Restrictions on Asset-Price Movements Under Rational Expectations: Theory and Evidence*

Ned Augenblick
UC Berkeley Haas School of Business
Economic Analysis & Policy Group

Eben Lazarus
MIT Sloan School of Management
Finance Group

JULY 3, 2018

Abstract

How restrictive is the assumption of rational expectations in asset markets? We provide two contributions to address this question. First, we derive restrictions on the admissible variation in asset prices in a general class of rational-expectations equilibria. The challenge in this task is that asset prices reflect both beliefs and preferences. We gain traction by considering market-implied, or risk-neutral, probabilities of future outcomes, and we provide a mapping between the variation in these probabilities and the minimum curvature of utility — or, more generally, the slope of the stochastic discount factor — required to rationalize the marginal investor's beliefs. Second, we implement these bounds empirically using S&P 500 index options. We find that very high utility curvature is required to rationalize the behavior of risk-neutral beliefs, and in some cases, no stochastic discount factor in the class we consider is capable of rationalizing these beliefs. This provides evidence of overreaction to new information relative to the rational benchmark. We show further that this overreaction is strongest for beliefs over prices at distant horizons, and that our findings cannot be explained by factors specific to the option market.

*Contact: ned@haas.berkeley.edu and elazarus@mit.edu. This paper is a revised version of Lazarus's job market paper, and he thanks John Campbell, Emmanuel Farhi, Matthew Rabin, and Jim Stock for their outstanding guidance and support. We are also grateful to Laura Blattner, Jaroslav Borovička, Gabriel Chodorow-Reich, Tim Christensen, Xavier Gabaix, Mark Gertler, Ben Golub, Gita Gopinath, Niels Gormsen, Sam Hanson, Bryan Kelly (discussant), David Laibson, Matteo Maggiori, Ian Martin, Benjamin Moll, José Luis Montiel Olea, Mikkel Plagborg-Møller, Simon Potter, Gianluca Rinaldi, Andrei Shleifer, Robert Solow, Jann Spiess, David Sraer, Stefanie Stantcheva, Jeremy Stein, Adi Sunderam, Neil Thakral, Linh Tô, Johan Walden, Chenzi Xu, and seminar participants at Harvard, UC Berkeley Haas, Chicago Booth, Yale SOM, Stanford GSB, Northwestern Kellogg, MIT Sloan, Duke Fuqua, LSE, the NBER Behavioral Finance Meeting, and the Chicago Fed Rookie Conference for advice, comments, and helpful discussions.

1. Introduction

The assumption that individuals have rational expectations (RE) over future outcomes is central to the positive and normative implications of many models in macroeconomics and finance. Testing the restrictiveness of this assumption is challenging, as individuals' beliefs are difficult to observe directly. A large body of previous literature has attempted to address this difficulty by turning to asset prices, which contain information on investor expectations of future cash flows. But asset prices also reflect agents' unobservable risk and time preferences. Past literature has accordingly appealed to restrictive assumptions on preferences, or on other features of the data-generating process for asset prices, to achieve identification in tests examining expectations.

The seminal work of Shiller (1981) serves as a useful benchmark: Shiller documents excess volatility in equity-index prices relative to a proxy for fundamental value, but this proxy is constructed under the assumption that discount rates over future cash flows are constant over time. This assumption is controversial in light of empirical evidence from recent asset-pricing literature,[1] which calls into question the conclusion that expectations are in fact overly volatile relative to RE. It would seem that changes in unobserved discount rates are capable of rationalizing any observed variation in asset prices, so can anything be said about the rationality of expectations?

In this paper, we show in a general theoretical framework that there are in fact certain bounds on asset-price movements that must hold in a broad class of rational-expectations equilibria, even when relaxing the identification assumptions used in much of the past literature. These bounds yield direct information on the restrictiveness of the rational-expectations assumption in the data: for any level of observed volatility in the asset prices we consider, our results give a precise lower bound for the curvature of utility required to rationalize the data. We show further that there are values of observed volatility that cannot possibly be rationalized by any amount of utility curvature, providing unambiguous evidence of excess volatility of beliefs in these cases.

As in Shiller's case, we focus our analysis on expectations over the future value of an equity index. But the key feature distinguishing our analysis from previous literature is that we consider the behavior of so-called *risk-neutral beliefs* over the underlying index's future price, rather than the behavior of the underlying index itself. The risk-neutral belief distribution can be calculated directly using option prices — options allow for bets over the future asset price, and thus the prices of these bets allow us to back out a probability distribution over this future price — so as is standard, we treat risk-neutral beliefs as observable. These risk-neutral beliefs represent the probability distribution that would be equal to a hypothetical risk-neutral agent's true (or *physical*) belief distribution about the future asset price, but risk-neutral beliefs are in general distorted relative to the marginal investor's physical beliefs in the case that the investor is risk-averse. Intuitively, the probability distribution we observe using asset prices will overweight states in which the marginal investor has low wealth (e.g., when the underlying asset has a low return), since the

[1] See, for example, Campbell (2003) and Cochrane (2011) for surveys discussing evidence on time variation in discount rates (or rationally expected returns), and the end of this section contains a full literature review.

investor will be willing to pay to insure against these high-marginal-utility states.

We show that statements about the "correct" amount of variation in risk-neutral beliefs under RE require less-restrictive assumptions than statements about variation of the index price itself. Previous analyses focusing on index-price variation require keeping track of some measure of the index's fundamental value; in contrast, we show that one can place restrictions on the intertemporal behavior of risk-neutral beliefs without any knowledge of the asset's fundamental value, or knowledge of the marginal investor's underlying physical beliefs. Aside from the maintained assumptions of RE and no arbitrage, our main results require only one general restriction on the stochastic discount factor (SDF), the random variable that determines an asset's ex-ante price by discounting the asset's random future cash flows: we assume that the SDF realization does not depend on the path of unobservable state variables realized between a given trading date and the option expiration date. We refer to this assumption as *conditional transition independence (CTI)*, and this assumption is met in many common macro-finance models. Further, we provide sufficient conditions under which our bounds are robust to mild violations of the CTI assumption.

To understand the economic intuition underlying our results, it is useful to consider the simple case in which we can directly observe an agent's beliefs (or equivalently in our context, the case in which the marginal trader is risk-neutral). Augenblick and Rabin (2018) show in this case that under RE, when the agent's beliefs about a given future outcome change, she must on average be getting closer to certainty about the ultimate realization of the event. If not, then she is, loosely, overreacting to new information relative to the rational (Bayesian) benchmark, or underweighting her prior relative to this new information.

For example, consider this agent's beliefs over whether a single binary outcome will occur at some future date T. She believes ex ante, as of date 0, that the outcome has a 10 percent chance of occurring. Then we observe her beliefs moving to 90% as of date 1, to 10% on date 2, to 90%, to 10%, and so on until date T, at which point her beliefs resolve to either 0% (if the outcome does not occur) or 100% (if it does). Observing one such stream of beliefs, we might simply conclude that she received extreme and alternating signals, forcing her to reverse her beliefs dramatically on a daily basis. But if we were to observe such a pattern repeatedly, we would instead conclude that she is systematically overreacting to new information; her beliefs are mean-reverting (or anti-persistent) in a predictable manner, which violates the martingale property of beliefs under Bayes' rule. Further, this matches the intuition from the previous paragraph: the fact that her beliefs are mean-reverting is equivalent to the fact that her beliefs are never moving closer to certainty, measured as distance from 0% or 100%, despite their large day-to-day changes.

