Estimating the Value of Information

Ohad Kadan Asaf Manela*

November 2016

Abstract

We derive a general expression for the value of information to a small price-taking investor in a dynamic environment and provide a framework for its estimation from index options. We apply this framework and estimate that a consumer-investor with commonly-used preference parameters would pay 1 to 4 percent of her wealth to preview and act on key macroeconomic indicators (GDP, unemployment, etc.). The value of information increases with the time discount factor, decreases with risk aversion, and increases with the elasticity of intertemporal substitution. Rational expectations (or lack thereof) play an important role in the value of information.

JEL classification: G10, G12, G13, D80, D82, D83

Keywords: value of information, informativeness, options prices, recovery.

*Both authors are at Washington University in St. Louis. Email: kadan@wustl.edu or amanela@wustl.edu. We thank Fang Liu, Semyon Malamud, Thomas Maurer, Alan Moreira, Ngoc-Khanh Tran, and seminar participants at IDC and WashU for helpful comments.

1 Introduction

How much would investors pay to receive valuable information slightly before it is released publicly? This question is far from hypothetical. Legions of analysts and computer programs are constantly collecting millions of pieces of information and producing thousands of reports on various aspects of the economy and individual firms. Understanding the private incentives to collect information is central to the study of the informational efficiency of capital markets (Fama, 1970; Grossman and Stiglitz, 1980). Quantifying the value of private information is also key for penalizing insider trading (DeMarzo, Fishman, and Hagerty, 1998). In this paper we present a general framework for evaluating informative signals from the point of view of a utility maximizing investor. We then illustrate our framework by estimating the values of key macroeconomic indicators and providing comparative statics for the determinants of the value of information.

Our thought exercise is as follows. We empathize with a risk-averse investor who can trade state-contingent securities that pay a dollar if a particular value of the market return realizes. Without any additional information, the investor would use her prior probabilities about each state to optimally choose her consumption and investments in the state-contingent securities. We offer this investor a peek into a posterior distribution updated to reflect an additional signal from a particular information source (e.g. GDP report). With this updated distribution, the investor can improve her consumption-investment choice. We estimate the share of wealth she is willing to forgo to observe such information, by proxying for the prior and posterior using probabilities recovered from the prices of S&P 500 options observed just before and after informational releases.

One may wonder why not simply measure announcement returns on the index (Fama, Fisher, Jensen, and Roll, 1969). While that statistic may provide an intuitive test whether additional price-relevant information is revealed on announcement dates, we would not learn much about the economic value of the announced information to an investor without a model of preferences that pins down how consumption and investment would change in response. For example, the value of information is small for a highly risk-averse investor because even with an additional signal she would remain mostly invested in a riskless portfolio. Similarly, an agent that strongly dislikes substituting her current consumption for future consumption would have little use for information that improves her investments. Moreover, the optimal financial instrument or trading strategy for utilizing a particular piece of information is not obvious. Quantifying the value of information instead requires a model of preferences, probabilities and opportunity sets.

To this end, we build on the dynamic stochastic recursive preferences setup of Epstein and Zin (1989) who generalize the Merton-Samuelson consumption and investment problem. We augment this setup by allowing a small price-taking agent to purchase a subscription to a signal, which she can observe just before it becomes available to the public. The small agent is then allowed to make consumption and investment decisions based on her private information without affecting prevailing market prices. We define the value of an information source as the fraction of wealth the agent is willing to give up in order to obtain a sequence of signals generated by this source.

We derive an estimable expression for the value of information, which is based on its expected effect on the value-to-consumption ratio. The value-to-consumption ratio reflects the expected ratio of lifetime utility to current consumption, where the expectation takes into account the different informative signals generated by the information source. This approach naturally leads to the application of the generalized method of moments (GMM) to estimate increases in the value-to-consumption ratios, which, in turn, determine the value of information. By conditioning on dates of particular information releases, we estimate the value of information associated with different information sources, building on the standard approach that uses large sample means to estimate population moments (Hansen and Singleton, 1982). We also provide a simpler closed-form first-order approximation for the value of information, which does not require numerical solution and proves to be quite accurate.

Our estimation approach takes the parameters of the recursive utility function (discount rate, risk aversion, and elasticity of inter-temporal substitution) as given. Then, to estimate the value of information associated with an information source we require two sets of inputs: state prices and their "physical" probabilities. State prices are relatively straightforward to extract from option prices using the approach introduced in Breeden and Litzenberger (1978). Recovering forward looking physical probabilities from state prices is an active topic of research, which is not the focus of the current paper, but a necessary step nonetheless.[1] We therefore take the following two distinct approaches: (i) parametric recovery that relies on assumptions about the growth rate of consumption and its relationship with the market return, (ii) nonparametric recovery as in Ross (2015). While the value of information we obtain in these two methods is somewhat different, all of our comparative statics results hold regardless.

We estimate the value of information from January 1996 to August 2015 for each of a set of key macroeconomic indicators including Unemployment, GDP and Jobless Claims reports, as well as FOMC Decisions. We set our benchmark preference parameters to standard expected utility levels, i.e., relative risk aversion and the inverse of the elasticity of intertemporal substitution (EIS) both equal to 5, and the time discount rate is set to $\beta = 0.998$. Our benchmark estimates indicate that our small agent is willing to pay between 1 and 4 percent of her wealth for the improved consumption/investment decisions associated with early observations of the leading macroeconomic indicators, regardless of using parametric or non-parametric recovery.

We then conduct a series of numerical comparative statics exercises by solving the model for different values of the underlying parameters, utilizing the flexibility of the Epstein-Zin preferences in separating between risk-aversion and the EIS. First, we find that the value of information is increasing with the time discount factor as the agent assigns more value to future periods. Second, we find that the value of information declines with risk aversion. Intuitively, more risk averse investors are reluctant to make risky investments, and the value of informative signals is lower for them. In fact, when both risk aversion and the EIS approach 1 (the *log* utility case), the value of informative signals approaches 100%. Namely, a *log* utility agent would be willing to pay nearly all of her wealth to obtain early informative signals. Comparative statics in the EIS are also intuitive: a higher EIS is associated with a higher willingness of the agent to defer consumption and save. Accordingly, we find that an increase in the EIS leads to a higher value of information.

Our estimation approach, which recovers the probabilities of states of the world from option prices, gives rise to a special role for rational expectations. Specifically, we refer to expectations

[1] The standard approach to calculating physical probabilities used realized moments of historical returns (see Jackwerth (2000) and Aït-Sahalia and Lo (2000)). This approach, however, is inappropriate here because we rely on the immediate response to changes in information upon announcement while historical moments are slow moving.

as being "rational" if the recovered probabilities we use to estimate the value of information are time-consistent in the sense that the law of total probably is satisfied. Deviations from rational expectations reflect that current probabilities of future states do not correctly and fully anticipate future realizations of information. We do not impose such rationality on the probabilities we recover, and find that the data often features significant deviations from rational expectations.

To illustrate the role of rational expectations in the value of information, we consider the particularly tractable case of log utility. Cabrales, Gossner, and Serrano (2013) focus on this case and show that for a log utility agent, faced with a static investment problem, the value of information equals the mean reduction in entropy that the information source can generate. We establish that in our dynamic consumption-investment framework, assuming that expectations are rational, the value of information generalizes to the present value of the expected future changes in entropy associated with the information. Importantly, we show that if expectations are not rational, the value of information in the log utility case may deviate substantially from the expected change in entropy, and we quantify these deviations empirically.[2]

We focus on the private value of information for investment. By considering a better-informed price-taking agent, we depart from the literature focusing on the public/social value of information, and instead focus explicitly on the private value of information.[3] Private information is pervasive yet hardly reflected by equity prices (Collin-Dufrense and Fos, 2015; Kacperczyk and Pagnotta, 2016). Recent work calibrates the willingness of a representative agent with recursive utility to pay for early resolution of uncertainty (Ai, 2007; Epstein, Farhi, and Strzalecki, 2014; Croce, Marchuk, and Schlag, 2016). Non-expected utility agents prefer early resolution of uncertainty if, and only if, their relative risk aversion exceeds the inverse of their EIS (Kreps and Porteus, 1978; Epstein and

[2] Cabrales, Gossner, and Serrano (2013) focus on the theoretical aspects of ranking different information sources, but do not provide a practical method for its estimation. Thus, our option implied estimation approach is useful if one would like to estimate such "change in entropy" rankings. In general, however, our approach to calculating the value of information differs from Cabrales, Gossner, and Serrano (2013) who study an upper bound on the value of information across all investors in a certain class. Instead, we use results from asset-pricing, which point us to a set of benchmark parameters for the Epstein-Zin preferences and use these preferences to calculate the value of information. This difference is key for practical reasons as our estimates show that the economic magnitude of the value of information and the ranking of different informational sources differ substantially between the log utility case and our benchmark parameters, which are commonly used to fit stylized asset pricing facts (e.g. Bansal and Yaron, 2004).

[3] The seminal contributions are Hirshleifer (1971) and Spence and Zeckhauser (1972). See Maurer and Tran (2015) for recent work on the social value of information.

Zin, 1989). The timing premium that Epstein, Farhi, and Strzalecki (2014) measure is entirely about this attitude of the agent toward uncertainty, even when she cannot do anything to alter its future consumption stream – the so called psychic value of information. Our thought exercise is quite different, and instead focuses on what Epstein, Farhi, and Strzalecki (2014) call the instrumental value of information. We allow the agent to alter her consumption process based on the information, without constraining it to match aggregate income in the economy. Thus, the value of information that we estimate captures both the instrumental value of information and, in the non-expected utility case, its timing premium. We therefore find it instructive to shut down the psychic channel in our benchmark estimates, and then show what happens as we deviate from the expected utility case.

Our work also relates to a literature studying information acquisition and markets for financial information. Recent such work studies media frenzies (Veldkamp, 2006), large price movements (Barlevy and Veronesi, 2000; Dow, Goldstein, and Guembel, forthcoming; García and Strobl, 2011; Mele and Sangiorgi, 2015), portfolio choice (Peress, 2004; Van Nieuwerburgh and Veldkamp, 2010), mutual funds behavior (García and Vanden, 2009; Kacperczyk, Van Nieuwerburgh, and Veldkamp, 2016), the home bias puzzle (Van Nieuwerburgh and Veldkamp, 2009), sell-side equity analysts (Kelly and Ljungqvist, 2012; Kadan, Michaely, and Moulton, 2016), and bank runs (He and Manela, 2016).[4]

Quantitative work in this field is rare, and has thus far relied on stronger assumptions. Savov (2014) is an exception, which calibrates a dynamic dispersed information model with constant relative risk aversion (CRRA) agents to study mutual fund performance. Manela (2014) estimates the value of diffusing information in a competitive noisy rational expectations equilibrium with constant absolute risk averse (CARA) investors and four periods. In that model, the value of information takes into account the equilibrium fraction of informed investors and the resulting partial informativeness of prices. Avdis (forthcoming) shows that in such dynamic equilibrium models, the stochastic supply of the risky asset itself has a significant informational role. Breon-Drish (2015) and Malamud (2015) suggest that moving beyond CARA utility can be important.

[4]See Veldkamp (2011) for a good review of this literature.

Here, we generalize to a neoclassical infinite horizon recursive preferences investor, with a more general state space and trading opportunities, but at the cost of abstracting from such information externalities and equilibrium effects.[5]

The paper proceeds as follows. We derive a general expression for the value of information in Section 2. We provide a framework for its estimation from options in Section 3. We apply this methodology to estimate the values of key macroeconomic indicators and present comparative statics results in Section 4. In Section 5 we discuss the role of rational expectations in the value of information by focusing on the log utility case. We conclude in Section 6. Proofs and derivation details are in Appendix A, while estimation details are in Appendix B.

2 Theoretical Framework

Our starting point is the standard dynamic setup with recursive preferences of Epstein and Zin (1989), in which consumption and investment are endogenous (unlike the endowment economy setup in Epstein, Farhi, and Strzalecki (2014)). We augment this setup by allowing a small price taking agent to obtain a private signal, which enables the agent to improve her decision making, yielding a better consumption stream. We use this tool to study the agent's willingness to pay for this signal – the value of private information. Since this setup is largely well-known, we focus on the main results needed for estimation, and refer the interested reader to the original papers including the corrections in Backus, Routledge, and Zin (2005) for details.

2.1 State Space and Preferences

Time is discrete with an infinite horizon. At each date $t > 0$ random state z_t is drawn from a finite set $\{1, ..., n\}$. We assume that the state evolution is Markovian and denote the transition probabilities by $p(z_{t+1}|z_t)$, where $\sum_{z_{t+1}} p(z_{t+1}|z_t) = 1$. To prevent pathologies we assume that any state can be reached from any other state, i.e, $0 < p(z_{t+1}|z_t) < 1$ for all $z_t, z_{t+1} \in \{1, ..., n\}$. In particular, the Markov matrix is ergodic.

[5] Recent work by Borovička (2015) shows that in models with recursive preferences agents with heterogeneous beliefs may survive in the long run, which suggests that equilibrium effects may be important for the value of information. Quantifying such effects would be an interesting avenue for future work.

When state z_t is realized there is trade in n Arrow-Debreu securities corresponding to the n states in time z_{t+1}. The price of a security paying 1 when state z_{t+1} is realized and zero otherwise given that the current state z_t is denoted by $q(z_{t+1}|z_t) > 0$ (the state price). We assume no arbitrage so that such positive state prices do exist.

Consider a small, price-taking agent whose preferences over consumption c_t are represented by a recursive utility function as in Epstein and Zin (1989),

$$U_t = J\left[c_t, \mu_t\left(U_{t+1}\right)\right],$$

where U_t is short-hand for utility starting at some date-t history, J is a time aggregator, and μ_t is a certainty equivalent function based on the conditional probabilities as of time t. The time aggregator J is the only determinant of consumption dynamics in a deterministic environment, while the certainty equivalent function μ is all that matters in a static problem with uncertainty.

