Uncertainty premia for small and large risks

Martin Puhl[a] Pavel Savor[b, *] Mungo Wilson[c]

[a] Oesterreichische Nationalbank and Vienna University of Technology, Austria

[b] Driehaus College of Business, DePaul University, USA

[c] Said Business School, Oxford University, and Oxford-Man Institute, United Kingdom

June 2024[†]

Abstract

We develop a model showing that the effect of smooth ambiguity aversion on large risks, those that are independent of the holding period, is of first-order importance, in contrast to risks that are proportional to the holding period. To test this hypothesis, we construct an ex-ante measure of the price of uncertainty based on changes in the option-implied concavity of preferences. As predicted by our model, we find that such concavity increases ahead of scheduled macroeconomic announcements, which represent large risks. We also provide an estimate of the coefficient of relative ambiguity aversion and show how uncertainty varies across different announcements. Our results suggest that the macroeconomic announcement premium arises at least partly because of an increase in the price of uncertainty. One implication is that a fundamental benefit of securities markets is that they break large risks into small ones by allowing frequent trading, thereby reducing discount rates.

JEL Classification: G12

Keywords: Uncertainty; price of risk; price of uncertainty; risk aversion; ambiguity aversion; options

*Corresponding author. psavor@depaul.edu.

†We thank Hengjie Ai, and seminar participants at the 2017 Western Finance Association Meeting, American University, City University of New York (Baruch), Rutgers University, and University of South Carolina (Moore) for their valuable comments. The expressed opinions are strictly those of the authors and do not necessarily reflect the official viewpoint of the OeNB or the Eurosystem.

1. Introduction

A number of studies find that the stock market enjoys high average excess returns on days of macroeconomic announcements (Savor and Wilson 2013, 2014; Lucca and Moench 2015; Cieslak, Morse, and Vissing-Jorgensen 2018; and Ai and Bansal 2018). Sharpe ratios are also much higher, often ten to twenty times greater than in normal periods. One interpretation for this result is that investors are averse to the risk, or more generally to the risk and ambiguity if investors are both risk- and ambiguity-averse, associated with announcements and consequently demand a higher compensation to bear this risk. However, existing papers do not definitively rule out the hypothesis that the origin of high announcement returns is either an increase in the quantity of risk or a series of news that were, on average, positive for investors.

Although Savor and Wilson (2013) show that the realized and option-implied variance of market returns around announcements does not increase commensurately to realized returns, and that other moments of the return distribution are more favorable to investors, there is no canonical way to measure the total quantity of risk, given that stock market returns are unconditionally neither normally nor lognormally distributed (see, e.g., Backus, Chernow, and Martin 2013). Similarly, since actual investor expectations are unobservable and, in any case, it is not necessarily clear what constitutes good news for investors, it is always arguable that investors on average received favorable surprises over the sample period under observation. What is needed to conclusively address this question is an ex-ante measure of the price of risk and uncertainty, the reward investors demand for risk and ambiguity associated with an asset, which we supply in this paper.

We develop the measure by first showing that for a representative investor with smooth ambiguity aversion preferences (Klibanoff, Marinacci, and Mukerji 2005, henceforth KMM) the concavity of the utility function increases ahead of scheduled announcements.[1] In the absence of scheduled announcements, ambiguity-averse agents with KMM preferences will not demand an additional premium for holding risky assets, provided they can rebalance their portfolios sufficiently frequently (Skiadas 2013). However, these same agents will demand an ambiguity premium in a period con-

[1] Because the concept of 'concavity' in the setting of ambiguity aversion is somewhat vague, we introduce and define a well-motivated concavity index in Section 3. Our concavity index is the elasticity of the marginal utility of wealth with respect to wealth, which is equal to risk aversion plus the negative of the wealth-weighted derivative with respect to wealth of the ratio of the risk- and ambiguity-adjusted state price density to the actual physical density.

taining a scheduled announcement. This will be true even if the amount of uncertainty about the distribution (i.e., ambiguity) is the same across announcement and non-announcement periods. The intuition underpinning this finding is that trading more frequently reduces the holding-period risk for non-announcement periods but not for announcement periods, and therefore the uncertainty around expected returns remains first-order important only for the latter. Based on this insight, our principal hypothesis is that investors' concavity of preferences is higher before scheduled announcements due to a higher price of uncertainty during these periods.

We confirm the prediction by using S&P 500 index options to infer our concavity index for a representative agent. Our measure relies on the idea in Ait-Sahalia and Lo (2000) that option-implied state-price densities can be used to construct a local (state-specific) measure of the concavity of preferences for the representative investor, which under the null of no ambiguity aversion equals risk aversion. Estimating one global measure of concavity, rather than many local ones, requires making further assumptions about the form of expected utility (under the null of zero ambiguity aversion). Assuming for example CRRA utility, the concavity index (henceforward, 'concavity') increases from 9.3 to between 11.3 and 15.1 (depending on the type of announcement) ahead of scheduled announcements. The difference between announcement and non-announcement days is economically meaningful and statistically significant, and holds for all announcements (employment, inflation, and Federal Reserve decisions) together as well as for each separately.

We also plot in Figure 1 the resulting implied local (non-parametric) concavities for each degree of option moneyness, as in the original study of Ait-Sahalia and Lo (2000), and do so separately for days immediately before scheduled macroeconomic announcements (blue line) and for 'regular' non-announcement days (black line). Each dot in the figure is a local estimate of concavity, with 95% confidence intervals shown as dotted lines of the same color. The main result is that there is a clear increase in implied concavity ahead of announcements over nearly the entire range of states, which is consistent with smooth ambiguity aversion. It is not consistent with expected utility, which implies constant risk aversion and no ambiguity aversion, and thus concavity cannot change ahead of announcements.

[FIGURE 1 ABOUT HERE.]

We next use the closed-form solution in Gollier (2011) of a model for KMM preferences with constant absolute risk aversion and constant relative ambiguity aversion to show how the announcement premium can be decomposed into ambiguity (the quantity of uncertainty about the true expected return) and ambiguity aversion, and risk and risk aversion.[2] We combine this insight with a simple estimate of the quantity of uncertainty to compute the coefficient of relative ambiguity aversion and explore its implications for the magnitude of the equity risk premium. Similar to the results based on implied concavity, our findings suggest that smooth ambiguity aversion represents an important driver of premia in equity markets. Under the assumption that relative ambiguity aversion does not change, we explore how uncertainty varies across different types of announcements, finding it is the highest for the Federal Reserve's monetary policy decisions and lowest for inflation announcements.

In a recent paper, Ai and Bansal (2018) provide a general characterization of the kinds of preferences that generate a positive announcement premium through an increase in the price of risk (holding constant the quantity and type of risk). This class of preferences includes, but is not limited to, smooth ambiguity aversion. It is possible that the entire general class of preferences identified by Ai and Bansal (2018) implies the same increase in option-implied concavity. For instance, if announcements provide information not only about current payoffs but also future ones, the preferences of investors with non-time-separable utility (e.g., Epstein-Zin preferences with elasticity of intertemporal substitution greater than one) will likely appear to become more concave ahead of announcements. Models combining learning and ambiguity aversion could also potentially generate increased concavity of preferences.[3] We leave the formal consideration of both these issues to future research. Finally, announcements could constitute a different type of risk (or ambiguity) that commands a higher risk (or ambiguity) price (as was assumed in the model in Savor and Wilson 2013).[4]

This paper has two principal objectives. First, we seek to provide deep intuition for why the price of uncertainty should increase ahead of scheduled announcements when agents have smooth

[2] We provide a formal model for this decomposition in Appendix B, which is based on our argument that investors will demand a lower ambiguity premium when they are able to trade more frequently.

[3] See Epstein and Schneider (2007), Epstein and Schneider (2008), Epstein and Ji (2013), Illeditsch (2011), and Illeditsch, Ganguli, and Condie (2021) for examples of such models.

[4] The results in Savor and Wilson (2014), who find that the CAPM holds only on announcement days, are however difficult to reconcile with this last possibility. See the discussion in that paper for more details.

ambiguity aversion preferences. Our arguments provide a novel perspective on investors' demand for liquidity, which in our model is valuable because more frequent trading converts a large risk into a series of smaller ones, in turn reducing premia investors demand to hold the affected assets. Second, we use our example to motivate an ex-ante measure of the price of uncertainty derived from option prices and find that as expected it indeed increases ahead of announcements. The stock market announcement premium, first documented by Savor and Wilson (2013), therefore arises at least partly due to an increase in the price of uncertainty ahead of macroeconomic announcements. An implication of our argument is that if news currently released in the form of scheduled announcements (such as FOMC monetary policy decisions) were instead released either more frequently or at random times, then the announcement premium might disappear.

The rest of the paper is organized as follows. Section 2 provides a sketch of the intuition and discusses the relevant literature; Section 3 outlines the formal arguments; Section 4 discusses alternative explanations; Section 5 explains our empirical methodology; Section 6 presents our results; Section 7 uses a simple model to estimate ambiguity and ambiguity aversion; and Section 8 concludes.

2. Why do ambiguity-averse investors care particularly about scheduled announcements?

Aversion to uncertainty (Knight 1921), as opposed to risk, has been a central theme of much research in economics and other fields. Ambiguity (which we use as a synonym for uncertainty) aversion can rationalize experimental evidence that is not consistent with expected utility theory. In experiments, agents reveal a strong preference for lotteries with a known distribution of payoffs over those with an unknown distribution (Ellsberg 1961 and Trautmann, Vieider, and Wakker 2011). Moreover, ambiguity aversion has the potential to explain a number of asset-pricing puzzles.[5] Smooth ambiguity aversion models such as KMM attempt to retain the tractability and spirit of expected utility models while allowing for aversion to uncertainty as well as risk. An interesting

[5] For example, Collard, Sheppard, Mukerji, and Tallon (2011) argue that smooth ambiguity aversion, which we describe below, can rationalize the otherwise puzzlingly high risk aversion required to explain the equity premium puzzle. Mankiw, in his discussion of Parker (2001), implicitly argues that some form of ambiguity aversion can rationalize the puzzlingly low rates of equity participation in the USA (Parker 2001, p. 335). It is likely that other puzzles, such as the home bias puzzle of French and Poterba (1991), can also to a certain extent be attributed to aversion to uncertainty in the presence of learning.

feature of such models is that when risks can be reduced by updating and acting more frequently (small risks), the effect of uncertainty is second order (Skiadas 2013). However, updating and acting more frequently does not affect the importance of uncertainty during periods containing a scheduled announcement (large or point-in-time risks). In our paper, we argue that this insight offers an opportunity to estimate the price of uncertainty associated with ambiguity aversion by focusing on differences between small and large risks (a distinction we explain further below).

The intuition behind our argument is straightforward. Consider the following generic problem in economics: an agent tries to choose an action α to maximize the expected value of an objective function $f(\widetilde{Y};\alpha,\widetilde{\theta})$ over some uncertain outcome $\widetilde{Y}$ dependent on α and with a probability distribution that in turn depends upon an uncertain parameter $\widetilde{\theta}$ with expectation $\overline{\theta}$. A second function $g(f)$ assigns weights to different expected values of $f(.)$ given different possible values of $\widetilde{\theta}$. Assuming the problem can be well approximated by a second-order Taylor expansion, the agent's value function simplifies to

$$V = \max_{\alpha} E\left[g(f(\widetilde{Y};\alpha,\widetilde{\theta}))\right] \approx V(\alpha^{*},\overline{\theta}) + \frac{1}{2}V''_{\overline{\theta}}Var[\widetilde{\theta}]. \tag{1}$$

Relative to the standard problem without uncertainty, welfare depends on two additional parameters: $V''_{\overline{\theta}}$, the concavity of the agent's welfare function, normally assumed to be negative if the agent dislikes uncertainty and often termed ambiguity aversion; and $Var[\widetilde{\theta}]$, the quantity of uncertainty about the true value of $\widetilde{\theta}$. An ambiguity premium, defined as $-V''_{\overline{\overline{\theta}}}Var[\widetilde{\theta}]$, is positive (negative) if V is strictly concave (convex) and zero if it is linear.

If an agent can adjust the action α periodically, the risk borne from exposure to $\widetilde{Y}$ in the time interval between each adjustment may be proportional to the length of the time interval. That is, by adjusting more frequently, the agent can reduce the quantity of risk borne between each adjustment. In such cases, we show in Appendix A that the agent's uncertainty $Var[\widetilde{\theta}]$ is proportional to the *square* of the time interval, and therefore, for sufficiently small time intervals, irrelevant to the agent's welfare. The argument above is essentially the same as the one in Skiadas (2013), who was the first to make this point.

However, in other cases, adjusting the action, no matter how frequently, will not materially reduce the risk or uncertainty borne from exposure to $\widetilde{Y}$ in a given time interval. This means that $Var[\widetilde{\theta}]$ will be independent of the time interval, and consequently that uncertainty will have a first-order effect on welfare, even for finite values of V_{θ}''. These cases involve events that the agent knows will occur at a particular time, in other words scheduled events. Such events are different from standard diffusions or Poisson jumps (or their discrete-time analogs). More specifically, we formally show in Appendix A that if preferences are of the second-order expected utility form of KMM, then they become more concave ahead of scheduled announcements relative to both Brownian and Poissonian ('small') risks.[6]

The key idea in this paper is thus that an impending scheduled event will increase the concavity (defined as the elasticity of marginal utility to wealth) of an ambiguity-averse agent's preferences across possible outcomes. Provided we can estimate the degree of concavity of an agent's preferences over payoffs in different states, this provides us with a test for the existence of ambiguity aversion. In the case of scheduled announcement premia, it also constitutes evidence of an increase in the price of uncertainty, as opposed to an increase in the quantity of risk or just good luck in a particular sample.

Until recently, the only formal model of ambiguity aversion in decision theory was that of Gilboa and Schmeidler's (1989) maxmin preferences, in which agents maximize expected utility over the most pessimistic prior distribution of returns in the support of distributions.[7] While Gilboa and Schmeidler's model does not say how this support is formed, most commentators agree that the support does not include all logically possible priors. The consensus in practice seems to be to estimate a confidence interval around the mean and choose the lower bound of this interval. (Such an approach is obviously problematic in small samples or in the presence of rare disaster-type risks.) This results in interesting 'first-order' phenomena such as non-participation in stocks and portfolio inertia, whereby portfolios are insensitive to new information over wide ranges and suddenly much more sensitive than smooth models. While potentially relevant, such phenomena

[6]If preferences are of the entropic variational preferences form of Hansen and Sargent (2011), the same finding holds. Preferences for non-entropic smooth divergence preferences (allowed for in Maccheroni, Marinacci, and Rustichini 2006) are more concave for scheduled announcements than for Brownian risks but not necessarily for Poissonian risks. We leave this special case to future work.

[7]For an application see Maenhout (2004).

tend to make these models relatively difficult to work with. Finally, maxmin preferences do not allow the separation of the degree of ambiguity from the degree of aversion to ambiguity, and so cannot address comparative statics issues such as: how will decisions change as agents' aversion to ambiguity changes, holding the degree of uncertainty fixed?[8]

We choose to work with the KMM formulation for several reasons. First, KMM preferences are a well-specified model of decision-making under uncertainty with obvious appeal for applied work, and which nest both expected utility and Gilboa-Schmeidler maxmin expected utility as limiting special cases. Second, the literature on ambiguity aversion, especially in a dynamic setting but also in the limit of continuous time, has not yet reached a consensus on the appropriate way to handle the troika of aversion to risk, aversion to uncertainty, and preference for the early resolution to uncertainty (Strzalecki 2013) or on how to take appropriate limits as time intervals go to zero (Hansen and Sargent 2011). Rather than wait for the consensus, we chose to demonstrate an important implication of a well-specified example. We believe our finding about option-implied concavity should be a stylized fact that other work should also explain.

Our results suggest an important role for trading in securities markets that has not received much attention so far. Since small risks carry no ambiguity premium under many classes of smooth ambiguity-averse preferences while large risks do, then breaking down large risks into smaller ones will lower equity and other asset premia by eliminating or reducing the ambiguity premium. (Importantly, the total quantity of risk does not change here, just the way it is distributed across time.) This possibility implies that the provision of liquidity, which allows investors to adjust their positions more frequently and thereby convert a large risk into a series of smaller ones, is one of the primary ways in which markets help reduce asset premia, and that such premia should be affected by liquidity conditions. That is, far from being an 'anti-social... fetish' as Keynes (1936) termed it, the demand for liquidity is a fundamental way of managing uncertainty successfully, in the process lowering the cost of capital.

In fact, a perhaps surprising implication of our model is that if the information released in scheduled announcements were instead released continuously or at random times, then the announcement premium, and therefore well over half of the equity premium, would simply disappear

[8] Epstein and Schneider (2010) provide a thought experiment in which this claim does not hold in a particular setting.

or at the very least be much reduced. Such an implication could in principle be verified experimentally, perhaps even in the field, which we regard as an important direction for future research.

A number of studies attempt to identify ambiguity aversion at work in the field of asset pricing. These broadly fall into two categories. First, some studies identify the dispersion in a set of beliefs (most frequently professional analysts' forecasts) as a proxy for uncertainty or ambiguity about the true distribution of returns (see, e.g., Anderson, Ghysels, and Juergens 2009; Drechsler 2013; Shi 2013; Ulrich 2013; and Antoniou, Harris, and Zhang 2014). We refer to these studies as the dispersed predictions set of studies. Second, other studies estimate a structural model of ambiguity aversion on a set of asset pricing moments (see, e.g., Ju and Miao 2012; Jeong, Kim, and Park 2014; Thimme and Völkert 2014; and Gallant, Jahan-Parva, and Liu 2014). We refer to these studies as structural models studies.[9]

Dispersed predictions are neither a necessary nor a sufficient indicator of aggregate uncertainty about the true distribution of asset returns. Ambiguity-averse agents with the same preferences and prior beliefs who receive the same information should reach the same conditional beliefs about the world, even though they are ambiguity-averse and there exists genuine uncertainty about the true distribution. Therefore, dispersion in beliefs is not necessary for ambiguity aversion. Non-ambiguity-averse agents who receive different information can reach different beliefs, and their predictions and actions need not fully reveal these beliefs, as shown in classic studies of noisy rational expectations models such as Grossman and Stiglitz (1980) or Diamond and Verrecchia (1981). In such models, when two expected-utility-maximizing agents receive different private signals, then they cannot necessarily infer the other's private signal from his action, and so dispersion in beliefs can persist even though no one is ambiguity-averse. Thus, heterogeneous forecasts or actions cannot readily proxy for uncertainty, and are not sufficient for ambiguity aversion.

Structural models proceed by fitting an assumed model of decision-making to the data. In the case of ambiguity aversion, they fall into two groups: models which assume Gilboa and Schmeidler's (1989) maxmin preferences and those which assume smooth ambiguity aversion (of which Gilboa-Schmeidler preferences are an extreme special case). The results of models which assume Gilboa-

[9]Brenner and Izhakian (2011, 2012) and Izhakian and Yermack (2014) use intraday stock returns to measure the degree of ambiguity as the variance of the probability of gain or loss.

Schmeidler preferences are clearly sensitive to the minimum (worst-case) distribution allowed in the support of prior distributions. Although the Gilboa-Schmeidler model itself is silent as to how this support is formed, most implementations treat it as the lower bound of some confidence interval (e.g., Gajdos, Hayashi, Tallon, and Vergnaud 2008 and Garlappi, Uppal, and Wang 2007). Indeed, smooth ambiguity aversion models were developed to address this potential difficulty in implementing the maxmin model, which is unable to separate beliefs (priors over distributions) from preferences (aversion to ambiguity).

3. Smooth ambiguity aversion and asset prices

KMM propose a model of preferences satisfying smooth ambiguity aversion. Given a prior distribution indexed by θ over future states of the world s, $\{p_{s|\theta}\}_{s=1}^{S}$ (or a continuous-state analog $f_\theta(s)$) and given prior probabilities $\{q_\theta\}_{\theta=1}^{n}$ (or a continuous analog) that each such distribution $\{p_{s|\theta}\}_{s=1}^{S}$ is the true distribution, agents try to maximize expected utility $E[\widetilde{u}]$ while accounting for uncertainty about the true distribution θ (Knight 1921). If the agent is ambiguity averse, she puts a lower weight on distributions that imply higher expected utility. Given a positive, increasing, concave, and differentiable function $\phi(x)$, the agent tries to maximize

$$\Sigma_{\theta=1}^{n} q_\theta \phi(\Sigma_{s=1}^{S} p_{s|\theta} u(x_s)) = E[\phi(E_{\widetilde{\theta}}[u(\widetilde{W})])], \tag{2}$$

subject to the usual constraints. A linear $\phi(.)$ implies that agents are standard expected utility maximizers, using the compound distribution $\Sigma_{\theta=1}^{n} q_\theta p_{s|\theta} = E[p_{s|\theta}] = p_s$. Broadly speaking, as the relative curvature of $\phi(.)$ increases, the preferences approach maxmin.[10] Note that the probabilities $p_{s|\theta}$ will depend on the time interval under consideration in a dynamic setting, just as they do in a binomial tree, for example.

Before proceeding, let us be clear about the terminology. A risk-averse agent is concerned only about risk. An ambiguity-averse agent is concerned about her uncertainty about the parameters of the risk (those that describe the distribution of a price change, such as the drift or variance.) When we discuss 'uncertainty' about a particular risk, we mean that an agent does not know,

[10] Gollier (2011) provides a very helpful discussion of such preferences and the additional assumptions required under which smooth ambiguity aversion implies that an increase in ambiguity, or ambiguity aversion, increases uncertainty premia.

for instance, the drift of a stochastic process. Such an agent may seek compensation not just for bearing the risk (due to her risk-aversion), but also for bearing uncertainty about the drift (due to her ambiguity-aversion).

A remarkable, and hitherto unremarked, aspect of KMM preferences is that they effectively become more concave (i.e., agents become more averse to uncertainty) ahead of certain events such as important scheduled announcements. That is, the price required by an investor per unit of uncertainty increases ahead of announcements and decreases again after. Therefore, smooth ambiguity aversion has the potential to explain the macroeconomic announcement premium (Savor and Wilson 2013), the finding that average returns around scheduled economic announcements are much higher than at other times, without corresponding increases in the variance or higher moments of returns.

The intuition for this effect is that most of the time, in the absence of scheduled announcements, risks are 'small' in the sense that uncertainty about the parameters associated with those risks is proportional to the square of the time interval under consideration and therefore second-order. By contrast, in a time interval containing a scheduled announcement, risks are 'large' in that uncertainty is independent of the time interval and therefore first-order. Smooth ambiguity-averse investors consequently demand a premium for bearing uncertainty ahead of announcements, but not at other times.