Our main contribution is to show how the above logic — that movement in beliefs must correspond on average to reduction in uncertainty — can be applied to the general case in which the assumption of risk neutrality is dropped. Our task becomes considerably more difficult in this case, as observable risk-neutral beliefs need not follow a martingale given their distortion relative to physical beliefs. But given that the distortion between risk-neutral and physical beliefs is indexed by the risk aversion of the marginal trader, we show in this case that the admissible gap

between risk-neutral belief movement and uncertainty reduction under RE can be bounded as a simple function of this risk-aversion value (or, more generally, the slope of the SDF across states). Further, this bound is tight, as there exists a data-generating process for the signals received by the marginal agent that yields risk-neutral belief variation arbitrarily close to the bound.

We then take our bounds to the data using S&P 500 index option prices obtained from OptionMetrics. We find that very high risk aversion is needed to rationalize the observed variation in risk-neutral beliefs over the future index value, and in our baseline estimation there is in fact *no* amount of risk aversion capable of rationalizing the data. Thus the marginal investor's beliefs are overly volatile relative to the RE benchmark, suggesting that many leading frameworks capable of explaining medium-to-low-frequency variation in asset prices have difficulty rationalizing medium-to-high-frequency variation in beliefs.

Given that we conduct our estimation using index options data, we must also consider whether idiosyncracies specific to this market could be responsible for some of our empirical findings.[2] In this case, our findings would still be indicative of some apparent inefficiency, but specific to the options market rather than with respect to macroeconomic beliefs more generally. For example, if bid-ask bounce induces spurious variation in measured risk-neutral beliefs, this could upwardly bias estimated excess movement in those beliefs. We construct our benchmark empirical tests with such issues in mind — for example, we use end-of-day prices to avoid intraday bid-ask bounce — but we also consider whether related market-specific issues could nonetheless affect our results and attempt to account for them explicitly in robustness tests. Our estimates are weakened only slightly in these tests, and we still find that very high risk aversion is needed to rationalize the data. Thus factors specific to the option market are not capable of accounting for our main results.

We further consider the features of the data that yield our empirical conclusions, leading to three main additional findings. First, we find that excess volatility is concentrated in trading periods relatively far from a given option maturity date: for the last two weeks of trading before expiration, the data can be rationalized with reasonable risk-aversion values, while this is no longer true at longer horizons from expiration. Thus our findings arise largely because beliefs about events in the somewhat-distant future appear to react too strongly to new information. Second, reconducting our estimation at different sampling frequencies, we find that the risk-aversion value required to rationalize the data is decreasing as one decreases the sampling frequency from daily to weekly to monthly. The month-to-month variation in risk-neutral beliefs is moderate enough to be explained by finite (but somewhat large) risk aversion, but this masks substantial volatility (and required risk aversion) at higher sampling frequencies. Third, conducting regressions of our belief-volatility measure on a range of macroeconomic statistics, we find that belief volatility has

[2]There is indeed evidence of idiosyncracies in this market; for example, put-selling and related strategies have high measured returns (Coval and Shumway, 2001), and Jackwerth (2000) argues that a pricing kernel backed out from options data is nonmonotonic in the index return. However, Broadie, Chernov, and Johannes (2009) and Santa-Clara and Yan (2010) suggest that the apparent mispricings are insignificant once peso-problem-type sampling uncertainty and disaster- or jump-risk premia are accounted for: "Option and stock returns may remain puzzling relative to consumption and dividends, but there is little evidence for mispricing relative to the underlying stock index" (Broadie, Chernov, and Johannes, 2009, p. 4496). We return to this issue in Section 5.

a strong positive relationship with measures of macroeconomic uncertainty and no relationship with measures of liquidity or limits to arbitrage in asset markets. This may be considered additional evidence against the possibility that factors specific to the option market are the main driver of our results, but this regression evidence is only suggestive and reduced-form.

Finally, we consider the robustness of our results to violations of our assumption of conditional transition independence for the SDF. The habit-formation model of Campbell and Cochrane (1999) violates this assumption and is also capable of matching important asset-pricing moments in the data, so we consider a calibrated version of the model as an instructive example for how CTI violations might affect our conclusions. We solve and simulate the model numerically,[3] yielding two findings: (a) even given the violation of CTI, the simulated risk-neutral beliefs still exhibit substantially less variation than observed in empirical data; (b) when we naïvely apply our theoretical bounds to estimate the risk aversion required to rationalize the simulated beliefs data, the bounds still yield conservative estimates of the model's implied risk-aversion values. We conclude that reasonably calibrated CTI violations have difficulty accounting for the excess volatility in beliefs in the data, and further that our bounds still apply approximately under mild violations of CTI, as formalized in an additional theoretical result.

Relation to previous literature and interpretation of results. In addition to Shiller (1979, 1981), we follow, among others, LeRoy and Porter (1981), De Bondt and Thaler (1985), and Campbell and Shiller (1987) in conducting empirical tests for excess volatility in asset prices relative to RE. LeRoy and LaCivita (1981) note that these tests are in general joint tests of (a) stationarity of the relevant data-generating process (e.g., for prices in the case of Shiller, 1981); (b) constant discount rates; and (c) rational expectations. Kleidon (1986) and Marsh and Merton (1986) emphasize the importance of possible non-stationarity in accounting for apparent excess volatility; meanwhile, much of the modern asset-pricing literature rationalizes observed price volatility by appealing to time variation in discount rates (again see Campbell, 2003, and Cochrane, 2011, for surveys). We build on the excess-volatility literature by showing that even without imposing any restrictions on the structure of the data-generating process, and imposing only mild restrictions on the variation in discount rates, RE nonetheless restricts the admissible variation in option prices in an empirically testable way.

There are two costs associated with our additional generality. First, we consider derivative prices rather than directly considering the behavior of the underlying index. In this way, our work is complementary to that of Giglio and Kelly (2018), who document excess volatility for claims on equity and currency volatility, inflation swaps, commodity futures, and credit default swaps. Cash flows for these asset classes are well approximated by low-dimensional linear factor models (as in Chamberlain and Rothschild, 1983) in which the unobserved factors follow vector autoregressions under the risk-neutral measure. This autoregressive structure generates restric-

[3]For our purposes, solving the model numerically requires solving for the joint distribution of t-period-ahead ($t \in \{1, 2, \ldots, T\}$) realizations of the equity return and the SDF. Given that this is a high-dimensional object, we apply projection methods to solve for this distribution, and these numerical methods may be of interest in their own right.

tions on the relative revisions to risk-neutral expectations, and therefore prices, at the short versus long end of the term structure for each given asset class. These restrictions are then found to be violated in the data, as long-maturity claims exhibit excess volatility relative to the values implied by the factor models estimated using short-end prices. Their framework differs from ours in that they achieve identification by parameterizing the data-generating process for cash flows, whereas we restrict the evolution of the SDF. Their autoregressive parameterization applies well to the term-structure-like claims they consider, but not to claims on the level of the equity index, to which our framework does apply. These differences in setting and estimation strategy thus allow the two frameworks to provide independent and complementary evidence for excess volatility in expectations, and both do so in a manner that accounts in principle for discount-rate variation.