We specialize to the widely used constant elasticity of substitution (CES) time aggregator

$$J[c, \mu] = \left\{(1-\beta)\, c^{1-\rho} + \beta \mu^{1-\rho}\right\}^{\frac{1}{1-\rho}},$$

where $\rho \geq 0$ can be interpreted as the inverse of the elasticity of intertemporal substitution (EIS), and the certainty equivalent function takes the homogeneous form

$$\mu(U) = \left(E_t U^{1-\gamma}\right)^{\frac{1}{1-\gamma}}, \tag{1}$$

where $\gamma > 0$ is the coefficient of relative risk aversion and $E_t(\cdot)$ is the conditional expectation operator. The case of expected utility with constant relative risk aversion is then a special case in which $\rho = \gamma$.

2.2 Information and Timing

For a fixed date t we assume the agent is allowed to purchase a stream of informative signals about future states of nature in dates $t+1$, $t+2, \ldots$. The standard way to model such signals is using the notion of an information structure (e.g. Cabrales, Gossner, and Serrano, 2013). An information structure α as a finite set of S_α signals, which are potentially correlated with future states of nature. If an agent buys information structure α, then in each period starting from date t she gets to observe a signal $s_t \in \{1, ..., S_\alpha\}$, which may be informative on the state realization in the next period z_{t+1}. The question we are asking is how much would an agent be willing to pay to privately observe such a stream of signals.

Formally, an information structure α is represented by a matrix of conditional probabilities $\alpha(s_t|z_{t+1})$ for all $s_t \in \{1, ..., S_\alpha\}$ and $z_{t+1} \in \{1, ..., n\}$ representing the probability of observing signal s_t given that the next period's signal is z_{t+1}. We restrict attention to imperfect information structures. That is, any signal may be observed with some positive probability given any realized state: $\alpha(s_t|z_{t+1}) > 0$ for all s_t and z_{t+1}. Imperfect information structures are natural to consider as typical signals observed in reality likely provide useful but imperfect predictions on future states. To simplify, we assume that the signal in time t is independent of all state realizations and signals up to time t conditional on the state in time $t+1$. That is, the information structure α consists of signals, which at time t depend directly only on the state in time $t+1$. In particular, signals may be serially correlated, however, all the information incorporated in a signal about next period's signal is already included in the next period's state realization.

The total probability of observing signal s_t given that the current state is z_t is given by

$$p_\alpha(s_t|z_t) = \sum_{z_{t+1}} \alpha(s_t|z_{t+1})\, p(z_{t+1}|z_t). \tag{2}$$

Since the state Markov chain is ergodic and the information structure α is imperfect we have that $p_\alpha(s_t|z_t) > 0$ for all s_t and z_t. Thus, the posterior probability for state z_{t+1} given that signal s_t is observed is

$$p_\alpha(z_{t+1}|s_t, z_t) = \frac{\alpha(s_t|z_{t+1})\, p(z_{t+1}|z_t)}{p_\alpha(s_t|z_t)} > 0. \tag{3}$$

As a special case, let α_0 denote an information structure that consists of one signal only, i.e,. $S_\alpha = 1$. Then, α_0 is completely uninformative since regardless of z_{t+1} the same signal will always be observed by the agent. In particular, $p_{\alpha_0}(z_{t+1}|s_t, z_t) = p(z_{t+1}|z_t)$. Thus, obtaining the information structure α_0 is equivalent to not obtaining information at all.

The timing of events in each period is as described in Figure 1. At the beginning of time t the state z_t is realized and revealed publicly. Next, the signal s_t (taken from information structure α) is realized and is revealed privately to the agent who bought the information structure. The agent then makes a decision based on z_t and s_t choosing how to allocate her wealth a_t between consumption c_t and a vector of investments in Arrow-Debreu securities corresponding to the n states of nature, denoted by $\mathbf{w}_{t+1}$.[6] Here $\mathbf{w}_{t+1} = (w_{1,t+1}, ..., w_{n,t+1})$, where $w_{i,t+1}$ is the fraction of $a_t - c_t$ invested in the Arrow-Debreu security that pays 1 dollar when $z_{t+1} = i \in \{1, ..., n\}$ is realized and zero otherwise. The key point in this timing choice is that a small agent, who purchased an information structure α, gets to choose her consumption/investment profile using the realized signal s_t, giving her an informational advantage compared to the public information available at time t. Similar to Grossman and Stiglitz (1980), Admati (1985), Ausubel (1990), Admati and Pfleiderer (1991), we assume that the informed trader is essentially infinitesimal, so that her trades do not affect prices.

2.3 Optimization

The agent's problem is the classic Merton-Samuelson consumption/portfolio-choice problem, albeit with the additional signal. Her dynamic problem is characterized by the Bellman equation

$$V(a, z, s) = \max_{c,\mathbf{w}} \left\{ (1-\beta) c^{1-\rho} + \beta E_t \left[V(a', z', s')^{1-\gamma} \right]^{\frac{1-\rho}{1-\gamma}} \right\}^{\frac{1}{1-\rho}}$$

subject to the wealth constraint, $a' = (a-c) \sum_{i=1}^{n} w_i R_i = (a-c) R_p$, where a denotes wealth, R_i is the gross random return of the ith Arrow-Debreu security, $R_i = \frac{1_{\{z_t=i\}}}{q_i}$, and R_p is the gross return on a portfolio $(w_1, w_2, \ldots, w_{n-1}, 1 - \sum_{i=1}^{n} w_i)$, of such securities. Note that the expectation here is taken over both states and signals and that to conserve notation we drop time subscripts when convenient and instead use "primes" to indicate next period's variables.

The linearity of the budget constraint and the homogeneity of the time and risk aggregators

[6]Throughout the paper we denote vectors in bold.

imply homogeneity of the value function:

$$V\left(a,z,s\right)=aL\left(z,s\right) \tag{4}$$

for some scaled value function L. This useful separation implies that $L\left(z,s\right)$ is the marginal utility of wealth (Backus, Routledge, and Zin, 2005).

The scaled Bellman equation can be written as

$$L\left(z,s\right)=\max_{b,w}\left\{\left(1-\beta\right)b^{1-\rho}+\left(1-b\right)^{1-\rho}\beta E_t\left[R_p'^{1-\gamma}L\left(z',s'\right)^{1-\gamma}\right]^{\frac{1-\rho}{1-\gamma}}\right\}^{\frac{1}{1-\rho}}, \tag{5}$$

where $b=b\left(z,s\right)\equiv\frac{c}{a}$ is the consumption-to-wealth ratio.

Plugging the first-order condition for the optimal b,

$$\beta E_t\left[R_p'^{1-\gamma}L\left(z',s'\right)^{1-\gamma}\right]^{\frac{1-\rho}{1-\gamma}}=\left(1-\beta\right)\left(b^{-1}-1\right)^{\rho}, \tag{6}$$

into (5) and simplifying gives a relationship between V, L and b:

$$V\left(a,z,s\right)=aL\left(z,s\right)=a\left(1-\beta\right)^{\frac{1}{1-\rho}}b^{\frac{-\rho}{1-\rho}}. \tag{7}$$

This familiar "solution" (Backus, Routledge, and Zin, 2005) of the value function is still implicit because it remains to solve for b. Isolating b in (6) and plugging the result into (5) gives a recursion for the value-to-wealth ratio $L(z,s)$,

$$L\left(z,s\right)=\left\{\left(1-\beta\right)^{\frac{1}{\rho}}+\beta^{\frac{1}{\rho}}E\left[R_p'^{1-\gamma}L\left(z',s'\right)^{1-\gamma}|z,s\right]^{\frac{1}{\rho}\frac{1-\rho}{1-\gamma}}\right\}^{\frac{\rho}{1-\rho}}, \tag{8}$$

which depends on its next-period values $L(z',s')$, the return-on-wealth R_p', and the parameters of the model. The return on wealth is optimally chosen by the agent, by picking the portfolio weight

$w_{z'}$ on the Arrow-Debreu security for each future state z', to satisfy the first-order condition

$$w_{z'} = \left[\frac{1}{\lambda}\left\{(1-\beta)\, b^{1-\rho} + (1-b)^{1-\rho}\, \beta E\left[R_p'^{1-\gamma} L\left(z', s'\right)^{1-\gamma} |z\right]^{\frac{1-\rho}{1-\gamma}}\right\}^{\frac{1}{1-\rho}-1}\right]^{\frac{1}{\gamma}} \tag{9}$$

$$\times \left[(1-b)^{1-\rho}\, \beta E\left[R_p'^{1-\gamma} L\left(z', s'\right)^{1-\gamma} |z\right]^{\frac{1-\rho}{1-\gamma}-1}\right]^{\frac{1}{\gamma}} \tag{10}$$

$$\times \left\{q\left(z'|z\right)^{\gamma-1} E\left[L\left(z', s'\right)^{1-\gamma} |z'\right] p\left(z'|z, s\right)\right\}^{\frac{1}{\gamma}}, \tag{11}$$

where we have used the fact that $R_i\left(z' = i|z\right) = \frac{1_{\{z'=i\}}}{q(z'|z)}$, and $R_p' = \sum_i w_i R_i' = \sum_i w_i \frac{1_{\{z'=i\}}}{q(z'|z)}$, or $R_p\left(z' = i|z\right) = \frac{w_i}{q(z'=i|z)}$. Using the fact that the terms on the first two lines of (9) do not depend on the future state z', and the weights must sum to one, we get

$$w_{z'} = \frac{\left\{q\left(z'|z\right)^{\gamma-1} E\left[L\left(z', s'\right)^{1-\gamma} |z'\right] p\left(z'|z, s\right)\right\}^{\frac{1}{\gamma}}}{\sum_k \left\{q\left(k|z\right)^{\gamma-1} E\left[L\left(k, s'\right)^{1-\gamma} |k\right] p\left(k|z, s\right)\right\}^{\frac{1}{\gamma}}}, \tag{12}$$

which pins down the optimal return on wealth in state z' conditional on the current state z and signal s, $R_p\left(z'|z, s\right) = \frac{w_{z'}}{q(z'|z)}$.

Plugging the optimal R_p' into (8) and rearranging we get the following recursive expression

$$L\left(z, s\right) = \left\{(1-\beta)^{\frac{1}{\rho}} + \beta^{\frac{1}{\rho}}\left[\sum_{z'} \left\{q\left(z'|z\right)^{\gamma-1} E\left[L\left(z', s'\right)^{1-\gamma} |z'\right] p\left(z'|z, s\right)\right\}^{\frac{1}{\gamma}}\right]^{\frac{\gamma}{\rho}\frac{1-\rho}{1-\gamma}}\right\}^{\frac{\rho}{1-\rho}}, \tag{13}$$

where the expectation on the right is over next period's signal s' conditional on a state z' realizing. Beyond the posterior distribution needed to evaluate this expectation, this is a functional equation in $L(z, s)$ that only depends on state-prices q, physical probabilities p, and preference parameters.

Scaling both sides of the solution to the value function (7) by consumption and rearranging gives the following relationship between the consumption-to-wealth ratio b and the log value-to-consumption ratio $v \equiv \log(V/c)$ (Hansen, Heaton, and Li, 2008):

$$b = (1-\beta)\, e^{(\rho-1)v}, \tag{14}$$

or in terms of the wealth-scaled value function in (7) gives

$$L(z,s) = (1-\beta)\, e^{\rho v(z,s)}. \tag{15}$$

Recursion (13) can therefore be expressed in terms of the log value-to-consumption ratio:

$$v(z,s;\alpha) = \frac{1}{1-\rho} \log \left\{ 1-\beta+\beta^{\frac{1}{\rho}} \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} E\left[e^{(1-\gamma)\rho v(z',s';\alpha)} | z' \right] p(z'|z,s) \right\}^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{\rho}\frac{1-\rho}{1-\gamma}} \right\}. \tag{16}$$

In the case in which the signal is drawn from the uninformative information structure, α_0, this recursion becomes

$$v(z;\alpha_0) = \frac{1}{1-\rho} \log \left\{ 1-\beta+\beta^{\frac{1}{\rho}} \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} e^{(1-\gamma)\rho v(z';\alpha_0)} p(z'|z) \right\}^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{\rho}\frac{1-\rho}{1-\gamma}} \right\}. \tag{17}$$

2.4 The Value of Information

Information structure α is potentially valuable for the agent as it allows her to make better informed investment and consumption decisions using a private signal. Clearly, the value of the information structure to the agent cannot be negative since information can always be ignored and thus one cannot be worse off by obtaining more information.

We define the value of information structure α in state $z \in \{1,...,n\}$ as the fraction of current wealth that the agent is willing to give up in order to obtain a sequence of signals generated by α starting with the current period and for all future periods.

Definition 1. The value of information structure α in state z, before observing the signal s, is the fraction of wealth Ω such that

$$\mu\left[V\left(a\left(1-\Omega(z;\alpha)\right)\right)|z;\alpha\right] = \mu\left[V(a)\,|z;\alpha_0\right], \tag{18}$$

where $\mu[\cdot]$ is the certainty equivalent function defined in (1).

The left hand side of (18) is the certainty equivalent of lifetime consumption to the agent if she gives up a fraction Ω of her current wealth to pay for information structure α. The right hand side is the certainty equivalent of lifetime consumption resulting from resorting to the uninformative information structure α_0 (at no cost).

In what follows, it is useful to work with the transformed value of information $\omega(z;\alpha) \equiv -\log(1-\Omega(z;\alpha))$, which can potentially take values on the entire real line. A first-order expansion of $\Omega(\omega) = 1 - e^{-\omega}$, reveals that $\Omega \approx \omega$ when these are close to zero.

