguity attitudes (inside uncertainty) and ambiguous financial reality (outside uncertainty) when developing and estimating financial models. We introduce a novel approach for systematically dealing with outside uncertainty in addition to inside uncertainty in a tractable way. By bounding the effects of ambiguous data features, we avoid the adverse consequences of outside uncertainty, such as strongly biased equity premiums and investment policies. In a real data application, we show that asset managers can be more reliably evaluated using our bounded-influence approach.

Keywords: Knightian Uncertainty, Model Risk, Ambiguity Aversion, Robust Econometrics, Portfolio Choice, Option Pricing

JEL Codes: C13, C15, D81, G11, G12

Models are ubiquitous in financial markets. They are used by financial institutions, regulators, and central banks, e.g., to quantify risk or to assess the effect of monetary policy. When developing and estimating models, market participants face uncertainty about the underlying data generating process (DGP). While risk is the randomness that can be understood and modeled, Knightian or model uncertainty denotes unmeasurable ambiguity about the DGP. The latter means that the true DGP contains unknown, time-varying features, which poses significant challenges for model builders in financial markets.

In this paper, we introduce a novel approach for dealing with ambiguity, when taking financial models to the data. We take the point of view of a model builder, who develops an economic model and would like to estimate the model outputs, taking uncertainty into account. Uncertainty can be subdivided into “inside” and “outside” uncertainty (Hansen, 2014). Inside uncertainty pertains to the model builder’s assumptions about preferences and attitudes that agents inside the model have towards ambiguity.¹ Outside uncertainty denotes the fact that model builders need to estimate the model parameters, such that the model accurately reflects financial reality. While a large finance literature exists on inside uncertainty, outside uncertainty has largely been neglected. What are the consequences of ambiguity for model builders, and how can we deal with outside uncertainty on top of inside uncertainty in a tractable way? For instance, what is a reliable optimal asset allocation estimated from a time series of stock returns that contains unknown, time-varying features, when the model parameters are unknown? Similarly, how large are the premiums that investors require for bearing risk and ambiguity?

Our approach for dealing with ambiguity is designed to bound the effects of ambiguous data features in financial time series on model outputs. In various settings, we show that the presence of ambiguous data features can imply large biases for estimated optimal decision rules and for key equilibrium variables. These biases arise whenever the adverse effects of ambiguous data features are not bounded during the estimation, irrespective of whether a model builder accounts for inside uncertainty by allowing for an aversion to ambiguity in

¹At least since Knight (1921) and Ellsberg (1961), agents are known to dislike ambiguity. Various theoretical approaches to specify ambiguity aversion in static as well as dynamic economies have been introduced; see, e.g., Gilboa and Schmeidler (1989), Epstein and Wang (1994), Hansen and Sargent (2001), Chen and Epstein (2002), Epstein and Schneider (2003, 2007), Klibanoff, Marinacci, and Mukerji (2005, 2009).

agents' decisions. Thus, incorporating inside uncertainty is not enough to obtain results that are robust under outside uncertainty. In contrast, the negative effects of ambiguous data features are virtually eliminated by our approach.

To facilitate the discussion, Table 1 presents an overview of how our paper is positioned relative to the existing literature. The literature on financial modeling under uncertainty can be grouped into four categories, depending on whether (i) the estimated model includes an aversion to ambiguity (i.e., the model builder addresses inside uncertainty) and (ii) the estimation method bounds the effects of ambiguous data features (i.e., the model builder addresses outside uncertainty). Quadrants A and B in Table 1 are based on the standard approach in the literature, in which outside uncertainty is not taken into account and the effects of ambiguous data features are not bounded during the estimation of the model. Quadrant A focuses on models of standard risk averse behavior. Motivated by unknown, time-varying data features, the models in Quadrant B allow the agents inside the model to account for ambiguity when making decisions. An aversion to ambiguity can arise from the agents' concern that a distortion of their reference belief², generated by unknown, time-varying features, could make a decision rule less effective than under the reference belief assumptions.³ The literature shows that ambiguity aversion can provide explanations for important asset pricing questions, including the equity premium, risk-free rate, and credit spread puzzles, the large skew of implied volatility smiles of equity index options, the predictability of bond returns, stock market non-participation, underdiversification, home-bias, and flight-to-quality effects; Epstein and Schneider (2010) provide a review of this literature. To study how ambiguity aversion affects optimal decisions and asset prices, this literature

²A reference belief is a description of the DGP, which the model builder considers to be the best description of the data. It is the model that the agent inside the model is assumed to use for her decision making.

³Epstein and Schneider (2007) and Epstein and Schneider (2008), for instance, emphasize that the true DGP can depend on small, unknown, time-varying features, which are too difficult to understand and model. Most models of ambiguity averse behavior imply cautious ex-ante decision rules, which explicitly account for the fact that agent's reference belief might be an imperfect description of reality. Such behavior is linked to a reduced dependence of the effectiveness of optimal decision rules on time-varying, ambiguous features.

typically takes the choice of a reference belief including all parameters as given.⁴

[Table 1 about here.]

In reality, model builders do not only need to specify how agents inside the model confront uncertainty, but they also have to estimate the agents' reference model, facing outside uncertainty. Since the true DGP includes unspecified, time-varying features, identifying an appropriate reference belief from observed financial data is a nontrivial task. Even when the reference belief is appropriate for the vast majority of the data, ambiguous data features can have an excessive influence on the selection of the reference belief parameters, when using common estimation techniques such as maximum likelihood (ML) or generalized method of moments (GMM) estimation. This property is problematic from the perspective of a robust evaluation of asset pricing models, which is the model builder's primary goal, as it can imply estimated decision rules and equilibrium variables that mostly depend on unknown, time-varying features, rather than on the reference belief and agents' attitudes towards risk or ambiguity.

Our approach to the estimation of models under uncertainty is designed to bound the adverse effects of ambiguous financial reality, by limiting the dependence of estimated decision rules and equilibrium variables on ambiguous time-varying data features. To that end, we use and extend results from the robust econometrics literature and apply them in the context of financial models with ambiguity. In contrast to our approach, many conventional methods, such as ML or GMM, imply an unbounded dependence. Our method is applicable to a broad set of models under minimal assumptions. Sufficient conditions are the choice of a (parametric) reference belief for the relevant state dynamics and the determination of the maximal distortion that the model builder is willing to tolerate for the relevant financial variables. While our method also applies to models without ambiguity, it is particularly intuitive when the model builder assumes agents in the model to be ambiguity averse. Indeed,

⁴In some cases, the empirical predictions of ambiguity aversion are derived by calibrating the reference belief parameters to match unconditional asset returns. Other authors have estimated the reference belief of dynamic models of ambiguity using standard econometric techniques, such as, e.g., a (pseudo) maximum likelihood or a generalized method of moments approach. For instance, Benigno and Nisticò (2012) estimate a VAR(1), including various macroeconomic variables and excess returns. Ulrich (2013) applies pseudo ML methods to estimate the structural dynamics of a yield curve model with ambiguity aversion, while Chen, Ju, and Miao (2014) rely on ML to estimate the asset return dynamics in presence of ambiguous predictability.

if a model builder accounts for the presence of unknown, time-varying features in agents' behavior in the model, these features should not be ignored during estimation. Our approach to model building is coherent in this sense.

Formally, our method produces the most efficient estimation of the reference belief, under the constraint of a bounded effect of ambiguous data features. Therefore, it is applicable without explicit assumptions about investors' attitudes to ambiguity and it is distinct from approaches that explicitly incorporate, e.g., a max-min attitude to ambiguity at the estimation level, such as Chamberlain (2000) and Ait-Sahalia and Brandt (2001). These approaches are by definition difficult to apply when it is not natural or feasible to select a single attitude to ambiguity for estimation purposes, as, for instance, in heterogeneous equilibrium economies or in incomplete derivative markets with ambiguity. Moreover, an equivalence between max-min estimation and max-min ambiguity averse behavior arises only under additional assumptions, which may be violated in reality, such as the identity of the sets of beliefs that can be justified by the model builder and by the agent inside the model; see Cerreia-Vioglio et al. (2013).⁵ In our approach, we adopt a weaker assumption, which only requires the model builder's parametric set of reference beliefs to contain the reference belief of the agents inside the model.

Quadrants C and D in Table 1 are directly linked to the bounded-influence approach proposed in this paper. They consider models of financial behavior based on expected utility and aversion to ambiguity, respectively. Our work naturally complements the vast literature that has tested key empirical asset pricing predictions under the conditions of Quadrants A and B. Comparing the results of financial models in Quadrants A and B with those in Quadrants C and D allows us to quantify the estimation error and economic implications for model builders, caused by a given model estimation approach, in presence and absence of ambiguous, time-varying data features in financial reality, both for ambiguity neutral and

⁵More precisely, max-min estimation and max-min ambiguity are equivalent if the subjective ambiguity averse preferences and the restricted set of objectively rational probabilistic beliefs that are justifiable, e.g., from physical or symmetry considerations, based on available information are consistent (Cerreia-Vioglio et al., 2013). In our context, this assumption would imply on the one side the identity of the sets of beliefs that can be justified by the model builder and by the agent inside the model, based on their information sets. On the other hand, it would imply that with respect to this set of common beliefs agents inside the model would have to satisfy an additional rationality axiom.

for ambiguity averse decision rules.

We show that accounting for outside uncertainty is crucial for the reliable estimation of key financial variables and that it is not sufficient to specify ambiguity averse preferences to obtain robust asset pricing results in presence of ambiguity. In the realistic setting where financial reality contains ambiguous, time-varying features, we show that the estimation of a reference belief can imply surprisingly large biases for the decision rules and equilibrium variables of financial models, whenever the influence of ambiguous data features is not bounded appropriately. Such biases can arise when using conventional estimation methods, even if the reference belief is a good model for the vast majority of the data. The biases are statistically and economically relevant, both for economies displaying pure risk averse behavior and for economies in which agents are risk and ambiguity averse. In contrast, they are virtually eliminated using our bounded-influence approach.

We quantify the economic value of our approach by Monte Carlo simulations and in a real-data application. In a financial reality in which agents are ambiguity averse and the DGP contains unknown, time-varying features — the most realistic scenario — we show that the identified biases in estimated optimal policies, when not using bounded-influence methods, are of the same order as those of neglecting ambiguity in investors’ decisions, when the reference belief parameters are known. More precisely, in a portfolio planning problem with ambiguous rare disasters, we find that optimal portfolios and out-of-sample utilities estimated by conventional methods can be severely biased. In general equilibrium, conventional methods can result in strongly biased risk and/or ambiguity premiums, while in incomplete derivative markets these methods tend to overestimate bid-ask spreads. These distortions are successfully eliminated using our bounded-influence approach. Finally, in a real-data asset allocation problem with ambiguous return predictability, we uncover optimal portfolios largely insensitive to unknown, time-varying data features and show that our approach can lead to different economic implications compared with classical methods. In summary, this evidence suggests that our bounded-influence approach successfully deals with outside uncertainty in financial reality. Compared to conventional methods, it reduces or even eliminates the biases caused by ambiguous, time-varying data features for estimated decisions rules and equilibrium variables of financial models.

While we quantify the economic value of our bounded-influence approach for specific applications, it is important to note that the approach itself and the intuition of our results are valid also for more general settings and questions. Indeed, many interesting financial and economic models of ambiguity have been studied in the literature, mostly under the assumptions of Quadrant B. In a portfolio and asset allocation context, Cao, Wang, and Zhang (2005) and Trojani and Vanini (2004) focus on endogenous stock market participation. Routledge and Zin (2009) clarify that ambiguity can reduce liquidity, while Easley and O’Hara (2009) explain by the non-participation of ambiguity averse rational traders the absence of trading during the 2007–2009 financial crisis. Epstein and Miao (2003) and Uppal and Wang (2003) demonstrate that diverse aversions to ambiguity can imply home biases and underdiversification, while Benigno and Nisticò (2012) link the U.S. home bias to a hedging demand against long-run real exchange rate ambiguity. Caballero and Krishnamurthy (2008) explain the flight-to-quality during financial crises and Miao and Wang (2011) address option exercise in real investment and exit problems. All of those papers abstract from the potential effects of ambiguous data features for the estimation of agents’ reference beliefs. Therefore, the predictions of these models such as, e.g., the emergence of market participation, can be understood mostly as qualitative features of models of ambiguity. Our results suggest that a bounded-influence approach is helpful to make quantitative predictions for these models, because in reality reference beliefs are unknown to the model builder and need to be estimated from data that contain ambiguous, time-varying features.

Our results on the potential biases created by ambiguous data features are relevant also in a more structural general equilibrium perspective. For instance, Maenhout (2004) and Trojani and Vanini (2002) find a large equity premium and a low risk-free rate, while Leippold and Trojani (2008) show that learning additionally explains excess volatility. Trojani and Sbuelz (2008) explicitly link time-varying ambiguity to countercyclical returns and volatilities. Gagliardini, Porchia, and Trojani (2009) capture deviations from the expectations hypothesis in bond markets, while Ulrich (2013) estimates the link between uncertainty aversion and inflation premiums. Boyarchenko (2012) shows that ambiguous accounting signals help explain credit default swap spreads and Drechsler (2013) demonstrates that uncertainty about fundamentals explains equity premiums, return volatilities, and option skews. Finally,

Ju and Miao (2012) generate procyclical (countercyclical) dividend-price ratios (equity premiums) with ambiguous consumption and dividend dynamics. All these papers estimate the parameters of the relevant reference belief using methods that have an unbounded sensitivity to ambiguous, time-varying features. While the concrete development of bounded-influence methods for complex model dynamics requires a case-by-case study, the approach proposed in this paper is applicable also to these settings.

Our results are not only relevant for academic work, but they also have direct implications for market participants. A key insight is that model builders in financial institutions should give at least as much care to the estimation of the model as to the development of robust models in which uncertainty is factored in. Following our approach provides economic value and allows for more reliable decisions under uncertainty. Regarding the recent regulatory focus on model risk, our results highlight the importance of outside uncertainty and the benefits of bounded-influence estimation.

1. A Bounded-Influence Approach to Deal with Ambiguous Time-Varying Features

1.1. Bounding the Sensitivity of Investor's Utility to Time-Varying Data Features

Consider a typical setting from the ambiguity literature, in which a model builder wants to estimate the optimal policy of an ambiguity averse agent. The model builder assumes that the agent inside the model has a reference belief P_{θ_0} with parameter $\theta_0 \in \Theta$ and a given specification of attitudes to ambiguity. In line with common practice, let $u := \{u_t\}$ denote dynamic distortions to the reference belief. The set \mathcal{U} of distortions defines the subjective set of likelihoods $P_{\theta_0, u}$ that the agent inside the model considers for her decision making. The set of likelihoods $P_{\theta_0, u}$ contains the agent's reference model P_{θ_0} . The optimal policy and the utility implied by the agent's optimal choice depend on the reference belief parameter and are denoted by $C(\theta_0)$ and $V(\theta_0)$, respectively. We assume that $C(\cdot)$ and $V(\cdot)$

are smooth (differentiable) functions of reference belief parameter θ . This assumption is typically satisfied by models of ambiguity aversion.