Second, rather than allowing for fully binary (rejection vs. non-rejection) empirical tests of RE models, our general framework instead allows for a mapping between the observed asset-price variation and the risk aversion required to rationalize the data. Our results may thus appear similar in spirit to those of Mehra and Prescott (1985), and more generally Hansen and Jagannathan (1991), who find that the SDF must be highly volatile to rationalize the observed excess returns for risky assets. Our results differ from theirs in two respects. First, we obtain our mapping using the second moment (i.e., the variation) of observed returns, while they use the first moment (or average) of returns.[4] More importantly, the Hansen–Jagannathan results may in principle be explained by features of the data-generating process for consumption or returns rather than high risk aversion per se; for example, models of rare disasters (e.g., Rietz, 1988; Barro, 2006; Gabaix, 2012; Wachter, 2013) can generate sufficient SDF volatility to rationalize the observed equity premium without requiring high risk aversion. But this is not the case for our results, as we obtain a relationship between local changes in the risk-neutral belief distribution and local risk aversion (or the slope of the SDF) at those points of the distribution. If we observe highly variable risk-neutral beliefs over the event that the S&P's 90-day return will be between 8% and 10%, we know that this cannot be attributable to disasters that affect the left tail of the return distribution; instead, we conclude either that risk aversion is very high or that there is a departure from RE.

Our approach is, however, related to that of Hansen and Jagannathan (1991) at a somewhat higher level: we maintain the spirit of their general semi-parametric setting, as we use a sufficient-statistic-type approach to recover structural parameters from observable data. In this way one may also relate our work to the sufficient-statistics literature in other fields; for example, Chetty (2006) derives an upper bound on utility curvature using labor-supply behavior, and Chetty (2009) provides a longer survey of the literature in settings different from ours. Within the asset-pricing literature, see, for example, Alvarez and Jermann (2005) and Martin (2017), among many others.

Briefly summarizing our relation to literature more closely related to the specifics of our empirical setting and findings, our results complement evidence on beliefs obtained from survey data, as, for example, in Greenwood and Shleifer (2014), Gennaioli, Ma, and Shleifer (2016), and

[4]One could instead map between variation in returns and the volatility of SDF volatility (or the heteroskedasticity of the SDF), but we consider the results from our mapping to be somewhat more intuitive than this alternative.

Manski (2017), as well as the results of Augenblick and Rabin (2018) for settings in which beliefs are directly observable. Another set of related literature endeavors to measure physical beliefs (rather than risk-neutral beliefs) indirectly using options data, for the purpose of examining either expectations or preferences; see Bates (2003) and Garcia, Ghysels, and Renault (2010) for surveys. As a recent example, Ross (2015) assumes a Markov process for transitions between return states and a transition-independence assumption on the SDF similar to but more restrictive than the one we use (see Borovička, Hansen, and Scheinkman, 2016, for a discussion), from which Perron-Frobenius theory allows him to back out a distribution of physical beliefs. Our approach differs from this set of literature in that we need not measure physical beliefs at all or know the true data-generating process for returns to conduct our tests, so we accordingly require less structure.[5]

Organization. The remainder of the paper is organized as follows. Section 2 introduces our theoretical framework and the intuition for our results by considering a simple two-state example, first stated with respect to directly observed beliefs and then extended to consider the effects of risk aversion. Section 3 then presents a general asset-pricing framework, and our theoretical bounds in this general case are collected in Section 4. Section 5 describes the data we use to conduct our empirical test and presents our estimation strategy and main empirical results, while Section 6 conducts additional empirical tests to consider the statistical and macroeconomic correlates of these results and their robustness. Section 7 contains a brief discussion of possible underlying theoretical channels for our results, and Section 8 concludes. Proofs are contained in Appendix A, and Appendix B contains additional technical detail from Sections 3–6.

2. Theoretical Framework: A Simple Example

We first consider a simple two-state example to introduce our framework and to clarify three issues: (a) the economics underlying the restriction on belief movement under RE in the risk-neutral case; (b) how risk aversion complicates this analysis; and (c) how we can nonetheless bound belief movement with risk aversion given certain identifying assumptions.[6] The three subsections below deal with each of these three issues in turn. Readers interested in the more general formal framework may wish to skip ahead to Section 3.

2.1. Example with Directly Observed Beliefs

Consider a discrete-time economy with time indexed by $t \in \{0, 1, 2, \ldots, T\}$. A representative agent consumes C_T as of terminal date T, and this consumption value is exogenously determined

[5]Similar parametric concerns apply to the option-anomalies literature discussed in Footnote 2. If, for example, one measures the physical distribution of returns using historical data (e.g., Jackwerth, 2000), then an absence of crashes in the sample will lead one to overestimate the returns to a put-selling strategy and to incorrectly infer the shape of the pricing kernel over return states. Linn, Shive, and Shumway (2018) argue that this is an empirically relevant concern.

[6]The example we use for (a) works through a basic version of the results in Augenblick and Rabin (2018), who provide the equivalent of Lemma 1 in a general context in which beliefs are directly observed. We differ in specializing to a financial-market context, and the bulk of the work of our paper is related to complications arising from (b) and (c).

by the terminal value of her wealth portfolio, which is stochastic. Assume for now that there are only two possible terminal consumption (or wealth) states: $C_T \in \{C_{\text{low}}, C_{\text{high}}\}$. We also assume for now that the agent's only consumption is in period T.

Each period, the agent receives information and forms beliefs about her terminal consumption value, and these beliefs will be our object of interest. Denote by π_t the agent's date-t subjective belief that the bad state C_{low} will be realized, and the good-state probability is accordingly $1 - \pi_t$. We need not keep track of the information structure for now; the agent simply receives some arbitrary signal each period with new information about the relative likelihood of the two terminal consumption states, and she updates her belief over time accordingly.

We assume that expectations are rational; that is, the agent's beliefs coincide period by period with the true conditional probability of realizing the bad state: $\pi_t = \text{Prob}(C_T = C_{\text{low}} \,|\, \mathcal{F}_t)$, where $\mathcal{F}_t$ denotes time-t conditioning information. This requires that the agent have a correctly specified prior π_0 and that beliefs are updated according to Bayes' rule in response to new information. (We postpone a full formal discussion of the requirements of RE to Section 3.) The belief π_t is a martingale (with respect to $\mathcal{F}_t$) under rationality, or $\pi_t = \mathbb{E}_t[\pi_{t+1}]$, where $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot|\mathcal{F}_t]$ is the conditional expectation.[7]

Assume for now that the agent is risk-neutral and does not discount future consumption, so ex-ante utility is $\mathbb{E}_0[C_T]$. An outside observer can observe Arrow-Debreu state prices for the two terminal consumption states (that is, the date-t price of a zero-net-supply security that pays off 1 unit of consumption if $C_T = C_{\text{low}}$ and 0 otherwise, and similarly for the high-consumption state).[8] With risk neutrality and no discounting, in equilibrium these state prices $q_t(C_{\text{low}})$ and $q_t(C_{\text{high}})$ are equal to the agent's actual subjective beliefs π_t and $1 - \pi_t$, respectively, following the principle that beliefs are directly observable from asset prices under risk neutrality.