Since the certainty equivalent of a constant is itself, and the value function is homogeneous, (18) can be written

$$\mu\left[L(z,s;\alpha)\,|z;\alpha\right] = e^{\omega(z;\alpha)}L(z;\alpha_0), \tag{19}$$

or using (15), in terms of the value-to-consumption ratio

$$\mu\left[e^{\rho v(z,s;\alpha)}|z;\alpha\right] = e^{\rho v(z;\alpha_0)+\omega(z;\alpha)}, \tag{20}$$

which is equivalent to

$$E\left[e^{(1-\gamma)\rho v(z,s;\alpha)}|z;\alpha\right] = e^{(1-\gamma)[\rho v(z;\alpha_0)+\omega(z;\alpha)]}. \tag{21}$$

Plugging (16) into (21), and using next period's version of (21) we get

$$E\left[e^{(\gamma-1)[\rho v(z;\alpha_0)+\omega(z;\alpha)]}\left\{1-\beta+\beta^{\frac{1}{\rho}}\left[\sum_{z'}\left\{q\left(z'|z\right)^{\gamma-1}e^{(1-\gamma)[\rho v(z';\alpha_0)+\omega(z';\alpha)]}p\left(z'|z,s\right)\right\}^{\frac{1}{\gamma}}\right]^{\frac{\gamma}{\rho}\frac{1-\rho}{1-\gamma}}\right\}^{\frac{\rho(1-\gamma)}{1-\rho}}|z\right] = 1, \tag{22}$$

which given $v(z;\alpha_0)$, q, p, and preference parameters β, γ, and ρ is a set of n equations in the n unknown $\omega(z)$'s, one for each state z.

The conditional population moments (22) provide the theoretical basis for the GMM estimation approach we develop below. This representation is particularly useful because the outer expectation over signals can be estimated using sample averages without taking a stance on the exact signal space and the conditional distribution of the signals. In Appendix A.1 we provide expressions for some limiting cases (e.g. log utility).

3 Estimation

In this section we develop a general procedure for estimating the value of information from (22). To this end we follow prior literature and take the parameters β, γ, and ρ as given. We then estimate the state-dependent value of information $\omega(z, \alpha)$, $z = 1, \ldots, n$ using variations in state prices q and state probabilities p, both estimated from options as we elaborate below. To distinguish between an informative information structure α and an uninformative information structure α_0 – we focus on informational events occurring at some dates t and observe price behavior prior to the information release (beginning of date t) and shortly after the information release (date $t + dt$).

The crux of our empirical approach builds on a long tradition that uses large sample means to estimate population moments (Hansen and Singleton, 1982). On every observation date t where a particular signal s_t occurs, we form three $n \times n$ data matrices: state prices $Q_t = \{q_{ijt}\}_{i,j=1,\ldots,n}$, where q_{ijt} is a shorthand for $q(z_{t+1} = j | z_t = i)$, prior physical probabilities $P_t = \{p_{ijt}\}_{i,j=1,\ldots,n}$, where $p_{ijt} = p(z_{t+1} = j | z_t = i)$, and posterior physical probabilities $P_t^*(s_t) = \left\{p_{ijt}^*(s_t)\right\}_{i,j=1,\ldots,n}$, where $p_{ij}^*(s_t) = p_\alpha(z_{t+1} = j | s_t, z_t = i)$. We use S&P 500 index options to estimate Q_t, and then recover the physical probabilities P_t using both parametric and recently developed non-parametric approaches. We estimate the posterior physical probabilities by re-estimating P_t at time $t+dt$ – right after the information associated with signal s_t is released. We then use the Generalized Method of Moments (GMM) to estimate the information values ω in (22). By conditioning on dates of particular information releases (e.g. Unemployment report), we estimate the value of information released on those dates, from the (non-linear) average change in the transition probabilities matrices.

Intuitively, the thought experiment considers a small agent who can trade state contingent securities with prices Q_t on each date t, but can purchase the information that markets will have at time $t + dt$, without affecting prices. By observing many realized updates to her information set, we back out her willingness to pay for the average signal. We provide more details of of the data construction and methodology next, before moving to the results in Section 4.

3.1 Estimating the State Price Matrix

Following Breeden and Litzenberger (1978), we estimate state prices from S&P 500 index (SPX) European options, where the state is captured by values of the index. Options data is from OptionMetrics, which provides for each trading day from January 4, 1996 to August 31, 2015, a panel of options prices with strike prices on each side of the current index value, and at various maturities (terms). This data also includes the closing price of the SPX, a term structure of zero coupon risk-free rates, and an estimate of the dividend yield. Moneyness as usual is defined as the option's strike price over the SPX spot price, and terms are measured in years.

We apply several standard filters to avoid relying on stale prices. We only use at or out of the money calls (moneyness at least 1) and puts (moneyness at most 1) that are 7 days or more to maturity, and which have a strictly positive trading volume.

Since the Breeden and Litzenberger (1978) method relies on second-order-differences, the smoothness of the option price in its strike is essential. For this purpose, we use the Carr and Wu (2010) approach to construct an implied volatility surface on each day, which parameterizes the implied variance dynamics as a mean-reverting lognormal process (LNV). Carr and Wu show that, under their assumptions, implied variance as a function of log moneyness k and time to maturity τ is determined by an easy to solve quadratic equation. The whole implied variance surface is then determined by six parameters, which are related to the stock price and implied variance dynamics.

Unlike Carr and Wu who use a Kalman filter and assume these parameters are highly persistent (a random walk), we estimate the volatility surface at time t using only time t options. We do so because our value of information estimates come from changes in the implied volatility surface upon news arrival, and we would not want an arbitrary assumption about parameter persistence to drive our results.

Figure 2 provides an example volatility surface. As expected we have more options close to the spot price with moneyness around 1 that expire within a year, than farther out of the money or with longer time to expiration. The mean R-squared over the entire sample is 0.93 and the median is 0.94.

Following Ross (2015), on each date t we discretize the state relative to the current spot price of the SPX denoted by S_t, so that log returns on the SPX between t and $t+1$, $r_{t+1} = \log \frac{S_{t+1}}{S_t}$, takes one of 11 possible equally-spaced values in $[-0.24, 0.24]$. We focus on one month transition matrices, so that, for example, each entry q_{ijt} in the state price transition matrix is the price in state i of a security paying \$1 if state j realizes within a month (i.e., at time $t+1$). The estimation of the posterior probabilities $p^*_{ijt}(s_t)$ is done at the end of the day on which the signal is publicly revealed. The prior probabilities, p_{ijt} are estimated one trading day before that. That is, we take dt to be 1 day.

Note that the state-price matrix Q_t consists of prices of Arrow-Debreu securities maturing in one month for the current state as well as other potential states that did not realize at time t. However, as pointed out by Ross (2015), the volatility surface, instead encodes the state prices corresponding only to the state that was realized at time t for various horizons. Ross overcomes this problem using a least squares projection utilizing the fact that the moneyness-term matrix can be constructed from powers of Q_t. We take a different route by assuming that, consistent with the Carr-Wu model, the implied volatility surface is only a function of moneyness and term, but does not depend on the spot price. This assumption allows us to calculate implied volatility for states other than the current one and construct Q_t directly. By relying on the Carr-Wu model we avoid most forms of arbitrage by construction.

We provide further details in Section B.1, but one can get an intuitive sense of the state price matrices we estimate from Figure 3, which shows the Q_t we estimate from the implied vol surface of Figure 2 for a particular day. The center row represents prices of securities paying \$1 if the state next month corresponds to each column. This is a matrix where each off-center row is a shifted version of the center one. The figure also shows for comparison the indirect approach of Ross (2015), which inverts Q_t using least squares, while imposing that Q_t is unimodal. Both approaches start with the same state-term prices derived from the fitted Carr and Wu (2010) volatility surface. Because we identify the value of information from day-to-day changes in Q_t, we only keep observations where all elements of Q_t are positive (and therefore admit no arbitrage) on consecutive trading days.

3.2 Recovering the Physical Probability Matrix

Unlike with state prices, the estimation (recovery) of forward looking physical probabilities is an active topic of research.[7] Traditional approaches to this problem use historical realizations to proxy for future probabilities (see Jackwerth (2000) and Aït-Sahalia and Lo (2000)). Such an approach, however, results in a slow-moving probability distribution, which is inappropriate for our purposes. Another approach has been to place sufficient parametric assumptions on marginal utility and distributional assumptions. Ross (2015) reignited this literature by deriving a recovery result, which allows one to estimate the physical probability of asset returns and the market's risk aversion, the kernel, from the state price transition process alone. Ross's theoretical result has been extended substantially (Carr and Yu, 2012; Walden, 2014), but also criticized (Dubynskiy and Goldstein, 2013; Tran and Xia, 2014; Borovička, Hansen, and Scheinkman, forthcoming). Empirically, Malamud (2015) finds that probabilities recovered from index options are informative about expected returns. Recovery is not the focus of the current paper, but a necessary step nonetheless. We therefore take the following two distinct approaches.

3.2.1 Parametric Recovery

The stochastic discount factor of a representative agent with Epstein and Zin (1989) preferences is such that

$$q\left(z_{t+1}|s_t, z_t\right) = \beta^{\frac{1-\gamma}{1-\rho}} e^{\frac{\rho-\gamma}{1-\rho} r_p(z_{t+1})} E\left[e^{-\rho\left(\frac{1-\gamma}{1-\rho}\right) g(z_{t+1}, s_{t+1})} | z_{t+1}\right] p_\alpha\left(z_{t+1}|s_t, z_t\right), \tag{23}$$

where $r_p\left(z_{t+1}\right) = \log R_p\left(z_{t+1}\right)$ is the log return on wealth, $g\left(z_{t+1}, s_{t+1}\right)$ is the consumption growth rate in state z_{t+1}, and the expectation is taken over signals s_{t+1}. Note that here we allow the state price q to to depend on the signal s_t (when it becomes public). The reason this is warranted is that the stochastic discount factor is determined by the trade of a representative investor rather than the "small investor" we considered in Section 2.1. There the prices of assets were given exogenously and not affected by the trades of the small agent, whereas here prices reflect all publicly available

[7] The "recovery" nomenculture is due to Dybvig and Rogers (1997) who study the recoverability (or identifiability) of preferences in a neoclassical investment problem.

information including the signal - once it is revealed.[8]

Our goal is to estimate $p_\alpha\left(z_{t+1}|s_t, z_t\right)$ from $q\left(z_{t+1}|s_t, z_t\right)$ using (23). To this end we make the following assumptions:

A1.1 The return on wealth R_p coincides with the return on the S&P 500 index.

A1.2 The consumption growth rate can be decomposed into $g\left(z_{t+1}, s_{t+1}\right) = g_z\left(z_{t+1}\right) + g_s\left(s_{t+1}\right)$.

A1.3 Expected exponentiated deviations are independent of z_{t+1}: $E\left[e^{-\rho\left(\frac{1-\gamma}{1-\rho}\right)g_s(s_{t+1})}|z_{t+1}\right] = G$ for some $G > 0$ (Note that consumption growth can still depend on the signal s_{t+1}).

A1.4 Return on wealth and consumption growth are proportional to each other: $R_p\left(z_{t+1}\right) = \Phi E\left[\frac{c_{t+1}}{c_t}|z_{t+1}\right]$, for some $\Phi > 0$.

Using the fact that the physical probabilities must sum to one, (23) along with Assumption A1.1-A1.4 we obtain

$$p^*_{ijt}\left(s_t\right) = \frac{e^{\gamma r_p(z_{t+1}=j)} q_{ijt}\left(s_t\right)}{\sum_k e^{\gamma r_p(z_{t+1}=k)} q_{ikt}\left(s_t\right)}. \tag{24}$$

This "exponential tilting" formula is a straightforward and intuitive mapping between state prices and posterior physical probabilities, governed only by the coefficient of relative risk aversion γ. Intuitively, the prices of Arrow securities paying in good states with high returns are relatively cheap, with the size of the wedge determined by γ.[9] As a special case, when the signal s_t is taken from the uninformative information structure α_0, (24) yields the prior physical probability matrix P_t.

3.2.2 Non-parametric Recovery a-la Ross (2015)

Our second recovery approach relies on a recovery result by Ross (2015), which builds on earlier work by Backus, Gregory, and Zin (1989), and Hansen and Scheinkman (2009). It makes an alternative, and somewhat weaker assumption, that state prices have the following form

[8]Note the market incompleteness introduced by the assumption that the signal can affect market prices, but only the realized state z determines which securities pay-off.

[9]Note that this mapping is identical to the one that results if the representative agent is a expected power utility maximizer with risk aversion γ.

$$q_{ijt}\left(s_t\right) = \delta \frac{m_{jt}}{m_{it}} p^*_{ijt}\left(s_t\right), \tag{25}$$

for some $\delta > 0$ and a vector $m \in \mathbb{R}^n$ with positive entries. Ross shows that the problem of recovering P_t from Q_t, can be posed as an eigenvalue decomposition problem

$$Q_t e_t = \delta e_t,$$

which by the Perron–Frobenius Theorem has only one eigenvector, e, with strictly positive entries (up to scale). The uniqueness of the decomposition implies that $m_{it} = 1/e_{it}$ for all $i = 1, ..., n$, and that

$$p^*_{ijt}\left(s_t\right) = \frac{1}{\delta} \frac{m_{it}}{m_{jt}} q_{ijt}\left(s_t\right) = \frac{1}{\delta} \frac{e_{jt}}{e_{it}} q_{ijt}\left(s_t\right). \tag{26}$$

As pointed out by Borovička, Hansen, and Scheinkman (forthcoming), assumption (25) is not innocuous, and in a more general setting, (26) recovers a probability measure that reflects long-term pricing, as opposed to transition probabilities, which may be quantitatively important (Alvarez and Jermann, 2005; Bakshi and Chabi-Yo, 2012).

Future literature could potentially relax these assumptions and find better ways of recovering the natural transition probabilities. Such recovery functions can be plugged-in to our framework, and used to recover more accurate values of information, although the comparative statics results below are likely to go through.