We now explain in detail how this works. We begin by defining the difference between small and large (point-in-time) risks, essentially explaining what is special about scheduled announcements. To keep things as simple as possible, we consider a short time interval, possibly containing a scheduled announcement, in which one of only two states can be realized, as in the basic theory of option replication with binomial trees.

3.1. Small risks and large risks

A risk is a state space $s \in (0, 1)$ and a random variable $B(\widetilde{s})$ which takes on different values in either state with a prior probability. The prior belief $Q = (Q(1), Q(0) = 1 - Q(1))$ is mapped into a drift parameter ρ over a given time interval h, where ρ is uncertain. To keep the analysis obvious and simple, we assume ρ is zero in expectation over all priors, but with a positive variance. In other words, there are some priors in the support of the decision maker in which ρ is positive and others in which ρ is negative. This is without loss of generality, since the true drift can be equal to the unconditionally expected drift over all priors plus ρ. The degree of uncertainty as to the true value of ρ is measured by $Var[\rho]$.

3.1.1. Small risks

A Brownian risk (a large risk of a small change) is then

$$\begin{array}{lll} B(1) = +\sqrt{h} & w.p. & \frac{1}{2}(1+\rho\sqrt{h}) \\ B(0) = -\sqrt{h} & w.p. & \frac{1}{2}(1-\rho\sqrt{h}) \end{array}. \tag{3}$$

Under the prior ρ, the expectation of $B(\widetilde{s})$ written $E_\rho[B(\widetilde{s})]$ (henceforward $E_\rho[\widetilde{B}]$) is ρh and its expected square $E_\rho[(\widetilde{B})^2] = h$. All higher moments are $o(h)$ and therefore negligible over short enough time intervals h.

A Poissonian risk, or compensated jump process, (a small risk of a large change) is given by

$$\begin{array}{lll} B(1) = 1-h & w.p. & (1+\rho)h \\ B(0) = -h & w.p. & (1-(1+\rho)h) \end{array}. \tag{4}$$

Under ρ, $E_\rho[\widetilde{B}] = \rho h$. Higher moments $E_\rho[(\widetilde{B})^{n>1}]$ all equal $(1+\rho)h$.

What both types of small risk share is that the expected value of the change in $\widetilde{B}$ is proportional to the time interval: over small enough time intervals, little change is expected.

3.1.2. Large risks

We now introduce the concept of a large, or point-in-time, risk, whose distribution is first-order independent of the length of the time interval h for which it is borne before adjustment. In a given time interval, but not every time interval, a large risk of the form $B(\widetilde{s})$ is given by

$$\begin{array}{lll} B(1) = \varepsilon & w.p. & \frac{1}{2}(1+\rho) \\ B(0) = -\varepsilon & w.p. & \frac{1}{2}(1-\rho) \end{array}. \tag{5}$$

Under ρ, $E_\rho[\widetilde{B}] = \rho\varepsilon$. Higher moments depend on whether the moment is odd or even. Given an indicator of oddness of n, $1_{n \text{ is } odd}(n)$, $E_\rho[(\widetilde{B})^n] = \varepsilon^n(\rho 1_{n \text{ is } odd}(n) + (1 - 1_{n \text{ is } odd}(n)))$.

It is of course possible to pick a time interval h^* which is equal to ε. Our point is that changing the time interval, so long as the large risk is contained inside it, has no first-order effect on the size of the risk $\widetilde{B}$. In fact, the size of the time interval containing this large risk can go to zero, just like a Brownian motion or Poisson jump, and still the risk will remain large. Hence our term 'point-in-time' risk: a point in time is always contained inside a time interval, regardless of how short the interval. Thus, our major insight is not that Skiadas' critique does not apply to fixed time intervals, but instead that it does not apply to point-in-time risks.

Large risks are not degenerate Poisson processes, because they are not independent of calendar time and so do not have independent increments. An example of a large risk is a scheduled announcement with major implications for security prices. More generally, a large risk is just a large probability of a large change. In our view, all such risks correspond to scheduled announcements, even if not regularly repeated, such as a widely expected lockdown announcement during a pandemic.

We show in Appendix A that by taking second-order approximations around expectations over all priors for Brownian and Poissonian risks and replicating the argument in Skiadas:

$$E[\phi(E_\rho[U(\widetilde{B})])] = \phi_0 + o(h), \tag{6}$$

where $\phi_0 = \phi\left(E\left[E_\rho[U(\widetilde{B})]\right]\right)$, i.e. the value of the function $\phi(.)$ evaluated at $E\left[E_\rho[U(\widetilde{B})]\right]$, and ϕ_0' and ϕ_0'' are the values of the same function's first and second derivatives evaluated at the same point.

However, we also show that Skiadas' argument does not apply to large risks, such as announcements scheduled to occur at a given point in time. In that case, we have:

$$\begin{aligned} E\left[\phi\left(E_\rho[U(\widetilde{B})]\right)\right] &= \phi_0 + \frac{1}{2}\phi_0'' Var\left[E_\rho[U(\widetilde{B})]\right] + ... \\ &= \phi_0 + \frac{1}{2}\phi_0'' c(\varepsilon)^2 Var[\widetilde{\rho}] + ... \end{aligned} \tag{7}$$

Uncertainty about the drift, $Var[\widetilde{\rho}]$, can only lower welfare for a smooth-ambiguity-averse investor of the KMM form for large risks ($\phi(.)$ is concave so ϕ_0'' is negative). Thus, to a first approximation, relative to standard risk-averse but non-ambiguity-averse agents, smooth ambiguity-averse agents should demand an additional premium for bearing large, point-in-time risks but not for small ones. It is also the case that higher-order derivatives of the function $\phi(.)$ should matter for large risks and only for large risks. For example, background risk, which affects risk and uncertainty premia just through third and higher derivatives of the utility function, should only affect ambiguity premia through large risks.

3.2. Testing the theory

We have shown that smooth ambiguity-averse investors demand additional premia ahead of large risks, such as scheduled announcements, and therefore that this class of preferences can explain the macroeconomic announcement premium (of course, there exist many other plausible explanations). We now derive a test of our theory, showing that the implication that the concavity of $\phi(.)$ only matters ahead of large risks such as announcements has testable implications in options markets. This will allow us to reject the null of linear $\phi(.)$ (expected utility) and enable us to propose an estimate of $\phi_0''(.)Var[\widetilde{\rho}]$, the ambiguity premium.

The basic idea is that because index option prices reveal state prices, hence risk-neutral probabilities of states (or in our case risk- and ambiguity-neutral probabilities), and hence a measure of the concavity of preferences for a representative investor, we can recover these preferences given option prices, some assumptions about preferences, and estimates of the physical probabilities of states. This insight, to our knowledge, was first used to estimate risk aversion in Ait-Sahalia and Lo (2000) and later by Bliss and Panigirtzoglou (2004), with a more recent contribution by Ross (2015). Here, we explain how this concavity index is constructed from state prices for two states, and how the concavity index should increase prior to large risks relative to its normal level.

3.2.1. State prices with smooth ambiguity aversion

Assuming initial wealth to be W_0, future wealth is uncertain and equals $W_s = W_0 + B(s)$ with $s \in [0, 1]$ as before. Assuming the existence of the necessary Arrow-Debreu securities, following the analysis in Gollier (2011) for complete markets, the state prices π_s can be shown to equal

$$\pi_s = \widehat{p}_s u'(W_s), \tag{8}$$

where $u'(W_s)$ is the marginal utility of wealth in state s, $\widehat{p}_s$ is the ambiguity-adjusted subjective probability of state s,

$$\widehat{p}_s = E[w_\rho p_{s|\rho}],$$

and

$$w_\rho = \frac{\phi'(E_\rho[u(\widetilde{W})])}{E\left[\phi'(E_\rho[u(\widetilde{W})])\right]}.$$

Here, $p_{s|\rho}$ is the subjective probability of state $s \in [0, 1]$ under the prior belief indexed by ρ and $E[.]$ the expectation over all priors ρ weighted by their subjective probability of being the correct prior (the compound expectation) and by the ambiguity weight w_p, which we now explain.

Because the function $\phi(.)$ is concave, its derivative ϕ' is decreasing, so distributions that imply higher expected utility receive a lower weight w_ρ than more pessimistic distributions. The weight w_ρ measures $\phi'(x)$ relative to its unconditional mean, so will sometimes be greater than one, sometimes less, and always positive. Under the null hypothesis of ambiguity neutrality, ϕ' is constant for all priors, $w_\rho = 1$ for all priors ρ, and then

$$\widehat{p}_s = E[p_{s|\rho}] \equiv p_s, \tag{9}$$

the unconditional subjective probability over the compound prior.

The essential idea is that, in the absence of ambiguity aversion, state prices vary across states either because of different physical probabilities or because of different marginal utilities of wealth in different states. As in Ait-Sahalia and Lo (2000) or Ross (2015), scaling state prices by physical probabilities allows us to recover marginal utilities across states and therefore to recover the degree to which marginal utility varies (locally) across states. Using this approach, we obtain an estimate of concavity or, equivalently, the elasticity of marginal utility to wealth. (For brevity of exposition, we use the term 'concavity,' but please note that in a setting with ambiguity aversion concavity is not the same as risk aversion.) Under smooth ambiguity aversion, however, the estimate of concavity should increase ahead of scheduled announcements or other large risks, as the second derivative of the ambiguity aversion function $\phi(.)$ becomes first-order to security prices. By estimating this across states using option prices, we can test our theory.

In this paper, our test is simply whether or not estimated option-implied concavity increases during announcement periods. We present evidence that it does. In Section 7, we provide a closed-form example of smooth ambiguity preferences under which we can decompose this increase into risk and risk aversion versus ambiguity and ambiguity aversion to obtain estimates of magnitudes.

Our concavity index is the elasticity of marginal utility of wealth with respect to wealth:

$$\begin{aligned} CI(s) &\equiv -\frac{W_s \frac{d}{dW}\left(\frac{\widehat{p}_s}{p_s} u'(W_s)\right)}{\left(\frac{\widehat{p}_\omega}{p_\omega} u'(W_\omega)\right)} = -\frac{W_s u''(W_s)}{u'(W_s)} - W_s \frac{\frac{d}{dW}\left(\frac{\widehat{p}_s}{p_s}\right)}{\frac{\widehat{p}_s}{p_s}} \\ &= \gamma(s) - W_s \frac{d}{dW_s} \ln\left(\frac{\widehat{p}_s}{p_s}\right). \end{aligned} \tag{10}$$

We show in Appendix A that $\widehat{p}_s = p_s$ in all states for small (i.e., Brownian and Poissonian risks) but not for large, point-in-time risks. Therefore, concavity is equal to risk aversion in the presence of small risks even for smooth ambiguity-averse investors, but not in the presence of large risks. In the latter case, using our two-state example at $\omega = 1$:

$$CI(1) = \gamma(1) - \frac{(W_0 + \varepsilon)}{\varepsilon} \frac{\frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho]}{1 + \frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho]} \tag{11}$$

and at $\omega = 0$:

$$CI(0) = \gamma(0) - \frac{(W_0 - \varepsilon)}{\varepsilon} \frac{\frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho]}{1 - \frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho]}. \tag{12}$$

Since $\phi_0'' < 0$ for ambiguity-averse agents, $CI(s) > \gamma(s)$. (Because ambiguity-distorted probabilities must be positive, the term $\frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho]$ is smaller than one in magnitude.) Large risks have a first-order effect on the representative investor's concavity and consequently on the price of uncertainty in equilibrium. Provided we can measure such concavity ex ante, we can evaluate whether the announcement premium is, at least in part, due to an increase in the price of uncertainty.

In summary, for investors with smooth ambiguity aversion the elasticity of marginal utility of wealth with respect to wealth, or concavity, increases ahead of point-in-time risks such as scheduled announcements. We use option prices in a later section to empirically evaluate this effect.

4. Alternative hypotheses

Agents with smooth ambiguity averse preferences confronting large risks is obviously not the only explanation for why concavity of investors' preferences increases ahead of scheduled announcements. In this section, we discuss alternative explanations in the context of a simple thought experiment which aims to outline how our model works.[11]

Consider a situation in which an investor values two different payoff risks, which we can think of as two different securities. For both securities, A and B, a payoff is realized over a period $[0, 1]$. This period is divided into N subintervals of length Δt each, where $\Delta t = 1/N$. As N becomes large, Δt becomes small. The investor can rebalance her holdings of each security and the riskless asset once in each subinterval. The payoffs on each security are independent from each other and over time, and over the whole period are distributed $N(\theta, \sigma^2)$, with uncertainty about the drift θ which is distributed $N(0, \sigma_\theta^2)$.

For security A, the payoff is realized over the whole period as either a continuous Brownian motion or as a series of Poisson jumps (or some combination of both). As shown in Appendix A, when Δt becomes small, an investor with finite smooth ambiguity-averse preferences will require no compensation for her uncertainty about θ and will behave to a first-order approximation like an expected utility investor who believes the true drift is $\theta = 0$. The parameter σ_θ^2 will not feature in her portfolio choice decision and therefore, in equilibrium, σ_θ^2 will not affect the excess return for security A. (As Δt goes to zero, this approximation becomes exact.)

For security B, the payoff is realized in a given subinterval of the period $[0, 1]$ with probability 1, no matter how small that subinterval. The subinterval timing is known in advance before date 0. As there cannot be news in the other intervals, there is no risk and therefore no risk premium. An investor with finite smooth ambiguity-averse preferences cannot reduce the risk she bears in

[11] We provide a more formal treatment in Section 7 and Appendix B, where we consider the model in Gollier (2011) to estimate ambiguity aversion.

the subinterval by trading more frequently. As a result, she will demand a premium for uncertainty about the drift θ, which will be increasing in σ_θ^2.

In the immediately preceding subinterval, the last time the investor can trade before the payoff is announced, she bears positive uncertainty and in the rest of the entire period she does not. Consequently, her utility function is more concave than in earlier or later subintervals, as we show in the paper. This implication is tested by looking at the option-implied concavity of preferences for a representative investor who holds the market.

If the investor holds both stocks (which together amount to the market), abstracting from rebalancing between stocks, the apparent premium realized on the market portfolio will be higher during the announcement interval and preferences will become more concave immediately prior to the announcement interval.

What other models and settings could generate this concavity result?

4.1. Recursive preferences

In a different setting, we can allow the investor to not only learn about the current payoff in the announcement interval, but also about future payoffs at the same time.[12] Imagine a representative investor consuming payoffs from a risky asset, so that the current consumption is required to equal the current payoff, as in Ai and Bansal (2018) (or indeed the original Lucas model). If such an investor has time-separable utility, then this additional source of risk will have no impact on investor welfare, and therefore no impact on the concavity of the investor's indirect utility function immediately prior to the announcement.[13] If preferences are non-time-separable, for example Epstein-Zin with elasticity of intertemporal substitution greater than one, then it is likely that preferences will appear to become more concave ahead of announcements.[14] (It is difficult to be

[12] This seems a particularly obvious channel for interest rate announcements, in which central banks reveal the likely path of future interest rates as well as the current rate.

[13] Ai and Bansal (2018) consider an endowment economy in which aggregate dividends are exogenously fixed to equal D_0, D_1, and D_2 at dates 0,1, and 2, where only $\widetilde{D}_2$ is random. They show that expected utility maximizers will not demand a risk premium to hold aggregate risk at date 0 if the uncertainty about the value of $\widetilde{D}_2$ is resolved at date 1 rather than, as is more standard, at date 2. The reason for this is that aggregate consumption at date 1 is required to equal D_1, which is deterministic. Epstein-Zin utility maximizers, on the other hand, will demand a risk premium in such a situation.

[14] This is the channel proposed in Savor and Wilson (2013), equation (8), for the announcement premium, but they do not discuss concavity of preferences.

certain because the quantity of risk to recursive investors also increases if they expect to learn not just about current but also future consumption.)

However, in our thought experiment above we have switched off this channel, by assuming the investor learns only about the current payoff from the announcement and nothing about future payoffs. Since no information ever arrives about future welfare, preferences for early or late resolution of uncertainty never switch on or off.

Ai, Bansal, Guo, and Yaron (2020) use the concept of generalized risk sensitivity in Ai and Bansal (2018) to explore whether investors have a preference for early resolution of uncertainty. They do so by making a series of assumptions regarding how investors learn in advance about how 'informative' FOMC announcements are going to be. Ai et al. (2020) then measure the average excess returns on a portfolio constructed to earn high (low) payoffs when investors learn that the FOMC announcement will be (un)informative. They conclude that the average returns on a claim on market variance in the period prior to an announcement present evidence of a preference for early resolution of uncertainty. Again, it seems quite probable that such a model would generate an increase in concavity ahead of announcements.

4.2. Other models of ambiguity aversion

We do not know whether maxmin preferences such as Gilboa-Schmeidler ambiguity aversion can lead to higher concavity. Skiadas (2013) shows that ambiguity aversion is second order when risks can be reduced by updating and acting more frequently (small risks in our paper), but he does not discuss how his results apply as smooth ambiguity aversion approaches the infinite limit. Deriving the findings in Skiadas (2013) for maxmin preferences is beyond the scope of our paper, but we suspect the answer is non-trivial.

On the one hand, there is a strong intuition that a Gilboa-Schmeidler investor who always relies on the same prior, the min, in both small- and large-risk situations will not change her behavior or her attitude to risk across those situations. On the other hand, our expression for the uncertainty premium in Equation (21) in Section 7 seems to imply that the difference between small- and large-risk behavior should be increasing in ambiguity aversion. Reconciling these two conflicting intuitions is a task for future research.

Epstein and Schneider (2003, 2007) and Epstein and Ji (2013) study the role of learning in settings, including dynamic settings, with these preferences. Epstein and Schneider (2008) show that such ambiguity-averse investors who anticipate a low-quality signal about the fundamental value of an investment will demand an ambiguity premium for holding the investment ahead of the announcement when the quality of the signal is uncertain. In their model, the uncertainty is about the precision of the signal, in contrast to a standard Bayesian framework where the precision is known. Illeditsch, Ganguli, and Condie (2021), building on the work of Illeditsch (2011), show a similar mechanism is at work when the correlation between the signal and the fundamental is uncertain, yielding a richer set of predictions than the Epstein and Schneider framework, but with both models implying an ambiguity premium in advance of a signal whose informativeness is uncertain. It seems highly likely that option-implied concavity of preferences would increase in such a situation. This remains an interesting question for future research.

5. Empirical methodology

5.1. Data

Our study uses S&P 500 equity index options traded on the Chicago Board Options Exchange (CBOE). The S&P 500 is one of the most actively traded indices in the world and represents a good proxy for the aggregate U.S. stock market. S&P 500 index options are European, cash-settled options and are available for a wide array of expiration dates and a dense grid of strike prices.

S&P 500 index options have four maturity types. The first and predominant type has monthly maturity intervals. The expiration date is always set to be the Saturday after the third Friday in a given month. The second type also offers monthly maturities, but uses the last trading day of the month as the expiration date. The third type offers weekly maturities with expiration dates occurring every Friday. The two latter types were introduced in 2007 and have lower trading volumes. The fourth type are options with daily maturities, which were introduced recently and have grown rapidly in popularity. Our study focuses on the first type, since those options historically have the most trading activity and exhibit highest liquidity. Maturities up to 12 months are available on a monthly basis. Longer maturities are offered at lower frequencies going up to five years.

Our sample covers the period starting in January 1996 and ending in December 2016, over which we collect 10,131,322 bid and ask prices. The average daily trading volume during this time period was about 400,000 contracts, and the most active contracts are typically those with shortest maturities. We obtain option prices (bid and ask quotes at the daily market close, converted to midpoint prices for the subsequent extraction of state-price densities), required data on the underlying, dividend yields, and risk-free rates from OptionMetrics.

All our tests use daily prices and returns. This choice is primarily driven by the availability of data, as intra-day option prices are very noisy (for example, due to low liquidity). Furthermore, it is not obvious how long it takes for information released through announcements to be incorporated into prices (many models predict that prices do not adjust instantaneously, but generally do not offer more guidance on the exact length of the period). We do not have a strong prior about the exact length of the decision interval, though we expect that it is not infinitesimal but is shorter than a whole trading day. If this is indeed true, our analysis based on daily prices and returns will be affected by noise (as the horizon studied is longer than the actual announcement period), which should create a bias against us finding any results.

We exclude from our analysis options that have a bid price below \$0.5, as these usually have prohibitively high spreads, low volume, and low informational value. We also exclude sets for a given day that contain fewer than eight available strike values for puts and calls, as well as deep-in-the-money options, since their trading volumes are very low. We replace the removed data points for deep in-the-money put options with the corresponding values for out-of-the-money calls using put-call parity, and vice versa for deep in-the-money calls. We further exclude options where implied volatility cannot be calculated, which typically happens because of violations of no-arbitrage conditions. Since option trades do not occur continuously for all maturities and strikes and to ensure that prices have information content, on a given day we study only options that have a positive trading volume. Option contracts with maturities of fewer than eight days often show erratic behavior and consequently we also omit these contracts. Finally, as the majority of available data is concentrated in maturities of up to three months (see Figure A.1 in the Internet Appendix), we focus our analysis on this maturity interval.

5.2. *Methodology*

A set of options with a common expiration date represents a snapshot of aggregated market beliefs, which allows us to extract the state-price density (SPD) for that maturity. We can then apply the idea in Ait-Sahalia and Lo (2000) to infer the concavity of the utility function from SPDs (together with the physical probability density). The resulting measure is the empirical counterpart of the concavity index (elasticity of marginal utility with respect to wealth) from Section 3. Under the null of no ambiguity aversion, it equals relative risk aversion (Ait-Sahalia and Lo 2000). However, if investors are ambiguity averse, it captures not only risk aversion but also ambiguity aversion. The estimation approach we adopt follows the methodology in Figlewski (2008) and also incorporates ideas by Breeden and Litzenberger (1978), Banz and Miller (1978), Shimko (1993), and Bliss and Panigirtzoglou (2004).

5.2.1. *Interpolation and smoothing*

Before we can extract a SPD that is suitable for further analysis, we clean the data to remove possible inconsistencies and eliminate extreme outliers with very low volume and high spreads. This approach prevents our estimated SPDs from displaying theoretically impossible features such as negative density values. Many direct smoothing approaches lead to erratic fluctuations near the at-the-money strikes. To address this issue, Shimko (1993) proposes an algorithm that converts option prices into implied volatilities based on the Black-Scholes-Merton model. The implied volatilities are then smoothed and interpolated, before prices for a dense set of equally-spaced strike prices are re-converted back to prices as a basis for the extraction of the SPDs.[15] In this study, we try to use as little smoothing and adjustment of the data as possible. However, in order to aggregate the data by degrees of moneyness, we need comparable data points that do not necessarily correspond to available strike levels and thus require the use of interpolation.