We keep track of two objects related to the agent's beliefs $\{\pi_t\}$ and discuss shortly how these objects are related under rationality. First, belief *movement* is defined to be the sum of squared changes in beliefs from arbitrary period t_1 to $t_2 > t_1$:

$$m_{t_1,t_2} \equiv \sum_{t=t_1+1}^{t_2} (\pi_t - \pi_{t-1})^2.$$

Movement is accordingly a formalized notion of belief volatility, equivalent to the discrete-time quadratic variation in the belief process. For the full path, we denote $m \equiv m_{0,T}$.

Second, we define the *uncertainty* of belief π_t as

$$u_t \equiv (1 - \pi_t)\pi_t,$$

and *uncertainty resolution* as

$$r_{t_1,t_2} \equiv u_{t_1} - u_{t_2}.$$

[7] To review why, we have that $\pi_t = \text{Prob}(C_T = C_{\text{low}} \,|\, \mathcal{F}_t) = \mathbb{E}_t[\mathbb{1}\{C_T = C_{\text{low}}\}] = \mathbb{E}_t[\mathbb{E}_{t+1}[\mathbb{1}\{C_T = C_{\text{low}}\}]] = \mathbb{E}_t[\text{Prob}(C_T = C_{\text{low}} \,|\, \mathcal{F}_{t+1})] = \mathbb{E}_t[\pi_{t+1}]$, where the third equality follows from the law of iterated expectations.

[8] These prices can be inferred from the prices of options on the terminal value of the wealth portfolio; see Section 3.

Uncertainty intuitively measures the belief's distance from certainty, and it is maximized at $\pi_t = 0.5$. It is equivalent to the time-t conditional variance of the Bernoulli random variable with realization $\mathbb{1}\{C_T = C_{\text{low}}\}$. Uncertainty resolution simply measures the decrease in uncertainty over time. For the full path, $r \equiv r_{0,T} = u_0 - u_T = u_0$, where the last equality holds because $u_T = 0$ for a fully resolving belief $\pi_T = 0$ or 1 as of the terminal period. Resolution over the entire path is accordingly equal to the ex-ante Bernoulli variance.

Movement and uncertainty resolution are restricted under rational expectations according to the following lemma, which applies a known fact about martingales to formalize a notion of the "correct" amount of belief volatility over any horizon under rationality.[9]

Lemma 1 (***Augenblick and Rabin, 2018***). *Under rational expectations, for any data-generating process, belief movement must equal uncertainty resolution in expectation for arbitrary periods t_1 to $t_2 > t_1$:*

$$\mathbb{E}[m_{t_1,t_2}] = \mathbb{E}[r_{t_1,t_2}].$$

Proof. This follows from an application of the fact that for any square-integrable (e.g., bounded) martingale $\{Y_t\}$, we have $\mathbb{E}_{t_1}[\sum_{t=t_1+1}^{t_2}(Y_t - Y_{t-1})^2] = \mathbb{E}_{t_1}[Y_{t_2}^2] - Y_{t_1}^2$. Rearranging and setting $Y_t = \pi_t$ yields the stated equality; see Appendix A for details. □

The restriction in this lemma reflects the intuition that if the agent's beliefs are moving ($m > 0$), it must be the case on average that she is learning something about the true terminal state ($r > 0$). If instead $\mathbb{E}[m] > \mathbb{E}[r]$, this corresponds to a case in which the agent is systematically overreacting to new information relative to the rational benchmark, as this requires beliefs to be predictably mean-reverting and therefore excessively volatile.[10] To see this intuitively, returning to the example in the Introduction in which the agent's belief oscillates back and forth between 0.1 and 0.9 until resolution, every change in beliefs $0.1 \to 0.9 \to 0.1$ yields positive movement (two-period movement in this case is $2 \times 0.8^2 = 1.28$) but no resolution of uncertainty. (The same is even true for one-day belief changes in this example, as there is no uncertainty resolved from 0.1 to 0.9 or vice versa.) If we were to observe this pattern over repeated samples, we would conclude from Lemma 1 that the agent is exhibiting excess belief movement.

To gain further intuition, Lemma 1 can be rewritten as $\mathbb{E}[\sum_{t=1}^{T}(1 - 2\pi_{t-1})(\pi_t - \pi_{t-1})] = 0$ for the full path. Thus positive revisions to beliefs ($\pi_t - \pi_{t-1} > 0$) coinciding with low initial beliefs ($1 - 2\pi_{t-1} > 0$, or $\pi_{t-1} < 0.5$) lead the statistic $\mathbb{E}[m - r]$ to be positive, and similarly for negative revisions with high initial beliefs. So it could be the case that $\mathbb{E}[\pi_t - \pi_{t-1}] = 0$ unconditionally, but a test based on the lemma would still reject rationality if such unconditional martingale behavior arose due to revisions in the opposite direction of the initial belief.

[9] Other applications of this fact can be found in the continuous-time volatility-estimation literature (e.g., Barndorff-Nielsen and Shephard, 2001; Andersen, Bollerslev, and Diebold, 2010).

[10] The converse holds for $\mathbb{E}[m] < \mathbb{E}[r]$, but we focus on the case of excess volatility given our empirical findings. Further, our "overreaction" terminology is shorthand and should be taken to encompass the possibility that the agent is in fact underweighting her prior relative to information from newly observed signals; see Section 7 for discussion.

2.2. Incorporating Risk Aversion

Our main contribution is to develop a nonparametric test of RE that allows for investor risk aversion, unlike the test implied by the restriction in Lemma 1. An econometrician with access to many observations of beliefs for the agent in the previous subsection could test for excess belief movement using that lemma, as long as the econometrician knew with certainty that the agent was risk-neutral and thus that measured beliefs coincided with the agent's true beliefs. But this subsection shows why such a test is invalid in the presence of risk aversion, and the next subsection then discusses how we can nonetheless bound belief movement in this more general case.

We now make two additional assumptions for the agent in the above example. First, to make clear the numeraire in which assets are priced each period, it will be useful to assume that the agent consumes the exogenous stream $C_t = \overline{C} = 1$ deterministically for all $t < T$, and we again focus on the realization of uncertainty over terminal consumption C_T. Second, and more importantly, assume now that the agent has time-separable log utility with no discounting, $\mathbb{E}_0 \sum_{t=0}^{T} \log(C_T)$, and therefore relative risk aversion $\gamma \equiv -\frac{C_t U''(C_t)}{U'(C_t)} = 1$ (where $U(C_t)$ is period utility), but that this is unknown to the econometrician. For exposition, we set the possible terminal consumption values to $C_{\text{low}} = 1/2$, $C_{\text{high}} = 2$, and assume that these are known.