Consider now a realistic scenario that model builders face in reality, namely that real data observations are generated by probability P_{θ_0, u^*} , where u^* is an unknown, possibly time-varying, distortion to the reference belief P_{θ_0} . Given that the distortion u^* is unknown, it may or may not be an element of the set of subjective distortions \mathcal{U} . Let $\tilde{\theta}(P_{\theta_0, u^*})$ be the asymptotic value of an estimator for the reference belief parameters applied to probability P_{θ_0, u^*} . Following a simple plug-in approach, the optimal policy and utility of a decision maker with a given degree of subjective ambiguity and a given attitude to ambiguity can be estimated by $C(\tilde{\theta}(P_{\theta_0, u^*}))$ and $V(\tilde{\theta}(P_{\theta_0, u^*}))$, respectively. In this way, we obtain a convenient description of the relation between the distorted probability P_{θ_0, u^*} , generating the observed data, and the model builder's estimate for the optimal policy and utility of the decision maker.

We assume that the reference belief P_{θ_0} is a plausible, even though approximate, description of the physical probability generating the observed data, i.e., we assume distortion u^* to be small. This is imposed by assuming that the distance between measures P_{θ_0, u^*} and P_{θ_0} is small. Such an assumption allows us to motivate the following first order functional expansion of the estimated optimal policy $C(\tilde{\theta}(P_{\theta_0, u^*}))$ in a neighborhood of the reference belief:

$$C(\tilde{\theta}(P_{\theta_0, u^*})) - C(\theta_0) = \frac{\partial C(\theta_0)}{\partial \theta'} \frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0} + o(|P_{\theta_0, u^*} - P_{\theta_0}|), \quad (1)$$

where $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}}$ is a (functional) Gâteaux derivative and $|\cdot|$ is a suitable norm on the vector space of finite signed measures; see, e.g., Von Mises (1947), Fernholz (1983) and Künsch (1984).

The first term on the right hand side of Equation (1) describes the linearized asymptotic effect of a small distortion u^* on the estimated optimal policy $C(\tilde{\theta}(P_{\theta_0, u^*}))$. This term is the product of two components: the sensitivity of the reference belief optimal policy to parameter θ , and the sensitivity of the estimator's asymptotic value to the unknown distur-

tion u^* . The first component, $\frac{\partial C(\theta_0)}{\partial \theta'}$, depends on how the model builder incorporates inside uncertainty. Once she has made assumptions about the agent's preferences and attitudes towards ambiguity, the term is known. The second component, $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$, is unknown because of outside uncertainty, i.e., the model builder does not know the latent distortion u^* that is part of the DGP. Importantly, even though u^* is small and difficult to detect, it can produce an arbitrarily large effect on the policy estimate $C(\tilde{\theta}(P_{\theta_0, u^*}))$, whenever the mapping $u^* \mapsto \frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$ defines an unbounded function of u^* . To limit this effect, outside uncertainty needs to be addressed by appropriately bounding $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$.

The quantity $\frac{\partial C(\theta_0)}{\partial \theta'}$ can be interpreted as a measure of the sensitivity of the estimated optimal policy with respect to an incorrect choice of parameter value θ for an agent's reference belief. The smaller $\frac{\partial C(\theta_0)}{\partial \theta'}$, the more robust is the optimal policy to an incorrect estimation of the reference belief parameter θ_0 . In contrast, $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$ measures the sensitivity of an estimation procedure to the fact that the reference belief P_{θ_0} is only an approximation of reality P_{θ_0, u^*} . The smaller $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$, the more robust is the estimator to unknown, time-varying features in the data. Note that since in Equation (1) only $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$ depends on u^* , the sensitivity of the estimated optimal policy and utility to outside uncertainty is fully characterized by the sensitivity of estimator $\tilde{\theta}$ to the distortion u^* .

Example. To understand the intuition behind the term $\frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$ in Equation (1), consider a simplified setting, in which we need to estimate the reference belief Sharpe ratio of a return time series, using an estimator $\tilde{\theta}$. The reference belief P_{θ_0} for returns is a normal distribution with parameters $\mu_0 = 0.05$ and $\sigma_0 = 0.2$. The true distribution P_{θ_0, u^*} is a normal distribution with parameters μ_0 and σ_0 , distorted by a small fraction α of returns that follow a different, unknown distribution G :

$$r_t \stackrel{iid}{\sim} (1 - \alpha)N(\mu_0, \sigma_0^2) + \alpha G . \quad (2)$$

We consider two possible estimators for the reference belief Sharpe ratio of returns, (i) the ratio of the sample mean and the sample standard deviation of returns, and (ii) the ratio of the sample median and the sample median absolute deviation of returns. Figure 1

presents the estimated Sharpe ratios under a contaminating distribution $G = N(\mu_0, 0.5^2)$, for different fractions of distortion $\alpha \in [0, 0.1]$. The results clearly show that the first estimation approach is more sensitive to deviations from the reference belief, i.e., it has a large sensitivity $\left. \frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \right|_{u^*=0}$. To illustrate, a small contamination of 3% of the data implies a reduction of about 10% in the estimated Sharpe ratio, whereas using the ratio of sample mean and sample standard deviation, the Sharpe ratio reduction using the ratio of sample median and sample median absolute deviation is less than 2.5%.

[Figure 1 about here.]

The intuition from the example holds in general. Whenever the goal is to estimate the reference belief of an ambiguity averse agent from ambiguous data, it is preferable to adopt estimators that have a low sensitivity to the DGP deviating from the reference belief. The goal of our bounded-influence method is to ensure that the difference $C(\tilde{\theta}(P_{\theta_0, u^*})) - C(\theta_0)$ in Equation (1) is bounded, independent of the particular form of unknown, time-varying features in the data.⁶ To attain this goal, we focus on estimators implying a bounded sensitivity to small distortions in the data, i.e., estimators for which $\left. \frac{\partial \tilde{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \right|_{u^*=0}$ is a bounded function of u^* .

1.2. Bounded Utility Sensitivity and Bounded Influence Estimators

The median or the well-known Huber (1964) estimator are examples of bounded-influence estimators. Such estimators aim at bounding the sensitivity of the resulting point estimates to arbitrary local deviations from a given reference belief. According to Equation (1), bounded-influence estimators imply a bounded sensitivity of estimated optimal policies and utilities to unknown, time-varying data features, thus addressing outside uncertainty.

For practical purposes, a general class of optimal bounded-influence estimators for ML, GMM or efficient method of moments (EMM)-type estimation settings can be defined. Such estimators are characterized by a bounded estimating function $\psi_c(\mathbf{s}(\cdot))$, which is obtained

⁶This is a very basic robustness requirement, because in all other cases a small distortion of the reference belief would cause an arbitrarily large bias in the point estimate for the optimal policy and utility of the decision maker inside the model.

by optimally weighting the unbounded ML, GMM or EMM estimating function $\mathbf{s}(\cdot)$ implied by the assumption of an exact reference belief; see Mancini, Ronchetti, and Trojani (2005) and Ortelli and Trojani (2005), among others. Let \mathbf{y} denote the time series of observations from the true DGP and $\boldsymbol{\theta}$ be the parameter vector of interest for the given reference belief. Bounded influence estimator $\widehat{\boldsymbol{\theta}}_{\text{rob}}$ is the ML, GMM or EMM-type estimator defined by the orthogonality condition $\mathbb{E}[\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}))] = \mathbf{0}$. The weighted estimating function $\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}))$ is defined by:

$$\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta})) = \mathbf{A}(\boldsymbol{\theta})(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}) - \boldsymbol{\tau}(\mathbf{y}, \boldsymbol{\theta}))\mathbf{w}(\mathbf{y}, \boldsymbol{\theta}), \quad (3)$$

where matrix \mathbf{A} and vector $\boldsymbol{\tau}$ are tuning parameters, while weight function \mathbf{w} is defined for $c \geq \sqrt{\dim(\mathbf{s}(\cdot))}$ by:

$$\mathbf{w}(\mathbf{y}, \boldsymbol{\theta}) = \min(\mathbf{1}, c \|\mathbf{A}(\boldsymbol{\theta})(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}) - \boldsymbol{\tau}(\mathbf{y}, \boldsymbol{\theta}))\|^{-1}), \quad (4)$$

with $\|\cdot\|$ denoting the Euclidean norm.⁷ The weighted estimating function $\boldsymbol{\psi}_c(\mathbf{s}(\cdot))$ is bounded and therefore ensures a bounded sensitivity of estimator $\widehat{\boldsymbol{\theta}}_{\text{rob}}$ if the reference belief deviates from the true DGP. The constant c can be chosen in a fully data-driven way. It controls the tradeoff between efficiency of the estimator in absence of outside uncertainty and estimator robustness when unknown, time-varying features are present in the DGP. Throughout this paper, we follow the common approach to choose c , such that the estimator achieves 95% efficiency, when the reference belief and the true DGP coincide.

Bounded-influence estimators of the above type can be defined for a fairly general class of reference beliefs, including, e.g., models with non-Gaussian returns, state-dependent volatility, jumps, or predictability features.

⁷Tuning matrix $\mathbf{A}(\boldsymbol{\theta})$ in Equation (3) ensures that the scaling condition

$$\mathbb{E}_{\boldsymbol{\theta}_0} [\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}_0))\boldsymbol{\psi}_c(\mathbf{s}(\mathbf{y}, \boldsymbol{\theta}_0))^\top] = \mathbf{I} \quad (5)$$

is satisfied, which implies that the norm of the self-standardized influence of the estimator is bounded by c . $\boldsymbol{\tau}$ is a nuisance parameter that adjusts the bias associated with a modification of estimating function $\mathbf{s}(\mathbf{y}, \boldsymbol{\theta})$ by weighting function \mathbf{w} . The internet Appendix provides various supplemental derivations and results. E.g., details on the updating algorithm to obtain bounded-influence parameter estimates according to Equation (3) are discussed in Section II. It is available on SSRN: <http://ssrn.com/abstract=2218212>.

2. Taking Ambiguity to Reality

In this section, we study the economic value of our bounded-influence approach for three concrete settings in which ambiguity is a natural concern. First, we address outside uncertainty when estimating the optimal portfolios of ambiguity averse agents in presence of unknown, time-varying features in the DGP.⁸ Second, we estimate the risk and ambiguity premiums in a jump-diffusion general equilibrium economy. Finally, we estimate bid and ask derivative prices of economies in which ambiguity generates a degree of market incompleteness.

2.1. Ambiguous Event Risk and Portfolio Choice

Unknown, time-varying features can influence different properties of the conditional distribution of asset returns, including, e.g., skewness, kurtosis, and event risk dynamics. As noted by Liu, Pan, and Wang (2005), event risk dynamics are very difficult to estimate precisely, because they depend on events that are rarely observed. Therefore, in such settings a misspecification of the reference belief dynamics for event risk is likely.

2.1.1. Time-Varying Event Risk: Inside Uncertainty

The model builder specifies a Merton (1976) jump-diffusion (JD) as the reference belief for market returns:

$$dS_t = \mu_S S_t dt + \sigma S_t dB_t + \left(e^{\xi^Y} - 1 \right) S_t dN_t, \quad (6)$$

with standard Brownian motion B , $\mu \in \mathbb{R}$, $\sigma > 0$, $\mu_S = \mu + \frac{1}{2}\sigma^2$, an iid jump size $\xi^Y \sim N(\mu_Y, \sigma_Y^2)$, and a Poisson process N with constant intensity λ_Y . The model builder addresses inside uncertainty by considering ambiguity with respect to the structure of event risk. For tractability, she specifies the agent's subjective attitudes to ambiguity with the max-min approach proposed by Liu, Pan, and Wang (2005).

The agent acts as a decision maker who invests a share π_t of her wealth into the JD stock

⁸The Internet Appendix helps building further intuition for inside and outside uncertainty in a simplified setting.

and the remainder at the risk-free rate r . She faces the following portfolio choice problem:

$$V(W_t, t) = \sup_{\{C_t, \pi_t\}} \inf_{\{a_t, b_t\}} E^{\{a, b\}} \left[\int_0^T e^{-\beta t} \left(\frac{C_t^{1-\gamma}}{1-\gamma} + \frac{1}{\Psi} H(a_t, b_t) \right) dt \right], \quad (7)$$

subject to the wealth dynamics

$$dW_t = [W_t(r + \pi_t(\mu_S - r)) - C_t] dt + \pi_t \sigma W_t dB_t + \pi_{t-} W_{t-} (e^{\xi^Y(b_t)} - 1) dN_t(a_t), \quad (8)$$

where C_t is the agent's consumption, a_t and b_t are the potentially time-varying parameters that describe a deviation from the reference model, and Ψ specifies the state dependent level of robustness sought. Time preference rate and risk aversion are given by β and γ , respectively. Function $H(\cdot)$ measures the discrepancy between the reference belief and alternative specifications for event risk at time t , parameterized by (a_t, b_t) :

$$H(a_t, b_t) = \lambda \left[1 + \left(a_t + \frac{1}{2} b_t^2 \sigma_Y^2 - 1 \right) e^{a_t} + d(1 + (e^{a_t + b_t^2 \sigma_Y^2} - 2) e^{a_t}) \right], \quad (9)$$

for some $d > 0$.⁹ The alternative specifications for event risk are implied by the Radon-Nikodym derivative for jump size and intensity:

$$dZ_t = \left(e^{a_t + b_t \xi_t^Y - b_t \mu_Y - \frac{1}{2} b_t^2 \sigma_Y^2} - 1 \right) Z_{t-} dN_{t-} - (e^{a_t} - 1) \lambda_Y Z_t dt. \quad (10)$$

This Radon-Nikodym process changes the jump intensity from λ_Y under the reference belief to $\lambda_Y^Z = \lambda_Y e^{a_t}$ for the distorted alternative. Similarly, the mean jump size $E \left[e^{\xi^Y} \right] = e^{\mu_Y + \frac{1}{2} \sigma_Y^2}$ under the reference belief is transformed to a mean jump size $E^{\{b\}} \left[e^{\xi^Y} \right] = e^{\mu_Y + \frac{1}{2} \sigma_Y^2} \cdot e^{b_t \sigma_Y^2}$ under the distorted alternative. The Internet Appendix derives the solution for the above portfolio choice problem when $\Psi(W, t) = \frac{\vartheta}{(1-\gamma)V(W, t)}$, where ϑ parameterizes the agent's attitudes to ambiguity.

As discussed in Section 1, a different degree of ambiguity aversion directly changes the

⁹ $H(\cdot)$ is the continuous time limit of an extended entropy measure; see Liu, Pan, and Wang (2005) for details.

sensitivity of the agent’s optimal policies and utilities to a small change in the value of the reference belief parameter. This effect can be measured by the derivative $\frac{\partial C(\theta_0)}{\partial \theta'}$ in Equation (1). Intuitively, a larger ambiguity aversion is linked to a more cautious optimal policy that is less sensitive to a variation of the reference belief parameters. This intuition is confirmed by the optimal policies in the economy with ambiguity about event risk. Indeed, the elasticities of portfolio weights with respect to all parameters, that are shown in Figure 2, are smaller, in absolute value, as the degree of ambiguity aversion, ϑ , increases.

[Figure 2 about here.]

For instance, while a one percentage point increase in μ increases the portfolio weight of an expected utility investor by more than 2%, the ambiguity averse investor only increases the allocation to the risky asset by about 1.75% ($\vartheta = 3$). Economically, this means that if μ is reduced from 0.08 to 0.06, the expected utility investor changes her allocation to risky assets from 0.53 to 0.29. In contrast, the ambiguity averse investor ($\vartheta = 3$), reduces the portfolio weight of risky assets by a smaller amount, from 0.46 to 0.27. Similar findings emerge for the elasticities with respect to all other reference belief parameters.