We evaluated several smoothing methods. Higher-order polynomials already provide a high degree of flexibility and are able to fit a large part of the volatility surface. However splines, as for example those used by Figlewski (2008), allow an even more accurate replication of the features

[15] Shimko (1993) stresses that the transformation does not assume the correctness of the Black-Scholes-Merton model, but merely uses the Black-Scholes-Merton formula as an algorithm that allows researchers to smooth and interpolate the given option set in a space that, though being theoretically inconsistent, preserves the volatility structure and information content comparably well.

of the curves. Cubic and fourth order splines run mandatorily through all k data points and fit the resulting curve by $k-1$ aggregated polynomials of third (or fourth) order. Apart from the function value itself, all derivatives up to the second (third) derivative are aligned at the transition points between the polynomials. The 6^{th} order polynomial imposes its curvature on the data points and leads partly to comparably large deviations. The cubic spline on the other hand has to run through all data points and also incorporates outliers that should be removed in the fit. In order to address these weaknesses, we also test smoothing splines, which allow the curve to not exactly go through all data points. The amount of acceptable deviation is defined by a smoothness parameter as proposed by Bliss (2002). Nevertheless, outliers in one area can still negatively affect the whole fitted curve.

In this section, we thus employ B-splines to reduce drastic outliers. B-splines are a generalization of Bezier curves and are closely related to smoothing splines. They also do not require that the curve runs through all data points. We choose a set of knot points that act as transition points between areas, in which the curve has different characteristics. As SPDs have different maturities, their width also varies. Therefore, the position of the knot points also has to vary with the maturity of the SPD. For a maturity of 1 year, we position the knot points at-the-money, at a moneyness level of 90 %, and at 110 %. For other maturities, the two outer knot points are scaled with the square root of time until maturity. For example, for a maturity of 3 months, the knot points would be set at $\sqrt{\frac{3}{12}} = \frac{1}{2}$ of the 10 % per annum distance, i.e. at moneyness levels of 95 % and 105 %. In between those knot points, a polynomial function is fitted with a least square algorithm.[16] We employ this smoothing technique at the level of the extraction of the risk- and ambiguity-neutral distribution and smooth the curve before it is transformed into the state-price density. As this method already enables us to extract suitable SPD estimates, we refrain from using smoothing in volatility space in our results. We only use the method suggested by Shimko (1993) as a comparison method. Figure A.2 in the Internet Appendix provides an illustration of the impact of the smoothing method we employ, which aims to keep the underlying shape with as little distortion as possible while removing random fluctuations in the data.

[16] To minimize the impact on the information content, we calibrate the process to employ a dense set of knots and allow only small deviations from the actual data points.

5.2.2. State-price density function extraction

Figlewski (2008), following Breeden and Litzenberger (1978), exploits the fact that the state-price density function $\pi(x)$ allows us to express the value of a security as the value of its expected payoff at maturity T under the risk- and ambiguity-neutral measure.[17] For a call option, this is the risk- and ambiguity-neutral probability-weighted difference between the stock price S_T and the strike K at maturity T:

$$C = \int_K^{\infty} e^{-rT} \left(S_T - K\right) \pi(S_T)\, dS_T. \tag{13}$$

C represents the call price, P the put price, S_t the price of the underlying asset at time t, K the strike price, r the risk-free rate, and T the maturity.[18] Correspondingly, we can write the price of a put option as:

$$P = \int_0^{K} e^{-rT} \left(K - S_T\right) \pi(S_T)\, dS_T. \tag{14}$$

As shown in Figlewski (2008), it is possible to extract $\pi(x)$ and $\Pi(x)$ (the state-price distribution function $\Pi(x) = \int_{-\infty}^{x} \pi(z)dz$) from the partial derivatives of option prices with respect to K

$$\frac{\partial C}{\partial K} = -e^{-rT} \int_K^{\infty} \pi(S_T)\, dS_T = -e^{-rT} \left[1 - \Pi(K)\right]. \tag{15}$$

Solving for $\Pi(K)$ and taking the second derivative yields:

$$\Pi(K) = e^{rT} \frac{\partial C}{\partial K} + 1 \qquad \text{and} \qquad \pi(K) = e^{rT} \frac{\partial^2 C}{\partial K^2}. \tag{16a,b}$$

For put options, the analogous expressions are given by:

$$\Pi(K) = e^{rT} \frac{\partial P}{\partial K} \qquad \text{and} \qquad \pi(K) = e^{rT} \frac{\partial^2 P}{\partial K^2}. \tag{17a,b}$$

[17] Figlewski (2008), like many others, uses the term 'risk-neutral density,' but in our setting with ambiguity aversion 'state-price density' is a more appropriate term.

[18] Please note that this expression for the option price is independent of the Black-Scholes-Merton formula and its assumptions and is based purely on the SPD-weighted integral of option payouts across all possible states at maturity T. This value is then discounted to $t = 0$ at the risk-free rate and not the cost of carry for S&P 500 futures (i.e., it does not include the dividend yield).

The state-price distribution function $\Pi(x)$ usually runs S-shaped from 0 to 1, and represents the probability that the index does not exceed a particular moneyness value at maturity. Its form is close to a cumulative normal distribution function. The SPD $\pi(x)$ is the first derivative of the state-price distribution function and resembles a normal distribution. It is however typically leptokurtic and skewed. We show the characteristic forms of $\Pi(x)$ and $\pi(x)$ in Figures A.3 and A.4 in the Internet Appendix.

Given a set of options with a common maturity spanning all realistic strike prices K, it is possible to estimate the SPD and state-price distribution functions using numerical derivatives. The denser the grid of strike prices that is provided by the option set and the smaller the bid-ask spreads, the more accurate the extracted functions.

5.2.3. Option-implied concavity of preferences

Given our estimates of the SPD $\pi(S_T)$ and physical densities $p(S_T)$, using the index level S_T as wealth, and letting $\lambda(S_T) = \pi(S_T)/p(S_T)$, we can estimate our concavity index using Equation (10). The ratio λ is proportional to the pricing kernel and only differs by a discount factor. The implied concavity is then:

$$CI(S_T) = -\frac{S_T \frac{d}{dS_T}\lambda(S_T)}{\lambda(S_T)}. \tag{18}$$

The assumption to equate wealth with the equity index is unproblematic for the representative agent who by definition invests in the market portfolio. Consequently $\frac{dW_s}{dS_T} = 1$ holds and can be ignored.

In our model with smooth ambiguity-averse investors, this concavity index will reflect both risk aversion and ambiguity aversion. However, as we argue above, ambiguity will matter only for large risks such as scheduled announcements, allowing us to distinguish the effect of ambiguity aversion from that of risk aversion.

6. Results

Savor and Wilson (2013) show that asset prices behave differently on days of pre-scheduled macroeconomic announcements about inflation, unemployment, and interest rates. In this section, we estimate the effect of such announcements on the physical probability densities of equity indices and on the option-implied state-price densities. We then show that option-implied concavity on days with macroeconomic announcements is higher than concavity for regular non-announcement (non-a) days.

As in Savor and Wilson (2013), we consider scheduled announcements for the Producer Price Index (PPI), unemployment (Emp), and Federal Open Markets Committee (FOMC) interest rate decisions. The dates of U.S. employment and inflation announcements come from the Bureau of Labor Statistics, and the dates of scheduled FOMC meetings directly from the Federal Reserve. Our sample includes 839 announcements in the period starting in January 1996 and ending in December 2016. We use the S&P 500 equity index and its index options as a proxy for the U.S. stock market. We define all days associated with an announcement as 'announcement days' and all other days without announcements as 'non-announcement days.'

6.1. Implied volatility and bid-ask spreads

Table 1 presents the summary statistics for the implied volatility surface of S&P 500 index options, separately for non-announcement days and different types of announcements days. Consistent with the hypothesis that announcement days are periods of elevated risk, implied volatility is higher on announcement days. This result is strongest for FOMC announcements and for options that are near expiration. Table 2 reports average bid-ask spreads of the options in our sample, which are significantly higher on all three types of announcement days. For FOMC and employment announcements (but not for PPI), spreads are also elevated on the day before the announcement. These findings provide further evidence that investors (and market makers) perceive announcement days as either riskier or more uncertain than 'regular' trading days.

[TABLES 1 AND 2 ABOUT HERE.]

6.2. Implied state-price density

Option markets offer insights into the aggregated beliefs of market participants. In particular, state-price densities extracted from liquid option prices offer a sensitive instrument to visualize market expectations. Here we examine the effect of announcement days on the implied SPDs. In Figure 2, we show the average SPD on non-announcement days (over our full sample), implied by options with maturities up to three months and an average maturity of about 44 days.[19] We compare this SPD to announcement-day SPDs, separately for PPI, employment, and FOMC, as well as the aggregate of all three. Once more, we document an increase in the width of the curve and more uncertainty reflected in the option market for FOMC, PPI, and employment announcement days. The effect is particularly pronounced for PPI and FOMC announcement days. For employment announcement days, the effect is of less strong, but increased skewness is still visible.

[FIGURE 2 ABOUT HERE.]

Table 3 provides the analysis of the effect of announcement days on the form of the SPD. We perform a simple t-test, a Kolmogorov-Smirnov test, and a Mann-Whitney test. In all tests we compare the set of SPD values at a moneyness level of 1.02 (maximum of the average SPD distribution) on announcement days to the set of values on non-announcement days.[20] The t-test examines whether the mean of the non-announcement day SPD values is higher than the mean of announcement-day SPD values. The Kolmogorov-Smirnov test compares the two sets of SPD values and studies whether the non-announcement day value distribution stochastically dominates the one on announcement days. The Mann-Whitney test determines whether the two distributions of SPD values, i.e. the one on non-announcement days and the one on announcement days, significantly differ by a positive location shift.[21]

[TABLE 3 ABOUT HERE.]

Our results show that there is significant impact of announcements on the market. With the exception of employment announcements, p-values are quite low and indicate that the SPD on

[19] In the aggregation of SPDs, we face the choice between either scaling the moneyness of the SPDs according to the maturity before aggregation or simply averaging the SPDs. The maturity scaling approach has the disadvantage of accepting the partly distorting effect of the scaling approximation and its implicit assumptions. This becomes particularly relevant for SPDs that show non-standard behavior. The direct approach has the disadvantage that we have to accept a certain degree of inaccuracy by averaging distributions with similar but not same maturities. We use the direct averaging approach in this paper, but we also show that our results are robust to employing maturity scaling. Please see Tables A.1 and A.2 in the Internet Appendix.

[20] We employ a definition of moneyness that compares the strike price to the current index value (i.e., a moneyness of 1 corresponds to the current S&P 500 index value). An alternative definition of moneyness would compare the strike price to the forward price. Since the main results in the paper rely on the comparison of implied and physical probability density, moneyness cancels out as long as its definition is consistent for both densities. Furthermore, the options we use in our tests have short maturities (average is 44 days), so that the difference between current and forward prices is relatively minor. Please refer to Figures A.5 and A.6 in the Internet Appendix showing the average probability density, SPD, and concavity over the 1996–2016 period based on forward instead of current moneyness.

[21] Please note that whenever we mention those tests in the results sections, we refer to one-sided tests.

announcement days is distinctly different than on regular non-announcement days. In particular, announcements lead to a measurable reduction of the SPD's maximum and to a widening of the distribution. We interpret this result as stemming from the higher uncertainty reflected in the option market on announcement days that leads to a widening of the SPD and a subsequent reduction in the mean maximum SPD value.

6.3. Physical probability density

Figure 3 plots the physical probability density of S&P 500 returns, aggregated over the sample period from 1996 to 2016. We use daily returns and show their frequency distribution comparing announcement days to 'regular' days without announcements.[22] Daily returns on non-announcement days are clearly concentrated in a narrow area above at-the-money. Returns strongly deviating from this area are rare and the observed physical probability density falls away very quickly for rising or falling moneyness.

[FIGURE 3 ABOUT HERE.]

The behavior of the physical density is different on days with announcements, with lower clustering around the center of the distribution and much higher mass on higher-return outcomes, as well as somewhat higher mass on lower-return outcomes. In fact, as the Kolmogorov-Smirnov tests in Table 4 show, the distribution of FOMC- and employment-day (as well as of all announcements together) market returns first-order stochastically dominates that of non-announcement day returns. More generally, Table 4 shows that announcement-day returns are economically and statistically significantly higher than returns on non-announcement days, in line with the findings in Savor and Wilson (2013). Announcement returns account for 57% of the total equity risk premium in the 1996–2016 period, with a Sharpe ratio that is about 9 times higher.

[TABLE 4 ABOUT HERE.]

6.4. Comparison of physical and implied state-price densities

Physical and implied probability distributions can differ considerably. Figure 4 compares the physical probability density to the average implied SPD. The implied curves, aggregated from

[22] More specifically, we group returns by the type of day, assign them to 0.01 intervals (i.e., 0.99–1, 1–1.01, etc.), count them, and then convert into the respective frequencies.

1996 to 2016, are distinctly wider and show that implied volatility is historically higher than realized volatility, a phenomenon usually referred to as the variance risk premium in the literature (in our model, ambiguity and ambiguity aversion represent another channel that may lead to this result).

[FIGURE 4 ABOUT HERE.]

The physical distributions used in the following sections are based on the aggregation of historical S&P 500 returns. All returns, irrespective of the type of day (announcement or non–announcement day) on which the period starts are aggregated and converted into a physical return distribution for a particular maturity for a particular day.[23]

In order to avoid distortion due to the scaling of 1-day returns, we calculate the physical distribution of observed returns directly based on average historical returns over a rolling period consistent with the maturity of the options used for the respective implied SPDs. This means that when we later compute the pricing kernel and the concavity index, the time horizon of the respective physical distribution used corresponds to the maturity of the SPD it is compared to.

6.5. Implied concavity

In this section, we investigate the effect of announcement days on the implied concavity of the utility function for a representative investor in the market, estimated using Equation (18). Figure A.7 in the Internet Appendix provides an illustrative example of the pricing kernel with maturity March 20, 1999 for the S&P 500 index and is based on the extracted SPD and the observed physical density used in this study.[24] From left to right the curve is downward-sloping as one would expect for a system with complete markets, risk- and/or ambiguity-averse investors, and common beliefs. However, around at-the money, the pricing kernel slopes upwards again and forms a maximum at low in-the-money levels. This indicates that investors attribute a higher marginal value to payouts in rising markets than in falling ones. This phenomenon is also replicated in other studies and is referred to as the Pricing Kernel Puzzle in the literature. For instance Ait-Sahalia and Lo (2000),

[23] Almost every considered time period contains more than one announcement day, making a clear-cut separation between returns on announcement and non-announcement days impractical.

[24] The particular date is chosen randomly, as the figure simply serves as an example of the typical form of the pricing kernel. Please note that the figure is based on the observed physical density, so that it reflects the unconditional historical distribution of returns. In addition to the general smoothing at the level of the state-price distribution as explained in Section 5.2.1, we also apply smoothing and interpolation at the level of the pricing kernel to present a continuous function.

Jackwerth (2000), Rosenberg and Engle (2002), and Bakshi and Chabi-Yo (2012) all find this anomaly in their tests.

As shown in Equation (18), this pricing kernel can be used to extract the form of the concavity index over moneyness, shown in Figure 5.[25] Note that the concavity index takes negative values at certain moneyness levels in this example. This is again due to the pricing kernel puzzle. As concavity is proportional to the negative slope of the pricing kernel, an upward-sloping kernel implies negative concavity. The effect that concavity is reduced to a minimum, usually closely below at-the-money, is present throughout the data. Whether this minimum turns negative or not on a given day, varies with market conditions.[26]

[FIGURE 5 ABOUT HERE.]

In order to show that market implied concavity differs significantly for days with and without announcements, we investigate each announcement type separately and compare the respective local (non–parametric) implied concavity to that on non-announcement days. We furthermore conduct t-tests, Kolmogorov-Smirnov tests, and Mann-Whitney tests to show that differences are statistically significant. Figures 6 through 8 show the concavity index values for FOMC, PPI, and employment announcement days, respectively. (Figure 1, which we discussed above, plots the concavity index values for all three announcement types together.) The charts plot concavity values aggregated per type of day via the arithmetic mean for the 1996–2016 period and are centered around the respective at-the-money value. Note that the time horizon used for the estimation is based on options with a mean maturity of 44 days or about 6 weeks. Consequently, a 1 % variation in moneyness corresponds to a significantly larger annualized price fluctuation. We also always present the concavity estimate for non-announcement days (i.e., days other than PPI, employment, or FOMC days) for comparison. We plot 95 % confidence intervals as dotted lines of the same color.

[FIGURES 6, 7, AND 8 ABOUT HERE.]

[25] We chose the range of this figure to highlight the main features of the observed concavity index.

[26] Hens and Reichlin (2013) propose various potential explanations for the pricing kernel puzzle, such as non-concave utility functions, incorrect beliefs, and incomplete markets. In our setup, ambiguity loving preferences would imply convexity. However, despite experimental evidence documenting that agents can be ambiguity loving in certain circumstances, this is not a widely held view in the literature. Thus, we do not attempt to resolve this puzzle or distinguish between different explanations.

We find a clearly visible increase in our concavity index ahead of scheduled announcements. On days of FOMC, PPI or employment announcements, the concavity index (CI) increases on average by more than 20 % relative to non-announcement days. These increases are distinctly visible only on announcement days. On post-announcement days, the implied concavity reverts back to its normal non-announcement-day level. Our results are not driven by day-of-the-week effects, as we find no difference in the concavity index on different days of the week.

6.5.1. Aggregated concavity estimates

We provide statistical analysis in which we aggregate concavity estimates over a predefined moneyness area, spanning the range from 0.98 to 1.02, to form one concavity value per day and maturity. To condense the information of the concavity index in this range into one comparable number, we fit a function based on a representative agent's utility function to it. As in Bliss and Panigirtzoglou (2004), we use the power utility function and the exponential utility function (Table 5). The associated respective Relative Risk Aversion (RRA) can be computed as γ_{pow} or $S_T\,\gamma_{exp}$ and is fitted to the empirically determined concavity index over the range considered (0.98–1.02, Equation 19).

$$C.I.(S_T) \qquad \sim \gamma_{pow} \qquad or \qquad \sim \gamma_{exp}\,S_T \tag{19}$$

[TABLE 5 ABOUT HERE.]

Each of the two utility functions can be characterized by a single parameter γ_{pow} or γ_{exp}, respectively. Bear in mind, however, that following Equation (10), the concavity index reflects not only relative risk aversion but also ambiguity aversion. Thus, one should not interpret the γ coefficients as estimates of relative risk aversion, at least not on announcement days. Our principal objective in this exercise is simply to aggregate concavity estimates over different moneyness levels and then test whether the resulting measure is higher on announcement days, as predicted by our model and consistent with a higher price of uncertainty on such days.

In Panel A of Tables 6 and 7, we compare the distributions of those parameters γ on announcement days to the respective values on non-announcement days. We focus on the area around at-the-money (0.98 to 1.02), as it provides the highest data quality and the largest amount of available data.[27] To assess the distribution of the γ parameters, we conduct t-tests, Kolmogorov-Smirnov tests, and Mann-Whitney tests. The t-test examines whether the mean of the non-announcement day γ values is higher than that on announcement days. The Kolmogorov-Smirnov test explores whether the announcement-day γ distribution stochastically dominates the non-announcement-day distribution. The Mann-Whitney tests assesses whether the distribution of announcement-day concavity differs by a positive location shift from that on non-announcement days.

[**TABLES 6 AND 7 ABOUT HERE.**]

Each of the three announcement types separately, as well as all announcements together, are associated with an increase in implied concavity before the announcement. For FOMC announcements, PPI announcements, employment announcements, and all announcements combined, p-values for the t-test and for the Mann-Whitney test are below 0.1 %. For the very demanding Kolmogorov-Smirnov test, p-values for all announcement types reach values between 0.0 % and 2.0 % on the single and aggregated level, and generally show the highest level of statistical significance. For employment announcements, Figure 8 indicates that the distribution of increases across moneyness levels is not as stable as for the other announcement types. Nevertheless, even for employment announcements, the results are still highly significant.[28]

In Tables A.5 and A.6 in the Internet Appendix, we examine market returns and the concavity index over a period beginning three days before the announcement and ending three days after. The key takeaway is that announcement-day returns are much higher than returns on other days, which exhibit no real pattern. The only exception is the high average return one day before FOMC announcements, which is the well-known pre-FOMC announcement drift in Lucca and Moench (2015).[29] With respect to the concavity index, the major difference is between announcement days and non-announcement days, with no strong patterns either before or after the announcement. The concavity index is higher one day before FOMC announcements compared

[27] Since the implied distributions are based on individual options, the informational value and precision of the estimation decrease the further one moves away from the at-the money region that has the highest liquidity and most trading volume. Tables A.3 and A.4 in the Internet Appendix present the results of robustness tests when we extend the moneyness interval. Our findings remain the same.

[28] Tables A.1 and A.2 in the Internet Appendix present the results based on the alternative approach of scaling moneyness according to maturity. For consistency with the no-scaling tests, we use a slightly longer moneyness interval from 0.97 to 1.03.

[29] The exact origin of this drift is subject to much debate, with different potential explanations that include those based on microstructure effects, explicit leaks, or other concurrent information releases.

to other non-announcement days, but is still significantly lower than on the FOMC announcement day itself. Overall, we conclude that announcement days indeed are special, even relative to non-announcement days shortly before or after the announcement.

The average maturity of options used to extract our SPDs is 44 days, which is likely considerably longer than the horizon over which markets impound information released through macroeconomic announcements into prices. Furthermore, the 44-day window will always cover multiple announcements. While both of these issues should make it more difficult for our tests to establish an increase in concavity on announcement days, a shorter window would allow us to more closely align our empirical tests with our theory and better control for the difference in physical distributions across different types of trading days. Unfortunately, short-maturity options suffer from data quality issues and are difficult to work with, which is why most papers in the literature exclude them from their samples (Bliss and Panigirtzoglou 2004; Panigirtzoglou and Skiadopoulos 2004; and Figlewski 2008), an approach that we also adopt for our analysis above.[30]

As an additional test, we relax some of the data filters (minimum maturity, minimum price, minimum trading volume, minimum set of strike values per set) used to construct our main sample, and focus on options with a maturity of five days or less. This significantly reduces the number of announcements we can use. Moreover, the quality of the extracted SPDs is lower than that of SPDs used in our main tests, and the resulting SPDs often exhibit erratic behavior. However, as Table A.7 in the Internet Appendix shows, even in this restricted sample the concavity index is higher on announcement days than on non-announcement days. The statistical significance decreases relative to our principal results, which is not surprising given the lower number of observations and reduced general data quality. The findings also hold for FOMC and PPI announcements separately. We are not able to perform the analysis for employment announcements due to a lack of usable observations.[31]

[30] Trading volumes for weekly options are low during our sample period. Daily options, which nowadays are actively traded, were introduced after our sample ends.