The econometrician can once again observe Arrow-Debreu state prices over time for the two terminal consumption states, but now these prices will not be equal to the agent's actual subjective beliefs. Assume the agent's unobservable rational prior beliefs are $\pi_0 = 0.3$ for the bad state and $1 - \pi_0 = 0.7$ for the good state. Optimality (with no discounting) implies that the state prices are

$$q_t(C_i) = \frac{U'(C_i)}{U'(C_t)} \pi_t \quad \text{for } i \in \{\text{low}, \text{high}\}, \tag{1}$$

so using the parameters assumed above, the time-0 state prices are $q_0(C_{\text{low}}) = \frac{2}{1}\frac{3}{10} = 0.6$ and $q_0(C_{\text{high}}) = \frac{1/2}{1}\frac{7}{10} = 0.35$.

We can now define the bad-state *risk-neutral belief* π_t^* by dividing the relevant state price by the sum of state prices:

$$\pi_t^* \equiv \frac{q_t(C_{\text{low}})}{q_t(C_{\text{low}}) + q_t(C_{\text{high}})} = \frac{U'(C_{\text{low}})}{\mathbb{E}_t[U'(C_T)]} \pi_t. \tag{2}$$

The high-state risk-neutral beliefs is similarly $\frac{q_t(C_{\text{high}})}{q_t(C_{\text{low}}) + q_t(C_{\text{high}})} = 1 - \pi_t^*$. So the two states' risk-neutral beliefs are both positive and sum to 1 by construction, implying they define a valid probability distribution. As is standard, we refer to them as risk-neutral beliefs because they coincide with actual subjective beliefs for a risk-neutral agent, as can be seen in the last expression in (2) (and as in the previous subsection). So the risk-neutral beliefs can be interpreted as the beliefs for an as-if risk-neutral agent.

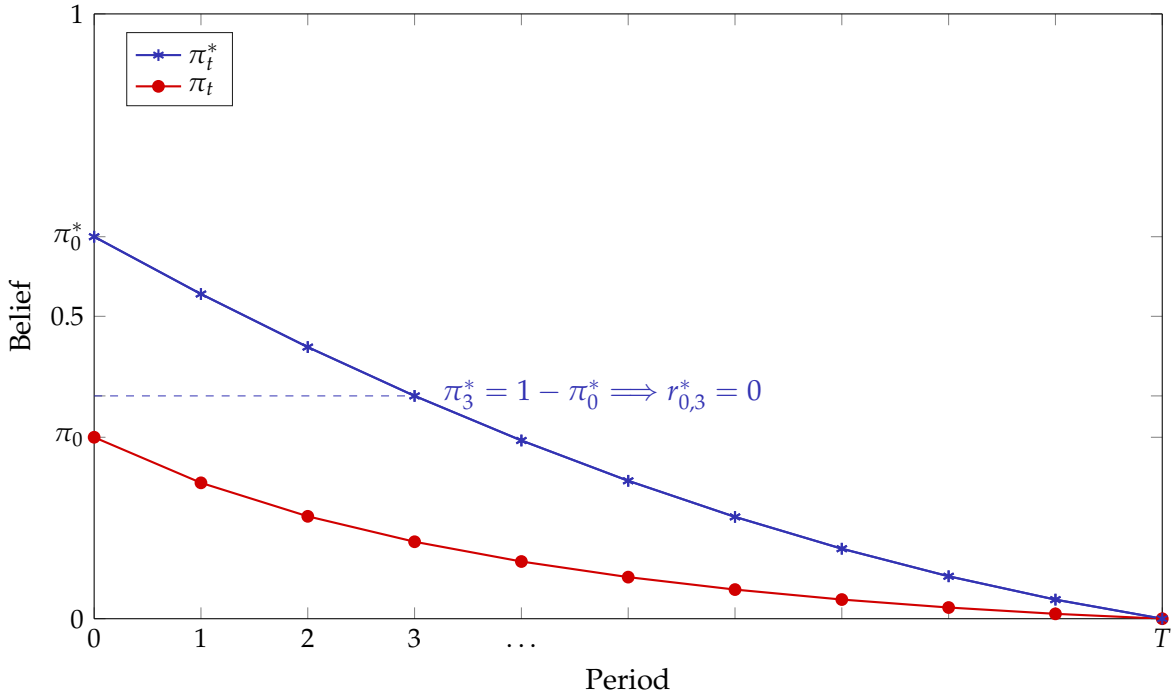
In period 0, the state prices found above yield risk-neutral beliefs $\pi_0^* = 0.63$, $1 - \pi_0^* = 0.37$. So the pseudo-belief distribution backed out from asset prices reflects a combination of both beliefs

and risk preferences; following the usual logic, the agent is willing to pay more than the actuarially fair value for bad-state consumption given her high marginal utility in that state (and vice versa), upwardly biasing the bad-state risk-neutral belief relative to the agent's subjective belief.

To see how this distortion in observed beliefs could affect inference regarding belief movement, consider the paths of beliefs plotted in Figure 1. The series in red circles corresponds to an example realization of the agent's actual subjective beliefs. π_0 in this example is equal to 0.3, the value assumed above. Given this prior for the bad state, a typical realization under rationality will involve beliefs eventually converging to 0 for this state; in this example path, the agent receives signals such that her subjective beliefs slowly and monotonically converge to 0 at date T.

Meanwhile, the observable risk-neutral prior in the figure is $\pi_0^* = 0.63$ as above, and risk-neutral beliefs must follow the agent's true beliefs to 0, as plotted in blue asterisks. Given that $\pi_0^* > 0.5$, the risk-neutral belief moving to 0 implies that eventually this belief must cross $1 - \pi_0^*$, and in this example it does so exactly at $t = 3$. Thus *risk-neutral uncertainty resolution* from $t_1 = 0$ to $t_2 = 3$, $r^*_{t_1,t_2} \equiv (1 - \pi^*_{t_1})\pi^*_{t_1} - (1 - \pi^*_{t_2})\pi^*_{t_2}$, is equal to 0; moving from a belief of 0.63 to 0.37 means the belief has moved no closer to certainty in either direction. But *risk-neutral belief movement*, $m^*_{t_1,t_2} \equiv \sum_{\tau=t_1+1}^{t_2}(\pi^*_\tau - \pi^*_{\tau-1})^2$, has of course been positive in the transition from $t_1 = 0$ to $t_2 = 3$.