In summary, our findings confirm the intuition that a bias in the point estimate of the reference belief parameters is likely to produce a smaller bias for the optimal policies and utilities of agents with a larger ambiguity aversion. If the model builder addresses inside uncertainty by incorporating subjective ambiguity, this implies a higher degree of robustness with respect to estimation error in the reference belief parameters.

2.1.2. Time-Varying Event Risk: Outside Uncertainty

The second component of the first term in Equation (1) measures the sensitivity of optimal portfolios to a bias in the point estimates for the reference belief parameters. The overall sensitivity of the optimal policy is bounded, if and only if $\frac{\partial \hat{\theta}(P_{\theta_0})}{\partial P_{\theta_0, u^*}} \Big|_{u^*=0}$ is a bounded function of u^* . To quantify the implications of our bounded-influence approach in this context, we assume that returns over a period of $T = 6$ years follow the reference belief in Equation (6),

distorted by small, unknown, time-varying features:

$$dS_t = \mu_S S_t dt + \sigma S_t dB_t + \left(e^{\chi_t^Y} - 1 \right) S_t dN_t \quad ; \quad t \in [0, H], \quad (11)$$

where χ_t^Y is an iid jump size drawn from a time-varying distribution $N(\mu_{Y,t}, \sigma_Y^2)$ and $\mu_{Y,t}$ is, without loss of generality, a piecewise constant function of time given by:

$$\mu_{Y,t} = \begin{cases} -0.04 & t \in [0, 0.21T) \\ -0.01 & t \in [0.21T, 0.54T) \\ 0.01 & t \in [0.54T, 0.58T) \\ 0.12 & t \in [0.58T, 0.75T) \\ -0.02 & t \in [0.75T, 0.83T) \\ -0.10 & t \in [0.83T, T]. \end{cases} \quad (12)$$

This specification of time-varying distortion to simulate outside uncertainty is consistent with the way the model builder accounts for inside uncertainty in the model. The choice of the specific form of the distortion in Equation (12) is for illustration purposes, and other examples of specifications of unknown, time-varying features produce very similar effects (see the Internet Appendix).

The unconditional average jump size implied by the dynamics in Equation (11) is $\mu_Y = -1\%$, as in the reference belief, but it conditionally oscillates around this value. For instance, in the first 1.25 years the average jump size is -4% , instead of -1% . Note that such a distortion of the reference belief is small statistically. Jumps occur infrequently, and certain jumps may both be compatible with conditional and unconditional jump size distributions. Due to the structure of the distortion, the time-varying jump size alters the tail behavior of returns, but not their unconditional mean. The top left panel of Figure 3 illustrates the density of returns for the model in Equation (11), together with the density implied by the reference belief dynamics for the true parameters $\mu = 0.08, \sigma = 0.15, \lambda_Y = 3$,

$\mu_Y = -0.01$, $\sigma_Y = 0.04$. This comparison suggests a small statistical discrepancy between the two distributions.¹⁰ These conclusions are confirmed when looking at more formal measures of statistical discrepancy, such as detection error probabilities (DEP).¹¹ The detection error probabilities per piecewise constant jump size in Table 2 show that a model builder is not able to statistically distinguish the reference belief and the true process in up to 45% of the cases, which is far above the common significance level of 5%. The specific form of distortion is also reasonable given the the worst case model considered by the agent inside the model. With five years of data, the true reference belief parameters, and preference parameters $\gamma = \vartheta = 3$, the agent’s worst case model has a detection error probability of 36.23% relative to the reference belief.

In summary, even though the tail behavior of the two distributions is different, it is difficult to detect the specific structure of the time-varying jump component in Equation (11), using sample sizes typically available in practice. In reality, it would be even more difficult to identify such time-varying distortions, because their structure is unknown to the model builder.

[Figure 3 and Table 2 about here.]

2.1.3. Inside Uncertainty, Outside Uncertainty, and Bounded-Influence Estimation

In this section, we quantify the biases in the model builder’s estimation of the agent’s utility $V(\tilde{\theta}) - V(\theta_0)$ for different assumptions about the agent’s subjective attitudes to ambiguity, implied by different estimators $\tilde{\theta}$, when data include an unknown, time-varying distortion (Equation (11)). In this way, we quantify the added value of bounded-influence estimation methods, when estimating the optimal policies and utilities of agents with different subjective ambiguity attitudes.

¹⁰A sample realization for the distorted process (top right panel) and the decomposition into clean data (bottom right panel) and time-varying jump component (bottom left panel) further highlight the difficulty to distinguish these two processes based on realistic sample sizes of returns.

¹¹Given the exact likelihood functions for model (11) and the reference belief, the detection error probability is large when model (11) and the reference belief are statistically hardly distinguishable; see, e.g., Hansen and Sargent (2008), among others.

The model builder observes return data generated by the ambiguous dynamics in Equation (11) and aims to estimate reference belief P_{θ_0} , and the agent’s optimal policies and utility. Following Table 1, we consider both standard estimators and bounded-influence estimators. For each, we quantify the arising biases in the model builder’s estimation for different degrees of subjective ambiguity. To estimate the reference belief parameters, we derive the likelihood function for the dynamics in Equation (6).¹² The first order condition for the ML estimator of this likelihood function implies an unbounded estimating function $\mathbf{s}(\mathbf{y}, \boldsymbol{\theta})$. Thus, unknown, time-varying features in the DGP can have an unbounded influence on ML estimates in this setting. We obtain a new bounded-influence estimator for the JD in Equation (6) by applying the general approach outlined in Section 1. The resulting estimator has a weighted estimating function implied by Equation (3).

To illustrate the relation between subjective ambiguity attitudes and bounded-influence estimation in presence of unknown, time-varying features in the DGP, we fix an estimation sample of five years of daily market returns, generated by the distorted jump diffusion in Equation (11). We then compute the one year out-of-sample utility implied by the estimated optimal policies, both for ambiguity-neutral and ambiguity-averse agents. We finally compare estimation results and relate the estimation accuracy to the sensitivities of the different estimators with respect to unknown, time-varying features in the data (Equation (1)), illustrating the bias in the model builder’s conclusions, if she neglects inside uncertainty, outside uncertainty, or both. Panel (a) of Table 3 presents the wealth equivalents of the average estimated ex ante in-sample utility and the ex post realized out-of-sample utility when in- and out-of-sample returns follow a distorted jump diffusion (Equation (11)). Panel (b) reports average estimated utilities for the case in which both in- and out-of-sample returns follow the reference belief. Note that in this case the ex-ante expected utility is always equal to the ex-post realized utility. In both panels we distinguish three cases: (i) parameters are known, (ii) parameters are estimated by ML, and (iii) parameters are estimated using the new bounded-influence estimator.

First, consider the hypothetical case in which the reference belief, including parameters, is known to the model builder. These results constitute the benchmark for reality, in which the

¹²We report the analytic expression for the likelihood function in the Internet Appendix.

model builder needs to estimate the reference model parameters. When the parameters are known, the ex-ante utilities do not depend on the data and are thus equal in Panels (a) and (b). The model builder finds that the expected utility agent (Quadrant A in Table 1) has, as expected, the highest certainty-equivalent wealth. The reduction in the certainty-equivalent wealth of ambiguity averse agents ($\vartheta > 0$) is small. Due to the more cautious policy, an ambiguity averse agent with $\vartheta = 3$ expects 0.1% less wealth than the expected utility maximizer. The ex-post realized utility is equal to the expected utility, when ambiguity is absent, i.e., the reference model is equal to the true DGP and returns in- and out-of-sample follow a clean JD (first row of Panel (b)). When returns follow the distorted JD from Equation (11), the agent's ex-post realized utility is uniformly lower than in the clean JD case (fourth row of Panel (a)). A subjective concern for misspecification of the jump component improves the agent's utility by 2%, which is the motivation of ambiguity averse behavior. Does a model builder get similar results when the reference model parameters need to be estimated?

When relaxing the assumption of known parameters, a striking impact from parameter estimation emerges. Panel (a) of Table 3 shows that the difficulty of identifying rare event components in the true DGP may significantly bias the model builder's assessment of the agent's utility. For instance, when using ML estimation, the estimated wealth of an expected utility maximizer (Quadrant A) is reduced by 17% compared to the known parameter case. While accounting for inside uncertainty (Quadrant B) reduces this bias, the agent's utility is still estimated 6% too low. In contrast, bounded-influence estimation uniformly reduces the estimation error for the model builder. If the agent inside the model is an expected utility maximizer, utility is underestimated by 5%, while the estimation error relative to the known parameter case is reduced to 3% for an ambiguity averse agent. Thus, a model builder who uses our bounded-influence estimation approach is able to estimate the agent's utilities much more accurately, independent of her assumptions about how the agent in the model addresses inside uncertainty. As expected, the added value of bounded-influence estimation is largest if the agent inside the model does not consider ambiguity in her decisions. Moreover, even when the agent inside the model is ambiguity averse ($\vartheta = 3$) and the correct realized utility is 953, the cost of neglecting outside uncertainty (realized utility 895 vs. 924) is larger than the

cost of neglecting inside uncertainty in the hypothetical case of known parameters (realized utility 933 vs. 953), further highlighting the importance of a bounded-influence estimation approach.

Our results also have important implications if the model builder would like to assess the added value of incorporating inside uncertainty in the optimal decision. If the model builder relies on ML estimates, the estimated added value of taking inside uncertainty into account can be significantly overstated. Indeed, the estimated utility for an expected utility maximizer would be improved from 773 to 895, i.e., by 17%, when incorporating a concern for ambiguity in her decisions. If the model builder estimates the added value of being ambiguity averse using bounded-influence estimation techniques, the gain is reduced to 5% (from 882 to 924), which is much closer to the true parameter benchmark of 2% (from 933 to 953). Finally, when a clean JD generates the data, bounded-influence estimators (third row of Panel (b)) do not significantly change the estimated utility, when compared to ML estimators (second row of Panel (b)), despite the latter being optimal in this setup. Thus, the cost of using bounded-influence estimators under ideal model assumptions is negligible in this JD setting.

It is important to note that the excessive sensitivity of results to ambiguity in the data is an intrinsic feature of policies that are estimated using non-robust (pseudo) ML methods. To underline the generality of the problem, we can reverse the direction of the time-varying mean jump size contamination, relative to the unconditional jump size of -1% in Equation (11). Also in this case, our bounded-influence approach outperforms the ML estimators, in terms of realized certainty-equivalent wealth. The Internet Appendix provides detailed results.

In summary, our results show that the bounded-influence approach produces the best results for a model builder facing a portfolio problem with event risk and ambiguity, generated by small unknown and time-varying jump components. In contrast, neglecting outside uncertainty and relying on non-robust estimators can imply dramatic losses in estimation accuracy, even when inside uncertainty is taken into account to derive portfolio rules.

[Table 3 about here.]

2.2. Event Risk and the Equilibrium Market Prices of Risk and Ambiguity

Bounded-influence estimation is not only important for estimating the optimal policies of ambiguity averse decision makers. In this subsection we consider a second setting to quantify the added value of our bounded-influence approach. Namely, we consider a general equilibrium version of the (partial equilibrium) rare event economy introduced above. In this setting, the goal of the model builder is to estimate ambiguity and risk premiums. In equilibrium, ambiguity has to be priced and reflected by the time-series and cross-sectional behavior of asset prices. Liu, Pan, and Wang (2005) show that ambiguity aversion about event risk can help reconcile the equity premium and some of the smirk patterns of index options. To obtain estimates for the market prices of risk and ambiguity, the model builder develops a general equilibrium model from which these quantities can be derived. Then, she estimates the premiums based on time series of option implied volatility smiles and the corresponding underlying returns.

2.2.1. Identifying the Market Price of Event Risk and Ambiguity: The Model

Studying the general equilibrium problem of a representative agent with preferences given by Equation (7) allows for the computation of equilibrium premiums for risk and ambiguity, as well as equilibrium option prices. Given the market clearing condition $\pi_t^* = 1$, the Internet Appendix derives the implicit solutions for the equilibrium worst case choices a^* and b^* , which together with the representative agent's first order condition for optimal portfolio choice imply the following equity premium in the JD economy:

$$\mu_S - r = \gamma\sigma^2 - \lambda_Y e^{a^*} \left(e^{\mu_Y(1-\gamma) + \frac{1}{2}(1+\gamma^2)\sigma_Y^2(1+2b^*)} - e^{-\gamma\mu_Y + \frac{1}{2}\sigma_Y^2\gamma^2(1+2b^*)} \right). \quad (13)$$

How does this equity premium compare to that of standard economies? Clearly, the Black-Scholes (BS) environment implies $\lambda_Y = 0$, while the JD setting without ambiguity emerges

for $a^* = b^* = 0$:

$$\mathbf{BS:} \quad \gamma\sigma^2, \tag{14}$$

$$\mathbf{JD:} \quad \gamma\sigma^2 - \lambda_Y \left(e^{\mu_Y(1-\gamma) + \frac{1}{2}(1+\gamma^2)\sigma_Y^2} - e^{-\gamma\mu_Y + \frac{1}{2}\sigma_Y^2\gamma^2} \right). \tag{15}$$

Given a choice of parameters for the reference belief, Table 4 shows the relation between the equity premium and the preference parameters governing risk and ambiguity aversion. For instance, for $\gamma = 3$ we see that moving from a Black-Scholes to a JD economy nearly doubles the equity premium from 6.75% to 11.08%, under the given parameter choice. When the representative agent is additionally ambiguity averse (i.e., $\vartheta > 0$), the equity premium increases by approximately another 4% for each increase in parameter ϑ .

[Table 4 about here.]

Given the market incompleteness arising from the presence of jumps in the asset return process, the model builder needs additional information, e.g., from a time series of option prices written on the underlying JD equity index, to decompose the latent structure of the equity premium into a risk and an ambiguity premium component. The Internet Appendix shows that equilibrium option prices have to satisfy the following fundamental differential equation:

$$\begin{aligned} rC_t = & \frac{\partial C}{\partial t} + \left(r - \lambda_Y e^{a^*} E^{\{b^*\}} \left[e^{(1-\gamma)\xi^Y} - e^{-\gamma\xi^Y} \right] \right) S_t \frac{\partial C}{\partial S} + \frac{1}{2} \frac{\partial^2 C}{\partial S^2} \sigma^2 S_t^2 \\ & + \lambda_Y e^{a^*} E^{\{b^*\}} \left[e^{-\gamma\xi^Y} (C_t - C_{t-}) \right], \end{aligned} \tag{16}$$

subject to the standard boundary conditions. The solution of this equation provides the equilibrium pricing function $C_t = C(S_t, \boldsymbol{\theta}, \gamma, \vartheta)$ in the economy with ambiguity aversion, where parameter $\boldsymbol{\theta}$ collects the parameters describing the stock return dynamics under the reference belief. Note that this pricing function depends only on the reference belief and the representative agent's preference parameters. Therefore, the cross section of option prices is independent of potentially unknown, time-varying features in the DGP of underlying returns.