3.3 Informational events and timing

Stock returns realized around pre-scheduled macroeconomic announcements, such as the employment report and the FOMC statements, account for virtually all of the market equity premium during our sample period (Ai and Bansal, 2016). By conditioning (22) on dates of particular information releases, we estimate the value of information released on those dates, from the change in the transition probabilities matrices. To fix ideas, consider an information release that occurs on the first day of each month. We then estimate state prices and probabilities on the day prior to the information (date t) and on the day immediately after the information release (date $t + dt$). We consider the information structure just before the information release as α_0 and the one just after the information release as α. Applying GMM on the moment conditions (22) we can repeat this process each month and estimate the value of information $\omega_i(\alpha)$ for each state $i = 1, ..., n$. The point is that using α is tantamount to making consumption/investment decisions which condition on the information release each month, whereas using α_0 reflects consumption/investment decisions which only rely on the probabilities prior to the information release.

Our data on macroeconomic indicators release dates comes from Bloomberg's Economic Calendar. We focus on the following key indicators:

1. **FOMC Decision**. The Federal Open Market Committee (FOMC) of the Federal Reserve Board of Governors Fed funds rate decision, which usually occurs eight times a year for 153 events during our sample period. Because Lucca and Moench (2015) document large average excess returns on U.S. equities in anticipation of FOMC decisions (a pre-FOMC announcement drift), we also report results for the trading day period to FOMC decisions. We restrict our sample to scheduled releases.[10]

2. **Unemployment.** The employment situation is reported monthly by the Bureau of Labor Statistics (BLS), usually on the first Friday of the month for 217 events during our sample period.

3. **GDP.** Gross domestic product is reported by the Bureau of Economic Analysis (BEA). Initial reports are for the previous quarter, which are then revised in each of the following two months. We have 220 such events during our sample period.

4. **Jobless Claims.** Initial jobless claims are reported each Thursday by the U.S. Department of Labor. We have 940 such events during our sample period.

5. **Mortgage Apps.** Mortgage Bankers' Association purchase applications index of applications at mortgage lenders. This leading indicator for single-family home sales and housing construction is reported each Wednesday. We have 598 such events during our sample period.

[10]By following the procedure of Lucca and Moench (2015) we exclude 6 unscheduled FOMC announcements, which seem to occur after sharp market drops (e.g. October 8, 2008). Including these hardly changes our estimates.

6. **Consumer Comfort.** The Bloomberg Consumer Comfort Index is a weekly, random-sample survey tracking Americans' views on the condition of the U.S. economy, their personal finances and the buying climate. The survey is formerly sponsored by ABC News since 1985. It is reported each Thursday. We have 601 such events during our sample period.

In the cases where events occur after trading hours, we designate the following trading day as the event day.

To get a better sense of our data, Figure 5 plots the mean prior and posterior risk-neutral probabilities of the future state (S&P 500 return in one month), estimated from options on the day before and on the day of Unemployment reports. We use these to recover the corresponding physical prior and posterior return distributions, here using the parametric recovery approach. As expected, for both the risk-neutral and the physical distributions, the distribution on average becomes more centered around its mean when such information arrives due to a reduction in uncertainty. The risk-neutral distribution is, however, relatively more pessimistic than the physical one because it also accounts for a risk premium.

3.4 GMM estimation

We now turn to the estimation of the value of information $\omega = (\omega_1, ..., \omega_n)$ for states $1, ..., n$ from (22). Let f denote the empirical counterpart to the bracketed term in (22), i.e.,

$$f_i(\omega; x_t) = e^{(\gamma-1)[\rho v_{it}(\alpha_0)+\omega_i]} \left\{ 1 - \beta + \beta^{\frac{1}{\rho}} \Sigma_{it}^{\frac{\gamma(1-\rho)}{\rho(1-\gamma)}} \right\}^{\frac{\rho(1-\gamma)}{1-\rho}} - 1, \tag{27}$$

where

$$\Sigma_{it} = \sum_k (p^*_{ikt})^{\frac{1}{\gamma}} (q_{ikt})^{-\frac{1-\gamma}{\gamma}} e^{\left(\frac{1-\gamma}{\gamma}\right)[\rho v_{kt}(\alpha_0)+\omega_k]}. \tag{28}$$

Here x_t denotes time varying data consisting of Q_t, P^*_t, and the no-information case log value-to-consumption ratio v_t (which depends on the prior probabilities P_t).[11] We estimate the vector v_t by numerically finding a fixed point to (17) separately on each day t. Note that $f_i(\omega; x_t)$ are simply the time t residuals associated with (22).

[11]For brevity we use here v_{it} as a shorthand for $v(z_t = i)$ given the information structure α_0.

We construct the GMM estimator of ω for each subsample of T observations by minimizing the quadratic form

$$\omega_T = \arg\min_{\omega} J_T(\omega), \tag{29}$$

where $J_T(\cdot)$ is the GMM objective function given by

$$J_T(\omega) = \frac{1}{2} g_T(\omega)' W g_T(\omega), \tag{30}$$

and the sample moments are

$$g_T(\omega) = \frac{1}{T}\sum_{t=1}^{T} f(\omega, x_t). \tag{31}$$

Calculation shows that the ij-th entry of the Jacobian of f is

$$\begin{aligned} Df_{ij}(\omega; x_t) &= (1-\gamma)\, e^{(\gamma-1)[\rho v_{it}(\alpha_0)+\omega_i]} \\ &\times \Bigg\{ \left[1-\beta+\beta^{\frac{1}{\rho}} \Sigma_{it}^{\frac{\gamma(1-\rho)}{\rho(1-\gamma)}}\right]^{\frac{\rho(1-\gamma)}{1-\rho}-1} \beta^{\frac{1}{\rho}} \Sigma_{it}^{\frac{\gamma(1-\rho)}{\rho(1-\gamma)}-1} \left(p^*_{ijt}\right)^{\frac{1}{\gamma}} (q_{ijt})^{-\frac{1-\gamma}{\gamma}} e^{\left(\frac{1-\gamma}{\gamma}\right)[\rho v_{jt}(\alpha_0)+\omega_j]} \\ &-1_{\{i=j\}} \left[1-\beta+\beta^{\frac{1}{\rho}} \Sigma_{it}^{\frac{\gamma(1-\rho)}{\rho(1-\gamma)}}\right]^{\frac{\rho(1-\gamma)}{1-\rho}} \Bigg\}. \end{aligned} \tag{32}$$

We report second-stage (efficient) GMM estimates, where the first stage uses an identity weighting matrix W, and the second stage uses the inverse of the Newey-West estimate of the covariance matrix with two lags V_T.[12] The standard errors we report are the square-root of the entries along the diagonal of

$$var(\omega_T) = \frac{1}{T}\left(Dg_T' V_T^{-1} Dg_T\right)^{-1} \tag{33}$$

where $Dg_T(\omega) \equiv \frac{\partial g_T'}{\partial \omega_T} = \frac{1}{T}\sum_{t=1}^{T} Df(\omega; x_t)$ is the mean Jacobian of f.

We estimate ω_T for different subsamples T associated with announcement days of the different macroeconomic indicators discussed in Section 3.3. We report the value of information as percent of wealth $\Omega_T = 1 - e^{-\omega_T}$, for the middle (current) state (estimates for other states are similar in magnitude given a set of parameters). We calculate standard errors using the Delta method.

[12] Theoretically, because the number of moments equals the number of estimated parameters in our setting, under the identifying assumption that $g(\omega_0) = 0$ at the true parameter vector ω_0, the weighting matrix should not matter. Reassuringly, our second-stage estimates are almost identical to first-stage ones.

Because $\Omega'(\omega) = e^{-\omega} = 1 - \Omega(\omega)$, the standard errors for $\Omega(\omega_T)$ are simply those of ω_T multiplied by $1 - \Omega(\omega_T)$.

3.5 A useful approximation

A first-order Taylor approximation of the moment equation (31) around $\omega = 0$, which is a reasonable guess for many events and parameter values gives

$$0 = g_T(\omega) \approx g_T(0) + Dg_T(0)\,\omega. \tag{34}$$

Solving for ω we get a first-order approximation for small values of information

$$\tilde{\omega}_T = -Dg_T(0)^{-1}\,g_T(0), \tag{35}$$

which requires no optimization and turns out to be quite accurate. Thus in addition to being a good initial guess for GMM estimation, for some applications it could suffice to simply calculate this approximation.

4 Estimation results

4.1 Benchmark estimates

Our benchmark parameters deliberately focus on the expected utility specification $\gamma = \rho$, where there is neither preference nor aversion to early resolution of uncertainty. Specifically we set the time discount rate to $\beta = 0.998$, and a moderate degree of relative risk aversion $\gamma = 5$.

It is well-known that agents have preference for early resolution of uncertainty if and only if $\gamma > \rho$ (Kreps and Porteus, 1978; Epstein and Zin, 1989; Epstein, Farhi, and Strzalecki, 2014). Such preference is entirely about the attitude of the agent toward uncertainty, even when she cannot do anything to alter its future consumption stream – the so called psychic value of information. Our thought exercise is quite different, and instead focuses on what Epstein, Farhi, and Strzalecki (2014) call the instrumental value of information. We therefore find it instructive to shutdown the psychic channel at first.

The first three columns of Table 2 report our benchmark estimates of the value of information Ω_T for the middle (current) state, using parametric recovery. We find that the agent is willing to pay between 1.5 and 3.2 percent of her wealth for the informational content of the leading macroeconomic indicators. Standard errors are small, and allow us to reject zero at the one percent significance level. The third column reports the first-order approximation for the value of information and yields very similar estimates.

We next consider the log utility case: $\gamma = \rho = 1$. This case offers an interesting point of view since Cabrales, Gossner, and Serrano (2013) show that this case sets an upper bound for the value of information in a static expected utility setup. Table 2 shows that such an agent is willing to pay nearly all of her wealth for the same information. The table shows that the percentage of wealth a log utility agent is willing to pay for macroeconomic signals ranges between 90% and 97%. The standard errors show that these estimates are indistinguishable from 100%. Thus, while the log utility case offers a tractable upper bound on the value of information, this bound seems somewhat lax when considering levels of risk aversion that are commonly used to fit stylized asset pricing facts (e.g. Bansal and Yaron, 2004).

The last three columns of Table 2 consider the case in which the $EIS = 1/\rho = 0.9$. Consistent with a substantial timing premium (Epstein, Farhi, and Strzalecki, 2014), we find that the value of these signals increases 6-fold relative to the expected utility benchmark.

Table 3 reports the results using the same parameter sets, but using Ross non-parametric recovery instead. We can see that the magnitude of the estimates is remarkably similar to the parametric recovery ones, suggesting that these are mostly driven by the preference parameter values instead. The recovery method does change the ordering of the information sources, but statistically, the values of information released on these events are similar.

4.2 Comparative statics and robustness

We next consider several comparative statics exercises around these benchmark parameters for Jobless Claims reports. The results are qualitatively the same for other events.

Figure 6a shows that increasing the coefficient of relative aversion γ while keeping it equal to the inverse of the EIS ρ, reduces the value of information. Intuitively, investment-relevant information is less useful to a more risk averse agent because her willingness to change her portfolio to take into account the information is limited. An agent with an extreme level of risk aversion would choose a near-riskless portfolio regardless of the signal.

The main advantage of the recursive utility framework is the separation between risk aversion and intertemporal substitution. Figure 6b illustrates the effect of increasing risk aversion weakens when the EIS is held fixed. With recursive utility, as we increase γ there is a counter effect coming from a stronger preference for early resolution of uncertainty, even if the agent's actual consumption and investment are less sensitive to the information.

Figure 6c shows that the effect on the value of information of increasing the EIS while holding risk aversion fixed is positive for values of EIS below one. As the EIS increases, the agent is more willing to substitute intertemporally and use the information to increase her future consumption. The preference for early resolution of uncertainty also works in the same direction. We have thus far been unable to estimate the value of information for values of EIS above one.

Figure 6d shows that, as expected, when the time discount factor β increases, and the agent attaches more value to future periods, the value of information increases. Relatedly, we chose to focus on the monthly horizon in our benchmark tests because options are relatively highly traded around this expiration. The horizon in our model is the time between the arrival of any two signals belonging to the stream of signals being valued. We should expect more frequent signals to be more valuable, and indeed this is what we find in Figure 7a.[13]

It is interesting to see how our results change as we change some of the choices we made in optimization and recovery.

In our benchmark estimates, we stop the search for a minimum of the GMM objective when $|\omega_{i+1} - \omega_i| \leq 10^{-4}$, that is when the change in the value of information is less than a basis point. Figure 7b shows that increasing the numerical precision even to 10^{-8} does not change our estimates.

Figure 7c shows, however, that as we increase the number of states n used to discretize the

[13] As we change the the horizon we make a corresponding change to the time discount factor β so that the annual discount rate remains the same (i.e. $\beta(\tau) = e^{-r \times \tau}$ for a fixed r).

state space, the estimated value of information increases somewhat. The estimates are nonetheless statistically indistinguishable from each other, though different from zero.

A second choice we had to make was the distance (in log returns) between each state. A larger value is better at capturing the tails of the distribution, especially in later periods, but also gives a more coarse estimate of the center of the distribution. Figure 7d shows that at least in the neighborhood of our benchmark estimates, this choice makes a small difference.

5 The Role of Rational Expectations

The analysis thus far does not make any assumption on the rationality of the recovered probability distribution in the following sense. The probabilities recovered at date t may or may not satisfy the law of total probability with respect to realized signals

$$p\left(z'|z\right)=\sum_{s}p_{\alpha}\left(s|z\right)p_{\alpha}\left(z'|z,s\right). \tag{36}$$

This potential departure from rational expectations has implications for our estimates of the value of information. To illustrate this point we consider the case of log utility, $\gamma=\rho=1$, for which the value of information is particularly tractable. Cabrales, Gossner, and Serrano (2013) have shown that in a one period investment choice model with log utility, the value of information coincides with the expected reduction in entropy. Our dynamic consumption/investment setup leads to a similar result when (36) holds, but differs substantially otherwise.