[31] A likely explanation is that employment announcements occur on the first Friday of a month (or the second if the first Friday falls on the first day of a month), which is always more than five days before standard option expiration dates.

6.5.2. *Bootstrap test*

In addition to the above statistical tests, we also conduct a bootstrap analysis in Panel B of Tables 6 and 7. Bootstrapping tests the reliability of model estimates by replicating a varied version of the underlying sample. One of its advantages in our setting is that it takes into account estimation error, which we need to do since γ is only an estimate and therefore analysis based on just its point estimates likely overstates p-values.

Analogous to Bliss and Panigirtzoglou (2004), we construct samples of the same size as the original sample by sampling with replacement. We consider the uncertainty of the estimation process starting at the physical distribution (in addition to the variation in daily γ estimates). We resample the returns of the physical distribution for each sample and calculate the resulting γ estimates, which are then used in sampling with replacement. In particular, we create 1,000 sets of resampled physical probability densities (with replacement) for each set of options. With the redrawn physical probability distributions of returns, we proceed through our estimation process to combine them with the extracted SPDs and create the implied concavity measures. These are then converted into the γ values for the power and exponential utility functions. We use these resulting γ estimates to create 1,000 bootstrap samples of the same size as the original sample by sampling with replacement. (In total, we use 1,000,000 bootstrap samples.) Finally, we compute estimates of the γ parameter for each bootstrapping sample and compare the distribution of these γ estimates to the empirical results in Panel A.

We conduct standard bootstrap hypothesis tests (see for example MacKinnon 2009) that confirm the statistical significance of the t-test, Kolmogorov-Smirnov, and Mann-Whitney test results for the difference in the concavity index between announcement and non-announcement days. Crucially, the bootstrapped p-values remain extremely low, highlighting again the clear difference between γ values on announcement and non-announcement days.

In summary, both the conducted statistical tests and the bootstrap analysis provide strong evidence for an increase in concavity around announcement days and therefore for the existence of a significant ambiguity premium. The difference between announcement and non-announcement days represents a lower bound on ambiguity aversion, since if non-announcement day risks are not all small, ambiguity aversion will also matter on non-announcement days.[32]

7. Calibration exercise

In this section, we calibrate a simple smooth ambiguity aversion model with both risk-aversion and ambiguity-aversion parameters. Our starting point is Gollier (2011), who provides a closed-form example of smooth ambiguity aversion. We provide all the details and derivations in Appendix B, including for a model with return uncertainty (which is the simpler version) and one with payoff uncertainty (the more complex version).

We begin with a representative investor who allocates her wealth at date 0, W_0, between a risk-free asset with net return R_f and the market portfolio with risky net return R_{MKT}. The investor believes market excess returns are normally distributed with mean $\widetilde{\theta}$ (unknown) and variance σ^2 (known). (The assumption that the variance of market returns is known is convenient, but also highly plausible. See Merton, 1980, for a discussion of the differences in estimating the mean and variance of returns.) She has CARA utility of wealth W_1 with absolute risk aversion A, given by

$$U(W_1) = -\frac{1}{A}\exp(-AW_1). \tag{20}$$

The investor has a prior distribution over θ, which is also normal: $\theta \sim N(\mu, \sigma_\theta^2)$. She has constant relative ambiguity aversion, specifically:

$$\phi(U) = -\frac{(-U)^{1+\eta}}{1+\eta},$$

where $\eta \geq 0$ is her coefficient of relative ambiguity aversion ($\eta = 0$ corresponds to standard expected utility).

In equilibirum, the expected market excess return is then:

$$\mu = A(\sigma^2 + (1+\eta)\sigma_\theta^2). \tag{21}$$

Applying our argument from Section 2 above (see explanation and formal proofs in Appendix

[32] Risks on non-announcement days can be large if investors cannot trade as rapidly as they would like or because the combination of preferences and risks is one of the exceptions to Skiadas' (2013) general finding.

B), we then show that the per-period expected returns in non-announcement periods, μ^{NA}, and announcement periods, μ^{A}, are:

$$\mu^{NA} = A\sigma^2 > 0 \tag{22}$$

and

$$\mu^{A} = \mu^{NA} + A(1+\eta)\sigma_\theta^2 > \mu^{NA}. \tag{23}$$

Note that the variance of returns over the same length of periods in the two regimes is the same, σ^2.

In this model the uncertainty premium is $A(1+\eta)\sigma_\theta^2$, uncertainty is measured by σ_θ^2 and the price of uncertainty is $A(1+\eta)$, the product of absolute risk aversion and relative ambiguity aversion plus one. The price of uncertainty is the increase in premium μ for an increase in uncertainty σ_θ^2, holding all else equal.

Setting $\eta = 0$ in Equation (23) shows that parameter uncertainty σ_θ^2 still drives up the risk premium for the investor, and moreover only ahead of announcements.[33] But this occurs only because the quantity of risk borne by the investor (as in a Bayesian model with learning, for example) has increased from σ^2 to $\sigma^2 + \sigma_\theta^2$. The risk premium per unit of risk, or price of risk, remains equal to A. Applying the reasoning in Section 2, one can show that since $\phi'' = 0$ when $\eta = 0$, there is no increase in concavity of preferences in this case.

7.1. Estimating ambiguity aversion

The major challenge in estimating any model with ambiguity aversion is how to measure uncertainty (σ_θ^2) as distinct from risk (σ^2). In our exercise here, we take a straightforward approach where we use the square of the standard error of the mean excess return as our proxy for σ_θ^2 (i.e., $\sigma_\theta^2 = \frac{\sigma^2}{N}$, where N is the number of observations). This measure directly takes into account the variance of the sampling distribution of the mean return and thus corresponds to σ_θ^2 in Gollier's model.

Our approach implicitly assumes that investors can reduce, and eventually eliminate, uncertainty through learning (as the sample size grows, the standard error will decrease). In our model there is no learning, but we still use standard errors, as they are a common measure of uncertainty

[33] We thank the referee for pointing this out.

in the literature. There is also a large literature that discusses the role of learning and its potential limitations in reducing uncertainty (e.g., if parameters of interest are not constant, as is likely in many real-world environments such as financial markets).[34] One takeaway from this prior work is that measuring uncertainty with the standard error is in many circumstances an overly simple approach. We understand and accept this point, with all the caveats it implies for our exercise here. We also note that reliance on the standard error alone and neglecting other data investors possess may bias upwards our estimates of uncertainty.

We begin by estimating the mean of the market excess return for announcement and non-announcement days separately. Over our sample period, the mean announcement-day excess return is 15.1 basis points (bps) while the mean non-announcement-day excess return is much lower at 1.6 bps. Crucially, the standard deviation of returns is roughly the same for both types of days at 1.2%, in line with the assumption in Gollier's model. We assume investors care about annual returns and therefore annualize the mean return estimates (and the associated standard errors).

We need to stress at this point that, unlike in simple expected utility models, investor horizon can matter when investors exhibit smooth ambiguity preferences. This is easily seen from the solution for the risky asset premium in Gollier's example: $\mu = A(\sigma^2 + (1+\eta)\sigma_\theta^2)$. This premium depends on σ_θ^2, which is just the squared standard error in our setup. Since the squared standard error is proportional to $\frac{1}{N}$ (where N is the number of observations), the importance of the ambiguity-aversion term $A(1+\eta)\sigma_\theta^2$ decreases relative to the risk-aversion term $A\sigma^2$ as investor horizon becomes shorter (and the number of individual observations is thus higher).[35]

Combining equations (22) and (23), we can solve for η as:

$$\eta = \frac{\mu^A - \mu^{NA}}{\mu^{NA}} \frac{\sigma^2}{\sigma_\theta^2} - 1. \tag{24}$$

[34] In the smooth setting, this point is implicit in the work of Ju and Miao (2012) and others (e.g., Collard, Mukerji, Sheppard, and Tallon 2018), who use a hidden Markov-switching model to prevent learning from resolving uncertainty.

[35] This may seem counter to the intuition in Merton (1980), who shows that only the length of the sample period (and not return frequency) matters when estimating expected returns. However, that conclusion is based on the fact that both means and associated standard errors vary in the same way with return frequency, so that their ratio does not change. For example, if investor horizon changes from monthly to annual, the mean return will be 12 times higher, but the standard error of the mean will also be 12 times higher. The key difference is that in Gollier's model with smooth ambiguity aversion the term that matters is the square of the standard error (which would be 144 times higher in the previous example).

To estimate μ^A (as the mean announcement-day excess return), we have 671 daily observations in our sample. Assuming there are 252 trading days in a typical year, for the annual expected excess return we get:

$$\frac{\sigma^2}{\sigma_\theta^2} = \frac{\sigma^2}{\sigma_{\bar{\mu}^A}^2} = \frac{\sigma^2}{\frac{\sigma^2}{(671/252)}} = 2.66. \tag{25}$$

From equation (24), we can then compute our η estimate, which equals 21.5. Applying this same approach to individual announcements, we get η estimates of 11.7 for FOMC days (mean excess return of 32.2 bps and 167 observations), 2.5 for PPI days (7.2 bps and 252 observations), and 5.3 for employment days (11.6 bps and 252 observations).[36]

What do these estimates mean? One way to answer this question is to explore their implications for equity premia. Assume the (annual) equity premium is 6% and the volatility is 15%. If investors do not exhibit ambiguity aversion, we can immediately from equation (22) get their coefficient of absolute risk aversion A, which is 2.67. The next step is to explore how much higher would the equity premium be if investors were also ambiguity averse. Using our η estimate of 21.5, equation (23) gives us a premium of 12.4%, roughly double its level with no ambiguity aversion.[37]

This is certainly a high number, but is at least of the same order of magnitude as the observed equity premium. Furthermore, the estimate of A we used is inflated, since it is computed assuming there is no ambiguity aversion. If the true A is lower, the 'extra' ambiguity premium would also be lower. For example, if we estimate A just based on non-announcement-day returns, we get 1.15. In that case, the equity premium for ambiguity-averse investors with $\eta = 21.5$ is 5.4%. With the caveat that these numbers are based on a very simple model of smooth ambiguity aversion, we conclude that our η estimates imply that smooth ambiguity aversion represents an important driver of expected excess returns in equity markets.

[36] Our η estimates differ across various announcements. One potential explanation for this is that our estimate of σ_θ^2, the square of the standard error of the mean excess return on a particular type of announcement day, does not accurately reflect expectations of investors. On the one hand, investors have access to much more data than just historical stock returns, which could reduce ambiguity below the level we estimate. On the other hand, parameters of interest to investors may not be constant, which would limit how much learning can reduce ambiguity.

[37] If we use returns over our 1996–2016 sample period to compute the standard error of the mean excess return, we get that $\sigma_\theta^2 = \frac{\sigma^2}{N} = \frac{0.15^2}{21} = 0.00107$.

7.2. Comparing uncertainty across announcements

If we make a (plausible) assumption that η is the same across all announcements, we can use equation (23) to examine how uncertainty varies across different announcements without relying on standard errors as our measure. Taking all announcements pooled together as the baseline, uncertainty is 126.7% higher on FOMC days, 58.5% lower on PPI days, and 25.9% lower on employment days. (Please note that we cannot compute uncertainty on non-announcement days, since the price of uncertainty on these days is zero in our model, given that investors can always reduce it to an arbitrarily low level by trading more frequently.) We can also compare our uncertainty measure based on standard errors with uncertainty estimates implied by announcement premia. If annualized market volatility on announcement days is 15%, equation (25) implies $\sigma_\theta = 9.2\%$ (on annualized basis). Setting $A = 1.15$ and $\eta = 21.5$, we get that $\sigma_\theta = 11.5\%$ (see equation (23)), which is quite close and potentially validates the standard-errors approach.

An alternative method for measuring uncertainty across announcements would use the concavity indices from Table 7. More specifically, in Appendix B we derive an expression for the concavity index under the Gollier model:

$$CI_s = AW_s - W_s \left(\frac{E\left[\omega_\theta N\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]}{E\left[\omega_\theta \Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]} - \frac{E\left[N\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]}{E\left[\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]} \right), \tag{26}$$

where s denotes the state (i.e., a given stock price level) and $\Phi(.)$ and $N(.)$ are the standard normal cumulative density and probability density functions, respectively. ω_θ is the weighting function implied by the Gollier model for each possible prior θ.

We take the concavity index estimate on non-announcement days as the risk aversion coefficient. We use excess returns and volatilities from Table 4 to measure expected returns and risk. We then evaluate equation (26) numerically and match it to the estimates in Table 7 to extract ambiguity aversion η and uncertainty σ_θ^2.[38] Over all announcement days, we obtain an η estimate of 19.3, which is encouragingly similar to the 21.5 estimate based on equity premia that we computed in the previous section. Again assuming this η is the same across all announcements, we compute

[38]The model in Gollier (2011) assumes CARA utility over wealth given a prior expected return on the risky asset, which corresponds to the concavity index in Table 7.

σ_θ by finding the level for which the concavity index most closely matches Table 7 numbers. The results show that uncertainty is almost the same on FOMC and employment days, with σ_θ of 6.7% in both cases.[39] Uncertainty is 9% lower on PPI days, with the important caveat that it is difficult to match the associated concavity index estimate (it is significantly higher than on other announcement days, but at the same time the excess return is lower, which is a difficult combination).

8. Conclusion

In this paper we argue that a distinction between large risks (risks that cannot be materially reduced by trading more frequently) and small risks (risks that can be so reduced) allows us to identify an ex-ante increase in the price of stock market uncertainty before scheduled announcements. We demonstrate how to estimate the concavity of preferences for a representative agent by combining the idea in Ait-Sahalia and Lo (2000) with more recently developed techniques for extracting SPDs. Our resulting empirical findings provide strong evidence for an increase in implied concavity ahead of macroeconomic announcements and thus for the existence of an uncertainty premium. Finally, we use a simple model of ambiguity aversion to directly estimate the coefficient of relative ambiguity aversion, whose magnitude suggests that smooth ambiguity aversion represents an important driver of expected stock returns.

These findings have important implications. They suggest that one key role of securities markets, which has not received much attention so far, is to facilitate breaking down large risks into smaller ones by allowing investors to trade frequently. Since small risks carry no ambiguity premium (or possibly premia implied by a more general class of non-standard preferences) while large risks do, trading lowers equity and other asset premia by eliminating or reducing the ambiguity premium. In other words, the provision of liquidity, which allows investors to adjust their positions more frequently and thereby convert a large risk into a series of smaller ones, is one of the primary ways in which markets help reduce the cost of capital.

Of course, there are alternative hypotheses that can explain our results. Perhaps the most obvious one is that a representative agent simply becomes more risk-averse ahead of announcements.

[39] It may seem surprising that uncertainty is so similar even thought concavity is about 12% higher on FOMC days than on PPI days. However, expected returns are also higher on FOMC days, which is an offsetting effect.

On its face, this idea seems very ad hoc, but if heterogeneous agents reallocate risks ahead of announcements, then it is conceivable that the risk aversion of the marginal holder of the market might be higher. Such a model would be much more complicated than our simple, tractable model of a representative agent with smooth ambiguity aversion, and to our knowledge no such model has ever been developed. Another potential explanation is that investors with recursive utility expect to learn about consumption further in the future, as in Ai and Bansal (2018).

A third possibility is that our estimates of physical probability are subject to small sample bias. However, this is more likely to be a problem in the tails of the distribution, as opposed to the center, and our results actually are the strongest for the center of the distribution.

The distinction between small and large risks may be relevant in other non-expected utility settings. Ai and Bansal (2018) provide a general characterization of preferences that give rise to an announcement premium holding constant the type and quantity of risk. It remains a question for future research whether all such preferences imply our documented increase in option-implied concavity, and whether they do so for the same or for different reasons.

References

Ai, H., Bansal, R., 2018. Risk preferences and the macro announcement premium. Econometrica 86, 1383–1430.

Ai, H., Bansal, R., Guo, H., Yaron, A., 2020. Identifying preference for early resolution from asset market data. Working paper.

Ait-Sahalia, Y., Lo, A., 2000. Nonparametric risk management and implied risk aversion. Journal of Econometrics 94, 9–51.

Anderson, E., Ghysels E., Juergens J. L., 2009. The impact of risk and uncertainty on expected returns. Journal of Financial Economics 94, 233–263.

Antoniou, C., Harris R., Zhang R., 2014. Ambiguity aversion and stock market participation: An empirical analysis. Working paper.

Backus, D., Chernov, M., Martin, I., 2011. Disasters implied by equity index options. Journal of Finance 66, 1969–2012.

Banz, R., Miller, M., 1978. Prices for state-contingent claims: some estimates and applications. Journal of Business 51, 653–672.

Bliss, R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. Journal of Banking and Finance 26, 381–422.

Bliss, R., Panigirtzoglou, N., 2004. Option implied risk aversion estimates. Journal of Finance 59, 407–446.

Breeden, D. T., Litzenberger, S., 1978. Prices of state-contingent claims implicit in option prices. Journal of Business 51, 621–651.

Brenner, M., Izhakian Y., 2011. Asset prices and ambiguity: Empirical evidence. Working paper.

Brenner, M., Izhakian Y., 2012. Pricing systematic ambiguity in capital markets. Working paper.

Cieslak, A., Morse, A., Vissing-Jorgensen, A., 2019. Stock returns over the FOMC cycle. Journal of Finance 74, 2201–2248.

Collard, F., Mukerji, S., Sheppard, K., Tallon J-M., 2013. Ambiguity and the historical equity premium. Working paper.

Diamond, D., Verrecchia, R. E., 1981. Information aggregation in a noisy rational expectations economy. Journal of Financial Economics 9, 221–235.

Drechsler, I., 2013. Uncertainty, Time-varying fear, and asset prices. Journal of Finance 68, 1843–1889.

Ellsberg, D., 1961. Risk, ambiguity, and the Savage axioms. Quarterly Journal of Economics 75, 643–669.

Epstein, L. G., Farhi, E., Stralezcki, T., 2014. How much would you pay to resolve long-run risk? American Economic Review 104, 2680–2697.

Epstein, L. G., Ji, S., 2013. Ambiguous volatility and asset pricing in continuous time. Review of Financial Studies 26, 1740–1786.

Epstein, L. G., Schneider, M., 2003. IID: Independently and indistinguishably distributed. Journal of Economic Theory 113, 32–50.

Epstein, L. G., Schneider, M., 2007. Learning under ambiguity. Review of Economic Studies 74, 1275–1303.

Epstein, L. G., Schneider, M., 2008. Ambiguity, information quality, and asset pricing. Journal of Finance 63, 197–228.

Epstein, L. G., Zin, S., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns. Econometrica 57, 937-969.

Gilboa, I., Schmeidler, D., 1989. Maximin expected utility with non-unique prior. Journal of Mathematical Economics 18, 141–153.

Figlewski, S., 2008. Estimating the implied risk neutral density for the U.S. market portfolio. In: *Volatility and Time Series Econometrics: Essays in honor of Robert F. Engle*, ed. Bollerslev, T., Oxford, Oxford University Press.

French, K. R., Poterba, J. M., 1991. Investor diversification and international equity markets. American Economic Review, Papers and Proceedings 81, 222–226.

Gajdos, T., Hayashi, T., Tallon, J.-M., Vergnaud, J.-C., 2008. Attitude toward imprecise information. Journal of Economic Theory 140, 27–65.

Gallant, R., Jahan-Parva, M. R., Liu, H., 2015. Measuring ambiguity aversion. Working paper.

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

Gilboa, I., Schmeidler, D., 1989. Maximin expected utility with non-unique prior. Journal of Mathematical Economics 18, 141–153.

Gollier, C., 2011. Portfolio choices and asset prices: The comparative statics of ambiguity aversion. Review of Economic Studies 78, 1329–1344.

Grossman, S., Stiglitz, J., 1980. On the impossibility of informationally efficient markets. American Economic Review 70, 393–408.

Hansen, L. P., Sargent T. J., 2011. Robust control and model uncertainty, American Economic Review 91, 60–66.

Hens, T., Reichlin, C., 2013. Three solutions to the pricing kernel puzzle. Review of Finance 17, 1065–1098.

Illeditsch, P., 2011. Ambiguous information, portfolio inertia, and excess volatility. Journal of Finance 66, 2213–2247.

Illeditsch, P., Ganguli, J., Condie, S., 2021. Information inertia. Journal of Finance 76, 443–479.

Izhakian, Y., Yermack, D., 2014. Risk, ambiguity, and the exercise of executive stock options. Working paper.

Jackwerth, J. C., 2000. Recovering risk aversion from option prices and realized returns. Review of Financial Studies 13, 433-451.

Jeong, D., Kim, H., Park, J. Y., 2015. Does ambiguity matter? Estimating asset pricing models with a multiple-priors recursive utility. Journal of Financial Economics 115, 361–382.

Ju, N., Miao, J., 2012. Ambiguity, learning, and asset returns. Econometrica 80, 559–591.

Keynes, J. M., 1936. The general theory of employment, interest and money. Macmillan, London, United Kingdom.

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

Knight, F. H., 1921. Risk, uncertainty and profit. Hart, Schaffner & Marx, Houghton Mifflin Company, Boston, MA.

Lucca, D., Moench, E., 2015. The Pre-FOMC announcement drift. Journal of Finance 76, 443-479.

Maccheroni, F., Marinacci M., Rustichini A., 2006. Ambiguity aversion, robustness and the variational representation of preferences. Econometrica 74, 1447–1498.

MacKinnon, J. G., 2009. Bootstrap hypothesis testing. In: Handbook of Computational Econometrics (edited by Belsley, D. A., and Kontoghiorghes, E. J.). John Wiley & Sons, Ltd.

Maenhout, P., 2004. Robust portfolio rules and asset pricing. Review of Financial Studies 17, 951–983.

Merton, R. C., 1980. On estimating the expected return on the market: An exploratory investigation. Journal of Financial Economics 8, 323–361.1383–1430.

Panigirtzoglou, N., Skiadopoulos, G., 2004. A new approach to modeling the dynamics of implied distributions: Theory and evidence from the S&P 500 options. Journal of Banking & Finance 28, 1499–1520.

Parker, J. A., 2001. The consumption risk of the stock market. Brookings papers on economic activity 2, 279–348.