In the Internet Appendix, we visualize how the equilibrium option implied volatility smile depends on the representative agent’s risk and ambiguity aversion parameters in the economy. The implied volatility smile has the typical skewed shape of index option implied volatilities, in which out-of-the-money put implied volatilities are systematically larger. Moreover, as highlighted by Liu, Pan, and Wang (2005), risk and ambiguity aversion affect both the level and the slope of the smile, allowing the model builder to use option-implied information to identify risk and ambiguity premiums.

2.2.2. Market Prices of Risk and Ambiguity: Bounded-Influence Estimation

The model builder observes the equilibrium interest rate r , a time series of index prices S_1, \dots, S_T and a time series of equilibrium option prices $C(\boldsymbol{\theta}, S_1), \dots, C(\boldsymbol{\theta}, S_T)$. Using this information, she estimates the general equilibrium parameters and, in particular, the equilibrium risk and ambiguity premiums in the economy. This task is complicated by the fact that the reference model dynamics in Equation (6) are potentially distorted by unknown, time-varying features. Note that the specific features of such potential model distortions do not appear in any of the equilibrium parameters of interest. Therefore, they should also not strongly influence the point estimates of the model builder studying the general equilibrium economy with ambiguity aversion.

To assess the effect of time-varying features in the DGP, we simulate the data that the model builder observes. We use again the distortion specified in Equation (11) to simulate stock returns. We compute option prices using the reference belief of the representative agent, including the true parameter values and $\gamma = 2$ and $\vartheta = 1$. The true equity premium is thus 0.0826.

Based on these simulated data, the model builder estimates (i) the parameters of the reference belief and (ii) the risk and ambiguity preference parameters, such that model prices match the observed data. Table 5 shows the accuracy of estimated ambiguity and risk aversion parameters for ML and bounded-influence estimates. Using these estimates, the model builder computes risk and ambiguity premiums. Figure 4 plots the finite sample distributions of estimated equity premiums for the JD model, as well as those of the corresponding risk and

ambiguity premiums, obtained when using (i) ML estimators or (ii) our bounded-influence estimation approach. The left panels of Figure 4 show results for the case when the DGP includes time-varying features, according to Equation (11), that are unknown to the model builder. The right panels show results for the case in which the true DGP and the reference model coincide (Equation (6)). The estimation with clean data shows that risk and ambiguity premiums are estimated without any significant bias by all methods. Moreover, the finite sample distribution of premiums estimated by classical and bounded-influence methods is virtually indistinguishable, confirming that the efficiency costs of using bounded-influence instead of ML methods in the ideal JD setting are negligible. The situation is dramatically different in presence of time-varying distortions. When the data contain unknown, time-varying components, Panel (a) shows that the equity premium is systematically underestimated by ML. While our bounded-influence method implies an unbiased estimation of the true equity premium of 8.26%, the classical estimator yields a premium that is over nine percentage points smaller on average. Therefore, non-robust estimation approaches can lead to significant misconceptions about the size of the equity premium if the true DGP contains even small, time-varying distortions compared with the reference belief. Panels (c) and (e) further decompose the estimated equity premium into the risk and the ambiguity premium, respectively. We find that under time-varying distortion (Equation (11)), both the (diffusive and jump) equilibrium risk compensation and the ambiguity premium are underestimated by non-robust estimators.

Overall, these findings confirm that our bounded influence approach can consistently estimate the general equilibrium structure of economies with ambiguity aversion, when unknown, time-varying rare event features are present in the DGP of observed asset returns, while approaches using unbounded estimation techniques can again lead to large errors in the results.

[Table 5 and Figure 4 about here.]

2.3. Ambiguity in Incomplete Markets: Good Deal Bounds Pricing

The third setting that we investigate in order to assess the added value of our bounded-influence approach is the determination of good deal bounds for option prices in incomplete markets with ambiguity. In an incomplete market, explicit specifications of ambiguity aversion allow to identify a single stochastic discount factor and a single market price of ambiguity that are consistent with the given specification of ambiguity attitudes. However, in general, when a concrete optimizing behavior by a representative agent cannot be assumed, ambiguity in an incomplete market is associated with a multiplicity of stochastic discount factors and market prices of ambiguity that are all consistent with the observed prices of the underlying. In such settings, the multiplicity of stochastic discount factors generates an interval of arbitrage-free derivative prices, which gives rise to natural bid and ask price definitions, as the lower and upper bounds of the price interval. To narrow the set of arbitrage-free derivative prices in incomplete markets, a number of papers have proposed to exclude stochastic discount factors that are unrealistic in terms of their implied Sharpe ratios; see, e.g., Cochrane and Saá-Requejo (2000) or Björk and Slinko (2006). Following this intuition, bid and ask derivative prices in ambiguous incomplete markets can be naturally defined, as the minimally and maximally acceptable prices consistent with a maximal degree of misspecification with respect to the given reference belief. In this section, we show how our bounded-influence approach of Section 1 can be exploited to obtain robust estimates of (bid and ask) option prices, even when we remain agnostic about the attitudes to ambiguity.

2.3.1. Bid and Ask Derivative Prices in an Incomplete Market with Ambiguity

We assume that the model builder would like to quote option prices in an economy in which she is uncertain about the true DGP for the option's underlying. Such quotes are required in various scenarios, e.g., for market making, the search for profitable trading opportunities, or the valuation of projects using the real options approach. As in the previous sections, the model builder's reference model for the underlying's returns is the JD in Equation (6). Even without ambiguity about the jump component, the market is incomplete and there exist

infinitely many arbitrage free prices for the option, unless prices of other traded options are available. The presence of ambiguity introduces an additional degree of incompleteness, because the model builder considers alternative dynamics that are associated with different market prices of ambiguity. For simplicity and to isolate the pricing effect of ambiguity, we assume, without loss of generality, a zero market price of jump risk.

The model builder considers a set of potential distortions to the dynamics in Equation (6), generated by the Radon-Nykodim derivatives in Equation (10) for parameters a, b such that $H(a, b) \leq \eta$, where the relative entropy $H(a, b)$ is defined in Equation (9). The constraint on $H(a, b)$ ensures a bound on ambiguous good-deals, by admitting only stochastic discount factors and market prices of ambiguity, implying a maximal entropy discrepancy $\eta \geq 0$ between candidate alternative models and the reference belief. Intuitively, η plays a similar role as the ambiguity aversion parameter ϑ in the robust portfolio choice problem considered above. Given the maximum relative entropy η between reference belief and alternatives, the bid and ask pricing bounds for a call option are given by

$$\inf_{a,b} C^{(a,b)}(S, K) \quad \text{subject to} \quad H(a, b) \leq \eta \quad (17)$$

and

$$\sup_{a,b} C^{(a,b)}(S, K) \quad \text{subject to} \quad H(a, b) \leq \eta, \quad (18)$$

where $C^{(a,b)}(\cdot)$ is the price of a call option with strike K and initial stock price S under the distorted JD with parameters a and b . Let (a^*, b^*) be the inf and sup solutions for the bid and the ask in Equation (17) and (18), respectively. The bid and ask options prices satisfy the following partial differential equation

$$rC_t = \frac{\partial C}{\partial t} + \left(r - \lambda_Y e^{a^*} E^{\{b^*\}} \left[e^{\xi^Y} - 1 \right] \right) S_t \frac{\partial C}{\partial S} + \frac{1}{2} \frac{\partial^2 C}{\partial S^2} \sigma^2 S_t^2 + \lambda_Y e^{a^*} E^{\{b^*\}} [C_t - C_{t-}], \quad (19)$$

subject to the usual boundary conditions. The benchmark price is obtained by setting $a^* = b^* = 0$ (i.e., $\eta = 0$), giving the reference model price, for which the closed-form solution

is well known from Merton (1976) .

Figure 5 depicts bid and ask prices for different levels of moneyness (Subplot (a)) and ambiguity (Subplot (b)), when the parameters of the reference model are known. The presence of ambiguity about the true DGP leads to significant bid-ask spreads. Interestingly, uncertainty causes the ask price to deviate more from the Merton price than the bid price, indicating that the model builder requires an additional compensation for the ambiguity in the hedge position when writing a call option. This is intuitive as a short call can theoretically cost the writer an infinite amount, while a short put has a maximum possible loss.

[Figure 5 about here.]

2.3.2. Bid-Ask Spreads: Dealing with Outside Uncertainty

When parameters need to be estimated, the sensitivity of classical estimates to time-varying features carries over to the pricing bounds. The general setting of Equation (1) is applicable also in this case, with the call price instead of consumption as the decision variable in the formula. Subplots (a), (c), and (e) of Figure 6 show that bounded-influence parameter estimates result in option prices close to the true model prices, even when unknown, time-varying features distort the dynamics of the underlying. For classical estimates, the true ask (bid) price is only included in a 95% (77%) confidence interval of the estimated price distribution. Moreover, the dispersion of the estimated distributions is on average 37% smaller for bounded-influence estimates based on their inter-quartile range. For comparison, subplots (b), (d), and (f) show bid, reference model, and ask prices in absence of time-varying unknown data features. Here, as expected, bounded-influence and classical parameter estimates both perform equally well: Estimated bid and ask prices are all essentially identical to the true model values.

[Figure 6 about here.]

The good performance of bounded-influence estimates is confirmed by Table 6, which reports (relative) bid-ask spreads. Estimated bid-ask spreads of non-robust estimators are signifi-

cantly larger and further deviate from the bid-ask spreads based on true reference belief parameters.

These findings have interesting economic implications, as in incomplete markets investors who trade options are typically required to maintain cash reserves to support a trade. Intuitively, these capital reserves need to be larger the larger the bid-ask spreads, because an investor who buys the option at the ask, but has to sell it again at the bid, incurs a cost equal to the bid-ask spread. Therefore, to be able to cover the cost of an unfavorable unwinding of the trade, investors need to hold an amount of capital proportional to the bid-ask spread; see Carr, Madan, and Alvarez (2011), among others. From an economic perspective, the results in Table 6 indicate that bounded-influence parameter estimates correctly imply significantly lower capital requirements. For instance, for out-of-the money options the capital charge implied by classical estimates can be four times larger than the capital required when estimating the parameters of the reference model using our bounded-influence approach. Even for in-the-money options, investors using classical estimates are required to hold at least 50% more capital. Consequently, costly capital reserves can be significantly overstated by non-robust parameter estimates. Similarly, favorable trading opportunities or investment projects may be missed when the bounds are estimated to be too wide due to unknown, time-varying features in the data. Overall, our bounded-influence approach also performs better in this third setting compared to classical estimates.

[Table 6 about here.]

3. Empirical Application: Ambiguous Predictability and Bounded-Influence Estimation

Motivated by the results and intuition of the previous sections, in which we analyzed the added economic value of our bounded-influence approach in different economies with ambiguity aversion, theoretically and by Monte Carlo simulation, this section studies a real data application to ambiguous stock return predictability. We take the role of a model builder, who would like to analyze the investment approach of an asset manager, such as a mutual or

hedge fund. In reality, the model builder could be an investor, searching for the best asset managers or a regulator, who would like to validate asset managers' promises regarding their investment performance.

We consider the realistic setting in which we do not know the manager's exact trading strategy and utility function, but from the fund's prospectus we know that it tries to exploit predictability features in stock returns. To answer the question whether this investment approach creates values, we build a model and analyze stock return predictability, taking both inside and outside uncertainty into account. Do our conclusions change if we neglect outside uncertainty? The following subsections proceed as follows. First, we apply our bounded-influence approach developed in Section 1 to estimate a reference model for stock returns, including predictability features. Next, we derive the asset manager's optimal portfolio decisions under ambiguous predictability for different assumptions about the asset manager's risk and ambiguity aversion. Finally, we evaluate the resulting realized wealth in- and out-of-sample.

3.1. Unknown, Time-Varying Components and Ambiguous Predictability Features

A large literature studies the predictability of stock returns by macroeconomic or financial indicators, such as the dividend-price ratio. The empirical literature tends to find some evidence of stock return predictability, especially at quarterly or yearly horizons. However, the strength of this evidence is mixed. Results significantly depend on the choice of the sample under investigation and it tends to be difficult to exploit estimated predictability features to produce additional economic value, e.g., in the context of real-time asset allocation; see, e.g., Lettau and Ludvigson (2010) and Welch and Goyal (2008) for a review of this literature.

The difficulty to identify predictability structures can be amplified by the presence of time-varying features in the underlying DGP. Lettau and Van Nieuwerburgh (2008), for instance, suggest that a permanent downward shift in the mean of the log dividend-price ratio during the mid-1990s can help reconcile some of the mixed results on stock return predictability. However, as noted by Lettau and Ludvigson (2010), such a structural break

hypothesis is hardly consistent with the more recent reversal to higher average dividend-price ratios, suggesting that the mid-1990s evidence was more likely associated with a temporary unusual period, rather than with a permanent structural shift in the overall structure of aggregate financial ratios.

Thus, from a broader perspective, the observed pattern in log dividend-price ratios is consistent with the presence of unknown, time-varying components in the DGP, in line with the intuition provided in the previous sections. Indeed, the Internet Appendix shows that a DGP including unknown distortions is a valid alternative hypothesis for the observed pattern in return and dividend-price ratio data, motivating an ambiguous modeling of predictability features. The following subsections study the relation between bounded-influence estimation and dynamic portfolio choice with ambiguous predictability in more detail.

3.2. Ambiguous Predictability and Bounded-Influence Estimation

We estimate a standard predictive VAR model using monthly value-weighted stock index (with and without dividends) and T-Bill returns, obtained from CRSP, for the period 1929–2013. We construct continuously compounded excess market returns and dividend-price ratios following Fama and French (1988)¹³. The sampling frequency is $\Delta = 1/12$, i.e., one month, x_t is the log dividend-price ratio, and $r_{t+\Delta} = \log S_{t+\Delta} - \log S_t - \Delta r^f$ is the market index excess return over the (annualized) T-Bill risk free rate r^f . The corresponding predictive system reads:

$$\begin{bmatrix} r_{t+\Delta} \\ x_{t+\Delta} \end{bmatrix} = \begin{bmatrix} d_1 \\ d_2 \end{bmatrix} + \begin{bmatrix} 0 & d_3 \\ 0 & d_4 \end{bmatrix} \cdot \begin{bmatrix} r_t \\ x_t \end{bmatrix} + \begin{bmatrix} v_{1,t+\Delta} \\ v_{2,t+\Delta} \end{bmatrix}. \quad (20)$$

¹³Dividend payments of the firms in the stock index are extracted as follows: In month t , the value of one dollar invested without reinvestment of dividends is $P(t) = \exp(r_{ND}(0) + r_{ND}(1) + \dots + r_{ND}(t))$, where $r_{ND}(t)$ denotes the continuously compounded return of the stock index without reinvestment of dividends. Then, dividend payments in month t can be computed as $D(t) = P(t-1) \exp(r_D(t)) - P(t)$, where $r_D(t)$ is the continuously compounded return of the stock index with dividends reinvested. Given the strong seasonality in dividends paid by the companies in the stock index, the dividend-price ratio in month t is defined as the sum of the dividends paid in months $t - 11$ to t , divided by the value of the stock index at time t .

Parameter d_3 captures potential stock return predictability features, while parameter d_4 measures the persistence of dividend-price ratios. $v_t = (v_{1,t}, v_{2,t})'$ defines a bivariate martingale difference process with finite second moments. Given the potential ambiguity affecting specification (20), we estimate the predictive VAR using (i) the standard ML estimator based on a Gaussian assumption for v_t and (ii) the corresponding bounded-influence estimator implied by the methodology from Section 1.