To see this point, in the log utility case the value-to-consumption recursions (16) and (17) specialize to

$$\begin{aligned} v\left(z,s;\alpha\right) &= \beta\log\beta+\beta E\left[\log p_{\alpha}\left(z'|z,s\right)-\log q\left(z'|z\right)+v\left(z',s';\alpha\right)|z,s;\alpha\right] && (37)\\ v\left(z;\alpha_{0}\right) &= \beta\log\beta+\beta E\left[\log p\left(z'|z\right)-\log q\left(z'|z\right)+v\left(z';\alpha_{0}\right)|z;\alpha_{0}\right], && (38)\end{aligned}$$

and (21) specializes to

$$\omega\left(z;\alpha\right)=E\left[v\left(z,s\right)|z;\alpha\right]-v\left(z;\alpha_{0}\right), \tag{39}$$

which shows that the value of information is the expected increase in the value-to-consumption ratio due to the signal s.

To obtain further intuition denote by $H(z,s;\alpha) \equiv -\sum_{z'} p_\alpha(z'|z,s)\log p_\alpha(z'|z,s)$ the entropy of the future state z' distribution given the current state z and signal s. Similarly, $H(z;\alpha_0) \equiv -\sum_{z'} p(z'|z)\log p(z'|z)$ is the unconditional entropy in state z. Recall that entropy is a measure of the dispersion of the probability distribution. Thus, $H(z,\alpha_0) - H(z,s;\alpha)$, the reduction in entropy associated with signal s, is a measure of the information in this signal.

We can express the value of information $\omega(z)$ by plugging (37) and (38) into (39) and using the one-period ahead version of (39) as

$$\begin{aligned}
\omega(z;\alpha) = {} & \beta\sum_s\sum_{z'}\left[\log p_\alpha(z'|z,s) - \log q(z'|z) + v(z';\alpha_0) + \omega(z';\alpha)\right] p_\alpha(z'|z,s)\, p_\alpha(s|z) \\
& - \beta\sum_{z'}\left[\log p(z'|z) - \log q(z'|z) + v(z';\alpha_0)\right] p(z'|z) \\
= {} & \beta\sum_s p_\alpha(s|z)\left[\sum_{z'} p_\alpha(z'|z,s)\log p_\alpha(z'|z,s) - \sum_{z'} p(z'|z)\log p(z'|z)\right] \\
& + \beta\sum_{z'}\left[\sum_s p_\alpha(s|z)\, p_\alpha(z'|z,s) - p(z'|z)\right]\left[v(z';\alpha_0) - \log q(z'|z)\right] \\
& + \beta\sum_{z'}\left[\sum_s p_\alpha(z'|z,s)\, p_\alpha(s|z)\right]\omega(z';\alpha).
\end{aligned} \tag{40}$$

If the rational expectations condition (36) is satisfied, then this expression boils down to

$$\omega(z;\alpha) = \beta I^p(z;\alpha) + \beta\sum_{z'}\omega(z';\alpha)\, p(z'|z), \tag{41}$$

where $I^p(z;\alpha) \equiv \sum_s\left[H(z;\alpha_0) - H(z,s;\alpha)\right] p_\alpha(s|z)$ is the expected reduction in entropy from the signal. This is a similar result to the one obtained in Cabrales, Gossner, and Serrano (2013). Unlike in their static investment model, here the entire future path of entropy reductions matters. Thus, the case of log utility provides a tractable benchmark in which the value of information is closely tied to the change in entropy. The asset pricing literature, however, typically considers risk aversion and EIS parameters that deviate significantly from this case. In those cases entropy does not capture the value of information and our estimation approach could prove more applicable.

More generally, our estimation approach does not impose (36) making it more robust, but also potentially less efficient if prior probabilities are actually rational in the population.[14] While it may seem like a natural assumption, (36) turns out to be a poor description of our data and so the value of information in the log utility case often departs from (41) substantially.

In general, (40) can be re-written as

$$\omega(z;\alpha) = \beta\left\{I(z;\alpha) - E\left[\Delta E\left[\log q(z'|z)|z,s;\alpha\right]|z\right] + E\left[\Delta E\left[v(z')|z,s;\alpha\right]|z\right] + E\left[E\left[\omega(z';\alpha)|z,s;\alpha\right]|z\right]\right\}, \tag{42}$$

where $\Delta E\left[\tilde{X}|z,s;\alpha\right] \equiv E\left[\tilde{X}|z,s;\alpha\right] - E\left[\tilde{X}|z;\alpha_0\right]$ is the change in expectation of random variable $\tilde{X}$ after observing a signal s. Our corresponding GMM sample moments replace the expectations over signals by sample averages, and can be expressed succinctly in vector notation

$$\omega = \beta\left[I^p(\alpha) + I^q(\alpha) + I^v(\alpha)\right] + \beta\left[\frac{1}{T}\sum_t P_{t+dt}\right]\omega, \tag{43}$$

where the entries for state z of the I vectors are

$$\begin{aligned}
I^p(z;\alpha) &\equiv \frac{1}{T}\sum_t\sum_{z'}\left[p_{t+dt}(z'|z)\log p_{t+dt}(z'|z) - p_t(z'|z)\log p_t(z'|z)\right]\\
I^q(z;\alpha) &\equiv -\frac{1}{T}\sum_t\sum_{z'}\left[p_{t+dt}(z'|z) - p_t(z'|z)\right]\log q_t(z'|z)\\
I^v(z;\alpha) &\equiv \frac{1}{T}\sum_t\sum_{z'}\left[p_{t+dt}(z'|z) - p_t(z'|z)\right]v_t(z').
\end{aligned}$$

To understand the economic meaning of this expression note that in the log utility case $p(z'|z) = w(z'|z)$ and $q(z'|z) = 1/R(z'|z)$ (see Equation (9)). Thus, in this case $I^p(z;\alpha) = \frac{1}{T}\sum_t \Delta E_t \log w(z'|z)$ and $I^q(z;\alpha) = \frac{1}{T}\sum_t \Delta E_t \log R(z'|z)$.

Intuitively, the value of information is large when the expected reduction in entropy due to the signal $I^p(\alpha)$ is large. This channel operates by concentrating the distribution of portfolio weights on the more likely states. As discussed above, if we imposed rational expectations, this would be the only surviving term.

[14] See Cochrane (2005, Ch. 16) for a discussion of the tradeoff between efficiency and robustness in similar contexts.

Without imposing (36), the value of information is lower when the expected reduction in state-price "entropy" $I^q(\alpha)$ is large. This term coincides with the expected change in the expected log return on a portfolio that places unit weights on each of the Arrow-Debreu securities, and captures expected changes in the investment opportunity set. Moreover, the signal can be valuable if it is expected to increase the expected future value-to-consumption ratio. Finally, when the time discount factor is large, and therefore B is large, so is the value of information.

The moment equation (43) can be solved in closed-form:

$$\omega = B\left[I^p\left(\alpha\right) + I^q\left(\alpha\right) + I^v\left(\alpha\right)\right] = \omega^p + \omega^q + \omega^v \tag{44}$$

where $\omega^p \equiv BI^p\left(\alpha\right)$, $\omega^q \equiv BI^q\left(\alpha\right)$, $\omega^v \equiv BI^v\left(\alpha\right)$, and $B \equiv \beta\left[Id - \beta\frac{1}{T}\sum_t P_{t+dt}\right]^{-1}$.

When the agent is myopic, the β multiplying the probabilities matrix P in B is zero. Imposing rational expectations on the average change in probabilities, $\frac{1}{T}\sum_t\left[p_{t+dt}\left(z'|z\right) - p_t\left(z'|z\right)\right] = 0$ and (44) specializes to the Cabrales, Gossner, and Serrano (2013) result. Unlike in their static investment model, here the entire future path of entropy reductions matters. The case of log utility provides a useful tractable benchmark in which the value of information is closely tied to the change in entropy. The asset pricing literature, however, typically considers risk aversion and EIS parameters that deviate significantly from this case. In those cases entropy does not capture the value of information and our estimation approach could prove more applicable.

The empirical decomposition into the three terms of ω is illuminating. Table 4 reports direct estimates of the value of information for the log utility case based on (44). These estimates are quite similar to the GMM estimates reported for the log utility case in Table 2. Based solely on the mean reduction in entropy of the physical probabilities ω^p, FOMC decisions are the most informative, followed by unemployment and GDP releases. Pre-FOMC is less informative than FOMC based on this metric alone. These results flip, however, once the mean change in the expected return on an equal-weighted portfolio of Arrow-Debreu securities ω^q is taken into account.

Intuitively, these results say that prior to FOMC decisions, state prices implied by options markets adjust such that trading on FOMC decisions is less valuable, than trading on the information released to options traders the prior day. This can be because of information leaks prior to FOMC decisions, or because of higher expected returns being required to compensate investors for the risk of holding the market on FOMC decision days. In any case, they would be missed if we imposed the rational expectations assumption (36).

6 Conclusion

We derive a general expression for the value of information to an investor in a dynamic environment with recursive utility and provide a framework for its estimation. We then estimate the values of key macroeconomic indicators that are implied by index options, to an early-informed price-taking investor. Comparative statics exercises show that time discounting, risk aversion and attitudes toward intertemporal substitution all play an important role in determining the value of information.

Our estimation approach facilitates a general structural estimation of the value of informational signals. Thus, we contribute to the literature by shifting from the traditional view that information is beneficial to a more practical level of quantifying the benefit in an informational signal. While we apply this approach to macroeconomic indicators, future research may use our methodology to study the value of information at the firm level, such as in the context of mergers and acquisitions, earnings releases, or analyst forecasts.

Estimates of the value of information could be applied to the pricing of information services and for the enforcement and litigation of insider trading. US and UK securities laws tie penalties to the profit gained or loss avoided as a result of insider trading and communication. For example, the Financial Conduct Authority's handbook stipulates it "will seek to deprive an individual of the financial benefit derived as a direct result of the market abuse (which may include the profit made or loss avoided) where it is practicable to quantify this." We estimate an *ex ante* measure of benefit from market abuse, which depends on the individual investor, her wealth, preferences, and investment opportunity set. Unlike commonly-used *ex post* measures of profit, our estimates do not depend on the random nature of securities markets or on the actions of other market participants.[15]

[15] See Financial Conduct Authority's Decision Procedure and Penalties manual, Section 6.5C for the UK, and 15 U.S.C. § 78u-1 for the US.

Appendix

A Proofs and derivations

A.1 Limiting cases for the value of information

Here, for completeness, we provide results for some limiting cases of interest.

The log value-to-consumption ratio for the no-information benchmark in state z depends on state pries q and natural transition probabilities p via the following recursions:

$$v(z) \equiv \log \frac{V}{c} = \begin{cases} \frac{1}{1-\rho} \log \left\{ 1 - \beta + \beta^{\frac{1}{\rho}} \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} e^{(1-\gamma)\rho v(z')} p(z'|z) \right\}^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{\rho} \frac{1-\rho}{1-\gamma}} \right\} & \text{if } \rho \neq 1, \gamma > 1 \\ \frac{1}{1-\gamma} \log \left\{ 1 - \beta + \beta^{\frac{1}{\gamma}} \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} e^{(1-\gamma)\gamma v(z')} p(z'|z) \right\}^{\frac{1}{\gamma}} \right] \right\} & \text{if } \rho = \gamma \neq 1 \\ \frac{1}{1-\rho} \log \left\{ 1 - \beta + \beta^{\frac{1}{\rho}} \exp\left(\frac{1-\rho}{\rho} \sum_{z'} \left[\log \frac{p(z'|z)}{q(z'|z)} + \rho v(z') \right] p(z'|z) \right) \right\} & \text{if } \rho \neq 1, \gamma = 1 \\ \beta \log \beta + \frac{\beta\gamma}{1-\gamma} \log \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} e^{(1-\gamma) v(z')} p(z'|z) \right\}^{\frac{1}{\gamma}} \right] & \text{if } \rho = 1, \gamma > 1 \\ \beta \log \beta + \beta \sum_{z'} \left\{ \log \frac{p(z'|z)}{q(z'|z)} + v(z') \right\} p(z'|z) & \text{if } \rho = \gamma = 1 \end{cases} \tag{45}$$

The corresponding equations given a signal s are

$$v(z,s) = \begin{cases} \frac{1}{1-\rho} \log \left\{ 1 - \beta + \beta^{\frac{1}{\rho}} \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} E\left[e^{(1-\gamma)\rho v(z',s')} | z' \right] p_\alpha(z'|z,s) \right\}^{\frac{1}{\gamma}} \right]^{\frac{\gamma}{\rho} \frac{1-\rho}{1-\gamma}} \right\} & \text{if } \rho \neq 1, \gamma > 1 \\ \frac{1}{1-\gamma} \log \left\{ 1 - \beta + \beta^{\frac{1}{\gamma}} \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} E\left[e^{(1-\gamma)\gamma v(z',s')} | z' \right] p_\alpha(z'|z,s) \right\}^{\frac{1}{\gamma}} \right] \right\} & \text{if } \rho = \gamma \neq 1 \\ \frac{1}{1-\rho} \log \left\{ 1 - \beta + \beta^{\frac{1}{\rho}} \exp\left(\frac{1-\rho}{\rho} \sum_{z'} \left[\log \frac{p_\alpha(z'|z,s)}{q(z'|z)} + E\left[\rho v(z',s';\alpha) | z' \right] \right] p_\alpha(z'|z,s) \right) \right\} & \text{if } \rho \neq 1, \gamma = 1 \\ \beta \log \beta + \frac{\beta\gamma}{1-\gamma} \log \left[\sum_{z'} \left\{ q(z'|z)^{\gamma-1} E\left[e^{(1-\gamma) v(z',s';\alpha)} | z' \right] p_\alpha(z'|z,s) \right\}^{\frac{1}{\gamma}} \right] & \text{if } \rho = 1, \gamma > 1 \\ \beta \log \beta + \beta \sum_{z'} \left\{ \log \frac{p_\alpha(z'|z,s)}{q(z'|z)} + E\left[v(z',s') | z' \right] \right\} p_\alpha(z'|z,s) & \text{if } \rho = \gamma = 1 \end{cases} \tag{46}$$