Figure 1: Example Subjective and Risk-Neutral Beliefs Under RE



Notes: $\pi_0 = 0.3$, $\pi_0^* = 0.63$, following example in text. In this example of a single realization of uncertainty over time, the agent's subjective bad-state belief (red circles) converges to 0 monotonically, and the observed risk-neutral belief (blue asterisks) must follow and eventually reach 0 as well, while being distorted upwards relative to the true belief for all $t < T$. Since $\pi_0^* > 0.5$, the risk-neutral belief crosses $1 - \pi_0^*$ on its way to 0; in this case, this happens exactly at $t = 3$. As of this date, risk-neutral uncertainty resolution $r^*_{0,3} = (1 - \pi_0^*)\pi_0^* - (1 - \pi_3^*)\pi_3^* = 0$, while risk-neutral belief movement $m^*_{0,3} = \sum_{t=1}^{3}(\pi_t^* - \pi_{t-1}^*)^2 > 0$, illustrating that it can be the case that $\mathbb{E}[m^*] > \mathbb{E}[r^*]$ for risk-neutral beliefs.

Thus even under full rationality, the distortion in risk-neutral relative to actual beliefs induced by risk aversion can cause movement to exceed uncertainty resolution on average in the observed data. So if we naïvely test for rationality using Lemma 1 on observed risk-neutral (rather than actual) beliefs, we may spuriously conclude that beliefs are excessively volatile.

2.3. Identification: Preview of the Main Result

We now show, in the context of the above example, how we can nonetheless bound belief movement for risk-neutral beliefs even in the presence of risk aversion. For simplicity, we focus on beliefs only for periods 0 and 1 (with $T > 1$), and we assume the econometrician observes many draws (π_0^*, π_1^*) generated by subjective beliefs (π_0, π_1) with $\pi_0 = 0.3$ as above.

We make one additional assumption for exposition (which we relax fully in Sections 3–4): rather than maintaining the information structure that (implicitly) generated the series in Figure 1, we instead assume that as of $t = 1$, the agent learns with equal probability either that the bad state will certainly not be realized (in which case $\pi_1 = 0$) or that the probability that the bad state will be realized is $\pi_1 = 0.6$. (It can be seen that the prior $\pi_0 = 0.3$ is ex-ante rational under this signal structure, since $\mathbb{E}_0[\pi_1] = 0.3$.) The remainder of the structure of the example above is unchanged.

This assumption implies that for the subjective belief, expected movement is $\mathbb{E}[m_{0,1}] = 0.5 \times (0 - 0.3)^2 + 0.5 \times (0.6 - 0.3)^2 = 0.09$, and expected uncertainty resolution is also $\mathbb{E}[r_{0,1}] = 0.3 \times (1 - 0.3) - [0.5 \times 0 + 0.5 \times 0.6 \times (1 - 0.6)] = 0.09$, illustrating Lemma 1. For risk-neutral beliefs, we still have $\pi_0^* = 0.63$, and using the calculations (1)–(2) along with the period-1 subjective beliefs above, $\pi_1^* = 0$ or $\pi_1^* = 0.86$ with equal probability. We thus have expected risk-neutral movement $\mathbb{E}[m_{0,1}^*] = 0.5 \times (0 - 0.63)^2 + 0.5 \times (0.86 - 0.63)^2 = 0.22$, while expected risk-neutral uncertainty resolution is $\mathbb{E}[r_{0,1}^*] = 0.63 \times (1 - 0.63) - [0.5 \times 0 + 0.5 \times 0.86 \times (1 - 0.86)] = 0.17$, illustrating again that it can be the case that $\mathbb{E}[m^* - r^*] > 0$.

We can, however, achieve identification of the risk-aversion value required for the observed data to be consistent with rationality, by exploiting the fact that the underlying subjective beliefs must still meet Lemma 1. Denote

$$\phi \equiv \frac{U'(C_{\text{low}})}{U'(C_{\text{high}})}. \tag{3}$$

This slope of marginal utilities or marginal rate of substitution across states will be our structural object of interest for now. Abusing notation slightly, denote by $\pi_t(\phi, \pi_t^*)$ the function mapping from ϕ and the risk-neutral probability to the associated subjective probability. Equation (2) can be inverted to yield that

$$\pi_t(\phi, \pi_t^*) = \frac{\pi_t^*}{\phi + (1 - \phi)\pi_t^*}. \tag{4}$$

Note that the mapping between π_t and ϕ is one-to-one for any given observed value π_t^*, and it is decreasing in ϕ: as risk aversion increases, the underlying bad-state subjective belief decreases with respect to the observed risk-neutral belief.

Lemma 1 yields that, for the subjective beliefs, $\mathbb{E}[m_{0,1} - r_{0,1}] = \mathbb{E}[(1 - 2\pi_0)(\pi_1 - \pi_0)] = 0$.

Since π_0 only takes on one value in this example (even over repeated draws), the restriction $\mathbb{E}[m_{0,1} - r_{0,1}] = 0$ is equivalent to a simple martingale restriction, $\mathbb{E}[\pi_1] = \pi_0$.[11] The econometrician can observe in the data that π_1^* takes on two values with equal probability: $\pi_h^* = 0.86$ or $\pi_\ell^* = 0$. Using this along with (4), the restriction $\mathbb{E}[\pi_1(\phi, \pi_1^*)] = \pi_0(\phi, \pi_0^*)$ becomes

$$\mathbb{E}[\pi_1(\phi, \pi_1^*)] = \frac{1}{2}\frac{\pi_h^*}{\phi + (1-\phi)\pi_h^*} = \frac{\pi_0^*}{\phi + (1-\phi)\pi_0^*} = \pi_0(\phi, \pi_0^*), \tag{5}$$

which can be solved to yield

$$\phi = \frac{\pi_0^* \pi_h^*}{\pi_h^*(1 + \pi_0^*) - 2\pi_0^*} = 4.$$

This is in fact equal to the true ratio $\frac{U'(C_{\text{low}})}{U'(C_{\text{high}})} = \frac{2}{1/2}$ in the current example, so we have achieved identification. Further, $U'(C_{\text{high}}) = U'(C_{\text{low}}) + U''(C_{\text{low}})(C_{\text{high}} - C_{\text{low}}) + \mathcal{O}\big((C_{\text{high}} - C_{\text{low}})^2\big)$ as $C_{\text{high}} \to C_{\text{low}}$ by a Taylor expansion, which can be rearranged to yield

$$\gamma(C_{\text{low}}) \equiv -\frac{C_{\text{low}} U''(C_{\text{low}})}{U'(C_{\text{low}})} = \frac{\phi - 1}{(C_{\text{high}} - C_{\text{low}})/C_{\text{low}}} \tag{6}$$

to a first order, which tells us that risk aversion depends on the ratio of marginal utilities across states relative to the percent consumption gap across states. Thus in the current example, we recover $\gamma = 1$, which is in fact exact given that relative risk aversion is constant by assumption.

Solving for ϕ above, we assumed that the econometrician had access to the data-generating process governing risk-neutral beliefs: in (5), we used that $\pi_1^* = \pi_h^*$ or π_ℓ^* with equal probability. While this is consistent with the structure of the repeated-experiment thought exercise considered here, in reality the true data-generating process is difficult to estimate and potentially infinite-dimensional; for example, the signal structure at $t+1$ could depend on the value π_t.