Figure 7 shows increasing window estimates for the predictability parameter d_3 obtained using the two methods, with confidence interval bounds of plus and minus one (asymptotic) standard error around each point estimate.¹⁴ The bounded-influence point estimate for d_3 is systematically larger, suggesting that bounding the effects of influential data points helps uncover a positive, even if time-varying, predictive relation between the dividend-price ratio and stock returns: While there is hardly any statistical evidence for predictability using the classical estimator, the bounded-influence estimator indicates the presence of predictability for certain sample periods. In line with, e.g., Chen (2009), stock returns become predictable mainly in the middle of our sample, with weaker evidence for predictability in recent years.

[Figure 7 about here.]

Estimated parameters using the full sample of data are presented in Panel (a) of Table 7. In addition to the larger predictability parameter, bounded-influence estimates indicate a stronger negative correlation between the VAR innovations, a smaller variance of the error terms, and a larger intercept in the return equation. These systematic differences between classical and bounded-influence point estimates suggest that indeed the findings provided by classical estimators might be highly dependent on the features of a limited fraction of observations.

The bounded-influence methodology provides simple ways to identify influential observations and to illustrate the amount of outside uncertainty revealed by the data, namely by inspecting Huber weights (see Equation (4)). Figure 8 plots excess returns and dividend

¹⁴In every month t , we estimate the model based on sample data from January 1929 up to the current month. Classical standard errors are not adjusted for potential small sample biases, which might imply an even weaker evidence in favor of predictability when using classical estimators and tests; see, e.g., Campbell and Yogo (2006) or Amihud, Hurvich, and Whang (2009).

yields over time, together with the corresponding Huber weight of each observation. Overall, a fraction of 7.6% of the observations has a Huber weight of less than one, i.e., those observations are particularly influential for classical parameter estimates. An important fraction of influential observations clusters during the Great Depression. In the postwar period, the stock market crash in 1987 has the largest influence, followed by observations having patterns similar to those of a sequence of isolated outliers. In particular, we see that influential observations often emerge as a result of large negative market returns. In the Internet Appendix we show that these patterns of influential observations match the effects of ambiguity within a Gaussian VAR model contaminated by a time-varying jump component, similar to the distortion we used in the simulation study in Section 2.

[Table 7 and Figure 8 about here.]

All in all, there appears to be some support in favor of return predictability with the Gaussian VAR model, at least when using our bounded-influence estimation approach. At the same time, the evidence of the presence of potentially influential observations suggests to treat the strict Gaussian VAR dynamics as ambiguous, i.e., to account for inside uncertainty when developing the model and for outside uncertainty when estimating the model.

3.3. Ambiguous Predictability and Bounded-Influence Portfolio Choice

What is the optimal policy for an asset manager in light of these ambiguous predictability features? Can an ambiguity averse asset manager profit from exploiting predictability? To answer these questions, we derive a robust version of the dynamic portfolio problem studied in Xia (2001). Given the evidence in favor of predictability above, we specify the reference model for the stock return dynamics as follows:

$$dS_t = \mu(X, t)S_t dt + \sigma S_t dB_t^{(1)} , \quad (21)$$

where $\sigma > 0$ and $B^{(1)}$ is a standard Brownian motion. The instantaneous expected return μ is related to a predictor variable X (the log dividend-price ratio):

$$\mu(X, t) = \alpha + \nu X_t, \quad (22)$$

with $\alpha, \nu \in \mathbb{R}$. Under the reference model, the predictor variable follows a persistent Gaussian process with dynamics:

$$dX_t = \kappa_X(\bar{x} - X_t)dt + b dB_t^{(2)}, \quad (23)$$

where $\kappa_X, \bar{x}, b > 0$ and $B^{(2)}$ is a second standard Brownian motion such that $dB_t^{(1)}dB_t^{(2)} = \rho dt$. This reference model is consistent with the discrete-time model of Barberis (2000), who assumes the dividend yield, as predictor, to follow an AR(1) process. We account for inside uncertainty by allowing the asset manager to be ambiguity averse with respect to predictability. Thus, the manager solves a corresponding robust asset allocation problem:

$$V(t, W, X) = \sup_{\{\pi_t\}} \inf_{\{u_t\}} E^a \left[\frac{W_T^{1-\gamma}}{1-\gamma} + \int_0^T \frac{1}{2\Psi} (\pi_t \sigma W_t u_t)^2 dt \right], \quad (24)$$

subject to the dynamic budget constraint

$$dW_t = [W_t(r + \pi_t([\bar{\mu} + \nu(X_t - \bar{x}]) - r)) + \pi_t^2 \sigma^2 W_t^2 u_t] dt + \pi_t \sigma W_t dB_t^{(1)},$$

where $\bar{\mu} = \alpha + \nu \bar{x}$ and $\Psi = \vartheta / ((1-\gamma)V(t, W, X))$, for $\gamma, \vartheta \geq 0$. This specification of ambiguity aversion in robust portfolio choice is borrowed from Maenhout (2006). The optimal portfolio rule π_t^* for an ambiguity averse investor with power utility follows similarly to the results of Section 2 and is given by

$$\pi_t^* = \frac{1}{\gamma + \vartheta} \left(\frac{\mu(X, t) - r}{\sigma^2} + (B(\tau) + C(\tau)(\mu(X, t) - r)) \frac{\nu \rho b}{\sigma} \right), \quad (25)$$

where functions $B(\cdot)$ and $C(\cdot)$ are solutions to a system of ordinary differential equations, given in detail in the Internet Appendix, and $\tau = T - t$ is the investment horizon.

Having specified the model including the asset manager’s attitude towards ambiguity, we estimate the reference model parameters using our bounded-influence approach to account for outside uncertainty. We also compute classical estimates for comparison. Parameters of the continuous-time reference model in Equations (21) and (23) are easily estimated using its exact discretization. The Internet Appendix shows that model (20), which we estimated in the previous subsection, is the exact discrete-time equivalent of the reference model in Equations (21) and (23), and derives the explicit link between continuous-time and discrete-time model parameters. Estimated continuous-time parameters implied by the discrete-time point estimates in Panel (a) are shown in Panel (b) of Table 7. Intuitively, the main implications obtained for the discrete-time estimates carry over to the continuous-time parameters. For instance, bounded-influence estimation results imply (i) a larger slope coefficient ν of log dividend-price ratios in the (predictive) equation for expected returns, (ii) a more negative correlation parameter ρ and (iii) lower volatilities of both stock returns and log dividend-price ratios.

Given the continuous-time estimates, we derive optimal portfolio policies using Equation (25). Panel (a) of Figure 9 depicts the resulting optimal stock allocations for different time horizons, based on the parameter estimates for the 1929–2013 sample. Our bounded-influence estimation approach yields significantly different results compared with classical ML estimates, which leads to distinct interpretations for the optimal portfolio of ambiguity averse agents. Bounded-influence parameter estimates imply an increased myopic and intertemporal hedging demand, relative to the weights based on the standard ML estimator. Both the stronger predictive relation and the more negative correlation between returns and the predictor variable implied by the bounded-influence estimates lead to larger allocations to equities. These differences are even more pronounced for standard CRRA asset managers, who do not take inside uncertainty into account in selecting their optimal policies. Note that this is in contrast to common perceptions that bounded-influence parameter estimates always lead to more cautious portfolios.

For each set of estimated parameters, myopic and hedging demands in Panels (b) and

(c) are both estimated to be smaller for the ambiguity averse agent, which is in line with the intuition that standard CRRA asset managers are more optimistic about future stock performance than ambiguity averse managers. Since a CRRA asset manager is more confident with the predictive power of the dividend yield, she also more rapidly increases the equity allocation, as a function of the investment horizon, due to a larger intertemporal hedging motive.¹⁵ However, the size of these estimated effects depends on the estimation approach. For instance, the absolute decrease in total demand for the risky asset when accounting for ambiguity aversion is much smaller when using classical rather than bounded-influence estimates.

[Figure 9 about here.]

3.4. Ambiguous Predictability: Economic Implications of Bounded-Influence Estimation

We conclude our analysis by investigating the economic implications of our bounded-influence approach. To that end we quantify the added economic value of exploiting predictability features for dynamic asset allocation under ambiguity, by studying the realized wealth of the asset manager in the economy with ambiguous return predictability.

3.4.1. In-sample Results

First, consider point estimates of classical and bounded-influence estimators from the whole sample, i.e, we take an in-sample perspective. Based on the estimates from Table 7, we implement the optimal portfolio policy in Equation (25) for an investment horizon of one year, using the current dividend-price ratio and risk-free interest rate in each month t as inputs. Panels (a) and (b) of Table 8 show time-series averages of the resulting optimal portfolio weights, for different levels of risk and ambiguity aversion, which confirm the patterns shown in Figure 9.

¹⁵Assuming an intermediate level of risk aversion in line with Barberis (2000), γ is assumed to be 3. Based on the average T-Bill return, the continuously compounded annual risk-free interest rate equals $r_f = 3.42\%$. Moreover, the initial log dividend-price ratio is set to be the average log dividend-price ratio over the sample period ($\log(3.66\%)$).

The economic value of in-sample optimal policies can be evaluated based on realized utility. To this end, we compute the wealth obtained from investing in the optimal portfolio in month $t + 12$. Panels (c) and (d) depict the time-series average realized wealth, based on classical and bounded-influence parameter estimates, respectively. We find that for every level of γ and ϑ , realized wealth according to bounded-influence parameter estimates is higher than the estimated wealth based on classical estimates. The difference in utility is largest for low levels of risk aversion (e.g., more than 2.33% for $\gamma = \vartheta = 2$), but it is still economically significant, with at least 0.5% per year, for very risk averse investors.

[Table 8 about here.]

How robust is the estimated realized wealth with respect to particular data features? To shed light on this important aspect, we study the sensitivity of estimated portfolio weights and certainty equivalent wealth along the lines of Dell’Aquila, Ronchetti, and Trojani (2003), who perturb the most influential observations, identified from estimated Huber weights, to investigate the robustness of the relevant decision variable. As a simple example, consider a perturbation of the October 1987 return, the most extreme postwar observation in the sample, by up to $\pm 30\%$. This corresponds to varying that month’s actual return of -28.52% between -37.08% and -19.96% . Table 9 shows the minimum and maximum of the time series means of portfolio weights, after perturbing the five most influential returns by $\pm 30\%$.

We find that for both ambiguity neutral ($\vartheta = 0$) and ambiguity averse ($\vartheta = 3$) agents, portfolio weights implied by classical estimates are very sensitive to such a perturbation of only approximately 0.5% of all observations. While bounded-influence estimates only imply a change in average weights of maximally 2.04% (1.71% for $\vartheta = 3$), classical parameter estimates induce a difference between weights of 6.03%, as a result of the $\pm 30\%$ perturbations (2.94% for $\vartheta = 3$): The sensitivity of optimal portfolio weights to this small perturbation is almost three times as large for classical parameter estimates.

Large portfolio weight sensitivities are linked to large sensitivities of the corresponding optimal utility or certainty equivalent wealth. Table 9 confirms this intuition. The estimated average wealth when optimal portfolios are estimated with classical parameter estimates is much more sensitive to a changes in a tiny fraction of observations. The minimum average

wealth using classical estimates is 1.28% smaller than the maximum. This difference is only 0.19% when we estimate the reference model with bounded-influence methods. Thus, conclusions about the profitability of trading rules are much more robust to small perturbations in the data when using our bounded-influence approach.

[Table 9 about here.]

3.4.2. Out-Of-Sample Realized Wealth

In reality, a model builder needs to estimate model parameters in real time to analyze the optimal investment policy. Therefore, in this subsection, we investigate realized wealth out-of-sample. In this more realistic setup, we use a rolling window approach and estimate the reference model in each month t based on a sample of the previous 25 years of data. This leads to an additional potential variation in the portfolio allocation. On the one hand, each new observation can contain valuable information about the reference model itself. On the other hand, a newly incorporated return observation may simply come from an unknown, time-varying feature of the underlying DGP, and might thus be very influential for classical ML estimates.

Table 10 compares the realized wealth for ambiguity neutral ($\vartheta = 0$) and ambiguity averse asset managers ($\vartheta > 0$), estimated using classical ML and bounded-influence estimators. In the first two rows, the reference model incorporates return predictability. In the last two rows, we assume instead a random walk process without predictability, allowing us to assess the profitability of exploiting potential predictability features compared with the random walk benchmark. The average realized wealth over the full sample is higher when incorporating predictability in the reference model, both based on classical and bounded-influence estimates. However, the added value of predictability is estimated to be larger when using classical estimates. While the estimated wealth increase for an ambiguity averse investor is 0.9% per year with classical estimates, it is only 0.5% using our bounded-influence approach. Thus, over the full sample, classical estimates tend to overstate the profitability of predictability features, due to a number of influential observations.

We investigate the issue further by comparing the profitability of the two trading strategies in different subsamples (Columns 3 to 8 in Table 10). In line with Figure 7, predictability is profitable mostly in the period from 1973 to 1990. For this subsample, both classical and bounded-influence estimators indicate that an asset manager accounting for predictability features outperforms asset managers, investing according to the random walk. The results in the first and last subsample depend on the estimation approach. While classical estimators slightly favor the predictability strategy, our bounded-influence approach clearly indicates a higher realized wealth when returns are assumed to follow a random walk. Thus, relying on our bounded-influence approach, which is robust to the presence of a few influential observations, can result in very different economic interpretations in general, and assessments of optimal portfolio behavior in particular, compared with classical estimation techniques.

[Table 10 about here.]

4. Conclusion

Model builders in financial markets develop economic models to learn about preferences, optimal decisions, and equilibrium quantities. These models will always be a simplification of the complex financial reality, which contains unknown, time-varying features. Thus, to be able to draw relevant conclusions, model builders have to account for both inside and outside uncertainty. The model builder does not only need to develop a theory, including agents' preferences towards risk and ambiguity, but she also needs to estimate model parameters, such that the model accurately reflects financial reality. In this paper, we show that this goal is only achieved if the model builder bounds the adverse effects of outside uncertainty when taking the model to the data. Standard estimation methods frequently used in the ambiguity literature do not have this property.

We propose a novel bounded-influence approach that successfully deals with outside uncertainty on top of inside uncertainty, when developing and estimating financial models. We find that such a bounded-influence approach is key for producing accurately estimated optimal policies and equilibrium variables that are robust to the presence of unknown, time-

varying features in financial reality. We motivate the necessity to apply a bounded-influence approach to obtain robust inference results theoretically, and quantify the economic value of our approach for three exemplified applications, using Monte Carlo simulation, and in a real-data asset allocation problem. First, in an economy with ambiguity about rare events, we show that the bias in estimated optimal policies when not using bounded-influence methods in presence of ambiguous time-varying data features is of the same order as that of neglecting ambiguous data features in investors' decisions, when the reference belief is known. Second, in general equilibrium, standard methods can imply severely biased risk and ambiguity premiums. Third, in incomplete derivative markets they materially overestimate bid-ask spreads. These distortions are successfully eliminated using our bounded-influence approach. Finally, we apply our approach to a real-data asset allocation problem with ambiguous predictability. Our method allows us to more accurately determine optimal portfolios with a realized wealth largely insensitive to unknown, time-varying data features. In summary, this evidence suggests that our bounded-influence approach provides more accurately estimated decisions rules and equilibrium variables of financial models, whereas conventional estimates can be substantially biased by ambiguous time-varying data features.