Limiting cases of the value of information moments (27) are

$$f_i(\omega;x_t)=\begin{cases} e^{(\gamma-1)[\rho v_{it}(\alpha_0)+\omega_i]}\left\{1-\beta+\beta^{\frac{1}{\rho}}\Sigma_{it}^{\frac{\gamma(1-\rho)}{\rho(1-\gamma)}}\right\}^{\frac{\rho(1-\gamma)}{1-\rho}}-1 & \text{if } \rho\neq 1,\gamma>1\\ e^{(\gamma-1)[\gamma v_{it}(\alpha_0)+\omega_i]}\left\{1-\beta+\beta^{\frac{1}{\gamma}}\Sigma_{it}\right\}^{\gamma}-1 & \text{if } \rho=\gamma\neq 1\\ \frac{\rho}{1-\rho}\log\left\{1-\beta+\beta^{\frac{1}{\rho}}\exp\left(\frac{1-\rho}{\rho}\Sigma_{it}\right)\right\}-\rho v_{it}(\alpha_0)-\omega_i & \text{if } \rho\neq 1,\gamma=1\\ e^{(\gamma-1)[v_{it}(\alpha_0)+\omega_i]}\beta^{\beta(1-\gamma)}\Sigma_{it}^{\beta\gamma}-1 & \text{if } \rho=1,\gamma>1\\ \beta\log\beta+\beta\Sigma_{it}-v_{it}(\alpha_0)-\omega_i & \text{if } \rho=\gamma=1 \end{cases} \tag{47}$$

where

$$\Sigma_{it}=\begin{cases} \sum_k\left(p^*_{ikt}\right)^{\frac{1}{\gamma}}\left(q_{ikt}\right)^{-\frac{1-\gamma}{\gamma}}e^{\left(\frac{1-\gamma}{\gamma}\right)[\rho v_{kt}(\alpha_0)+\omega_k]} & \text{if } \gamma>1\\ \sum_k\left[\log\frac{p^*_{ikt}}{q_{ikt}}+\rho v_{kt}(\alpha_0)+\omega_k\right]p^*_{ikt} & \text{if } \gamma=1 \end{cases} \tag{48}$$

B Estimation details

B.1 State prices using the Carr-Wu implied volatility surface

Carr and Wu show that, under the assumptions of their LNV model, implied variance as a function of log moneyness k and time to maturity τ is determined by an easy to solve quadratic equation:

$$\begin{aligned}&\frac{w_t^2}{4}e^{-2\eta_t\tau}\tau^2 I_t^4(k,\tau)+\left[1+\kappa_t\tau+w_t^2e^{-2\eta_t\tau}\tau-\rho_t s_t w_t e^{-\eta_t\tau}\tau\right]I_t^2(k,\tau)\\ &-\left[s_t^2+\kappa_t\theta_t\tau+2\rho_t s_t w_t e^{-\eta_t\tau}k+w_t^2e^{-2\eta_t\tau}k^2\right]=0.\end{aligned} \tag{49}$$

The whole implied variance surface is then determined by the six coefficients $(\kappa_t, w_t, \eta_t, \theta_t, s_t, \rho_t)$, which are related to the stock price and implied variance dynamics. It turns out that these parameters are underidentified, so we fix $\kappa_t = 1$. To keep the variables to their respective domains, we follow Carr and Wu and optimize over the transformed variables $\left(\tilde{w}_t, \tilde{\eta}_t, \tilde{\theta}_t, \tilde{s}_t, \tilde{\rho}_t\right)$, which are defined on the entire real line. The first four are logs of the original parameters, and $\rho_t = \frac{e^{\tilde{\rho}_t}-1}{e^{\tilde{\rho}_t}+1} \in [-1, 1]$.

Let $h\left(k, \tau; \tilde{w}_t, \tilde{\eta}_t, \tilde{\theta}_t, \tilde{s}_t, \tilde{\rho}_t\right)$ denote the LNV model implied volatility, i.e. the square root of the sole real positive root of (49). Given a sample of options i on date t, with moneyness k_{it}, terms τ_{it}, and Black-Scholes implied volatilities y_{it}, we estimate the five parameters using weighted nonlinear least squares,

$$y_{it} = h\left(k_{it}, \tau_{it}; \tilde{w}_t, \tilde{\eta}_t, \tilde{\theta}_t, \tilde{s}_t, \tilde{\rho}_t\right) + \varepsilon_{it}, \tag{50}$$

with weights $e^{-\kappa_i^2/2}$, so that at the money options get the most weight.

Assuming the implied volatility function only depends on the relative return (moneyness) and term, but not the index level S, day t implied volatility given any spot price S, strike K, and term τ is

$$\sigma_t\left(S, K, \tau\right) = h\left(\log\left(K/S\right), \tau; \tilde{w}_t, \tilde{\eta}_t, \tilde{\theta}_t, \tilde{s}_t, \tilde{\rho}_t\right). \tag{51}$$

Call option premia on date t are

$$C_t\left(S, K, \tau\right) = BSCall\left(S, K, \sigma_t\left(S, K, \tau\right), r_\tau, \tau, \delta\right), \tag{52}$$

where r_τ is the annualized continuously compounded risk free rate over the horizon τ and δ is the dividend yield.

The usual Breeden-Litzenberger formula for state prices assumes that states are evenly-spaced over the range of possible spot prices (and strikes), with each two states separated by ΔK. But since, following Ross (2015), we maintain a constant log return difference Δk between any two consecutive states, we use the following modified expression for state prices

$$\begin{aligned} q_t\left(S, K, \tau\right) &= \frac{\left[C_t\left(S, K\left(k + \Delta k\right), \tau\right) - C_t\left(S, K, \tau\right)\right] - \frac{K(k+\Delta k)-K(k)}{K(k)-K(k-\Delta k)}\left[C_t\left(S, K, \tau\right) - C_t\left(S, K\left(k - \Delta k\right), \tau\right)\right]}{K\left(k + \Delta k\right) - K\left(k\right)} \\ &= \frac{C_t\left(S, Ke^{\Delta k}, \tau\right) - \left(1 + e^{\Delta k}\right) C_t\left(S, K, \tau\right) + e^{\Delta k} C_t\left(S, Ke^{-\Delta k}, \tau\right)}{\left(e^{\Delta k} - 1\right) K}. \end{aligned}$$

It can be shown that this portfolio of call options generates exactly \$1 if state $S = K$ occurs at τ and zero otherwise as desired.

On each day t, we discretize the state space relative to the current spot price S_t. With a constant log return difference Δk between states, the possible states are $S_{jt+1} = S_t e^{r_j} = S_t e^{(j-c)\times\Delta k}$, where c is the middle row's index (if $n = 11$, $c = 6$). Finally, each entry in row i and column j of the time

t state price transition matrix over horizon τ is given by $q_{tij} = q_t\left(S_i, S_j, \tau\right),$ for $i, j = 1 \ldots n$.

References

Admati, Anat R, 1985, A noisy rational expectations equilibrium for multi-asset securities markets, *Econometrica* pp. 629–657.

Admati, Anat R., and Paul Pfleiderer, 1991, Sunshine trading and financial market equilibrium, *Review of Financial Studies* 4, 443–481.

Ai, Hengjie, 2007, Information quality and long-run risks: Asset pricing and welfare implications, Wokring paper.

——— , and Ravi Bansal, 2016, Risk preferences and the macro announcement premium, Wokring paper.

Aït-Sahalia, Yacine, and Andrew W. Lo, 2000, Nonparametric risk management and implied risk aversion, *Journal of Econometrics* 94, 9–51.

Alvarez, Fernando, and Urban J Jermann, 2005, Using asset prices to measure the persistence of the marginal utility of wealth, *Econometrica* 73, 1977–2016.

Ausubel, Lawrence M, 1990, Insider trading in a rational expectations economy, *American Economic Review* pp. 1022–1041.

Avdis, Efstathios, forthcoming, Information trade-offs in dynamic financial markets, *Journal of Financial Economics*.

Backus, David K., Allan W. Gregory, and Stanley E. Zin, 1989, Risk premiums in the term structure: Evidence from artificial economies, *Journal of Monetary Economics* 24, 371–399.

Backus, David K, Bryan R Routledge, and Stanley E Zin, 2005, Exotic preferences for macroeconomists, in *NBER Macroeconomics Annual 2004, Volume 19* . pp. 319–414 (MIT Press).

Bakshi, Gurdip, and Fousseni Chabi-Yo, 2012, Variance bounds on the permanent and transitory components of stochastic discount factors, *Journal of Financial Economics* 105, 191–208.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *Journal of Finance* 59, 1481–1509.

Barlevy, Gadi, and Pietro Veronesi, 2000, Information acquisition in financial markets, *Review of Economic Studies* 67, 79–90.

Borovička, Jaroslav, Lars P Hansen, and José A Scheinkman, forthcoming, Misspecified recovery, *Journal of Finance*.

Borovička, Jaroslav, 2015, Survival and long-run dynamics with heterogeneous beliefs under recursive preferences, Working paper New York University.

Breeden, Douglas T, and Robert H Litzenberger, 1978, Prices of state-contingent claims implicit in option prices, *Journal of business* pp. 621–651.

Breon-Drish, Bradyn, 2015, On existence and uniqueness of equilibrium in a class of noisy rational expectations models, *Review of Economic Studies* p. rdv012.

Cabrales, Antonio, Olivier Gossner, and Roberto Serrano, 2013, Entropy and the value of information for investors., *American Economic Review* 103, 360–377.

Carr, Peter, and Liuren Wu, 2010, A new simple approach for constructing implied volatility surfaces, Working paper.

Carr, Peter, and Jiming Yu, 2012, Risk, return and ross recovery, *Journal of Derivatives* 20, 38–59.

Cochrane, John H., 2005, *Asset Pricing* (Princeton University Press) revised edn.

Collin-Dufrense, Pierre, and Vyacheslav Fos, 2015, Do prices reveal the presence of informed trading?, *Journal of Finance* 70, 1555–1582.

Croce, M. Max, Tatyana Marchuk, and Christian Schlag, 2016, The leading premium, Working paper.

DeMarzo, Peter M., Michael J. Fishman, and Kathleen M. Hagerty, 1998, The optimal enforcement of insider trading regulations, *Journal of Political Economy* 106, 602–632.

Dow, James, Itay Goldstein, and Alexander Guembel, forthcoming, Incentives for information production in markets where prices affect real investment, *Journal of the European Economic Association*.

Dubynskiy, Sergey, and Robert S Goldstein, 2013, Recovering drifts and preference parameters from financial derivatives, Working paper.

Dybvig, Philip H., and L.C.G. Rogers, 1997, Recovery of preferences from observed wealth in a single realization, *Review of Financial Studies* 10, 151–174.

Epstein, Larry G., Emmanuel Farhi, and Tomasz Strzalecki, 2014, How much would you pay to resolve long-run risk?, *American Economic Review* 104, 2680–2697.

Epstein, Larry G, and Stanley E Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica: Journal of the Econometric Society* pp. 937–969.

Fama, Eugene F., 1970, Efficient capital markets: A review of theory and empirical work, *Journal of Finance* 25, 383–417.

——— , Lawrence Fisher, Michael C. Jensen, and Richard Roll, 1969, The adjustment of stock prices to new information, *International Economic Review* 10, 1–21.

García, Diego, and Günter Strobl, 2011, Relative wealth concerns and complementarities in information acquisition, *Review of Financial Studies* 24, 169–207.

García, Diego, and Joel M. Vanden, 2009, Information acquisition and mutual funds, *Journal of Economic Theory* 144, 1965–1995.

Grossman, Sanford J., and Joseph E. Stiglitz, 1980, On the impossibility of informationally efficient markets, *American Economic Review* 70, 393–408.

Hansen, Lars Peter, John C Heaton, and Nan Li, 2008, Consumption strikes back? measuring long-run risk, *Journal of Political Economy* 116, 260–302.

Hansen, Lars Peter, and José A Scheinkman, 2009, Long-term risk: An operator approach, *Econometrica* 77, 177–234.

Hansen, Lars Peter, and Kenneth J. Singleton, 1982, Generalized instrumental variables estimation of nonlinear rational expectations models, *Econometrica* 50, 1269–1286.

He, Zhiguo, and Asaf Manela, 2016, Information acquisition in rumor-based bank runs, *Journal of Finance* 71, 1113–1158.

Hirshleifer, Jack, 1971, The private and social value of information and the reward to inventive activity, *American Economic Review* 61, 561–574.

Jackwerth, Jens Carsten, 2000, Recovering risk aversion from option prices and realized returns, *Review of Financial Studies* 13, 433–451.

Kacperczyk, Marcin T., and Emiliano Pagnotta, 2016, Chasing private information, Working paper.

Kacperczyk, Marcin T., Stijn Van Nieuwerburgh, and Laura Veldkamp, 2016, A rational theory of mutual funds' attention allocation, *Econometrica* 84, 571–626.

Kadan, Ohad, Roni Michaely, and Pamela C Moulton, 2016, Trading in the presence of short-lived private information: Evidence from analyst recommendation changes, Working paper.

Kelly, Bryan, and Alexander Ljungqvist, 2012, Testing asymmetric-information asset pricing models, *Review of Financial Studies* 25, 1366–1413.

Kreps, David M, and Evan L Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica* pp. 185–200.

Lucca, David O., and Emanuel Moench, 2015, The pre-fomc announcement drift, *Journal of Finance* 70, 329–371.

Malamud, Semyon, 2015, Noisy recovery, Working paper.

Manela, Asaf, 2014, The value of diffusing information, *Journal of Financial Economics* 111, 181–199.

Maurer, Thomas A., and Ngoc-Khanh Tran, 2015, The hirshleifer effect revisited: The case of disagreement, Working paper.

Mele, Antonio, and Francesco Sangiorgi, 2015, Uncertainty, information acquisition, and price swings in asset markets, *Review of Economic Studies*.