When generalizing the intuition from this example in the remainder of the paper, we accordingly take a conservative approach and prove a general bound on the minimum value of ϕ required to rationalize excess belief movement $\mathbb{E}[m^* - r^*]$, which holds under *all* possible data-generating processes. This bound, stated without proof for now, can be written in the context of this example as

$$\mathbb{E}[m^* - r^*] \leqslant {\pi_0^*}^2 \left(1 - \frac{1}{\pi_0^* + \phi(1 - \pi_0^*)}\right). \tag{7}$$

Intuitively, following Figure 1, the degree of admissible excess movement depends on the deviation of π_0^* from π_0, as encoded in $\phi = U'(C_{\text{low}})/U'(C_{\text{high}})$; we postpone a detailed discussion to Section 4.1. In our numerical example, one-period excess belief movement is 0.05, as calculated above (3). Using this on the left side of (7) along with $\pi_0^* = 0.63$, we obtain a lower bound for ϕ of 1.4. This illustrates the conservatism of the bound, as the true value of ϕ is 4. Similarly, applying (6) to the bound for ϕ of 1.4, we obtain a bound for γ of 0.14, as compared to its true value $\gamma = 1$.

[11] The example here simply works through identification in a stripped-down case, but we will in general use the more powerful restriction $\mathbb{E}[m - r] = 0$ to obtain a closed-form bound for ϕ given the observed values m^* and r^*.

This example further clarifies an important identification restriction: we have implicitly assumed that the value $\phi = U'(C_{\text{low}})/U'(C_{\text{high}})$ is constant over time, and this assumption will be maintained (and made explicit) in deriving the bound (7) below. Section 3.2 discusses the assumption in general contexts in greater detail, but it follows naturally in the current example from the assumption of time-separable utility and fixed state-contingent consumption values. This illustrates the manner in which restrictions on risk-neutral belief variation require weaker assumptions than restrictions on the underlying asset price: an Arrow-Debreu state price (and associated risk-neutral belief) depends on marginal utility in a single state, whereas the price of a consumption claim depends on the probability-weighted sum of marginal utilities over *all* states. Assuming constant discount rates allows for identification in the latter context (e.g., Shiller, 1981), but we need not make this assumption when working with risk-neutral beliefs.

To see directly how our framework allows for more generality than the constant-discount-rates framework, we can change our numerical example above slightly. Assume now that the deterministic consumption stream for $t < T$ is given by $(C_0, C_1, C_2, C_3, \ldots, C_{T-1}) = (1, 1/2, 1, 1/2, \ldots)$ but that π_t is constant at $\pi_t = \pi_0 = 0.5$ for $t < T$, and all other aspects of the example are unchanged. This induces time variation in the price of a consumption claim, which (again assuming no discounting) is given in equilibrium by $P_t(C_T) = \mathbb{E}_t\left[\frac{U'(C_T)}{U'(C_t)}C_T\right]$: we obtain $P_t(C_T) = 1$ for t even and $P_t(C_T) = 1/2$ for t odd. So we have extreme price variation despite *no* variation in expected terminal cash flows, as prices are changing entirely due to changes in discount rates. A Shiller-type variance-ratio test under the assumption of constant discount rates would thus spuriously reject the null of RE. Meanwhile, because the mapping between π_t and π_t^* is one-to-one for a given ϕ in (4), measured risk-neutral beliefs would be *constant* for $t < T$ in this case, allowing our bound to rationalize the data with $\phi = 1$, its minimal possible value, in (7).

Appendix B.1 discusses the relationship between risk-neutral beliefs and discount rates in greater detail and formalism. In particular, we make clear what forms of discount-rate variation are admissible under the assumption that ϕ is constant; in the example in the previous paragraph, all discount-rate variation arises from changes in the risk-free rate, but the appendix discusses cases in which the risk premium on the consumption claim may be time-varying as well.

3. General Theoretical Framework

We now consider a general many-state framework and show how our analysis applies in this case. Section 3.1 sets up and defines notation for a standard asset-pricing framework, and Section 3.2 presents and discusses the restriction on the SDF we use to derive our volatility bounds.

3.1. Preliminaries: Pricing and Beliefs

Probability space, market index, and options. We work in discrete time, and consider a discrete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t\in\mathbb{N}}$, so that time is indexed by $t \in \{0, 1, 2, \ldots\}$. A realization of the elementary state is denoted by $\omega \in \Omega$.

We will be concerned with the ex-dividend value of the market index, $V_t^m \colon \Omega \to \mathbb{R}_+$, on some option expiration date T. (The superscript m will generally refer to objects tied to the market index, and the subscript t to $\mathcal{F}_t$-adapted processes. When later considering empirical implementation, we will extend the notation to allow for a panel-data environment with multiple option expiration dates.) A European call option on the market index with strike price K has payoff $X_{T,K}^m = \max\{V_T^m - K, 0\}$, and we denote its time-t price as $q_{t,K}^m$. Assume without loss of generality that these option prices are observable for some set of strike prices $\mathcal{K} \subseteq \mathbb{R}_+$ beginning at date 0.

These option prices will be of interest for inferring a distribution over the change in value of the market index from 0 to T. For notation, we say that *return state* $s \in \mathcal{S} \subset \mathbb{R}_+$ is realized for the market index as of date T if

$$R_T^m \equiv \frac{V_T^m}{V_0^m} = s, \tag{8}$$

and the set of return states is accordingly a set of discrete values that the market return can take under Ω; for example, $\mathcal{S}$ could be $\{0, 0.01, \ldots, 0.99, 1, 1.01, \ldots, s_{\max}\}$, where $s = 1$ corresponds to a gross return of 1 (or a net return of 0).[12] The measure $\mathbb{P} \colon \mathcal{F} \to [0,1]$ governs the *objective* or *physical* probabilities of these return states.

Stochastic discount factor. Considering now the pricing of arbitrary assets, the absence of arbitrage implies the existence of a strictly positive *stochastic discount factor (SDF)* or *pricing kernel* process $\{M_t\}$ (i.e., $M_t \colon \Omega \to \mathbb{R}_{++}$) such that the price S_t of a claim to an arbitrary state-contingent payoff X_T is given by

$$S_t(X_T) = \mathbb{E}_t\left[\frac{M_T}{M_t} X_T\right], \tag{9}$$

where again $\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot|\mathcal{F}_t]$, and we can initialize $M_0 = 1$.[13]

In a Lucas (1978)–type economy, with a representative agent with consumption process $\{C_t\}$ and time-separable consumption utility with time discount factor β, the SDF evolves according to

$$\frac{M_{t+1}}{M_t} = \beta \frac{U'(C_{t+1})}{U'(C_t)} \tag{10}$$

by the agent's Euler equation, so the SDF can accordingly be interpreted as aggregate marginal utility. But the representation (9) is valid regardless of the existence of such a representative agent.

Risk-neutral measure. We define the *risk-neutral measure* $\mathbb{P}^*$ with respect to the objective measure $\mathbb{P}$ according to the Radon-Nikodym derivative

$$\left.\frac{d\mathbb{P}^*}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \frac{M_T / M_t}{\mathbb{E}_t[M_T / M_t]}. \tag{11}$$

[12] We could extend the analysis to continuous state spaces with additional technicalities, but do not do so given that empirical implementation requires discretization (and similarly for our discrete-time treatment) and that our theoretical results are more easily understood for probabilities than for densities. We can of course define $\mathcal{S}$ as finely as desired.