The approach itself and the intuition of our results are not limited to specific applications, but are applicable to a broad variety of dynamic financial settings and questions. Our results highlight that model builders in general should care at least as much about the estimation as about the development of models with ambiguity — accounting for inside uncertainty is useful, but not sufficient. Thus, our approach can be helpful to produce a number of new insights and interpretations for the growing literature on ambiguity aversion in finance, which has so far largely neglected outside uncertainty when studying the consequences of ambiguity for asset pricing.

Our paper also has direct implications for market participants. Model builders in financial institutions should account for outside uncertainty to obtain reliable estimates for their models, e.g., in risk management or portfolio allocation, allowing for more robust decisions in presence of uncertainty. Our findings are particularly relevant for regulators, who increasingly focus on the model risk involved in banks' operations. Which institutions already sufficiently address inside and outside uncertainty? What are the consequences for systemic

vs. idiosyncratic risk, and what are the implications for capital requirements? Applying our approach to models with ambiguity, in order to quantify the economic implications of both inside and outside uncertainty provides numerous avenues for future research.

References

- Aït-Sahalia, Y., and M.W. Brandt (2001), Variable selection for portfolio choice, *The Journal of Finance* 56, 1297–1351.
- Amihud, Y., C. M. Hurvich, and Y. Whang (2009), Multiple-predictor regressions: Hypothesis testing, *The Review of Financial Studies* 22, 413–434.
- Barberis, N. (2000), Investing for the Long Run when Returns are Predictable, *The Journal of Finance* 55, 225–264.
- Benigno, Pierpaolo, and Salvatore Nisticò (2012), International portfolio allocation under model uncertainty, *American Economic Journal: Macroeconomics* 4, 144–189.
- Björk, T., and I. Slinko (2006), Towards a General Theory of Good-Deal Bounds, *Review of Finance* 10, 221–260.
- Boyarchenko, Nina (2012), Ambiguity shifts and the 2007–2008 crisis, *Journal of Monetary Economics* 59, 493–507.
- Caballero, R.J., and A. Krishnamurthy (2008), Collective Risk Management in a Flight to Quality Episode, *The Journal of Finance* 63, 2195–2230.
- Campbell, J.Y., and M. Yogo (2006), Efficient Tests of Stock Return Predictability, *Journal of Financial Economics* 81, 27–60.
- Cao, H.H., T. Wang, and H.H. Zhang (2005), Model Uncertainty, Limited Market Participation, and Asset Prices, *The Review of Financial Studies* 18, 1219–1251.
- Carr, Peter, Dilip B. Madan, and Juan Jose Vicente Alvarez (2011), Markets, profits, capital, leverage and return, *The Journal of Risk* 14, 95–122.
- Cerreia-Vioglio, Simone, Fabio Maccheroni, Massimo Marinacci, and Luigi Montrucchio (2013), Ambiguity and Robust Statistics, *Journal of Economic Theory* 148, 974–1049.
- Chamberlain, G. (2000), Econometrics and decision theory, *Journal of Econometrics* 95, 255–283.

- Chen, Hui, Nengjiu Ju, and Jianjun Miao (2014), Dynamic asset allocation with ambiguous return predictability, *Review of Economic Dynamics* forthcoming.
- Chen, L. (2009), On the Reversal of Return and Dividend Growth Predictability: A Tale of Two Periods, *Journal of Financial Economics* 92, 128–151.
- Chen, Z., and L. Epstein (2002), Ambiguity and asset returns in continuous time, *Econometrica* 70, 1403–1443.
- Cochrane, J.H., and J. Saá-Requejo (2000), Beyond arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets, *Journal of Political Economy* 108, 79–119.
- Dell’Aquila, R., E. Ronchetti, and F. Trojani (2003), Robust GMM Analysis of Models for the Short Rate Process, *Journal of Empirical Finance* 10, 373–397.
- Drechsler, Itamar (2013), Uncertainty, Time-Varying Fear, and Asset Prices, *The Journal of Finance* 68, 1843–1889.
- Easley, David, and Maureen O’Hara (2009), Ambiguity and nonparticipation: The role of regulation, *The Review of Financial Studies* 22, 1817–1843.
- Ellsberg, D. (1961), Risk, Ambiguity, and the Savage Axioms, *The Quarterly Journal of Economics* 75, 643–669.
- Epstein, L.G., and J. Miao (2003), Two-Person Dynamic Equilibrium under Ambiguity, *Journal of Economic Dynamics and Control* 27, 1253–1288.
- Epstein, L.G., and M. Schneider (2003), Recursive Multiple Priors, *Journal of Economic Theory* 113, 32–50.
- Epstein, L.G., and M. Schneider (2007), Learning Under Ambiguity, *Review of Economic Studies* 74, 1275–1303.
- Epstein, L.G., and M. Schneider (2008), Ambiguity, Information Quality, and Asset Pricing, *The Journal of Finance* 63, 197–228.
- Epstein, L.G., and M. Schneider (2010), Ambiguity and asset markets, *Annual Review of Financial Economics* 2, 315–346.
- Epstein, L.G., and T. Wang (1994), Intertemporal Asset Pricing Under Knightian Uncertainty, *Econometrica* 62, 283–322.
- Fama, E.F., and K.R. French (1988), Dividend Yields and Expected Stock Returns, *Journal of Financial Economics* 22, 3–25.

- Fernholz, Luisa Turrin (1983), *Von Mises Calculus for Statistical Functionals* (Springer, New York, USA).
- Gagliardini, Patrick, Paolo Porchia, and Fabio Trojani (2009), Ambiguity aversion and the term structure of interest rates, *The Review of Financial Studies* 22, 4157–4188.
- Gilboa, I., and D. Schmeidler (1989), Maxmin Expected Utility with Non-unique Prior, *Journal of Mathematical Economics* 18, 141–153.
- Hansen, Lars Peter (2014), Uncertainty Outside and Inside Economic Models, *Nobel Lecture: Journal of Political Economy* forthcoming.
- Hansen, L.P., and T.J. Sargent (2001), Robust Control and Model Uncertainty, *American Economic Review* 91, 60–66.
- Hansen, L.P., and T.J. Sargent (2008), *Robustness* (Princeton University Press, Princeton).
- Huber, Peter J. (1964), Robust estimation of a location parameter, *The Annals of Mathematical Statistics* 35, 73–101.
- Ju, Nengjiu, and Jianjun Miao (2012), Ambiguity, learning, and asset returns, *Econometrica* 102, 559–591.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2005), A Smooth Model of Decision Making Under Ambiguity, *Econometrica* 73, 1849–1892.
- Klibanoff, P., M. Marinacci, and S. Mukerji (2009), Recursive Smooth Ambiguity Preferences, *Journal of Economic Theory* 144, 930–976.
- Knight, Frank (1921), *Risk, Uncertainty and Profit* (Houghton-Mifflin, Boston, USA).
- Künsch, H. (1984), Infinitesimal robustness for autoregressive processes, *Annals of Statistics* 12, 843–863.
- Leippold, M., and F. Trojani (2008), Learning and Asset Prices under Ambiguous Information, *The Review of Financial Studies* 21, 2565–2597.
- Lettau, M., and S. Van Nieuwerburgh (2008), Reconciling the Return Predictability Evidence, *The Review of Financial Studies* 21, 1607–1652.
- Lettau, Martin, and Sydney C. Ludvigson (2010), Measuring and modeling variation in the risk-return tradeoff, in Ait-Sahalia, Yacine, and Lars P. Hansen, eds., *Handbook of Financial Econometrics*, volume 1, chapter 11 (Elsevier, Amsterdam, The Netherlands).

- Liu, Jun, Jun Pan, and Tan Wang (2005), An equilibrium model for rare-event premia and its implication for option smirks, *The Review of Financial Studies* 18, 131–163.
- Maenhout, P.J. (2004), Robust Portfolio Rules and Asset Pricing, *The Review of Financial Studies* 17, 951–983.
- Maenhout, P.J. (2006), Robust Portfolio Rules and Detection-Error Probabilities for a Mean-Reverting Risk Premium, *Journal of Economic Theory* 128, 136–163.
- Mancini, L., E. Ronchetti, and F. Trojani (2005), Optimal Conditionally Unbiased Bounded-Influence Inference in Dynamic Location and Scale Models, *Journal of the American Statistical Association* 100, 628–641.
- Merton, R. C. (1976), Option pricing when underlying stock returns are discontinuous, *Journal of Economic Dynamics and Control* 3, 125–144.
- Miao, Jianjun, and Neng Wang (2011), Risk, uncertainty, and option exercise, *Journal of Economic Dynamics and Control* 35, 442–461.
- Ortelli, Claudio, and Fabio Trojani (2005), Robust efficient method of moments, *Journal of Econometrics* 127, 69–97.
- Routledge, B.R., and S.E. Zin (2009), Model Uncertainty and Liquidity, *Review of Economic Dynamics* 12, 543–566.
- Trojani, F., and A. Sbuelz (2008), Asset Prices with Locally-Constrained-Entropy Recursive Multiple Priors Utility, *Journal of Economic Dynamics and Control* 32, 3695–3717.
- Trojani, F., and P. Vanini (2002), A Note on Robustness in Merton’s Model of Intertemporal Consumption and Portfolio Choice, *Journal of Economic Dynamics and Control* 28, 1079–1113.
- Trojani, F., and P. Vanini (2004), Robustness and Ambiguity Aversion in General Equilibrium, *Review of Finance* 8, 279–324.
- Ulrich, Maxim (2013), Inflation Ambiguity and the Term Structure of Arbitrage-free US Government Bonds, *Journal of Monetary Economics* 60, 295–309.
- Uppal, R., and T. Wang (2003), Model Misspecification and Underdiversification, *The Journal of Finance* 58, 2465–2486.
- Von Mises, R. (1947), On the Asymptotic Distribution of Differentiable Statistical Functions, *Annals of Mathematical Statistics* 18, 309–348.

Welch, Ivo, and Amit Goyal (2008), A Comprehensive Look at the Empirical Performance of Equity Premium Prediction, *The Review of Financial Studies* 21, 1455–1508.

Xia, Y. (2001), Learning About Predictability: The Effects of Parameter Uncertainty on Dynamic Asset Allocation, *The Journal of Finance* 56, 205–246.




Figure 1. Estimated Sharpe Ratios. This figure shows the Sharpe ratio estimated either using the ratio of sample mean and standard deviation (blue, solid line) or the ratio of median and median absolute deviation (red, dashed line). The black, dotted line shows the true Sharpe ratio under the reference belief. True returns are distributed: $r_t \sim (1 - \kappa_t)N(0.05, 0.2) + \kappa_t N(0.05, 0.5)$. The estimates are based on 2 million simulated returns.






Figure 2. Parameter elasticities. This figure shows parameter elasticities of portfolio weight as a function of the ambiguity aversion parameter ϑ for the various parameters in the solution to the JD portfolio problem in Equation (7). The figure shows the percentage change in the optimal portfolio weight, given a percentage change in one of the parameters of the JD model. $\vartheta = 0$ corresponds to the classical portfolio problem and $\vartheta > 0$ reflects the investor's aversion to ambiguity; $\gamma = 3$, $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.


Panel A. Unconditional return densities



Panel B. Sample distorted process



Panel C. Sample time-varying component



Panel D. Sample clean process




Figure 3. Densities implied by the distorted JD and the clean JD. Panel (a) depicts densities implied by the distorted JD of Equation (11) and the corresponding JD model. The Figure shows the return densities of the true DGP and the agent's reference model based on 1,000,000 observations generated from the stock process specified in (11); $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$. A sample realization (six years) of the approximate JD model is shown in Panel (b). The components of this process, a time-varying jump size distortion and a clean JD model, are shown in Panels (c) and (d), respectively.







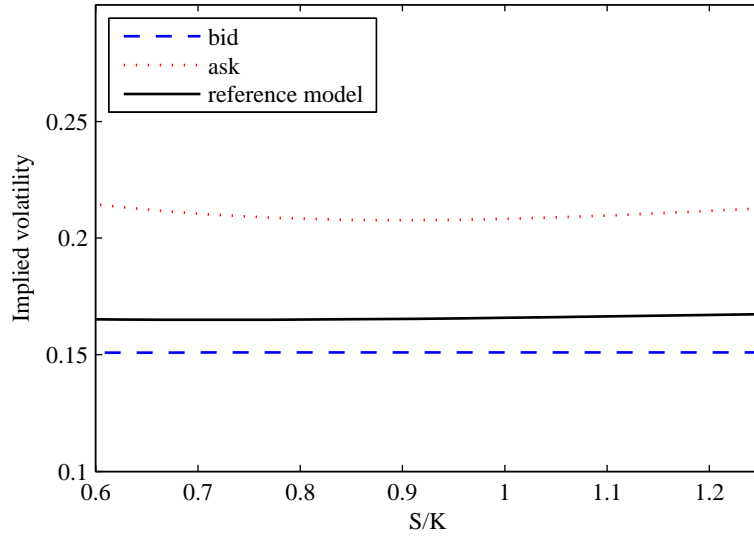
Panel A. Equity premium; distorted data**Panel B.** Equity premium; clean data**Panel C.** Risk premium; distorted data**Panel D.** Risk premium; clean data**Panel E.** Ambiguity premium; distorted data**Panel F.** Ambiguity premium; clean data

Figure 4. Equity, risk, and ambiguity premiums. This figure shows the total equity premium for distorted and clean data according to Equations (13) and (15). Panels (a) and (b) show the total equity premium of an economy with ambiguity aversion. The vertical line indicates the true premium ignoring parameter estimation. The blue, solid lines show estimated premiums based on classical parameter estimates, and the red, dashed lines give the distributions corresponding to robust estimates. Panels (c)–(f) decompose the total premium into risk and ambiguity components: Panels (c) and (d) show the risk premiums corresponding to a jump-diffusion economy and Panels (e) and (f) report the difference between total premiums and the risk premiums, which corresponds to the ambiguity premiums. The three plots on the left are obtained from distorted data in accordance with Equation (6) and equilibrium option prices computed from the partial differential equation in (16). The true preference parameters are $\gamma = 2$ and $\vartheta = 1$; the jump-diffusion is characterized by $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.

Panel A. Bid and ask prices as a function of moneyness



Panel B. Bid and ask prices as a function of ambiguity




Figure 5. Bid and ask option prices under ambiguity. This figure shows bid and ask option prices under ambiguity (quoted as Black-Scholes implied volatilities). Panel (a) shows bid and ask prices as a function of moneyness (S/K). The entropy bound η in Equations (17) and (18) equals 0.000612 corresponding to the entropy between the reference model and a realistic alternative model with jump standard deviation of $2\sigma_Y = 0.08$. The time to maturity of the option is 1 year, $S(0) = 100$, and $r = 3\%$. In Panel (b) bid and ask quotes for different levels of ambiguity are depicted. The strike price equals $K = 110$. For comparison, the solid line in both panels correspond to the price of the option without ambiguity.