Pastor, L., Veronesi, P., 2013. Political uncertainty and risk premia. Journal of Financial Economics 110, 520–545.

Rosenberg, J. V., Engle, R. F., 2002. Empirical pricing kernels. Journal of Financial Economics 64, 341–372.

Ross, S., 2015. The recovery theorem. Journal of Finance 70, 615–648.

Savor, P., Wilson, M., 2013. How much do investors care about macroeconomic risk? Evidence from scheduled economic announcements. Journal of Financial and Quantitative Analysis 48, 343–375.

Savor, P., Wilson, M., 2014. Asset pricing: A tale of two days. Journal of Financial Economics 113, 171–201.

Shi, Z., 2013. Time-varying ambiguity and asset pricing puzzles. Working paper.

Shimko, D., 1993. The bounds of probability. Risk 6, 33–37.

Sichert, T., 2020. The pricing kernel is U-shaped. Working paper.

Skiadas, C., 2013. Smooth ambiguity aversion toward small risks and continuous-time recursive utility. Journal of Political Economy 121, 775–792.

Strzalecki, T., 2013. Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion. Econometrica 81, 1039–1074.

Thimme, J., Völkert, C., 2015. Ambiguity in the cross-section of expected returns: An empirical assessment. Journal of Business & Economic Statistics 33, 418–429.

Trautmann, S. T., Vieider, F. M., Wakker, P.P., 2011. Preference reversals for ambiguity aversion. Management Science 57, 1320–1333.

Ulrich, M., 2013. Inflation ambiguity and the term structure of U.S. government bonds. Journal of Monetary Economics 60, 295–309.

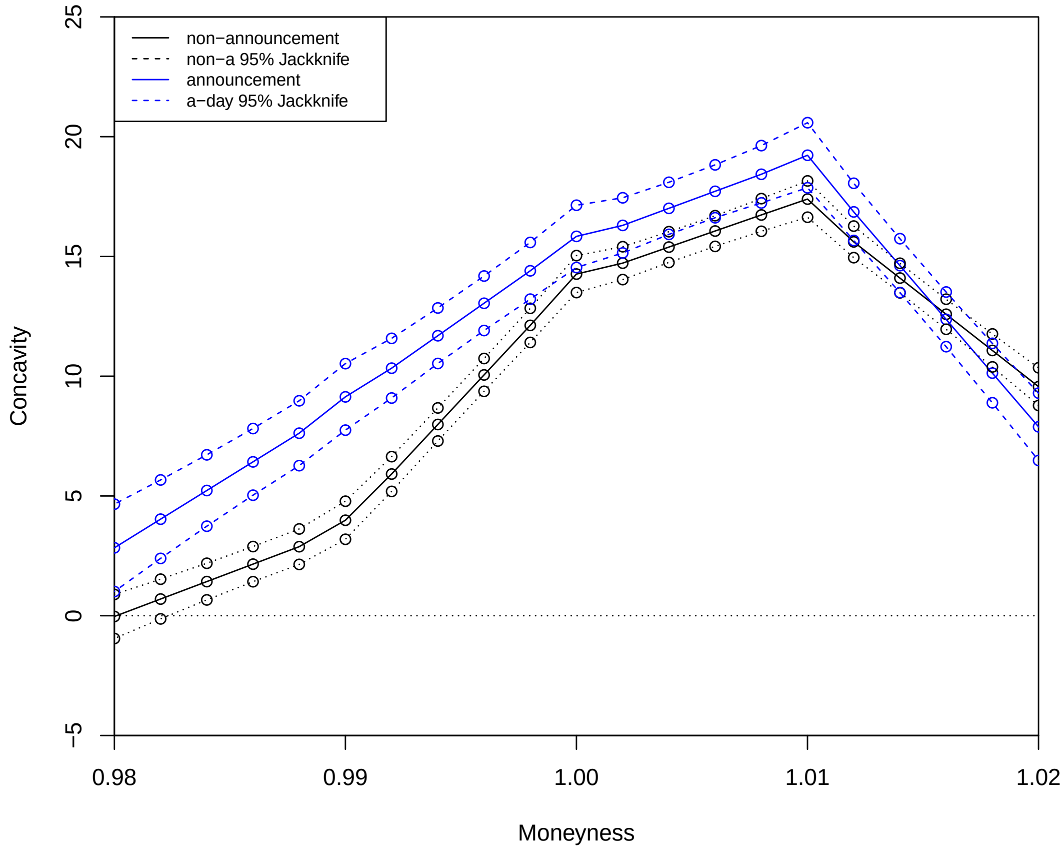

Figure 1: This figure shows the implied local (non-parametric) concavity on non-announcement and announcement days (FOMC, PPI or employment), together with 95% Jackknife error bands. The underlying state-price densities (SPDs) are extracted for each set of S&P 500 options based on the respective market parameters valid at the time and are compared to the physical probability density reflecting the historical average return over a period consistent with the SPD. The resulting ratio of the SPD and the physical probability density is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). The concavity curves are finally aggregated per type of day via the arithmetic mean centered around the respective at-the-money value. The sample period is 1996–2016.

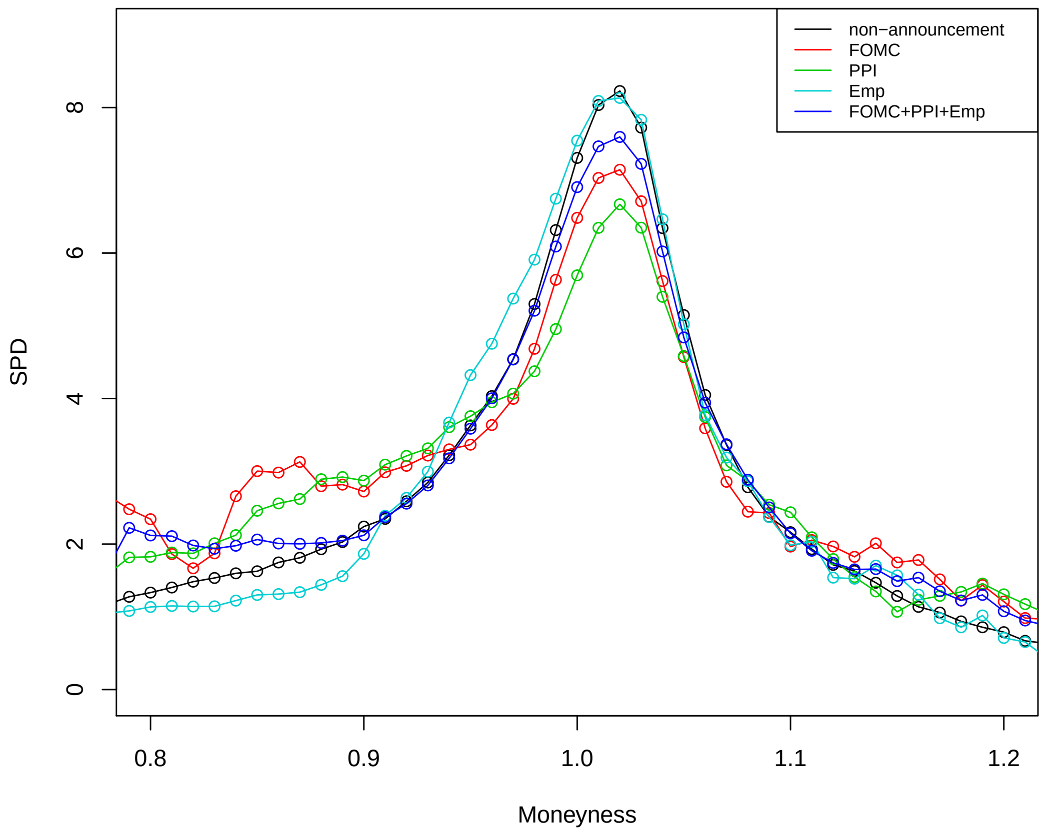

Figure 2: The figure shows the implied state-price densities (SPDs) for announcement days (FOMC, PPI, unemployment, and all three types jointly) and non-announcement days in the moneyness range 0.8 to 1.2. The SPDs are extracted based on the respective market parameters at the time and are subsequently aggregated via the arithmetic mean. The sample period is 1996–2016.

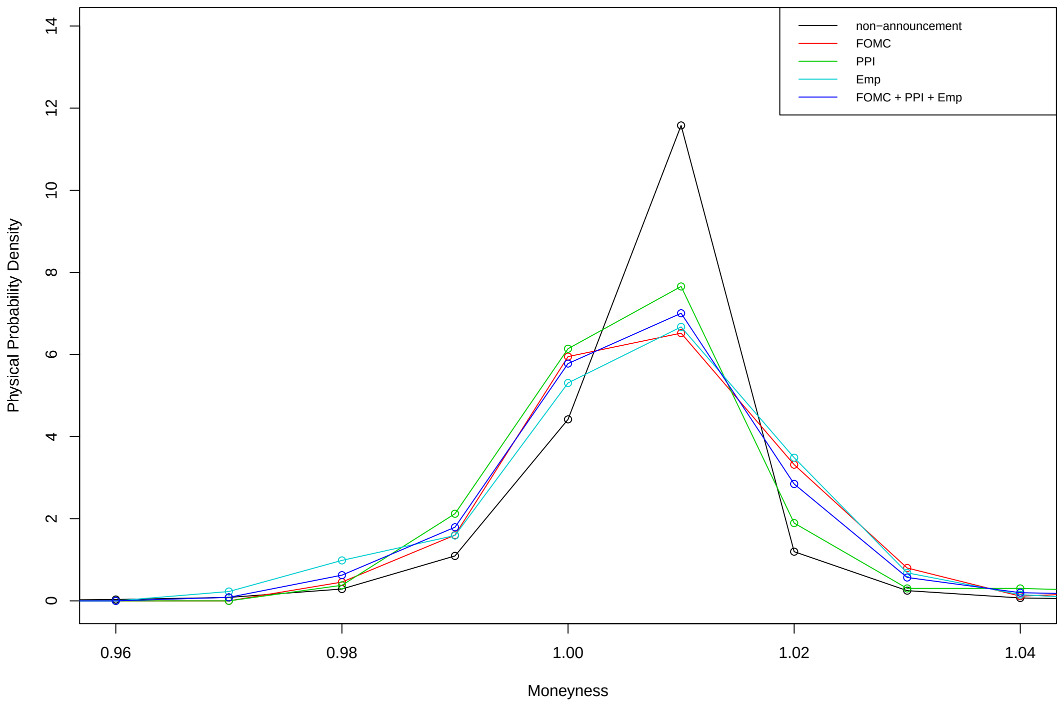

Figure 3: This figure shows the physical probability density of S&P 500 returns over a 1-day horizon. We sort daily returns by type of trading day (announcement and non-announcement) and compute their frequencies in 0.01 step moneyness buckets, which are then converted into the plotted probability density. The sample period is 1996–2016.

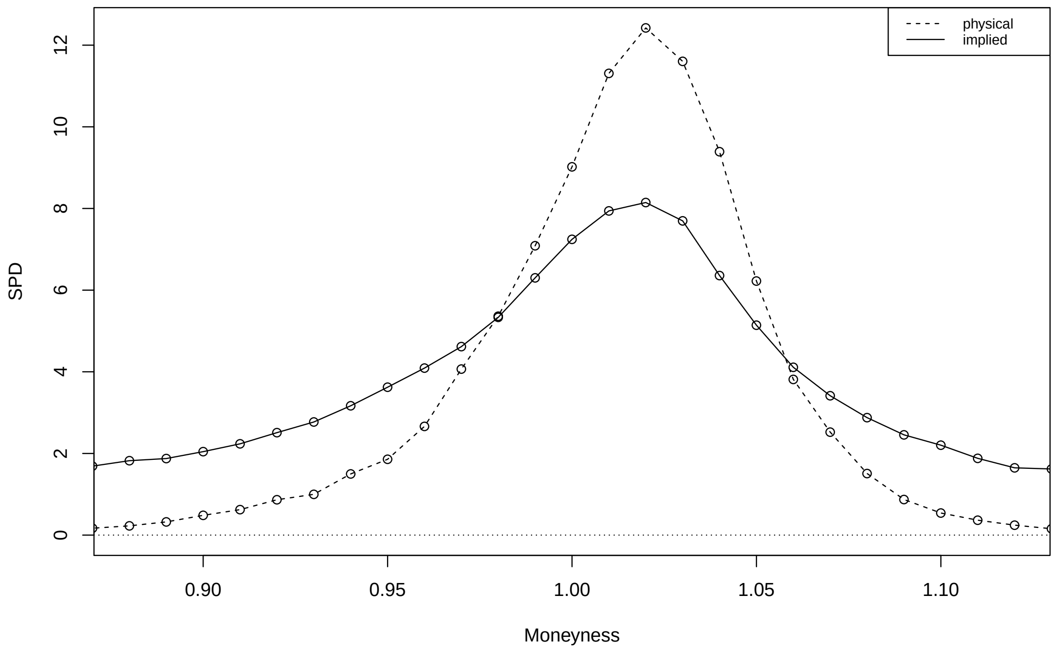

Figure 4: This figure plots the average physical probability density vs. the implied state-price density (SPD) over the 1996–2016 period. The employed physical probability densities of observed returns are based on average historical returns consistent with the maturity of the respective implied SPDs (on average 44 days).

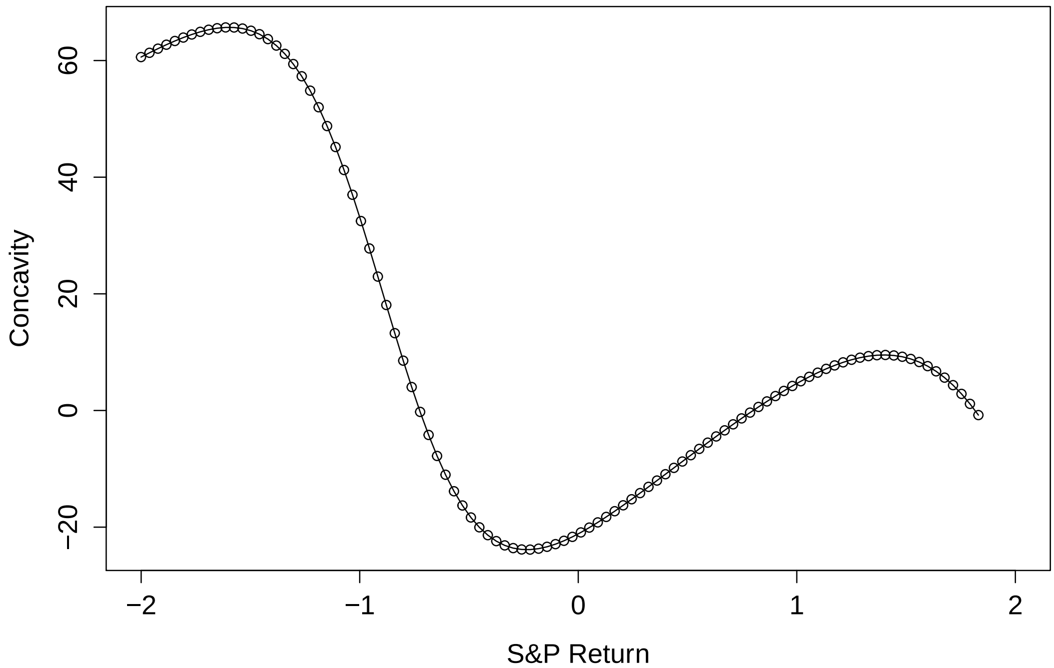

Figure 5: This figure shows the typical form of the concavity index for the S&P 500 over ω in a range of plus/minus two standard deviations (i.e., the graph is centered around the expected return plus error bands). The illustrative example shows the results for March 9, 1999 (S&P 500 index level of 1279.84) using options that expire on March 20, 1999.

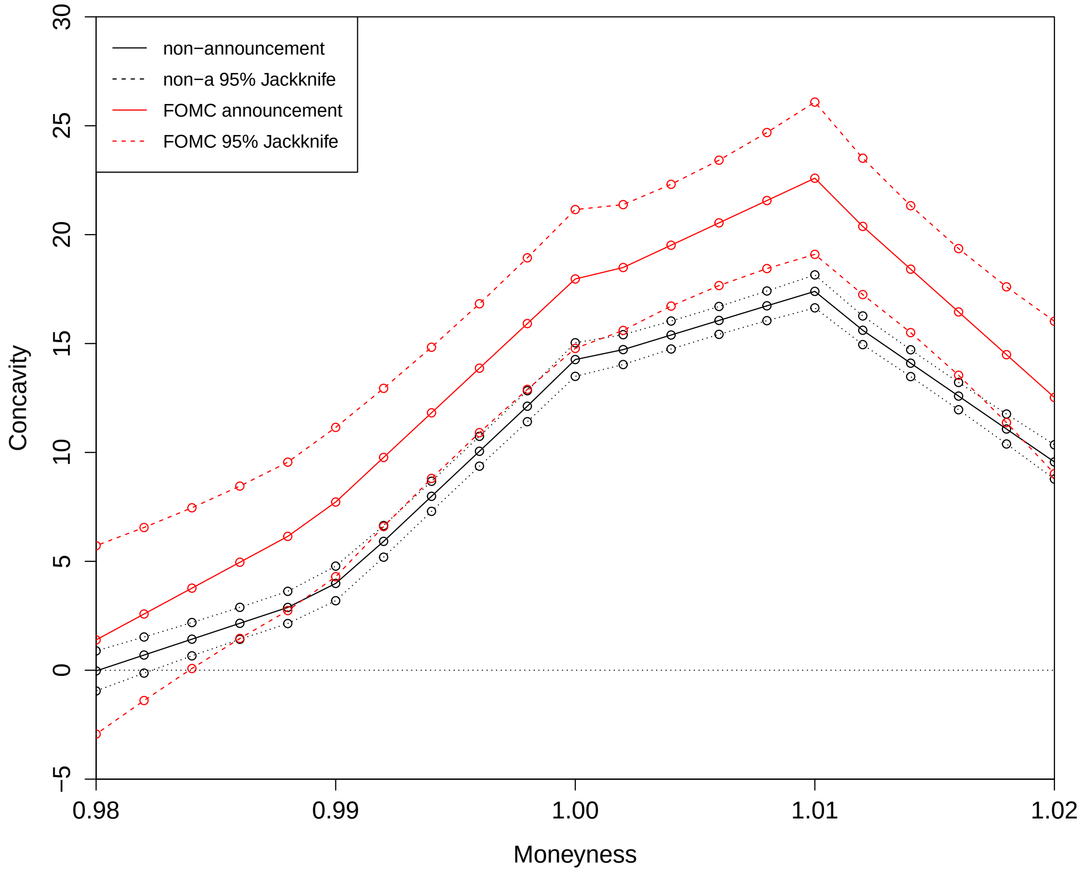

Figure 6: This figure shows the implied local (non-parametric) concavity on non-announcement and FOMC announcement days, together with 95% Jackknife error bands. The underlying state-price densities (SPDs) are extracted for each set of S&P 500 options based on the respective market parameters valid at the time and are compared to the physical probability density reflecting the historical average return over a period consistent with the SPD. The resulting ratio of the SPD and the physical probability density is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). The concavity curves are finally aggregated per type of day via the arithmetic mean centered around the respective at-the-money value. The sample period is 1996–2016.

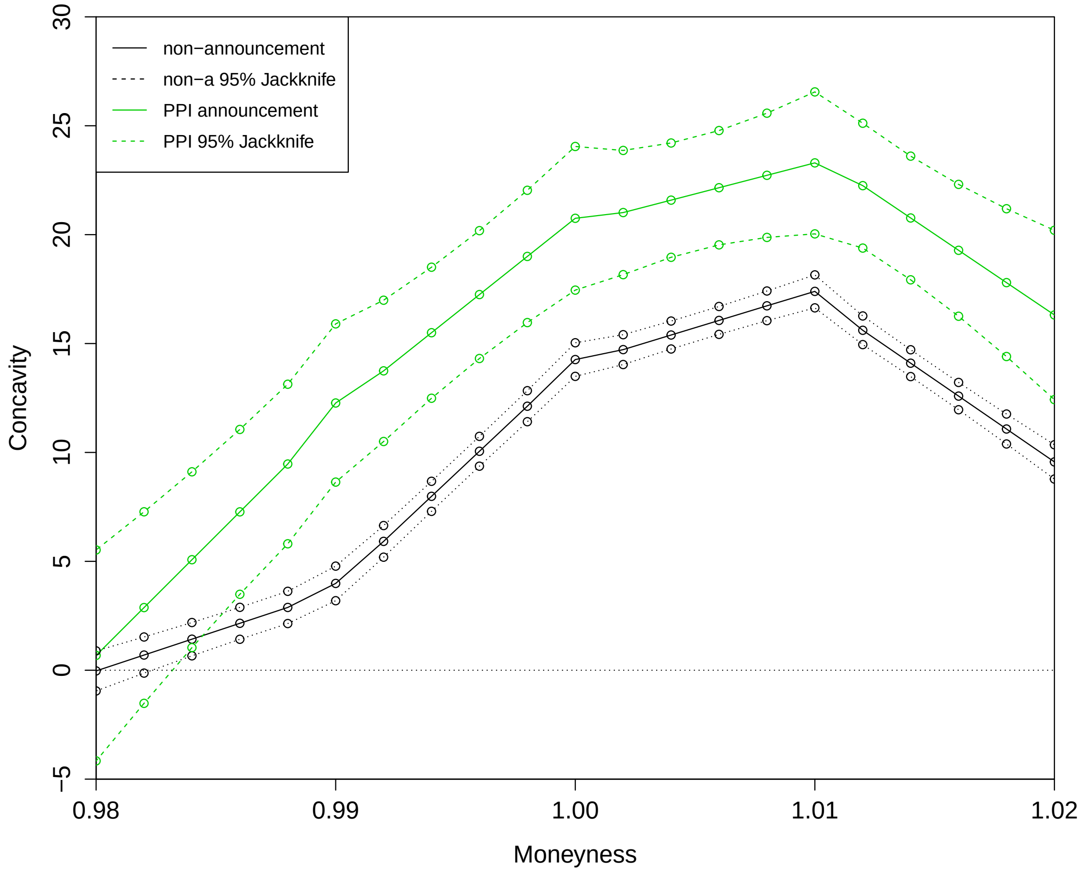

Figure 7: This figure shows the implied local (non-parametric) concavity on non-announcement and PPI announcement days, together with 95% Jackknife error bands. The underlying state-price densities (SPDs) are extracted for each set of S&P 500 options based on the respective market parameters valid at the time and are compared to the physical probability density reflecting the historical average return over a period consistent with the SPD. The resulting ratio of the SPD and the physical probability density is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). The concavity curves are finally aggregated per type of day via the arithmetic mean centered around the respective at-the-money value. The sample period is 1996–2016.