Merton, Robert C., 1969, Lifetime portfolio selection under uncertainty: The continuous-time case, *Review of Economics and Statistics* 51, 247–257.

Peress, Joël, 2004, Wealth, information acquisition, and portfolio choice, *Review of Financial Studies* 17, 879–914.

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

Samuelson, Paul A., 1969, Lifetime portfolio selection by dynamic stochastic programming, *Review of Economics and Statistics* 51, 239–246.

Savov, Alexi, 2014, The price of skill: Performance evaluation by households, *Journal of Financial Economics* 112, 213 – 231.

Spence, Michael, and Richard Zeckhauser, 1972, The effect of the timing of consumption decisions and the resolution of lotteries on the choice of lotteries, *Econometrica* pp. 401–403.

Tran, Ngoc-Khanh, and Shixiang Xia, 2014, Specified recovery, Working paper.

Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2009, Information immobility and the home bias puzzle, *Journal of Finance* 64, 1187–1215.

Van Nieuwerburgh, Stijn, and Laura Veldkamp, 2010, Information acquisition and under-diversification, *Review of Economic Studies* 77, 779–805.

Veldkamp, Laura, 2006, Media frenzies in markets for financial information, *American Economic Review* 96, 577–601.

——— , 2011, *Information Choice in Macroeconomics and Finance* (Princeton University Press: Princeton, NJ).

Walden, Johan, 2014, Recovery with unbounded diffusion processes, Working paper.

Table 1: **Event Summary Statistics**

	Power U Recovery				Ross Recovery				
Event	$\Delta E[r^e]$	$\Delta\sigma[r^e]$	ΔSR	I^p	$\Delta E[r^e]$	$\Delta\sigma[r^e]$	ΔSR	I^p	obs
All	0.00	0.00	0.01	0.01	0.00	0.00	0.01	0.01	4887
	(0.07)	(0.06)	(0.11)	(0.07)	(0.28)	(0.08)	(0.36)	(0.08)	
GDP	-0.17	-0.07	-0.50	-0.47	-0.02	-0.08	-0.11	-0.45	222
	(-0.87)	(-0.89)	(-0.94)	(-0.95)	(-0.91)	(-1.02)	(-1.02)	(-1.05)	
Unemployment	-0.58	-0.21	-2.24	-1.22	-0.07	-0.24	-0.52	-1.32	218
	(-2.81)	(-2.46)	(-3.96)	(-2.47)	(-2.35)	(-3.00)	(-3.97)	(-2.93)	
Jobless Claims	-0.01	-0.04	-0.16	-0.33	0.01	-0.03	0.03	-0.26	940
	(-0.13)	(-0.88)	(-0.54)	(-1.25)	(0.87)	(-0.74)	(0.57)	(-1.10)	
Pre-FOMC	0.04	0.03	0.48	0.33	0.02	0.02	0.06	0.23	147
	(0.13)	(0.18)	(0.54)	(0.45)	(0.81)	(0.14)	(0.40)	(0.35)	
FOMC	-0.48	-0.24	-0.99	-1.06	-0.02	-0.22	-0.11	-1.04	147
	(-1.60)	(-1.90)	(-1.33)	(-1.63)	(-0.79)	(-2.02)	(-0.81)	(-1.79)	
Morgtage App.	-0.07	-0.06	0.00	-0.22	-0.00	-0.07	-0.05	-0.34	596
	(-0.41)	(-0.84)	(0.01)	(-0.58)	(-0.33)	(-1.14)	(-0.79)	(-1.00)	
Consumer Comf.	-0.07	-0.04	-0.56	-0.30	-0.03	-0.08	-0.18	-0.49	600
	(-0.47)	(-0.54)	(-1.48)	(-0.79)	(-2.32)	(-1.34)	(-2.73)	(-1.48)	

Reported are subsample mean changes (Δ) multiplied by 100 of the expectation ($E[r^e]$), volatility ($\sigma[r^e]$), and Sharpe ratio (SR) of the excess return on the market implied by physical probabilities recovered from SPX options using either parametric recovery assuming expected utility and constant relative risk aversion $\gamma = 5$, or non-parametric Ross recovery. We also report the mean change in physical probability entropy (I^p) on the event day. The macroeconomic events are announcement days of GDP, unemployment, initial jobless claims, FOMC rate decision (FOMC), the day before (Pre-FOMC), mortgage applications, and Bloomberg's consumer comfort index. t-statistics are in parentheses.

Table 2: **Estimated Value of Information as Percent of Wealth: Power U Recovery**

	$RRA = 5 = 1/EIS$			$RRA = 1 = 1/EIS$			$RRA = 5,\ EIS = 0.90$			
Event	$\hat{\Omega}$	$se(\hat{\Omega})$	$\tilde{\Omega}$	$\hat{\Omega}$	$se(\hat{\Omega})$	$\tilde{\Omega}$	$\hat{\Omega}$	$se(\hat{\Omega})$	$\tilde{\Omega}$	obs
GDP	1.23	(0.51)	1.23	94.00	(3.90)	94.00	7.49	(2.99)	7.49	222
Unemployment	1.13	(0.41)	1.13	89.94	(4.72)	89.94	6.83	(2.43)	6.83	218
Jobless Claims	1.51	(0.27)	1.51	93.79	(2.30)	93.79	9.13	(1.52)	9.13	940
Pre-FOMC	1.53	(0.98)	1.53	95.77	(5.18)	95.94	9.42	(5.93)	9.42	147
FOMC	0.84	(0.68)	0.84	95.73	(3.96)	95.19	5.66	(4.21)	5.66	147
Morgtage App.	2.49	(0.59)	2.49	96.51	(2.18)	96.51	14.67	(3.29)	14.67	596
Consumer Comf.	2.10	(0.46)	2.10	96.54	(2.04)	96.54	12.23	(2.47)	12.23	600

Reported are GMM estimates for the state-dependent value of information as percent of wealth $\hat{\Omega}$ for the center (current) state of the 11 states. The first set of results uses our benchmark parameters, $\beta = 0.998$, $\gamma = \rho = 5$. The second set is the log utility limiting case. The third set of results keeps risk aversion at $\gamma = 5$ but increases the elasticity of intertemporal substitution (EIS) to 0.9. Each row is estimated separately using only announcement days of GDP, unemployment, initial jobless claims, FOMC rate decision (FOMC), the day before (Pre-FOMC), mortgage applications, and Bloomberg's consumer comfort index. Newey-West standard errors in parentheses correct for autocorrelation in errors with two lags. $\tilde{\Omega}$ is a first-order approximation around zero that requires no numerical optimization.

Table 3: **Estimated Value of Information as Percent of Wealth: Ross Recovery**

	$RRA = 5 = 1/EIS$			$RRA = 1 = 1/EIS$			$RRA = 5,\ EIS = 0.90$			
Event	$\hat{\Omega}$	$se(\hat{\Omega})$	$\tilde{\Omega}$	$\hat{\Omega}$	$se(\hat{\Omega})$	$\tilde{\Omega}$	$\hat{\Omega}$	$se(\hat{\Omega})$	$\tilde{\Omega}$	obs
GDP	2.72	(0.58)	2.72	97.40	(2.48)	97.40	15.38	(3.04)	15.38	222
Unemployment	3.08	(0.47)	3.08	99.16	(0.45)	97.35	17.29	(2.38)	17.29	218
Jobless Claims	2.75	(0.30)	2.75	97.82	(1.34)	97.82	15.61	(1.59)	15.61	940
Pre-FOMC	3.26	(1.20)	3.26	98.78	(1.60)	98.39	18.54	(6.39)	18.54	147
FOMC	3.19	(0.91)	3.19	98.79	(1.06)	96.57	18.00	(4.81)	18.00	147
Morgtage App.	4.01	(0.74)	4.01	98.20	(1.18)	98.20	21.75	(3.67)	21.75	596
Consumer Comf.	3.79	(0.52)	3.79	99.24	(0.72)	99.24	20.46	(2.50)	20.46	600

Reported are GMM estimates for the state-dependent value of information as percent of wealth $\hat{\Omega}$ for the center (current) state of the 11 states. The first set of results uses our benchmark parameters, $\beta = 0.998$, $\gamma = \rho = 5$. The second set is the log utility limiting case. The third set of results keeps risk aversion at $\gamma = 5$ but increases the elasticity of intertemporal substitution (EIS) to 0.9. Each row is estimated separately using only announcement days of GDP, unemployment, initial jobless claims, FOMC rate decision (FOMC), the day before (Pre-FOMC), mortgage applications, and Bloomberg's consumer comfort index. Newey-West standard errors in parentheses correct for autocorrelation in errors with two lags. $\tilde{\Omega}$ is a first-order approximation around zero that requires no numerical optimization.

Table 4: **Decomposition of the Value of Information for Log Utility**

Event	$\hat{\Omega} = 1 - e^{-\omega}$	$\hat{\omega} = \omega^p + \omega^q + \omega^v$	$\hat{\omega}^p$	$\hat{\omega}^q$	$\hat{\omega}^v$	obs
GDP	94.00	281.40	706.02	-416.37	-8.25	222
Unemployment	89.94	229.70	887.42	-610.51	-47.21	218
Jobless Claims	93.79	277.87	518.78	-219.89	-21.02	940
Pre-FOMC	95.94	320.50	751.52	-415.47	-15.55	147
FOMC	95.19	303.48	1237.83	-942.26	7.91	147
Morgtage App.	96.51	335.46	698.75	-359.68	-3.61	596
Consumer Comf.	96.54	336.52	697.12	-319.99	-40.62	600

Reported are GMM estimates for the state-dependent value of information as percent of wealth $\hat{\Omega} = 1 - e^{-\hat{\omega}}$ for the center (current) state of the 11 states. We focus on the log utility limiting case because ω in this case can be decomposed into three additive terms: the mean reduction in entropy (ω^p), the mean increase in risk-neutral entropy (ω^q), the mean increase in the expected value-to-consumption ratio (ω^v). Each row is estimated separately using only announcement days of GDP, unemployment, initial jobless claims, FOMC rate decision (FOMC), the day before (Pre-FOMC), mortgage applications, and Bloomberg's consumer comfort index.

Order of Events During Time t			
State z_t realized	Investor observes signal s_t	Investor chooses consumption c_t and investment portfolio weights w_{t+1}	Signal s_t becomes public and prices adjust

Figure 1: **Model Timeline**

Figure 2: **Implied Volatility Surface, January 12, 2012**

An example fitted implied volatility surface using the Carr and Wu (2010) LNV parameterization. Dots represent observations. We estimate the parameters of the volatility surface on each day t using only time t at and out of the money call and put SPX options.

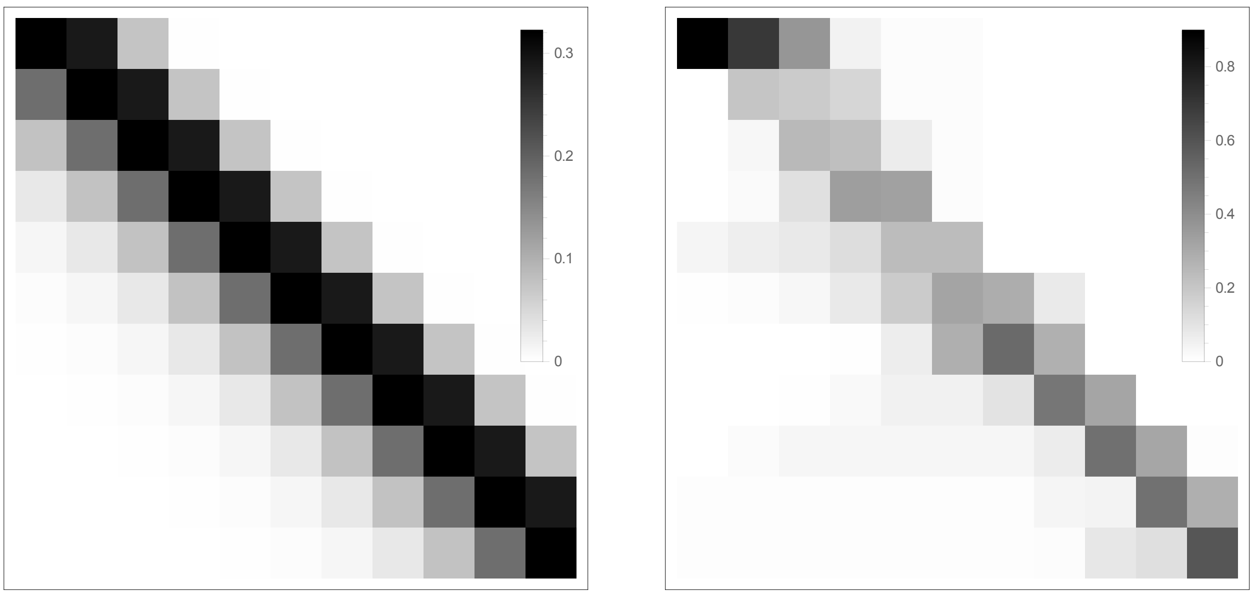

(a) Our Direct Approach

(b) Indirect Approach (Ross, 2015)

Figure 3: **State Price Transition Matrix, Q, January 12, 2012**

An example state price transition matrix Q. Panel (a) is our direct approach used in the paper. Panel (b) shows for comparison the indirect approach of Ross (2015), which inverts Q using least squares, while imposing that Q is unimodal. Both approaches start with the same state-term prices function $q\,(S, K, \tau)$ derived from the fitted Carr and Wu (2010) volatility surface. We estimate the parameters of the volatility surface on each day t using only time t at or out of the money call and put SPX options.