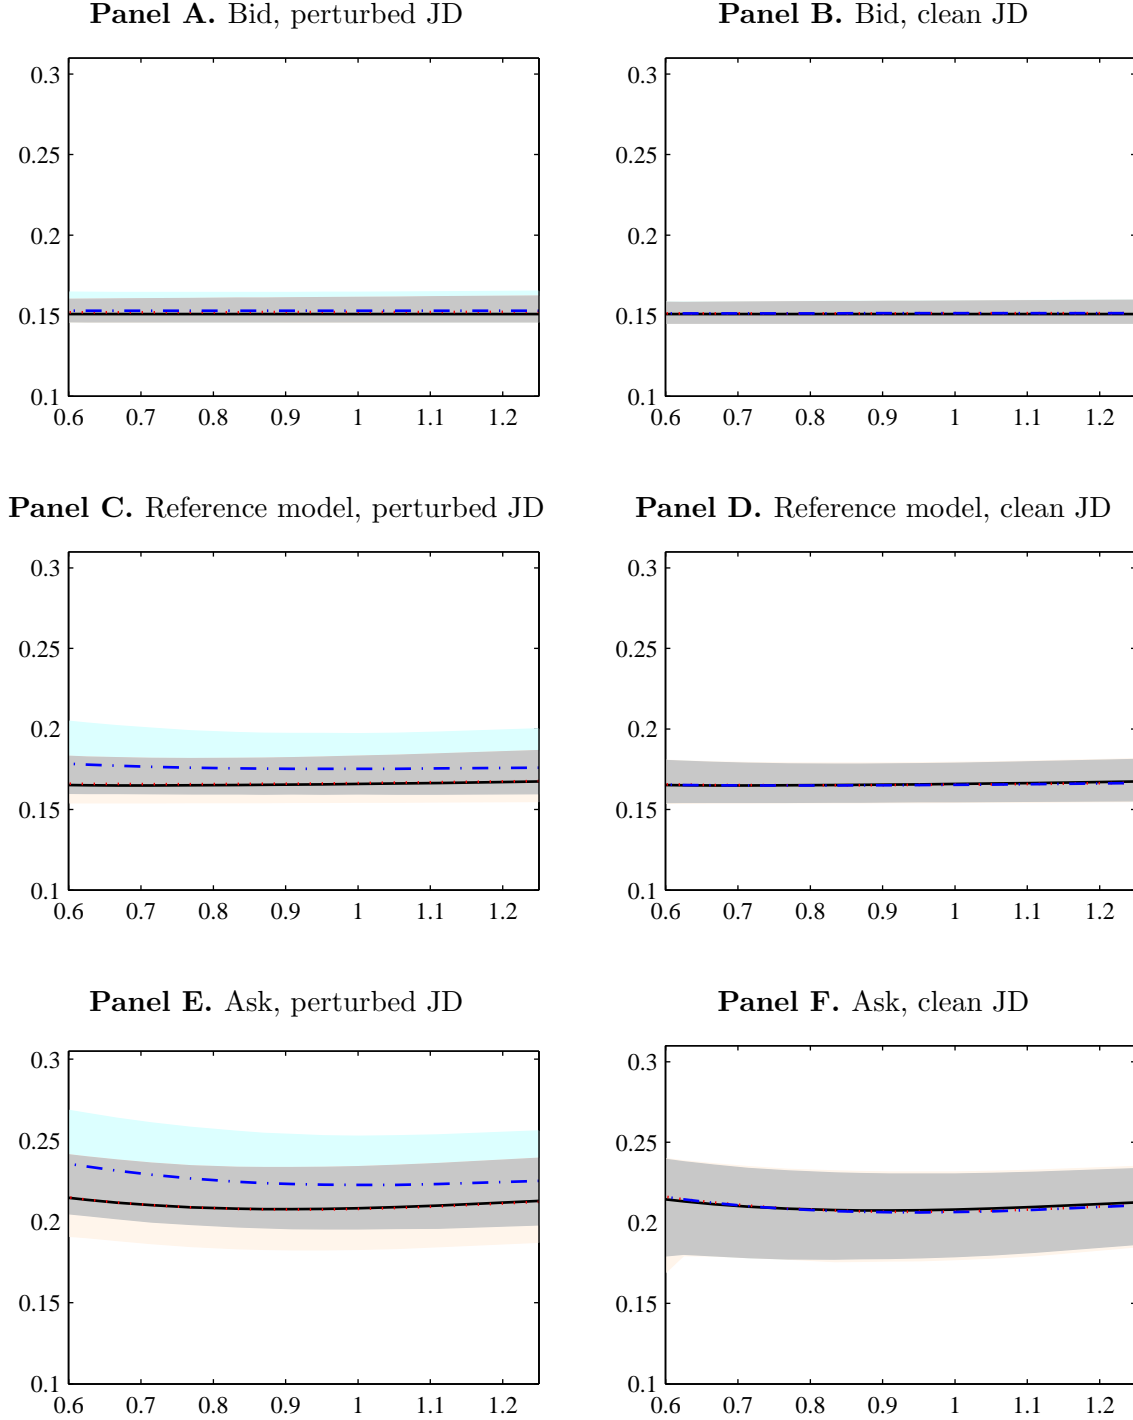


Figure 6. Estimated bid and ask option prices. This figure shows bid and ask prices (quoted as Black-Scholes implied volatilities) as a function of moneyness (S/K) for estimated parameters. In Panels (a), (c), and (e), estimates are obtained based on data from a perturbed jump diffusion (Equation (11)). Panels (b), (d), and (f) show bid and ask quotes for estimates from a clean jump diffusion. The red, dotted and the blue, dashed lines denote the robust and classical median estimate, respectively. The shaded areas mark the 5% and 95% quantiles. For comparison, the black, solid line in all panels corresponds to the price of the option based on true parameters. The time to maturity of the option is 1 year, $S(0) = 100$, and $r = 3\%$. Without loss of generality, we set the entropy bound $\eta = 0.000612$, which corresponds to the entropy between the reference model and a realistic alternative model with jump standard deviation of $2\sigma_Y = 0.08$.




Figure 7. Evolution of predictability parameter point estimates. This figure shows d_3 point estimates together with 68% confidence intervals (\pm one standard error) over time. In every month t , model (20) is estimated based on sample data from January 1929 up to the current month. The red, dotted line corresponds to robust point estimates ($c_{Huber} = 6$); the union of the top and middle shaded area depicts the robust confidence intervals. The solid, blue line represents the classical point estimates; the classical confidence interval is given by the union of the bottom and middle shaded areas.

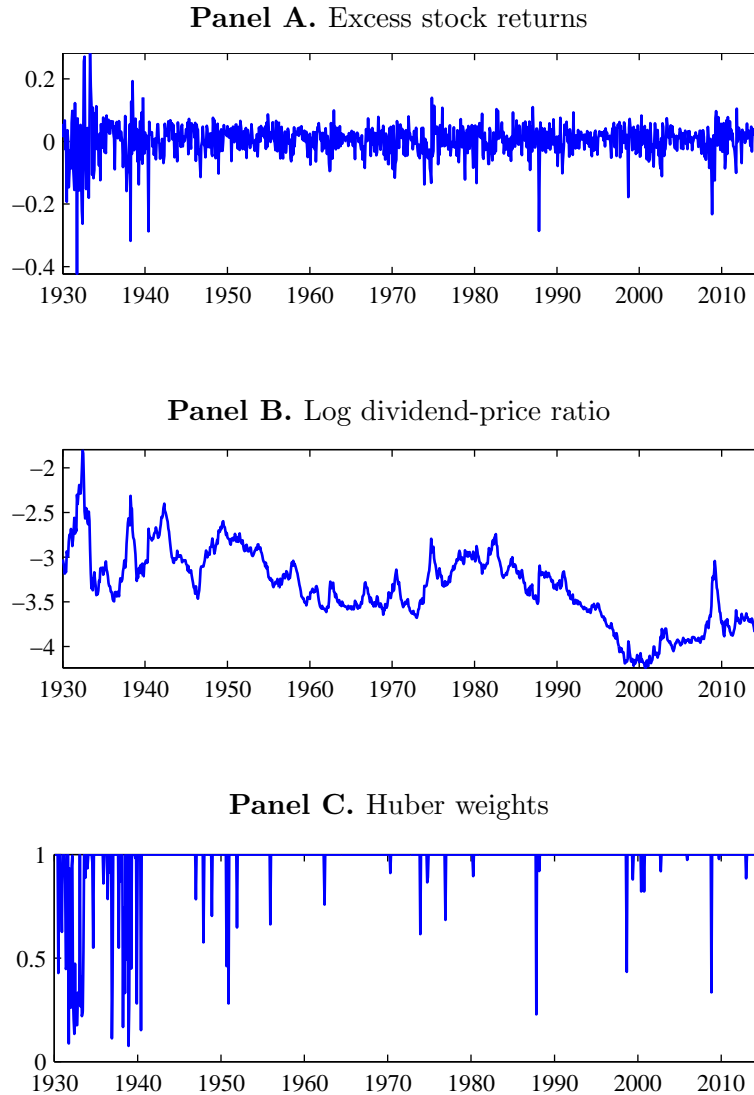


Figure 8. Excess stock returns, dividend-price ratios, and corresponding Huber weights. Panels (a) and (b) show excess stock returns and the dividend-price ratio of the CRSP value-weighted market portfolio for the period 1929–2013, respectively. The Huber weights corresponding to the observations are depicted in Panel (c). The econometric constant c_{Huber} used in the estimation equals 6.

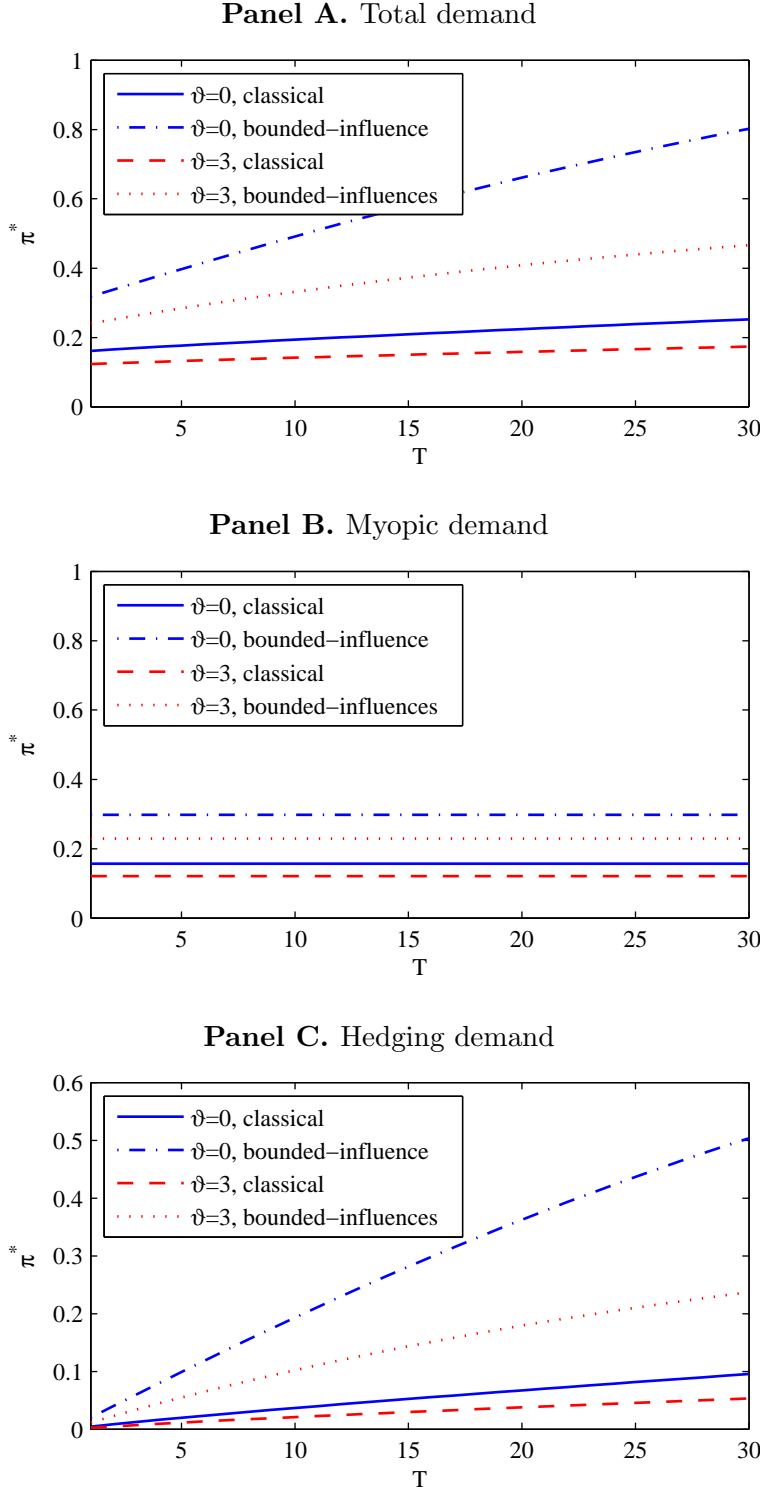


Figure 9. Portfolio allocation as a function of the investment horizon. Panel (a) shows the optimal allocation into equities, π^* , of classical and uncertainty averse investors as a function of the investment horizon for both classical and robust parameter estimates. The myopic and hedging demand components of total demand are shown in Panels (b) and (c), respectively. The continuously compounded annual risk-free interest rate, r_f , is equal to 3.59%, the coefficient of risk aversion $\gamma = 3$, and the initial log dividend-price ratio equals $\log(3.66\%)$.