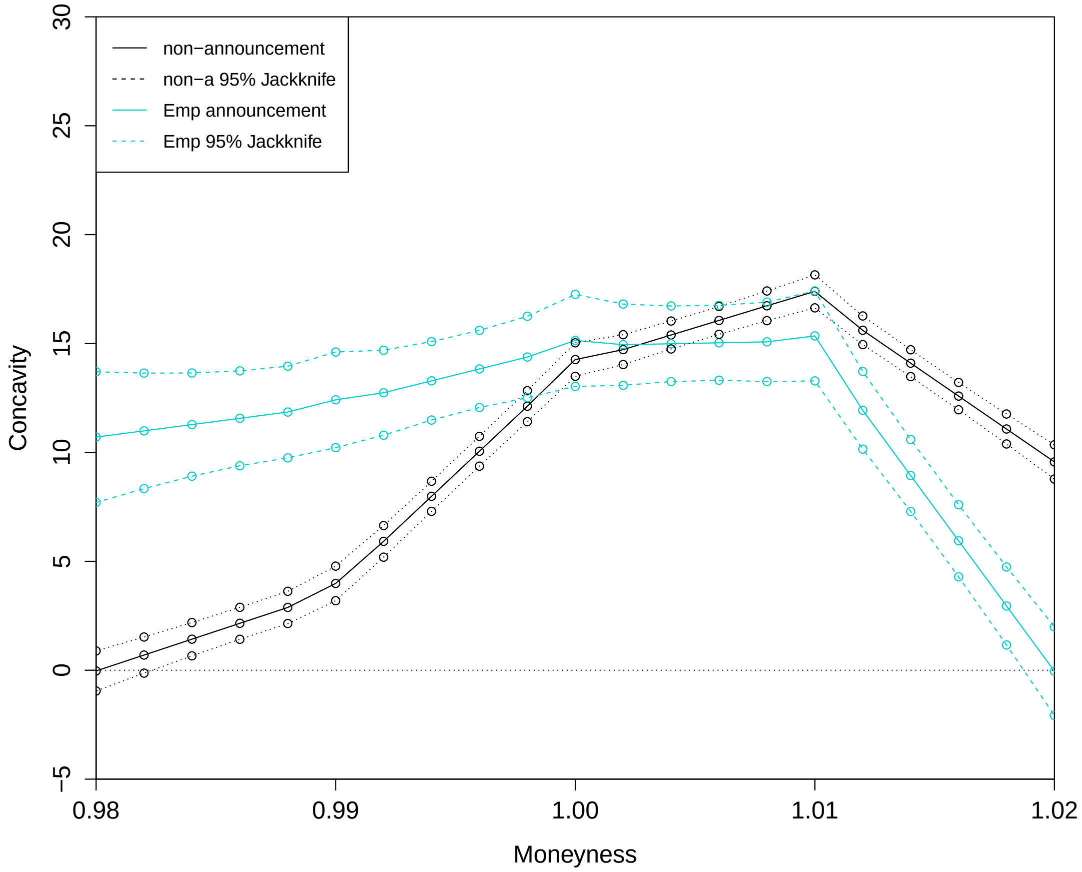

Figure 8: This figure shows the implied local (non-parametric) concavity on non-announcement and employment announcement days, together with 95% Jackknife error bands. The underlying state-price densities (SPDs) are extracted for each set of S&P 500 options based on the respective market parameters valid at the time and are compared to the physical probability density reflecting the historical average return over a period consistent with the SPD. The resulting ratio of the SPD and the physical probability density is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). The concavity curves are finally aggregated per type of day via the arithmetic mean centered around the respective at-the-money value. The sample period is 1996–2016.

Table 1
Implied Volatility Surface

T= 0-3 months

Moneyness	0.97–0.99	0.99–1.01	1.01–1.03
Non-Ann.	0.187	0.172	0.161
FOMC	0.192	0.190	0.193
PPI	0.189	0.173	0.164
Emp	0.186	0.172	0.162
All Ann.	0.188	0.176	0.168

T= 4–6 months

Moneyness	0.97–0.99	0.99–1.01	1.01–1.03
Non-Ann.	0.189	0.180	0.174
FOMC	0.189	0.182	0.182
PPI	0.191	0.181	0.174
Emp	0.186	0.178	0.172
All Ann.	0.189	0.180	0.175

T= 7–12 months

Moneyness	0.97–0.99	0.99–1.01	1.01–1.03
Non-Ann.	0.194	0.187	0.185
FOMC	0.195	0.185	0.190
PPI	0.192	0.186	0.186
Emp	0.195	0.191	0.188
All Ann.	0.194	0.187	0.188

This table presents the mean implied volatility for different areas of the volatility surface separated by type of announcement day. The table aggregates the available option contracts within pre-defined bands of moneyness and maturity T and uses the mid-price implied volatility based on OptionMetrics data. The sample period is 1996–2016.

Table 2
Bid/Ask Spreads

Panel A: Announcement Days

Type	Mean Spread	Std. Err.	t-test	KS	MW
Non-Ann.	12.31%	0.02%	-	-	-
FOMC	12.74%	0.07%	0.0%	0.0%	0.0%
PPI	12.54%	0.06%	0.0%	0.0%	0.0%
Emp.	12.42%	0.06%	3.3%	0.0%	0.0%
All Ann.	12.55%	0.04%	0.0%	0.0%	0.0%

Panel B: One Day Prior to Announcements

Type	Mean Spread	Std. Err.	t-test	KS	MW
Non-Ann.	12.30%	0.01%	-	-	-
FOMC	12.39%	0.07%	10.1%	0.0%	0.1%
PPI	12.23%	0.06%	83.8%	57.1%	100%
Emp.	12.51%	0.06%	0.1%	3.0%	4.1%
All Ann.	12.41%	0.04%	0.3%	15.4%	5.9%

This table shows the average bid-ask spread, defined as (ask price - bid price)/bid price, of the options included in our sample. Panel A focuses on announcement days and Panel B on the trading day prior to different announcements. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The sample period is 1996–2016. The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Table 3
State-Price Density

Type	Mean	Std. Err.	Skewness	Exc. Kurt.	t-test	KS	MW
Non-Ann.	8.25	0.08	-0.73	-0.27	-	-	-
FOMC	7.06	0.35	-0.40	-1.16	0.1%	0.4%	0.1%
PPI	6.18	0.30	-0.23	-1.19	0.0%	0.0%	0.0%
Emp.	8.16	0.21	-0.75	-0.34	34.0%	71.0%	42.0%
All Ann.	7.56	0.13	-0.52	-0.55	0.0%	0.0%	0.0%

This table presents the mean and standard error of the implied state-price density at its maximum on announcement and non-announcement days averaged across our 1996–2016 sample period. The table further reports the skewness and excess kurtosis. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the non-announcement-day average is higher (t-test), the non-announcement-day distribution stochastically dominates (KS), and the non-announcement-day distribution differs by a positive location shift (MW) compared to that on announcement days.

Table 4
Daily S&P 500 Returns

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	1.6 bp	1.7 bp	-	-	-
FOMC	32.2 bp	9.3 bp	0.1%	0.2%	0.2%
PPI	7.2 bp	7.9 bp	24.4%	14.6%	24.1%
Emp.	11.6 bp	7.7 bp	10.2%	0.3%	3.9%
All Ann.	15.1 bp	4.8 bp	0.4%	0.0%	0.2%

This table presents the mean and standard error of daily S&P 500 returns on announcement and non-announcement days. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days. The sample period is 1996–2016.

Table 5
Representative Agent Utility Functions

$$RRA = -\frac{S_T U''(S_T)}{U'(S_T)}$$

	Power	Exponential
$U(S_T)$	$\frac{S_T^{(1-\gamma)}-1}{1-\gamma}$	$-\frac{e^{-\gamma S_T}}{\gamma}$
$U(S_T)'$	$S_T^{-\gamma}$	$e^{-\gamma S_T}$
RRA	γ	γS_T

This table provides an overview of the representative agent's utility functions used to condense the concavity index into a single comparable measure. We consider the power utility and exponential utility function and calculate the associated relative risk aversion (RRA).

Table 6

Panel A: Concavity Index - Power Utility Function

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	9.34	0.17	-	-	-
FOMC	12.55	0.86	0.0%	2.0%	0.1%
PPI	15.06	0.87	0.0%	0.0%	0.0%
Emp	11.32	0.54	0.0%	0.0%	0.1%
All Ann.	11.33	0.31	0.0%	0.0%	0.0%

This table presents the mean and standard error of the γ parameters resulting from the fit of a power utility function to the concavity index on announcement and non-announcement days. For this process, we calculate the ratio of the option-implied state-price densities (SPD) and physical probability densities of S&P 500 returns that is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). We use a moneyness range of 0.98 to 1.02. The sample period is 1996–2016. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Panel B: Distribution of Bootstrap Estimates of γ
Power Utility Function

Type	t-test	KS	MW
FOMC	0.8%	2.8%	0.9%
PPI	0.0%	0.1%	0.0%
Emp	0.8%	2.5%	1.7%
All Ann.	0.0%	0.0%	0.0%

Based on 1,000,000 bootstrap samples (with replacement) of the same size as the original sample, we compute the distribution of γ parameters of the fit of a power utility function to the concavity index on announcement and non-announcement days. The bootstrapping approach starts with the resampling of the physical return probability distribution and subsequently aggregates it with the SPD to compute concavity and γ parameter values, which are further resampled with replacement. The table presents the bootstrapped p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Table 7
Panel A: Concavity Index - Exponential Utility Function

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	9.40	0.17	-	-	-
FOMC	12.62	0.85	0.0%	1.7%	0.1%
PPI	15.13	0.87	0.0%	0.0%	0.0%
Emp	11.30	0.54	0.0%	0.1%	0.1%
All Ann.	11.38	0.31	0.0%	0.0%	0.0%

This table presents the mean and standard error of the γ parameters of the fit of a exponential utility function to the concavity index on announcement and non-announcement days. For this process, we calculate the ratio of the option-implied state-price densities (SPD) and physical probability densities of S&P 500 returns that is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). We use a moneyness range of 0.98 to 1.02. The sample period is 1996–2016. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Panel B: Distribution of Bootstrap Estimates of γ
Exponential Utility Function

Type	t-test	KS	MW
FOMC	0.6%	2.0%	0.7%
PPI	0.0%	0.0%	0.0%
Emp	1.0%	3.2%	2.3%
All Ann.	0.0%	0.0%	0.0%

Based on 1,000,000 bootstrap samples (with replacement) of the same size as the original sample, we compute the distribution of γ parameters of the fit of an exponential utility function to the concavity index on announcement and non-announcement days. The bootstrapping approach starts with the resampling of the physical return probability distribution and subsequently aggregates it with the SPD to compute concavity and γ parameter values, which are further resampled with replacement. The table presents the bootstrapped p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Internet Appendix

Appendix A

In this appendix, we explain how smooth ambiguity aversion is first-order for large, or point-in-time, risks, unlike what Skiadas shows for small risks such as Brownian motions or Poisson jumps.

1.1. Expected utility under prior ρ

Given a prior ρ, expected utility $E_\rho[U(\widetilde{B})]$ can under standard assumptions be expressed as a Taylor expansion using the derivatives of $U(\widetilde{B})$ around $B = 0$ where $U(0) = U_0$ etc.:

$$\begin{aligned} E_\rho[U(\widetilde{B})] &= U_0 + U_0' E_\rho[\widetilde{B}] + \frac{1}{2} U_0'' E_\rho[(\widetilde{B})^2] + \frac{1}{3!} U_0''' E_\rho[(\widetilde{B})^3] + ... \qquad (A.1) \\ &= U_0 + \sum_{n=1}^{\infty} \frac{U_0^n}{n!} E_\rho[(\widetilde{B})^n]. \end{aligned}$$

Clearly for Brownian risks this approximates, for small time intervals, to

$$\begin{aligned} E_\rho[U(\widetilde{B})] &= U_0 + \left(U_0' \rho + \frac{1}{2} U_0'' \right) h + o(h) \qquad (A.2) \\ &= \left(U_0 + \frac{1}{2} U_0'' h \right) + U_0' h \rho + o(h) \\ &\approx U_0 + b_1 h + c_1 h \rho, \qquad (A.3) \end{aligned}$$

where the constant terms b_1 and c_1 are respectively defined as $b_1 = \frac{1}{2} U_0''$ and $c_1 = U_0'$.

For Poissonian risks this becomes

$$\begin{aligned} E_\rho[U(\widetilde{B})] &= U_0 + \quad U_0' \rho + (1 + \rho) \sum_{n=2}^{\infty} \frac{U_0^n}{n!} \Bigg) h + o(h) \qquad (A.4) \\ &= \quad U_0 + h \sum_{n=2}^{\infty} \frac{U_0^n}{n!} \Bigg) + \left(U_0' + \sum_{n=2}^{\infty} \frac{U_0^n}{n!} \right) h \rho + o(h) \\ &\approx U_0 + b_2 h + c_2 h \rho. \qquad (A.5) \end{aligned}$$

Similarly, the constants b_2 and c_2 are

$$b_2 = \sum_{n=2}^{\infty} \frac{U_0^n}{n!} \qquad \text{and} \qquad c_2 = c_1 + b_2.$$

For large risks this becomes

$$E_\rho[U(\widetilde{B})] = U_0 + \sum_{n=1}^{\infty}\left(\frac{U_0^n \varepsilon^n}{n!}\right)(\rho 1_{n \text{ is } odd}(n) + (1 - 1_{n \text{ is } odd}(n))) \quad \text{(A.6)}$$

$$= U_0 + \sum_{n=1}^{\infty}\left(\frac{U_0^n \varepsilon^n}{n!}\right)(1 - 1_{n \text{ is } odd}(n)) + \left(\sum_{n=1}^{\infty}\left(\frac{U_0^n \varepsilon^n}{n!}\right)1_{n \text{ is } odd}(n)\right)\rho \quad \text{(A.7)}$$

$$= U_0 + b(\varepsilon) + c(\varepsilon)\rho. \quad \text{(A.8)}$$

where the constant terms $b(\varepsilon)$ and $c(\varepsilon)$ are

$$b(\varepsilon) = \sum_{n \text{ even}}^{\infty}\left(\frac{U_0^n \varepsilon^n}{n!}\right) \qquad \text{and} \qquad c(\varepsilon) = \sum_{n \text{ odd}}^{\infty}\left(\frac{U_0^n \varepsilon^n}{n!}\right).$$

The key point to note here is that the uncertainty parameter ρ enters multiplicatively with h, the time interval, in the small risk cases but not in the case of large risk. Also for standard utility functions such as CARA or CRRA $b(\varepsilon) < 0$ while $c(\varepsilon) > 0$.

1.2. KMM smooth ambiguity aversion and each type of risk

The last step is to approximate $\phi(E_\rho[U(\widetilde{B}))$ and to take its expectation over the compound prior $E[\widetilde{\rho}]$.

Again we use Taylor expansions around the expected values of expected utility over the compound prior. Define $\phi_0 = \phi\left(E\left[E_\rho[U(\widetilde{B})]\right]\right)$, i.e. the value of the function $\phi(.)$ evaluated at $E\left[E_\rho[U(\widetilde{B})]\right]$, and define ϕ_0' and ϕ_0'' as the values of the same function's first and second derivatives evaluated at the same point. Then a second-order approximation of expected welfare given smooth ambiguity aversion is:

$$E\left[\phi\left(E_\rho[U(\widetilde{B})]\right)\right]$$

$$= \phi_0 + \phi_0' E\left[E_\rho[U(\widetilde{B})] - E\left[E_\rho[U(\widetilde{B})]\right]\right] + \frac{1}{2}\phi_0'' E\left[\left(E_\rho[U(\widetilde{B})] - E\left[E_\rho[U(\widetilde{B})]\right]\right)^2\right] + ...$$

$$= \phi_0 + \frac{1}{2}\phi_0'' Var\left[E_\rho[U(\widetilde{B})]\right] + ...$$

And we can then define the importance of ambiguity aversion $\phi(.)$ for small and large risks.

For Brownian risks:

$$\begin{aligned} E\left[\phi\left(E_\rho[U(\widetilde{B})]\right)\right] &= \phi_0 + \frac{1}{2}\phi_0^{''} Var\left[E_\rho[U(\widetilde{B})]\right] \\ &= \phi_0 + \frac{1}{2}\phi_0^{''} c_1^2 h^2 Var[\widetilde{\rho}] \\ &= \phi_0 + o(h). \end{aligned}$$

For Poissonian risks:

$$\begin{aligned} E\left[\phi\left(E_\rho[U(\widetilde{B})]\right)\right] &= \phi_0 + \frac{1}{2}\phi_0^{''} Var\left[E_\rho[U(\widetilde{B})]\right] \\ &= \phi_0 + \frac{1}{2}\phi_0^{''} c_2^2 h^2 Var[\widetilde{\rho}] \\ &= \phi_0 + o(h). \end{aligned}$$

For large risks this becomes:

$$\begin{aligned} E\left[\phi\left(E_\rho[U(\widetilde{B})]\right)\right] &= \phi_0 + \frac{1}{2}\phi_0^{''} Var\left[E_\rho[U(\widetilde{B})]\right] + ... \\ &= \phi_0 + \frac{1}{2}\phi_0^{''} c(\varepsilon)^2 Var[\widetilde{\rho}] + ... \end{aligned}$$

Comparison of these three equations shows that only large risks can affect welfare for an ambiguity-averse agent over small enough time intervals. The term $\frac{1}{2}\phi_0^{''} Var\left[E_\rho[U(\widetilde{B})]\right]$ (which is negative, since $\phi(.)$ is concave) in particular decreases expected welfare for large risks, where the concavity of $\phi(.)$ has a first-order effect on expected utility.

We now derive our concavity index using results from Section 2.

The expressions for expected utility under a given prior ρ above can be used to derive expressions for w_ρ for each type of risk. Thus, for Brownian risks

$$w_\rho = \frac{\phi'(U_0 + b_1 h + c_1 h\rho)}{E[\phi'(U_0 + b_1 h + c_1 h\rho)]}.$$

Again approximating $\phi'(.)$ around the compound expectation of its argument, and imposing that the compound expectation of ρ, $E[\rho] = 0$ gives

$$w_\rho = \frac{\phi_0' + \phi_0'' c_1 h\rho}{\phi_0'} = 1 + \frac{\phi_0''}{\phi_0'} c_1 h\rho.$$

For Poissonian risks an exactly parallel argument gives

$$w_\rho = 1 + \frac{\phi_0''}{\phi_0'} c_2 h\rho$$

and finally for large risks these weights are independent of the time interval:

$$w_\rho = 1 + \frac{\phi_0''}{\phi_0'} c(\varepsilon)\rho.$$

Then for Brownian risks

$$\widehat{p}_1 = E\left[\left(1 + \frac{\phi_0''}{\phi_0'} c_1 h\rho\right) \frac{1}{2}\left(1 + \rho\sqrt{h}\right)\right] = \frac{1}{2} + o(h) = p_1 = \widehat{p}_0 = p_0$$

and similarly for Poissonian risks

$$\widehat{p}_1 = E\left[\left(1 + \frac{\phi_0''}{\phi_0'} c_2 h\rho\right) (1+\rho)h)\right] = h + o(h) = p_1$$

and

$$\widehat{p}_0 = E\left[\left(1 + \frac{\phi_0''}{\phi_0'} c_2 h\rho\right) (1 - (1+\rho)h)\right] = 1 - h + o(h) = p_0.$$

While for large risks

$$\begin{aligned}
\widehat{p}_1 &= E\left[\left(1 + \frac{\phi_0''}{\phi_0'} c(\varepsilon)\rho\right) \frac{1}{2}(1+\rho)\right] \\
&= p_1 + \frac{1}{2}\frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho]
\end{aligned}$$

and

$$\widehat{p}_0 = p_0 - \frac{1}{2}\frac{\phi_0''}{\phi_0'} c(\varepsilon) Var[\rho].$$

The ambiguity-weighting derived in Gollier has no first-order effect on risk- and ambiguity-neutral probabilities, and therefore state prices, for small risks, but does have a first-order effect on such probabilities for large risks, increasing the state price for the 'bad' state $s = 0$ and lowering the state price for the 'good' state $s = 1$.

1.2.1. Concavity index for large risks

From Gollier's result about state prices we have

$$\frac{d}{dW}\left(\frac{\widehat{p}_s}{p_s}u'(W_s)\right) = u''(W_s)\frac{\widehat{p}_s}{p_s} + u'(W_s)\frac{d}{dW}\left(\frac{\widehat{p}_s}{p_s}\right) \tag{A.9}$$

and thus we can compute a concavity index $CI(\omega)$ as follows:

$$\begin{aligned} CI(s) &\equiv -\frac{W_s\frac{d}{dW}\left(\frac{\widehat{p}_s}{p_s}u'(W_s)\right)}{\left(\frac{\widehat{p}_\omega}{p_\omega}u'(W_\omega)\right)} = -\frac{W_s u''(W_s)}{u'(W_s)} - W_s\frac{\frac{d}{dW}\left(\frac{\widehat{p}_s}{p_s}\right)}{\frac{\widehat{p}_s}{p_s}} \\ &= \gamma(s) - W_s\frac{d}{dW_s}\ln\left(\frac{\widehat{p}_s}{p_s}\right). \end{aligned} \tag{A.10}$$

Clearly for Brownian and Poissonian risks the concavity index is just relative risk aversion, $\gamma(s)$, which is unaffected by ambiguity aversion ϕ_0'' since $\ln\left(\frac{\widehat{p}_s}{p_s}\right) = 0$ for all s and depends only on the underlying utility function $u(W)$. In other words, the concavity of preferences for ambiguity-averse representative agents facing small risks (Brownian or Poissonian) is identical to the concavity of preferences for expected utility (non-ambiguity-averse) agents using the compound prior. This is an implication of the central result in Skiadas (2013).

Skiadas' (2013) irrelevance results also extend to all variational preferences for the case of Brownian motions and to all entropic preferences (a subset of variational preferences), including Poisson risks.[40] Thus, as in the previous case, expected utility approximates such cases and consequently concavity is expected utility concavity.