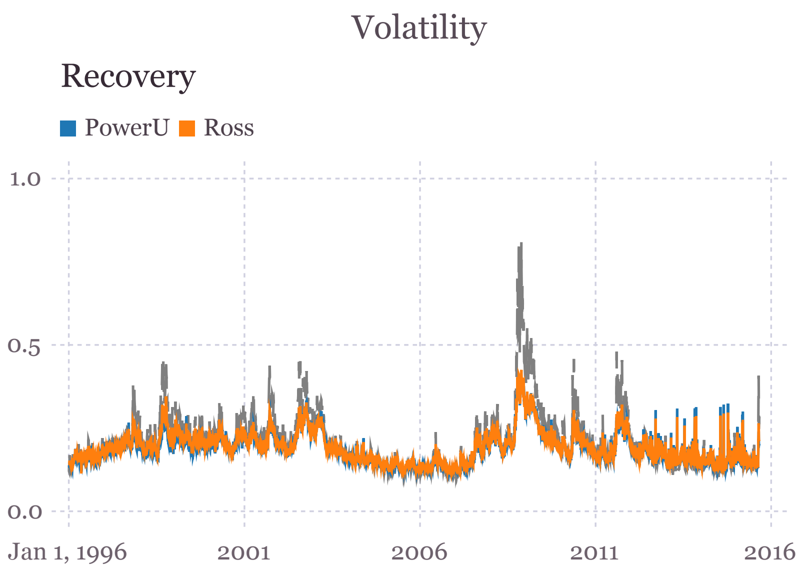

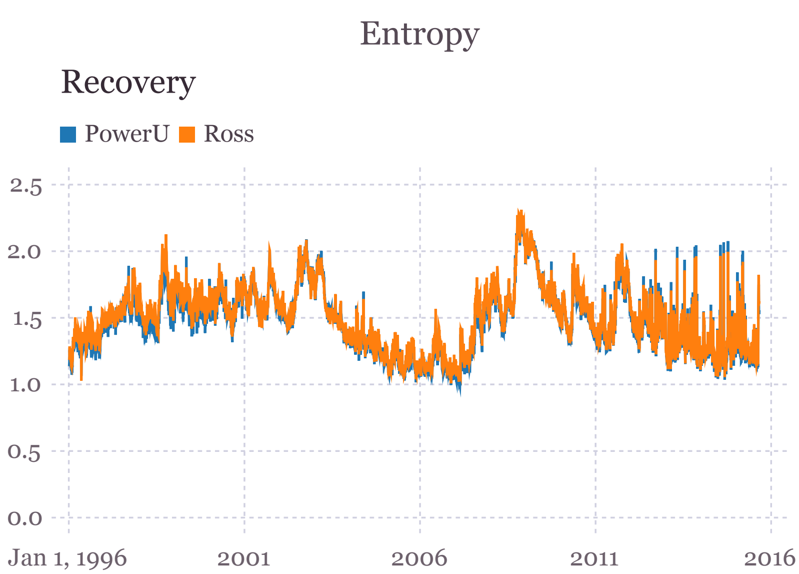

Figure 4: **Conditional Moments Recovered from SPX Options**
Plotted are the daily 30-day volatility and entropy of the excess return on the market, as implied by physical probabilities recovered from SPX options using either parametric recovery that assumes an expected utility representative agent and constant relative risk aversion $\gamma = 5$ (PowerU), or non-parametric Ross recovery. The sample is January 5, 1996 to August 31, 2015. The volatility plot also shows the VIX index (dashed line) for comparison.

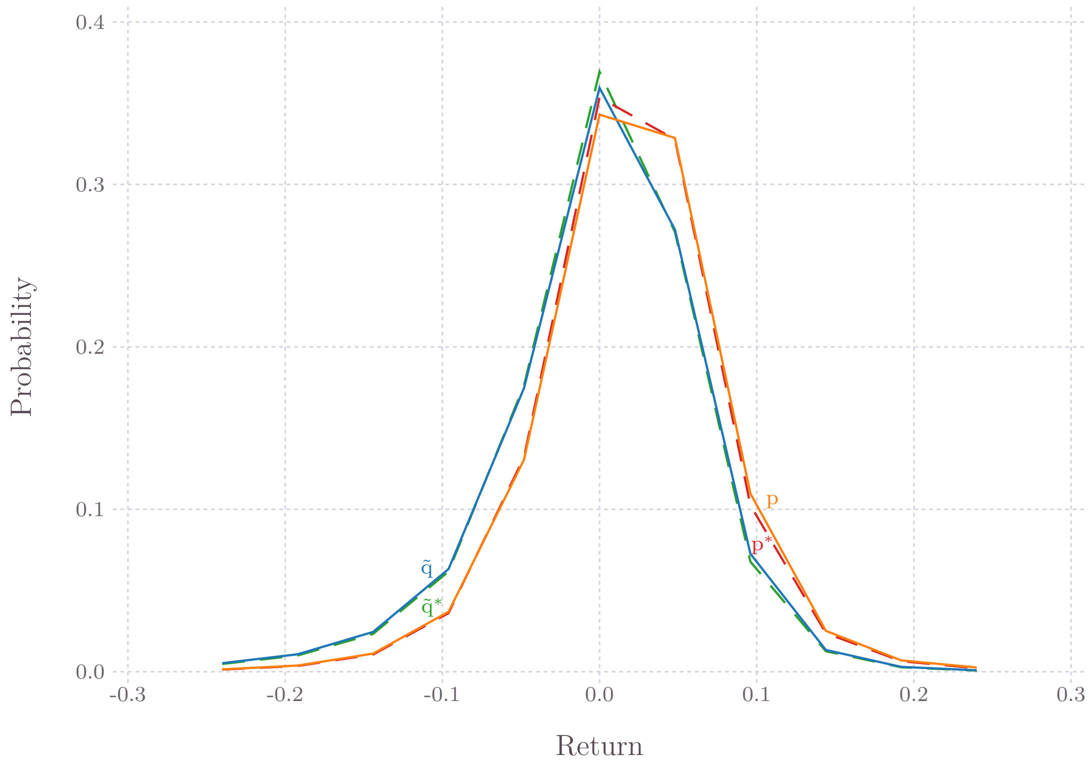

Figure 5: **Prior vs. Posterior Probabilities around Unemployment Releases**
Plotted are the mean prior $\tilde{q}_j = q_j / \sum_i q_i$ and corresponding posterior $\tilde{q}^*_j$ risk-neutral probabilities of the future state (S&P 500 return in one month), estimated from options on the day before and on the day of Unemployment reports. We use these to recover the corresponding physical prior p_j and posterior p^*_j return distributions, here using the parametric recovery approach.

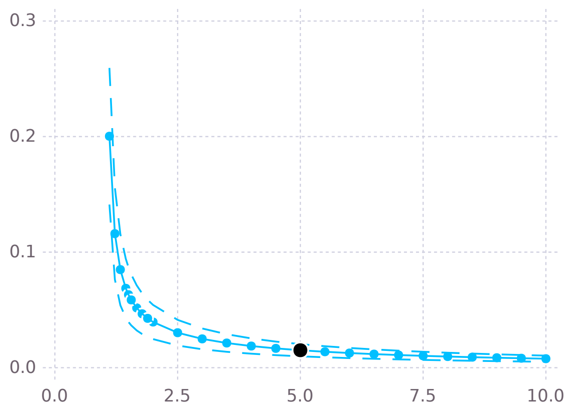

(a) Relative risk aversion $\gamma = \rho$

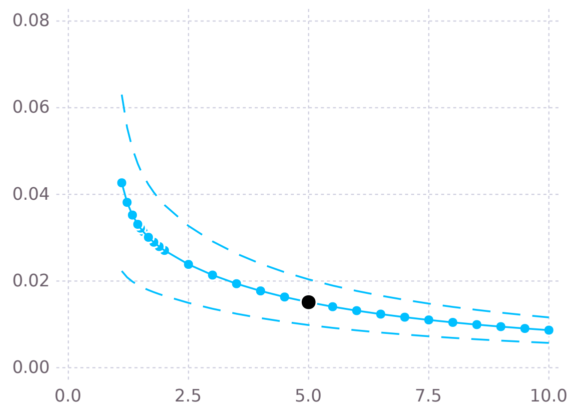

(b) Relative risk aversion γ

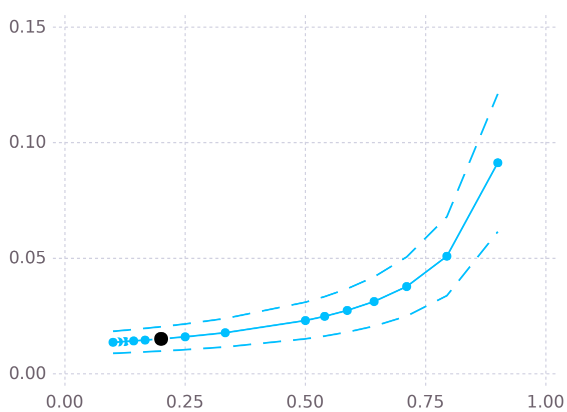

(c) Elasticity of intertemporal substitution $1/\rho$

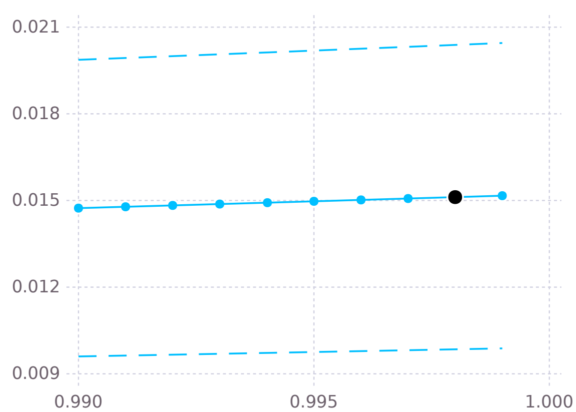

(d) Time discount factor β

Figure 6: **Comparative Statics: Value of Information on Jobless Claims**
Plotted are comparative statics of the value of information as a fraction of wealth around our benchmark parameters: $\beta = 0.998$, $RRA = \gamma = 5$, $EIS = 1/\rho = 1/5$, using parametric recovery that assumes an expected utility representative agent and constant relative risk aversion $\gamma = 5$. The benchmark estimate is circled in each plot. Dashed lines are the 95 percent confidence interval using Newey-West standard errors with two lags.

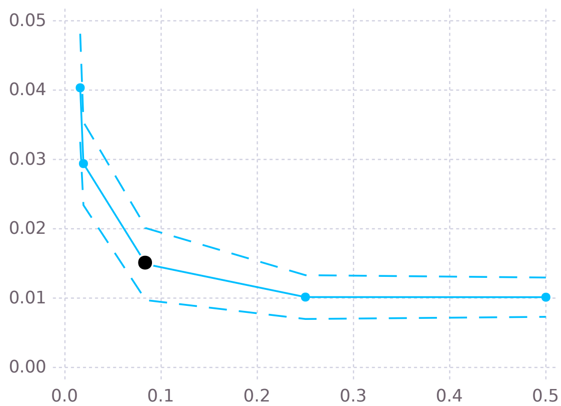

(a) Horizon in years

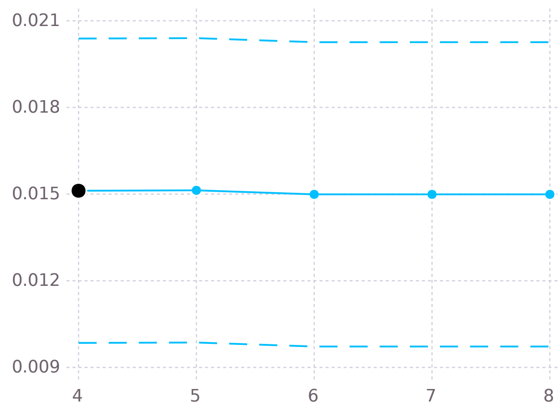

(b) Optimization precision

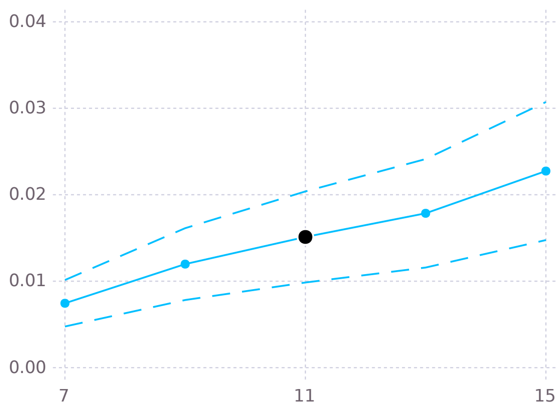

(c) Number of states, n

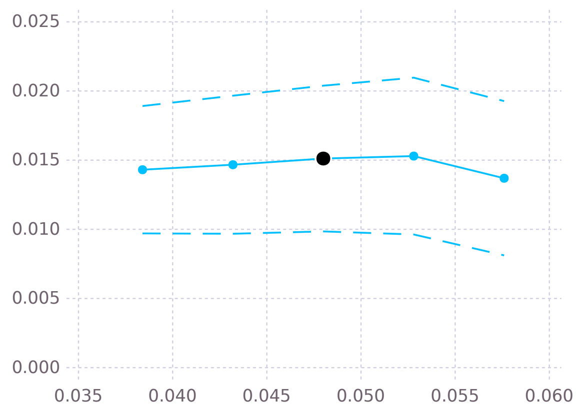

(d) State spacing, dk

Figure 7: **Comparative Statics: Value of Information on Jobless Claims**

Plotted is the value of information as a fraction of wealth for various alternative optimization and recovery specification parameters. The benchmark estimate is circled in each plot. The horizon, or time between state transitions, is measured in years. Optimization precision is the defined as the power x such that we stop the search for a minimum of the GMM objective when $|\omega_{i+1} - \omega_i| \leq 10^{-x}$. On each date t we discretize the state relative to the current SPX closing price, so that log returns (or equivalently log moneyness) take one of n possible equally-spaced states r_j, where the space between each state is dk. All plotted estimates use our benchmark parameters: $\beta = 0.998$, $RRA = \gamma = 5$, $EIS = 1/\rho = 1/5$, and parametric recovery that assumes an expected utility representative agent and constant relative risk aversion $\gamma = 5$. Dashed lines are the 95 percent confidence interval using Newey-West standard errors with two lags.