Table 1
Related literature

This table relates our paper to the existing literature on robust modeling and robust econometrics. The literature is grouped into four quadrants, depending on whether inside uncertainty and/or outside uncertainty is taken into account.

| | | Agent inside the model has subjective aversion to ambiguity | |
|--|-----|--|---|
| | | no | yes |
| Effect of ambiguous data features is bounded in estimation | no | classical literature A classical modeling and estimation; no uncertainty | robust modeling literature B robust modeling using classical estimation or calibration |
| | yes | robust econometrics literature C classical modeling using robust estimation | our paper D robust modeling and estimation |

Table 2
Detection error probabilities

This table shows the detection error probabilities (DEP) comparing the JD reference model with the assumed true data generating process for each piecewise constant jump size $\mu_{Y,t}$. The other parameters are $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$. The number of simulations to estimate the DEP equals 20,000.

| $\mu_{Y,t}$ | -0.04 | -0.01 | 0.01 | 0.12 | -0.02 | -0.1 |
|-------------|--------|-------|--------|-------|--------|-------|
| DEP | 26.22% | n/a | 44.20% | 3.46% | 44.70% | 8.15% |

Table 3
Expected and realized utility in the jump-diffusion case

This table shows the simulation based average *ex-post* realized utilities in wealth equivalents of an expected utility maximizer ($\vartheta = 0$) and her robust counterpart. Panel (a) depicts expected and realized utilities when the data contain small, time-varying distortions to the jump component in- and out-of-sample (Equation (12)) for the three cases of knowing the parameters, estimating them with maximum likelihood, and using robust estimates. Panel (b) shows the same quantities when the data follow a clean JD process. Portfolio weights for realized utilities are computed based on five years of daily data generated by the stock process specified in Equation (11). The econometric constant used in the robust estimation is chosen to ensure 95% efficiency in case the true data generating process is a clean jump-diffusion ($c_{Huber} = 40$); $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$. The number of simulations equals 10,000.

| | Degree of ambiguity aversion | | | | | |
|---|------------------------------|-------------------|-------------------|-------------------|-------------------|-------------------|
| | $\vartheta = 0$ | $\vartheta = 0.2$ | $\vartheta = 0.6$ | $\vartheta = 1.0$ | $\vartheta = 2.0$ | $\vartheta = 3.0$ |
| Panel (a): In- and out-of-sample distorted JD | | | | | | |
| Ex-ante expected utility | | | | | | |
| True parameters | 1050 | 1050 | 1050 | 1050 | 1050 | 1049 |
| (P)MLE | 1034 | 1039 | 1043 | 1045 | 1047 | 1048 |
| Robust estimates | 1033 | 1037 | 1041 | 1043 | 1045 | 1046 |
| Ex-post realized utility | | | | | | |
| True parameters | 933 | 938 | 942 | 945 | 949 | 953 |
| (P)MLE | 773 | 811 | 843 | 860 | 883 | 895 |
| Robust estimates | 882 | 891 | 901 | 907 | 916 | 924 |
| Panel (b): In- and out-of-sample JD | | | | | | |
| Ex-ante expected utility = ex-post realized utility | | | | | | |
| True parameters | 1050 | 1050 | 1050 | 1050 | 1050 | 1049 |
| (P)MLE | 1034 | 1038 | 1042 | 1043 | 1045 | 1046 |
| Robust estimates | 1034 | 1038 | 1042 | 1043 | 1045 | 1046 |

Table 4
Comparative statics for the equity premium

This table shows the equity premium as a function of risk aversion γ and ambiguity aversion ϑ in a JD economy (Equation ((15))). For comparison, the equity premium implied by a Black-Scholes (BS) world is also given as a function of γ (Equation ((14))). The stock price process is given by a JD model with by $\mu = 0.08$, $\sigma = 0.15$, $\lambda_Y = 3$, $\mu_Y = -0.01$, $\sigma_Y = 0.04$.

| | | BS | JD (with ambiguity aversion) | | | | | |
|-------------|---|--------|------------------------------|--------|--------|--------|--------|--------|
| ϑ | | | 0 | 1 | 2 | 3 | 4 | 5 |
| γ | 0 | 0.0000 | 0.0275 | 0.0316 | 0.0359 | 0.0402 | 0.0446 | 0.0490 |
| | 2 | 0.0450 | 0.0780 | 0.0826 | 0.0874 | 0.0924 | 0.0975 | 0.1028 |
| | 3 | 0.0675 | 0.1108 | 0.1559 | 0.2048 | 0.2579 | 0.3156 | 0.3780 |
| | 4 | 0.0900 | 0.1489 | 0.3098 | 0.5059 | 0.7461 | 1.0368 | 1.3786 |
| | 5 | 0.1125 | 0.1926 | 0.6201 | 1.2907 | 2.3698 | 3.9401 | 5.7840 |

Table 5
Accuracy of estimated ambiguity and risk aversion parameters

This table shows the accuracy of estimated ambiguity and risk aversion parameters based on equilibrium option prices for maximum likelihood and bounded-influence parameter estimates. Estimates are obtained from 5 years of daily data following a JD process. Columns 1 and 2 show the root mean squared error (RMSE) and mean absolute deviation (MAD) of estimated ambiguity and risk aversion parameters compared to the true values underlying equilibrium option prices (without loss of generality we chose $\vartheta = 1$ and $\gamma = 2$). Columns 3 and 4 show the same quantities when the data contain small, time-varying distortions to the jump component (Equation (11)).

| | | JD | | Distorted JD | |
|------|-------------|--------|------------------|--------------|------------------|
| | | (P)MLE | Robust estimates | (P)MLE | Robust estimates |
| RMSE | γ | 0.730 | 0.751 | 1.263 | 0.828 |
| | ϑ | 0.583 | 0.596 | 0.721 | 0.657 |
| MAD | γ | 0.599 | 0.618 | 1.117 | 0.675 |
| | ϑ | 0.483 | 0.496 | 0.676 | 0.530 |

Table 6
Bid-ask spreads due to ambiguity

Panel (a) shows median bid-ask spreads for estimates obtained based on a perturbed jump diffusion (Equation (11)). Bid and ask prices are computed from Equations (17) and (18), respectively. Relative bid-ask spreads in Panel (b) are computed as ask minus bid divided by the price of the option without ambiguity. In each panel we consider five parameter cases, namely true parameters (line 1) as well as classical and robust parameter estimates obtained from a clean jump diffusion (lines 4 and 5) as well as a perturbed (Equation (11)) jump diffusion (lines 2 and 3). The time to maturity of the option is one year, $S(0) = 100$, and $r = 3\%$. The entropy bound is $\eta = 0.00061$, which corresponds to the entropy between the reference model and a realistic alternative model with jump standard deviation of $2\sigma_Y = 0.08$. The number of simulations equals 10,000.

| K | 80 | 90 | 95 | 100 | 105 | 110 | 115 | 120 | 130 | 140 | 150 | 160 | 170 |
|--------------------------------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|-------|--------|
| Median bid-ask spread | | | | | | | | | | | | | |
| True parameters | 0.812 | 1.638 | 1.988 | 2.208 | 2.266 | 2.173 | 1.966 | 1.691 | 1.110 | 0.647 | 0.349 | 0.180 | 0.091 |
| Distorted, classical | 1.019 | 2.005 | 2.425 | 2.701 | 2.799 | 2.719 | 2.507 | 2.207 | 1.535 | 0.963 | 0.570 | 0.325 | 0.183 |
| Distorted, robust | 0.789 | 1.588 | 1.926 | 2.141 | 2.200 | 2.113 | 1.915 | 1.653 | 1.090 | 0.637 | 0.346 | 0.180 | 0.092 |
| Clean, classical | 0.774 | 1.577 | 1.922 | 2.142 | 2.206 | 2.123 | 1.929 | 1.666 | 1.105 | 0.651 | 0.356 | 0.187 | 0.096 |
| Clean, robust | 0.780 | 1.590 | 1.937 | 2.157 | 2.224 | 2.141 | 1.942 | 1.679 | 1.111 | 0.655 | 0.358 | 0.188 | 0.097 |
| Median relative bid-ask spread | | | | | | | | | | | | | |
| True parameters | 0.036 | 0.113 | 0.181 | 0.273 | 0.394 | 0.548 | 0.743 | 0.987 | 1.667 | 2.737 | 4.477 | 7.401 | 12.480 |
| Distorted, classical | 0.045 | 0.136 | 0.214 | 0.319 | 0.455 | 0.626 | 0.836 | 1.088 | 1.753 | 2.729 | 4.203 | 6.483 | 10.086 |
| Distorted, robust | 0.035 | 0.110 | 0.175 | 0.264 | 0.381 | 0.530 | 0.718 | 0.954 | 1.613 | 2.639 | 4.276 | 6.956 | 11.467 |
| Clean, classical | 0.034 | 0.109 | 0.175 | 0.265 | 0.384 | 0.537 | 0.730 | 0.973 | 1.653 | 2.718 | 4.434 | 7.269 | 12.116 |
| Clean, robust | 0.034 | 0.110 | 0.176 | 0.267 | 0.386 | 0.540 | 0.734 | 0.977 | 1.654 | 2.719 | 4.433 | 7.261 | 12.103 |

Table 7
Parameter estimates

Panel (a) shows parameter estimates for the discrete time VAR model in Equation (20). The econometric constant c_{Huber} used in the robust estimation equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model. Asymptotic standard errors are shown in parenthesis. Panel (b) depicts the corresponding continuous time parameter estimates and standard errors (Equations (21) and (23)).

| Panel (a): Discrete time (1929–2013) | | | | | | | |
|--|----------|----------|----------|------------|------------|------------|------------------|
| | d_1 | d_2 | d_3 | d_4 | $Var(v_1)$ | $Var(v_2)$ | $Corr(v_1, v_2)$ |
| Classical | 0.0117 | -0.0356 | 0.0025 | 0.9895 | 0.0029 | 0.0037 | -0.9492 |
| | (0.0102) | (0.0119) | (0.0033) | (0.0038) | (0.0001) | (0.0001) | (0.021) |
| Bounded-influence | 0.0245 | -0.0338 | 0.0058 | 0.9905 | 0.0020 | 0.0026 | -0.9547 |
| | (0.0126) | (0.0142) | (0.0038) | (0.0043) | (0.0001) | (0.0001) | (0.0403) |
| Panel (b): Continuous time (1929–2013) | | | | | | | |
| | σ | α | ν | κ_X | \bar{x} | b | ρ |
| Classical | 0.1875 | 0.1551 | 0.0300 | 0.1261 | -3.4056 | 0.2117 | -0.9493 |
| | (0.0018) | (0.1232) | (0.0396) | (0.0003) | (32.1229) | (0.0024) | (0.0209) |
| Bounded-influence | 0.1564 | 0.3050 | 0.0697 | 0.1143 | -3.5670 | 0.1767 | -0.9550 |
| | (0.0031) | (0.1531) | (0.0462) | (0.0004) | (33.5156) | (0.0035) | (0.04) |

Table 8
Portfolio weights and in-sample utility

Portfolio allocation as function of risk aversion (γ) and ambiguity aversion (ϑ). Panels (a) and (b) show the time series mean of the optimal allocation into equities for both classical and robust parameter estimates, respectively. In each month t , π^* is computed with an investment horizon of one year using the parameter estimates from Table 7 and the current dividend-price ratio. Panels (c) and (d) show the average realized wealth from these investment policies. Results are based on monthly CRSP data from 1929 to 2013.

| Panel (a): Portfolio weights with classical estimates | | | | | | | |
|---|----|--------|--------|--------|--------|--------|--------|
| ϑ | 0 | 1 | 2 | 3 | 4 | 5 | |
| γ | 2 | 78.58% | 52.83% | 39.52% | 31.57% | 26.28% | 22.51% |
| | 5 | 32.15% | 26.68% | 22.81% | 19.91% | 17.67% | 15.88% |
| | 10 | 16.13% | 14.62% | 13.38% | 12.33% | 11.43% | 10.65% |
| | 15 | 10.76% | 10.07% | 9.46% | 8.93% | 8.45% | 8.01% |
| | 20 | 8.08% | 7.68% | 7.32% | 7.00% | 6.70% | 6.42% |
| Panel (b): Portfolio weights with bounded-influence estimates | | | | | | | |
| ϑ | 0 | 1 | 2 | 3 | 4 | 5 | |
| γ | 2 | 93.30% | 84.77% | 72.48% | 59.97% | 50.19% | 43.00% |
| | 5 | 62.44% | 52.00% | 44.35% | 38.62% | 34.19% | 30.67% |
| | 10 | 31.82% | 28.75% | 26.22% | 24.10% | 22.30% | 20.75% |
| | 15 | 21.27% | 19.85% | 18.61% | 17.52% | 16.54% | 15.67% |
| | 20 | 15.97% | 15.16% | 14.42% | 13.76% | 13.15% | 12.60% |
| Panel (c): Realized wealth with classical estimates | | | | | | | |
| ϑ | 0 | 1 | 2 | 3 | 4 | 5 | |
| γ | 2 | 109.15 | 107.40 | 106.43 | 105.85 | 105.47 | 105.19 |
| | 5 | 105.89 | 105.50 | 105.21 | 105.00 | 104.84 | 104.71 |
| | 10 | 104.72 | 104.62 | 104.52 | 104.45 | 104.38 | 104.33 |
| | 15 | 104.33 | 104.28 | 104.24 | 104.20 | 104.16 | 104.13 |
| | 20 | 104.14 | 104.11 | 104.08 | 104.06 | 104.04 | 104.02 |
| Panel (d): Realized wealth with bounded-influence estimates | | | | | | | |
| ϑ | 0 | 1 | 2 | 3 | 4 | 5 | |
| γ | 2 | 109.89 | 109.60 | 109.05 | 108.23 | 107.51 | 106.97 |
| | 5 | 108.39 | 107.65 | 107.07 | 106.63 | 106.27 | 105.99 |
| | 10 | 106.09 | 105.84 | 105.64 | 105.47 | 105.33 | 105.20 |
| | 15 | 105.24 | 105.13 | 105.03 | 104.94 | 104.87 | 104.80 |
| | 20 | 104.82 | 104.76 | 104.70 | 104.65 | 104.60 | 104.55 |

Table 9
Sensitivity of portfolio weights to return perturbations

This table shows the effect of perturbing the 5 most influential observations by up to $\pm 30\%$ on the optimal portfolio weights and the corresponding realized wealths. Weight sensitivities are reported for the standard CRRA agent ($\vartheta = 0$) and her robust counterpart ($\vartheta = 3$) for classical and bounded-influence parameter estimates. The continuously compounded annual risk-free interest rate to convert the discrete time parameters to continuous-time, r_f , is equal to 3.42%, the coefficient of risk aversion $\gamma = 5$, and the initial log dividend-price ratio equals $\log(3.66\%)$. The econometric constant c_{Huber} used in the estimation of CRSP market returns between 1929 and 2013 equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model.

| | Portfolio weights | | | Realized wealth | | |
|------------------------------------|-------------------|--------|-------|-----------------|--------|--------|
| | Min | Max | Range | Min | Max | Range |
| Classical, $\vartheta = 0$ | 30.69% | 36.72% | 6.03% | 104.25 | 105.60 | 1.3545 |
| Bounded-influence, $\vartheta = 0$ | 62.44% | 64.49% | 2.04% | 106.11 | 106.31 | 0.1973 |
| Classical, $\vartheta = 3$ | 19.25% | 22.20% | 2.94% | 103.99 | 104.83 | 0.8495 |
| Bounded-influence, $\vartheta = 3$ | 38.62% | 40.33% | 1.71% | 105.23 | 105.42 | 0.1874 |

Table 10
Realized out-of-sample utility

This table shows the implied out-of-sample utility (wealth equivalent) for the full sample and for different subsamples. Utility is reported for the standard CRRA agent ($\vartheta = 0$) and her robust counterpart ($\vartheta = 3$) for classical and bounded-influence parameter estimates. The coefficient of risk aversion γ equals 5. The portfolio weight π^* is then either computed based on a random walk model (no predictability) or based on the optimal policy laid out in Equation (25) (in case of predictability), both using an investment horizon of $T = 1$ year and parameter estimates incorporating information up to time t and the current dividend-price ratio. The data extends from 1929-2013; the initial sample size for the estimation is 25 years. The econometric constant c_{Huber} used in the estimation equals 6 to achieve 95% efficiency if the data actually have been generated by the reference model.

| ϑ | Full sample | | 1954-1972 | | 1973-1990 | | 1990-2013 | |
|-----------------------------------|-------------|--------|-----------|--------|-----------|--------|-----------|--------|
| | 0 | 3 | 0 | 3 | 0 | 3 | 0 | 3 |
| Predictability, Classical | 107.16 | 106.47 | 106.90 | 105.68 | 111.02 | 110.59 | 104.31 | 103.84 |
| Predictability, Bounded-influence | 108.01 | 107.11 | 108.68 | 106.87 | 111.11 | 110.64 | 105.06 | 104.52 |
| RW, Classical | 105.94 | 105.52 | 106.74 | 105.58 | 107.47 | 107.68 | 104.15 | 103.77 |
| RW, Bounded-influence | 107.28 | 106.58 | 109.39 | 107.78 | 106.82 | 107.27 | 106.05 | 105.14 |