For large risks, concavity of preferences is affected by ambiguity aversion. We first need to

[40] Non-entropic divergence preferences and Poisson risks are one of the exceptions to Skiadas (2013) general finding and are not first-order equivalent to expected utility.

define derivatives with respect to wealth for discrete states. The natural definition is

$$p'_s = \frac{p_1 - p_0}{W(1) - W(0)} = \frac{p_1 - p_0}{B(1) - B(0)} = \frac{p_0 - p_1}{B(0) - B(1)}$$

for both $s = 0$ and 1, and likewise

$$\widehat{p}'_s = \frac{\widehat{p}_1 - \widehat{p}_0}{B(1) - B(0)} = \frac{\widehat{p}_0 - \widehat{p}_1}{B(0) - B(1)}.$$

Then

$$CI(s) = \gamma(s) + (W_0 + B(s)) \left(\frac{1}{p_s} p'_s - \frac{1}{\widehat{p}_s} \widehat{p}'_s \right). \tag{A.11}$$

Since, $p_0 = p_1$ are equal under the unconditional compound expectation (recall $E[\rho] = 0$)

$$CI(s) = \gamma(s) - (W_0 + B(s)) \frac{\widehat{p}'_s}{\widehat{p}_s}$$

and consequently our concavity index, evaluated for each state price, is at $\omega = 1$:

$$\begin{aligned} CI(1) &= \gamma(1) - (W_0 + \varepsilon) \frac{\frac{\frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]}{2\varepsilon}}{\frac{1}{2} + \frac{1}{2} \frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]} \\ &= \gamma(1) - \frac{(W_0 + \varepsilon)}{\varepsilon} \frac{\frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]}{1 + \frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]} \end{aligned} \tag{A.12}$$

and at $\omega = 0$:

$$\begin{aligned} CI(0) &= \gamma(0) - (W_0 - \varepsilon) \frac{\frac{\frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]}{2\varepsilon}}{\frac{1}{2} - \frac{1}{2} \frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]} \\ &= \gamma(0) - \frac{(W_0 - \varepsilon)}{\varepsilon} \frac{\frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]}{1 - \frac{\phi''_0}{\phi'_0} c(\varepsilon) Var[\rho]}. \end{aligned} \tag{A.13}$$

Appendix B

We explain here in more detail our arguments in Section 7 using Gollier's (2011) closed-form solution for smooth ambiguity aversion preferences.

1.1. Return uncertainty

We begin with a simple model where investor faces uncertainty about returns. In the next section, we relate uncertainty about returns to underlying uncertainty about payoffs. We also show there how uncertainty becomes second-order for small risks by specifically considering the time interval, an issue from which we abstract here.

There is a representative investor who allocates her wealth at date 0, W_0, between a risk-free asset with net return R_f and the market portfolio with risky net return R_{MKT}. With a share of wealth α allocated to the market, her date 1 wealth is

$$W_1 = W_0\left[(1+R_f) + \alpha(R_{MKT} - R_f)\right]. \tag{B.1}$$

The investor believes market excess returns are normally distributed with mean $\widetilde{\theta}$ (unknown) and variance σ^2 (known). (The assumption that the variance of market returns is known is convenient, but also highly plausible. See Merton, 1980, for a discussion of the differences in estimating the mean and variance of returns.) She has CARA utility of wealth W_1 with absolute risk aversion A, given by

$$U(W_1) = -\frac{1}{A}\exp(-AW_1) \tag{B.2}$$

so that given an expected return of θ her expected utility is

$$E_\theta[U(W_1;\alpha)] = -\frac{1}{A}\exp\left(-AW_0((1+R_f)+\alpha\theta) + \frac{A^2}{2}\alpha^2\sigma^2\right). \tag{B.3}$$

The investor has a prior distribution over θ which is also normal $\theta \sim N(\mu, \sigma_\theta^2)$. She has constant relative ambiguity aversion, specifically:

$$\phi(U) = -\frac{(-U)^{1+\eta}}{1+\eta},$$

where $\eta \geq 0$ is her coefficient of relative ambiguity aversion ($\eta = 0$ corresponds to standard expected utility). Then her welfare under smooth ambiguity aversion is:

$$V(\alpha) = E[\phi(U(W_1;\alpha) \propto \exp\left(-A(1+\eta)W_0((1+R_f)+\alpha\mu) + \frac{A^2}{2}(1+\eta)\alpha^2\sigma^2 + \frac{A^2}{2}(1+\eta)^2\alpha^2\sigma_\theta^2\right). \tag{B.4}$$

The first order condition is then

$$\mu = A(\sigma^2 + (1+\eta)\sigma_\theta^2)\alpha^* \tag{B.5}$$

so that in equilibrium, with $\alpha^* = 1$,

$$\mu = A(\sigma^2 + (1+\eta)\sigma_\theta^2). \tag{B.6}$$

Applying our argument from Section 2, in short periods of time without an important scheduled announcement, σ_θ^2 is the prior variance of ρh, where h is the time interval. $\theta = \rho h$ and its variance are therefore $o(h)$ and we can treat them as zero. For short periods of time containing an important announcement, σ_θ^2 is the prior variance of $\rho\varepsilon$, invariant to the time interval. (Our more detailed model of payoff risk in the next section of this appendix lays out the argument step by step.) Denoting the per-period expected returns in non-announcement periods and announcement periods as μ^{NA} and μ^A, respectively:

$$\mu^{NA} = A\sigma^2 > 0 \tag{B.7}$$

while

$$\mu^A = \mu^{NA} + A(1+\eta)\sigma_\theta^2 > \mu^{NA}. \tag{B.8}$$

Note that the variance of returns over the same length of periods in the two regimes is the same, σ^2.

1.2. Payoff Uncertainty

Our setup differs from the one in the previous section in two respects. First, we make assumptions about payoff risk, as opposed to return risk, and derive expressions for expected excess returns to show how uncertainty premia are first-order for large risks and second-order for small ones. Second, we make specific assumptions about the types of risk corresponding to our discussion in Section 2.

We start with a representative investor who allocates her wealth at date t, W_t, between a risk-free asset with net return $R_{f,t+\Delta t}$ and a risky asset with payoff $X_{t+\Delta t}$ and value P_t. Her objective is maximizing expected utility of wealth at time $t + \Delta t$, with the assumption that Δt is a short period of time.

In the absence of an announcement scheduled to take place between t and $t + \Delta t$, the risky payoff is given by:

$$X_{t+\Delta t} = X_t(\varepsilon_{B,t+\Delta t} + J_{t+\Delta t}),$$

where

$$\varepsilon_{B,t+\Delta t} \sim N(\theta \Delta t, \sigma^2 \Delta t)$$

is the discrete-time aggregation of a Brownian motion and

$$J_{t+\Delta t} = \begin{matrix} \varepsilon_P \text{ with probability } \lambda \Delta t \\ 0 \text{ with probability } 1 - \lambda \Delta t \end{matrix} ,$$

where

$$\varepsilon_P \sim N(\theta, \sigma^2)$$

is the approximate discrete-time aggregation of a Poisson jump for small time intervals (small enough that the probability of more than one such jump is negligible). It follows

$$E_t^{\theta}[X_{t+\Delta t}] = X_t(1 + \lambda)\theta \Delta t$$

and

$$Var_t[X_{t+\Delta t}] = X_t(1 + \lambda)\sigma^2 \Delta t.$$

If an announcement is scheduled to take place between t and $t + \Delta t$, the risky payoff is given by:

$$X_{t+\Delta t} = X_t \varepsilon_{A,t+\Delta t}$$

where the announcement news is independent of the time interval, as discussed:

$$\varepsilon_{A,t+\Delta t} \sim N(\theta, \sigma^2).$$

All three shocks are independent of each other. We make the strong assumption that when an announcement is scheduled to arrive, no other shock can occur. This greatly simplifies the exposition by eliminating nuisance interaction terms and also strikes us as empirically realistic.

Thus, when an announcement is scheduled to take place:

$$E_t^{\theta}[X_{t+\Delta t}] = X_t \theta$$

and

$$Var_t[X_{t+\Delta t}] = X_t^2 \sigma^2.$$

With a share of wealth α_t allocated to the market, investor's date $t + \Delta t$ wealth is

$$W_{t+\Delta t} = W_t \left[(1 + R_{f,t+\Delta t}) + \alpha_t \left(\frac{X_{t+\Delta t}}{P_t} - R_{f,t+\Delta t} \right) \right].$$

For ease of notation, we set $X_t = 1$ and $R_{f,t+\Delta t} = 0$. The excess return on the risky asset is then either $(\varepsilon_{B,t+\Delta t} + J_{t+\Delta t})/P_t$ in non-announcement periods or else $\varepsilon_{A,t+\Delta t}/P_t$ in announcement periods. Neither of these assumptions affects our claims in any way.

In the case of all three risks, the mean parameter θ is the same and is unknown, and variance σ^2 is the same and known. The investor has a prior distribution over θ, which is also normal: $\theta \sim N(\mu, \sigma_\theta^2)$.

Finally, define $I_{t+\Delta t}$ as an indicator variable that equals one if an announcement is scheduled in the period between t and $t + \Delta t$ and equals zero otherwise.

The investor has CARA utility of wealth $W_{t+\Delta t}$ with absolute risk aversion A, given by

$$U(W_{t+\Delta t}) = -\frac{1}{A}\exp(-AW_{t+\Delta t})$$

so that given an expected return of θ her expected utility, $E_t^{\theta}[U(W_{t+\Delta t};\alpha_t)]$, is

$$-\frac{1}{A}\exp\left(-AW_t\left(1+\alpha_t\frac{(1-I_{t+\Delta t})(1+\lambda)\theta\Delta t+I_{t+\Delta t}\theta}{P_t}\right)+W_t^2\frac{A^2}{2}\alpha_t^2\frac{(1-I_{t+\Delta t})(1+\lambda)\sigma^2\Delta t+I_{t+\Delta t}\sigma^2}{P_t^2}\right).$$

She has constant relative ambiguity aversion, specifically:

$$\phi(U) = -\frac{(-U)^{1+\eta}}{1+\eta},$$

where $\eta \geq 0$ is her coefficient of relative ambiguity aversion ($\eta = 0$ corresponds to standard expected utility). Her welfare under smooth ambiguity aversion is:

$$\begin{aligned}
V(\alpha_t) &= E_t[\phi(U(W_{t+\Delta t};\alpha_t) = -\frac{1}{A(1+\eta)}\exp\left(-(1+\eta)\,E_t\left[E_t^{\theta}[U]\right]+\frac{1}{2}(1+\eta)^2 Var_t\left[E_t^{\theta}[U]\right]\right) \\
&= -\frac{1}{A(1+\eta)}\exp\left[-(1+\eta)AW_t\left(1+\alpha_t\frac{(1-I_{t+\Delta t})(1+\lambda)\mu\Delta t+I_{t+\Delta t}\mu}{P_t}\right)\right] \\
&\quad -(1+\eta)AW_t\left(1+\alpha_t\frac{(1-I_{t+\Delta t})(1+\lambda)\mu\Delta t+I_{t+\Delta t}\mu}{P_t}\right) \\
&\quad +(1+\eta)W_t^2\frac{A^2}{2}\alpha_t^2\left(\frac{(1-I_{t+\Delta t})(1+\lambda)\sigma^2\Delta t+I_{t+\Delta t}\sigma^2}{P_t^2}\right) \\
&\quad \left[+\frac{1}{2}(1+\eta)^2A^2W_t^2\alpha_t^2\left(\frac{((1-I_{t+\Delta t})(1+\lambda)\Delta t+I_{t+\Delta t})^2\sigma_\theta^2}{P_t^2}\right)\right].
\end{aligned}$$

The first-order condition is then

$$\begin{aligned}
&\frac{((1+\lambda)\Delta t+I_{t+\Delta t})\mu}{P_t} \\
= \;& AW_t\left(\frac{((1-I_{t+\Delta t})(1+\lambda)\Delta t+I_{t+\Delta t})\sigma^2}{P_t^2}+(1+\eta)\frac{((1-I_{t+\Delta t})(1+\lambda)\Delta t+I_{t+\Delta t})^2\sigma_\theta^2}{P_t^2}\right)\alpha_t^*
\end{aligned}$$

so that in equilibrium, with $\alpha^* = 1$ and $P_t = W_t$,

$$\begin{aligned} & ((1 - I_{t+\Delta t})(1+\lambda)\Delta t + I_{t+\Delta t})\mu \\ = \; & A\left(((1 - I_{t+\Delta t})(1+\lambda)\Delta t + I_{t+\Delta t})\sigma^2 + (1+\eta)((1 - I_{t+\Delta t})(1+\lambda)\Delta t + I_{t+\Delta t})^2\sigma_\theta^2\right). \end{aligned}$$

Consider first a period in which there is no scheduled announcement, so that $I_{t+\Delta t} = 0$, this simplifies to

$$\begin{aligned} (1+\lambda)\Delta t\mu &= A\left((1+\lambda)\Delta t)\sigma^2 + (1+\eta)(1+\lambda)^2(\Delta t)^2\sigma_\theta^2\right) \\ &= A(1+\lambda)\Delta t)\sigma^2 + o(\Delta t). \end{aligned}$$

Cancelling the common term $(1+\lambda)\Delta t$ from both sides, we get

$$\mu \approx A\sigma^2.$$

Therefore, for short time intervals ambiguity aversion plays no role.

When $I_{t+\Delta t} = 1$, again eliminating terms which are $o(\Delta t)$,

$$\mu \approx A\sigma^2 + A(1+\eta)\sigma_\theta^2,$$

where now the second term is the ambiguity premium and is not second order.

We note that uncertainty matters even for an expected utility investor on announcement days, as the expected excess return is still increasing in uncertainty σ_θ^2 when $\eta = 0$. That is, parameter uncertainty matters even to a non-ambiguity-averse investor, but only for risks which are not proportional to the time interval. (There is assumed to be no learning in our model: θ is an independent draw from $N(\mu, \sigma_\theta^2)$ for each interval Δt.)

Gollier's weighting function (adapting the formula from Appendix A to a continuous distribution) is then

$$
\begin{aligned}
\omega_\theta &= \frac{\phi'}{E[\phi']} = \frac{\phi'(E^\theta[U(W_{t+\Delta t})]}{E[\phi'(E^\theta[U(W_{t+\Delta t})]]} \\
&= \frac{\exp\left(-A\eta(W_t+\theta)+0.5A^2\eta\sigma^2\right)}{\exp\left(-A\eta(W_t+\mu)+0.5A^2\eta\sigma^2+\frac{1}{2}A^2\eta^2\sigma_\theta^2\right)} \\
&= \exp\left(-A\eta(\theta-\mu)-\frac{1}{2}A^2\eta^2\sigma_\theta^2\right)
\end{aligned}
$$

which clearly equals one when $\eta = 0$. Thus, while there is an increase in uncertainty premia in this model for announcements when $\eta = 0$, there is no increase in concavity relative to small risks, since risk- and ambiguity-adjusted densities will be the same as standard risk-neutral densities. Note that when σ_θ goes to zero, the weight ω_θ does not, but the probability of θ differing from μ does, so their product does too.

We can now use ω_θ to derive an expression for the concavity index (elasticity of marginal utility of wealth with respect to wealth). Let $\Phi(.)$ be the standard normal cumulative density function. Then

$$
p_s = E_t\left[\Phi\left(\frac{W_s - E_t^\theta[W_{t+\Delta t}]}{Var[W_{t+\Delta t}]^{\frac{1}{2}}}\right)\right] = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty}\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\exp\left(-\left(\frac{\theta-\mu}{\sigma_\theta}\right)^2\right)d\theta
$$

and

$$
\begin{aligned}
\widehat{p_s} &= E_t\left[w_\theta\Phi\left(\frac{W_s - E_t^\theta[W_{t+\Delta t}]}{Var[W_{t+\Delta t}]^{\frac{1}{2}}}\right)\right] \\
&= \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{+\infty}\exp\left(-A\eta(\theta-\mu)-\frac{1}{2}A^2\eta^2\sigma_\theta^2\right)\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\exp\left(-\left(\frac{\theta-\mu}{\sigma_\theta}\right)^2\right)d\theta.
\end{aligned}
$$

Next we want to calculate, pointwise for each possible stock price W_s:

$$
\begin{aligned}
& -\frac{W_s}{\left(\frac{\widehat{p}_s}{p_s}\exp(-AW_s)\right)}\frac{d}{dW_s}\left(\frac{\widehat{p}_s}{p_s}\exp(-AW_s)\right) \\
= \quad & AW_s \\
& -W_s\frac{d}{dW_s}\ln\left(\frac{\int_{-\infty}^{+\infty}\exp\left(-A\eta(\theta-\mu)-\frac{1}{2}A^2\eta^2\sigma_\theta^2\right)\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\exp\left(-\left(\frac{\theta-\mu}{\sigma_\theta}\right)^2\right)d\theta}{\int_{-\infty}^{+\infty}\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\exp\left(-\left(\frac{\theta-\mu}{\sigma_\theta}\right)^2\right)d\theta}\right) \\
= \quad & AW_s - W_s\left(\frac{E\left[\omega_\theta N\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]}{E\left[\omega_\theta\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]}-\frac{E\left[N\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]}{E\left[\Phi\left(\frac{W_s-(W_t+\theta)}{\sigma}\right)\right]}\right).
\end{aligned}
$$

The last expression is our option-implied concavity ratio derived from the Gollier model, where $N(.)$ is the standard normal probability density function.

Appendix C

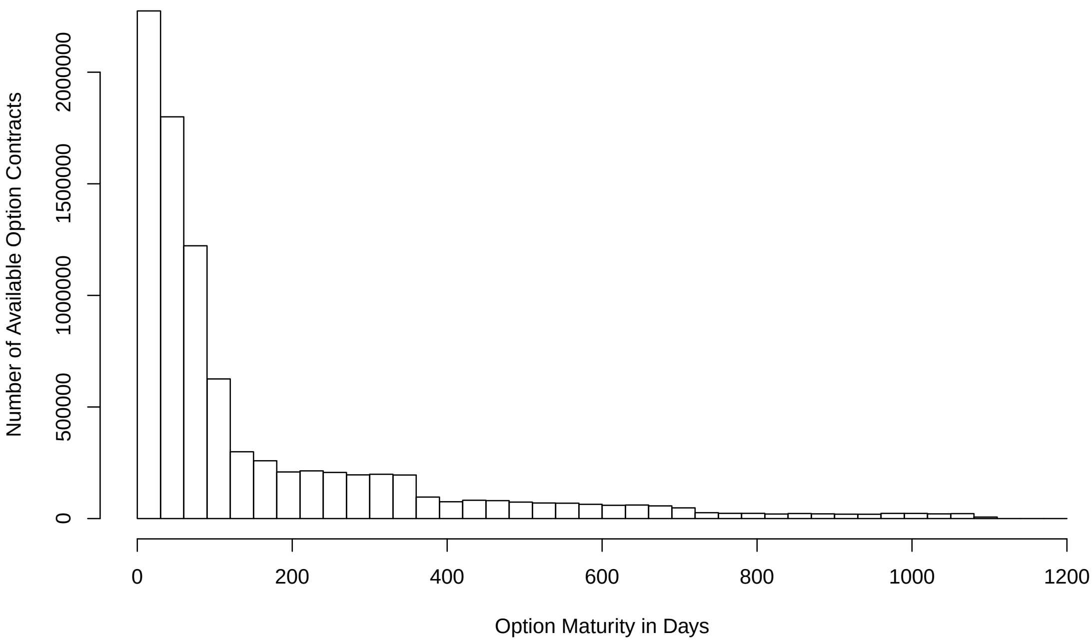

Figure A.1: Data availability over maturity. The histogram plots the number of S&P 500 option contracts in our 1996–2016 sample sorted by maturity. The majority of option contracts are available in the maturity range of up to three months.

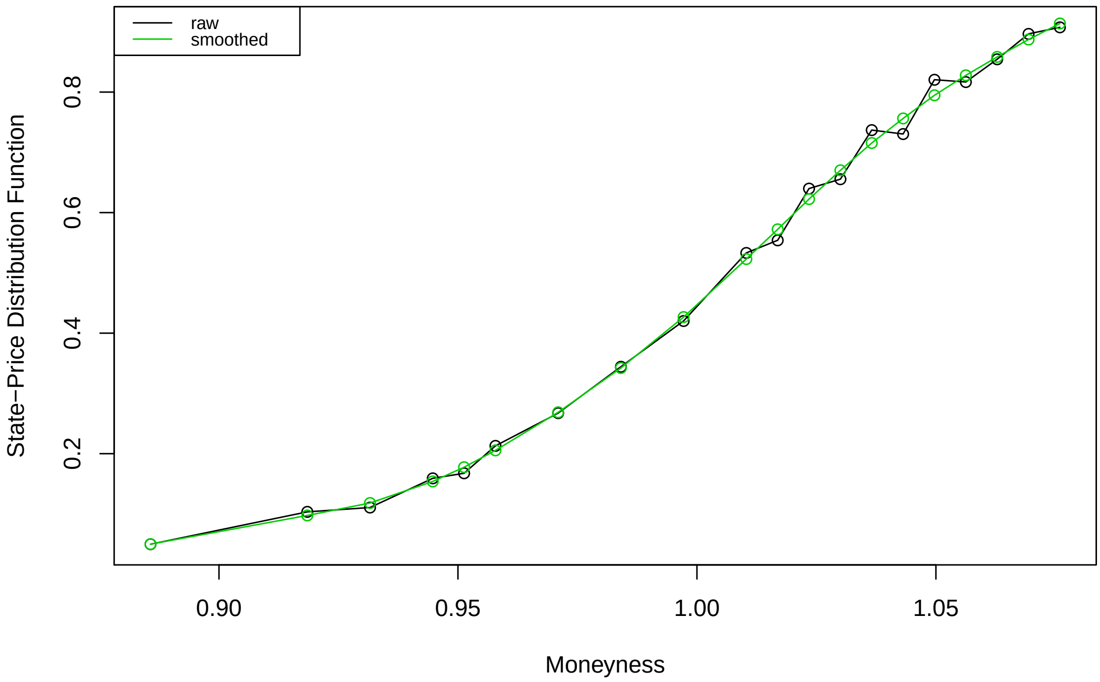

Figure A.2: This figure shows the raw vs. the smoothed distribution function with an illustrative example taken from April 07, 1997 (S&P 500 index level of 762.13) with a maturity of May 17, 1997. It visualizes the impact of the smoothing approach based on B-splines at the level of the state-price distribution employed in the paper to remove non-informative fluctuations.

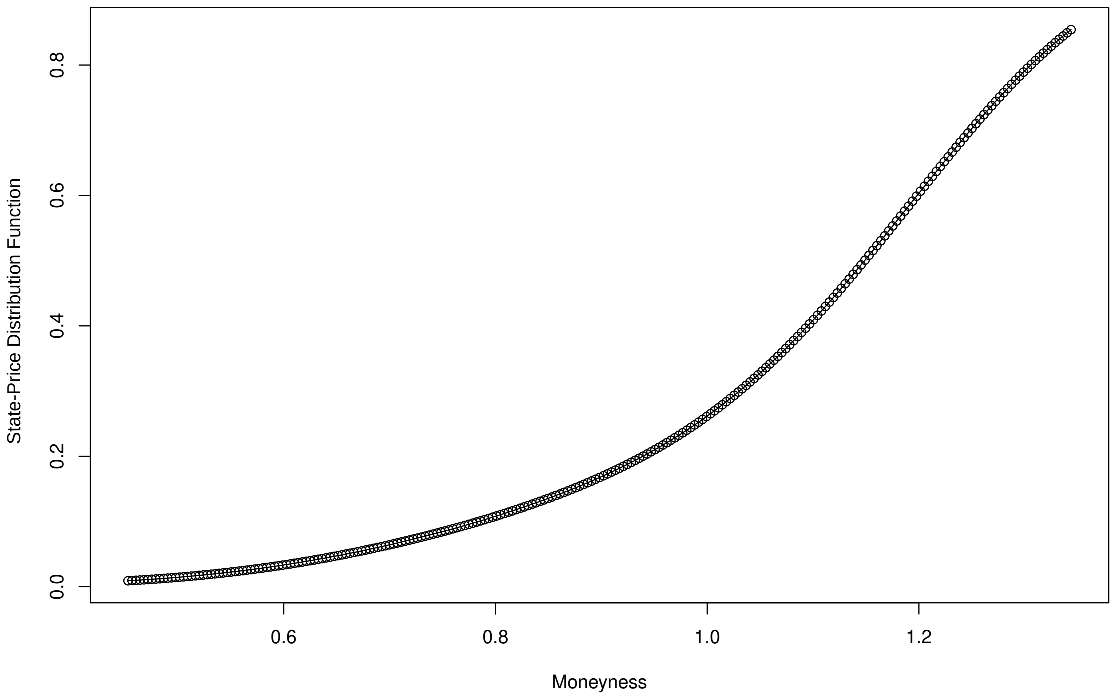

Figure A.3: This figure shows the typical form of the state-price distribution function for the S&P 500. The state-price distribution function is extracted from the partial derivatives of option prices of a particular day with common maturity with respect to the strike price K. The typical S-shape ranging from 0 to an ultimate value of 1 is clearly visible. The form of the state-price distribution determines the state-price density function as its first derivative. The purely illustrative example shows the curve for October 2, 2006 (S&P 500 index level of 1331.32) using options that expire on June 16, 2007. Option prices, index data, dividend yields, and risk free rates used in the extraction are sourced from OptionMetrics.

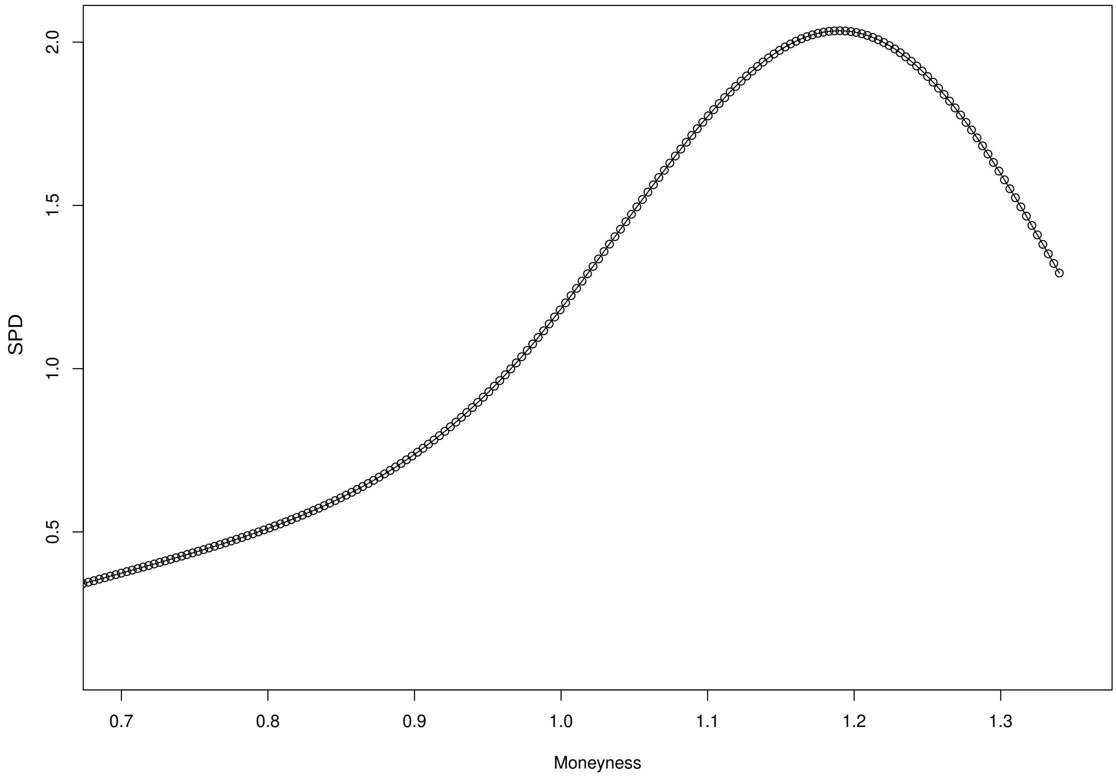

Figure A.4: This figure shows the typical form of the state-price density function (the first derivative of the state-price distribution function) for the S&P 500 that is close to a leptokurtic and skewed normal distribution. The purely illustrative example shows the SPD for October 2, 2006 (S&P 500 index level of 1331.32) using options that expire on June 16, 2007. Option prices, index data, dividend yields, and risk free rates used in the extraction are sourced from OptionMetrics.

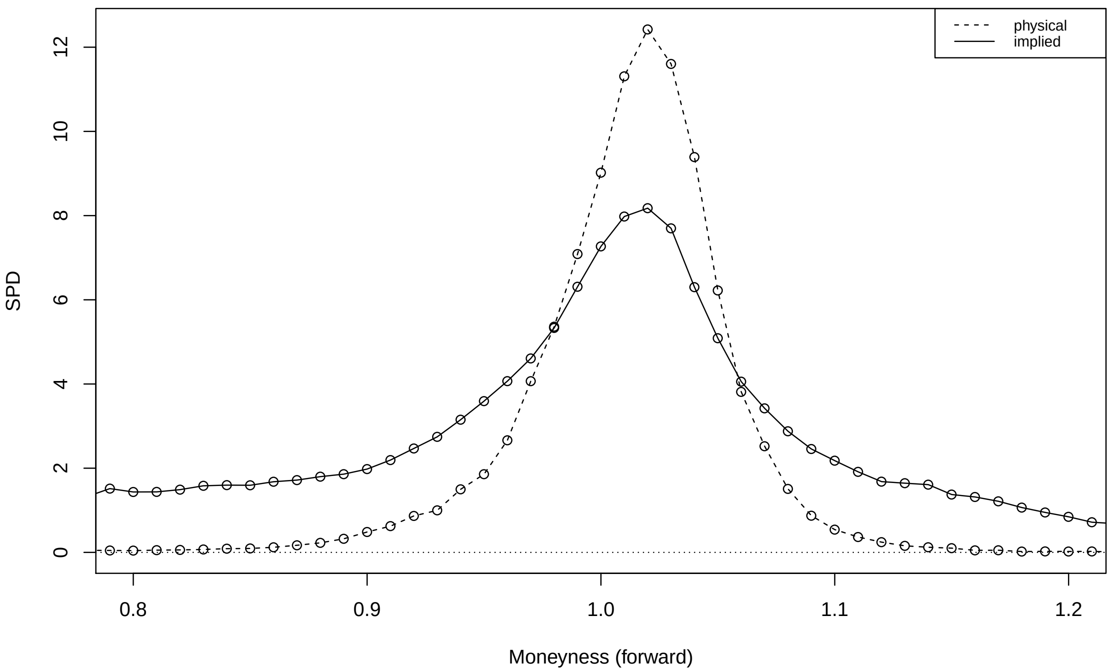

Figure A.5: This figure plots the average physical probability density vs. the implied state-price density (SPD) over the 1996–2016 period. The employed physical densities of observed returns are based on average historical returns consistent with the maturity of the respective implied SPDs (on average 44 days). In comparison to Figure 4, this figure plots the distributions in relation to moneyness with the futures/forward price assigned to a moneyness of 1.

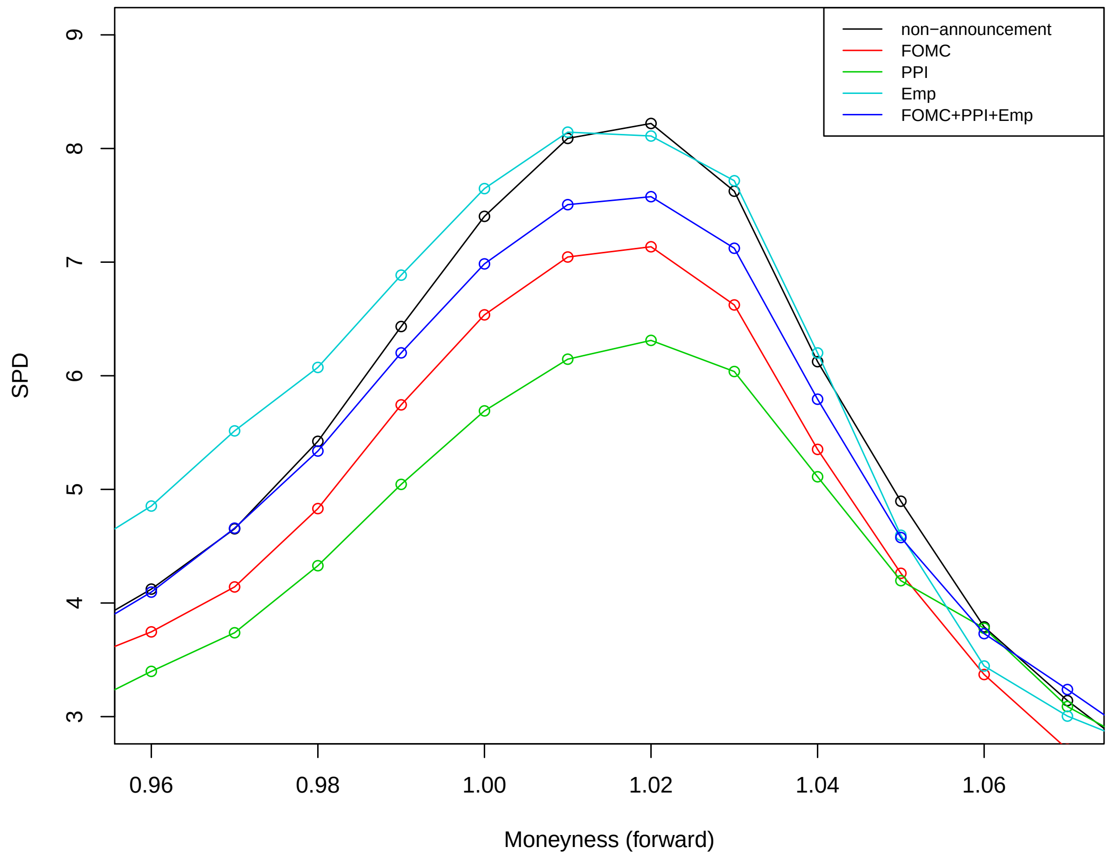

Figure A.6: The figure shows the implied state-price densities (SPDs) for announcement days (FOMC, PPI, unemployment, and all three types jointly) and non-announcement days. The SPDs are extracted based on the respective market parameters at the time and are subsequently aggregated via the arithmetic mean. The figure plots the distributions in relation to moneyness with the futures/forward price assigned to a moneyness of 1 (in contrast to current price moneyness used in Figure 2). The sample period is 1996–2016.

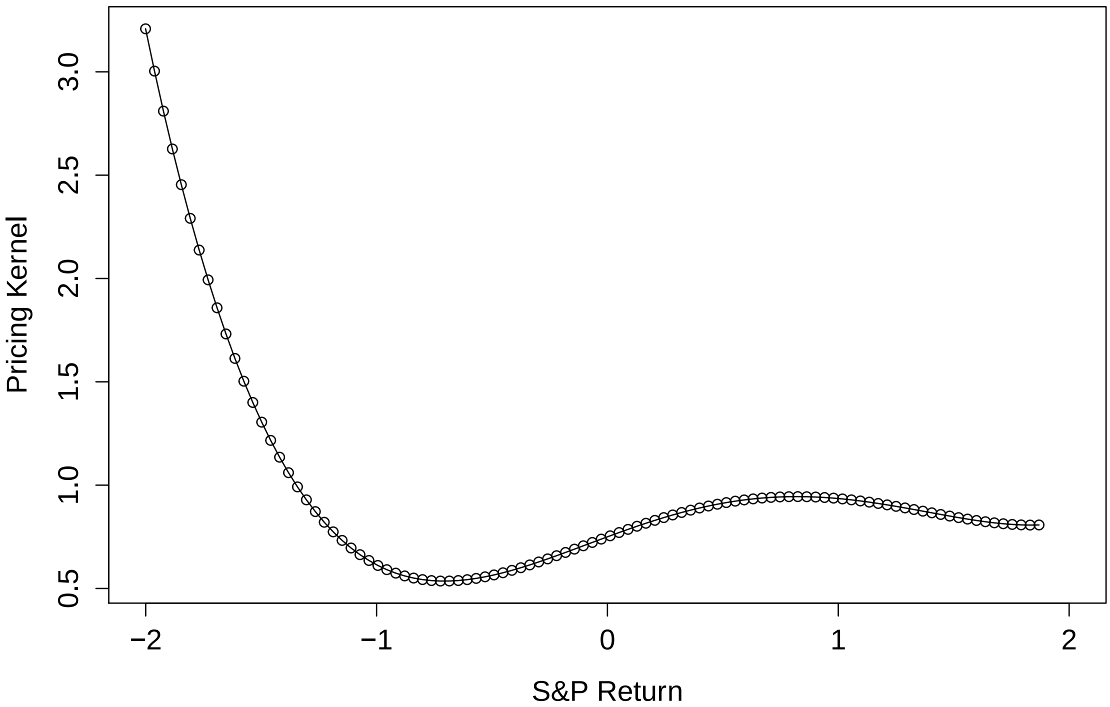

Figure A.7: This figure shows the typical form of the pricing kernel for the S&P 500 over ω in a range of plus/minus two standard deviations (i.e., the graph is centered around the expected return plus error bands). The illustrative example shows the results for March 9, 1999 (S&P 500 index level of 1279.84) using options that expire on March 20, 1999.

Table A.1
State-Price Density (with maturity scaling)

Type	Mean	Std. Err.	Skewness	Exc. Kurt.	t-test	KS	MW
Non-Ann.	8.17	0.07	-0.68	-0.24	-	-	-
FOMC	7.02	0.35	-0.41	-1.13	0.1%	0.1%	0.2%
PPI	6.18	0.29	-0.26	-1.14	0.0%	0.0%	0.0%
Emp	8.21	0.21	-0.71	-0.40	57.7%	56.6%	75.7%
All Ann.	7.54	0.13	-0.52	-0.55	0.0%	0.0%	0.0%

This table presents the mean and standard error of the implied state-price density at its maximum on announcement and non-announcement days averaged across our 1996–2016 sample period. Before the averaging, all SPDs are scaled to reflect a common maturity of 45 days. The table further reports the skewness and excess kurtosis. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the non-announcement-day average is higher (t-test), the non-announcement-day distribution stochastically dominates (KS), and the non-announcement-day distribution differs by a positive location shift (MW) compared to that on announcement days.

Table A.2
Concavity Index (with maturity scaling)

Panel A: Power Utility Function (0.97–1.03)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	10.61	0.21	0.0%	0.0%	0.0%
FOMC	14.71	1.06	0.0%	1.0%	0.0%
PPI	17.06	1.01	0.0%	0.0%	0.0%
Emp	12.17	0.60	0.7%	3.0%	1.0%
All Ann.	14.20	0.49	0.0%	0.0%	0.0%

Panel B: Exponential Utility Function (0.97–1.03)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	10.71	0.21	0.0%	0.0%	0.0%
FOMC	14.79	1.06	0.0%	0.8%	0.0%
PPI	17.16	1.01	0.0%	0.0%	0.0%
Emp	12.21	0.60	0.9%	4.2%	1.3%
All Ann.	14.27	0.48	0.0%	0.0%	0.0%

This table presents the mean and standard error of the γ parameters of the fit of a power (Panel A) and exponential (Panel B) utility function to the concavity index on announcement and non-announcement days averaged across our 1996–2016 sample period. For this process, we calculate the ratio of the option-implied state-price densities (SPD) and physical probability densities of S&P 500 returns that is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). Before the averaging, all SPDs are scaled to reflect a common maturity of 45 days. We use a moneyness range of 0.97 to 1.03. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Table A.3
Concavity Index

Panel A: Power Utility Function (0.98–1.08)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	2.47	0.19	-	-	-
FOMC	4.66	0.89	0.9%	2.6%	1.8%
PPI	6.42	1.04	0.0%	0.0%	0.1%
Emp	4.27	0.56	0.1%	0.3%	0.2%
All Ann.	4.95	0.45	0.0%	0.0%	0.0%

Panel B: Exponential Utility Function (0.98–1.08)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	2.38	0.19	-	-	-
FOMC	4.52	0.89	1.0%	3.5%	2.0%
PPI	6.20	1.04	0.0%	0.0%	0.1%
Emp	4.13	0.56	0.2%	0.4%	0.2%
All Ann.	4.79	0.45	0.0%	0.0%	0.0%

This table presents the mean and standard error of the γ parameters of the fit of a power (Panel A) and exponential (Panel B) utility function to the concavity index on announcement and non-announcement days. For this process, we calculate the ratio of the option-implied state-price densities (SPD) and physical probability densities of S&P 500 returns that is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). We use an extended moneyness range of 0.98 to 1.08. The sample period is 1996–2016. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Table A.4
Concavity Index

Panel A: Power Utility Function (0.92–1.08)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	5.69	0.12	-	-	-
FOMC	6.17	0.62	22.9%	58.0%	49.8%
PPI	9.05	0.69	0.0%	0.0%	0.0%
Emp	7.12	0.35	0.0%	0.0%	0.0%
All Ann.	7.44	0.30	0.0%	0.0%	0.0%

Panel B: Exponential Utility Function (0.92–1.08)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	5.48	0.13	-	-	-
FOMC	6.00	0.64	21.5%	52.3%	47.6%
PPI	8.82	0.71	0.0%	0.0%	0.0%
Emp	6.97	0.36	0.0%	0.0%	0.0%
All Ann.	7.26	0.30	0.0%	0.0%	0.0%

This table presents the mean and standard error of the γ parameters of the fit of a power (Panel A) and exponential (Panel B) utility function to the concavity index on announcement and non-announcement days. For this process, we calculate the ratio of the option-implied state-price densities (SPD) and physical probability densities of S&P 500 returns that is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). We use an extended moneyness range of 0.92 to 1.08. The sample period is 1996–2016. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Table A.5
S&P 500 Return around Announcement Days

	FOMC	PPI	Emp	All Ann.
3 days prior	-11.9	2.0	8.2	-0.6
2 days prior	1.7	-4.4	17.0	4.8
1 day prior	15.3	-6.2	-4.8	1.4
Ann. day	32.2	7.2	11.6	17.0
1 day past	-1.6	-0.9	-4.0	-2.2
2 days past	-2.2	8.6	-8.9	-0.9
3 days past	16.6	20.5	-3.7	11.1

This table presents the average S&P 500 daily excess returns around different types of announcement days. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The sample period is 1996–2016.

Table A.6
Concavity Index Around Announcement Days

Panel A: Power Utility Function

	FOMC	PPI.	Emp	All Ann.
3 days prior	7.69	11.93	9.28	9.64
2 days prior	8.61	12.76	9.48	10.28
1 day prior	10.22	12.83	9.61	10.89
Ann. Days	12.55	15.06	11.32	11.33
1 day after	8.90	11.72	10.04	10.22
2 days after	9.31	12.38	10.01	10.57
3 days after	9.26	12.14	10.20	10.53

Panel B: Exponential Utility Function

	FOMC	PPI.	Emp	All Ann.
3 days prior	7.79	11.97	9.30	9.68
2 days prior	8.69	12.81	9.48	10.33
1 day prior	10.31	12.89	9.59	10.93
Ann. Days	12.62	15.13	11.30	11.38
1 day after	8.98	11.83	10.02	10.28
2 days after	9.38	12.50	10.00	10.63
3 days after	9.31	12.25	10.19	10.58

This table shows the γ parameter of the fit of a power (Panel A) and exponential (Panel B) utility function to the concavity index on different types of announcement days and its behavior up to 3 days prior and after the respective announcement day. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The sample period is 1996–2016.

Table A.7
Panel A: Concavity Index - Power Utility Function ($\leq$ 5 days)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	14.34	2.08	-	-	-
FOMC	19.04	2.78	10.2%	22.8%	13.4%
PPI	17.86	1.52	9.2%	11.8%	15.3%
All Ann.	18.30	1.36	6.1%	5.8%	8.8%

Based on option sets with maturity of 5 days or less, this table presents the mean and standard error of the γ parameters resulting from the fit of a power utility function to the concavity index on announcement and non-announcement days. For this process, we calculate the ratio of the option-implied state-price densities (SPD) and physical probability densities of S&P 500 returns that is subsequently converted into the implied concavity (Ait-Sahalia and Lo 2000). We use a moneyness range of 0.98 to 1.02. The sample period is 1996–2016. The announcement days include days when the FOMC is releasing its decision (FOMC), when PPI numbers are released (PPI), when employment numbers are released (Emp), and all three announcements together (All Ann.). The last three columns show the p-values of a t-test, Kolmogorov-Smirnov (KS) test, and Mann-Whitney (MW) test, which test the hypotheses that the announcement-day average is higher (t-test), the announcement-day distribution stochastically dominates (KS), and the announcement-day distribution differs by a positive location shift (MW) compared to that on non-announcement days.

Panel B: Concavity Index - Exponential Utility Function ($\leq$ 5 days)

Type	Mean	Std. Err.	t-test	KS	MW
Non-Ann.	13.89	2.07	-	-	-
FOMC	18.55	2.77	10.2%	22.8%	13.4%
PPI	17.31	1.53	9.8%	11.8%	16.4%
All Ann.	17.78	1.37	6.4%	5.8%	9.3%

Analogous to Panel A and based on option sets with maturity of 5 days or less, this table presents the mean and standard error of the γ parameters, here resulting from the fit of an exponential utility function.