[13] Following Campbell (2017), we say that there is absence of arbitrage if (i) $S_t(X_T) \geqslant 0$ for all tradable payoffs X_T such that $X_T \geqslant 0$ almost surely, and (ii) $S_t(X_T) > 0$ for all tradable payoffs such that $X_T > 0$ with positive probability. See Campbell (2017) for a proof that absence of arbitrage implies a strictly positive SDF process (and vice versa).

Equation (9) yields that the $(T-t)$-period gross risk-free rate is given by $R^f_{t,T} \equiv 1/S_t(1_T) = 1/\mathbb{E}_t[M_T/M_t]$, where 1_T refers to one unit of the numeraire delivered at T. Using this along with the change of measure in (11), we can rewrite the basic pricing equation (9) as

$$S_t(X_T) = \frac{1}{R^f_{t,T}} \mathbb{E}^*_t[X_T], \tag{12}$$

as is standard, and where $\mathbb{E}^*_t[\cdot]$ is the conditional expectation under $\mathbb{P}^*$. Thus the price of the random payoff X_T is equal to the expectation of the payoff under $\mathbb{P}^*$, discounted at the $(T-t)$-period risk-free rate, so the change of measure to $\mathbb{P}^*$ incorporates the risk adjustment required to value payoffs at the relevant horizon.[14]

Return-state probabilities. We turn now to the implications of risk-neutral pricing for the market index. The time-t objective probability that the index realizes return state s at date T is

$$\mathbb{P}_t(R^m_T = s) = \sum_{\omega\,:\,R^m_T(\omega)=s} \mathbb{P}_t(\omega),$$

where $\mathbb{P}_t(\cdot) \equiv \mathbb{P}(\cdot|\mathcal{F}_t)$ is the conditional probability. Using this and (11), the associated risk-neutral probability is

$$\mathbb{P}^*_t(R^m_T = s) = \frac{\mathbb{E}_t[M_T/M_t \mid R^m_T = s]}{\mathbb{E}_t[M_T/M_t]}\, \mathbb{P}_t(R^m_T = s). \tag{13}$$

The risk-neutral pricing equation (12) can then be used to show that the date-t schedule of option prices $\{q^m_{t,K}\}_K$ reveals the set of risk-neutral probabilities $\{\mathbb{P}^*_t(R^m_T = s)\}_s$. Assume that the set of return states $\mathcal{S} = \{s_1, s_2, \ldots, s_J\} = \{s_j\}_{j=1,\ldots,J}$ is ordered such that $s_1 < s_2 < \cdots < s_J$, and assume for notational simplicity that the set of traded option strike prices $\mathcal{K}$ coincides with the set of possible date-T index values; that is, $\mathcal{K} = \{K_1, K_2, \ldots, K_J\}$, where $K_j = V^m_0 s_j$. (We will see that this can be relaxed.) We can then back out the risk-neutral probabilities of interest from option prices as follows:

$$\mathbb{P}^*_t(R^m_T = s_j) = R^f_{t,T} \left[\frac{q^m_{t,K_{j+1}} - q^m_{t,K_j}}{K_{j+1} - K_j} - \frac{q^m_{t,K_j} - q^m_{t,K_{j-1}}}{K_j - K_{j-1}} \right]. \tag{14}$$

Appendix A contains a brief derivation of this result, which follows directly from a discrete-state application of the classic result of Breeden and Litzenberger (1978). We see from this expression that we need not have the set of strikes $\mathcal{K}$ coincide with the full set of possible date-T index values

[14] This risk-neutral measure is defined for some particular option expiration date T; given that T is arbitrary, one can in fact interpret (11) as defining a set of risk-neutral measures $\{\mathbb{P}^{*_T}\}_{T\in\mathbb{N}}$, as will be implicitly used in our empirical implementation given that we have multiple option expiration dates. Further, $\mathbb{P}^*$ as defined in (11) is sometimes referred to as the *T-forward measure* rather than the risk-neutral measure; references include Jamshidian (1989), Geman, El Karoui, and Rochet (1995), and Hansen and Scheinkman (2017). Payoffs are discounted under $\mathbb{P}^*$ using a zero-coupon bond maturing at T (rather than an account accumulating short-term risk-free returns), and T-maturity forward prices $f_{t,T}$ (or expectations of prices at T) are martingales under $\mathbb{P}^*$, e.g., $f^m_{t,T} = \mathbb{E}^*_t[V^m_T]$. It is thus a natural *equivalent martingale measure* for use in considering options over the future market-index value, as in our case.

to back out the risk-neutral probability $\mathbb{P}_t^*(R_T^m = s_j)$: we must simply have strikes at $V_0^m s_{j-1}$, $V_0^m s_j$, and $V_0^m s_{j+1}$ to back it out exactly, and strikes near those values to obtain an approximation.

The fact that date-t option prices reveal the risk-neutral probabilities for the date-T return states without requiring us to account for the value of processes between t and T (e.g., one-period risk-free rates) motivates our use of the particular risk-neutral measure defined in (11).

Beliefs. To this point we have not taken a stance on the underlying structure of the economy responsible for generating prices and risk-neutral probabilities; the analysis above follows fully from the representation (9), which requires only the absence of arbitrage. We now specialize our exposition by viewing prices as being generated by some marginal trader observing public signals. None of what follows requires this to be the case, but it is useful for simple interpretation of our results.[15] One might intuitively think of this agent as corresponding to "the market."

The agent observes a finite vector of signals $\theta_t \in \Theta$ each period, with $\mathcal{F}_t = \sigma(\theta_\tau, 0 \leqslant \tau \leqslant t)$. The information provided by the date-t signals about date-T return states is described by the likelihood function or *signal-generating process* $\mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R_T^m)$, where $\mathbb{P}$ is the same physical measure as defined above. This general formulation implies that the conditional signal distribution can depend arbitrarily on the history $\theta^{t-1} \equiv (\theta_0, \theta_1, \ldots, \theta_{t-1})$, and the signals are informative about the relative likelihood of return states $s_j \neq s_k$ as long as $\mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R_T^m = s_j) \neq \mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R_T^m = s_k)$.

The agent's time-t *subjective belief distribution* over the return states for the market at T is denoted by $\Pi_{t,T} = \{\pi_t(R_T^m = s)\}_{s \in \mathcal{S}}$. The agent brings beliefs $\Pi_{t,T}$ into period $t+1$, observes signals θ_{t+1}, and forms new beliefs $\Pi_{t+1,T}$. In order to derive testable restrictions on rational-expectations price processes that can be taken to the data, we maintain the assumption that the agent has RE over the return states at T, defined as follows in a standard manner (e.g., Muth, 1961, and more specifically as first described in Lucas and Prescott, 1971, Lucas, 1972, and Green, 1975).

Definition 1 (*RE*). An agent has *rational expectations* over return states at T if and only if both:

(i) The agent's date-0 priors coincide with the objective probabilities:

$$\pi_0(R_T^m = s) = \mathbb{P}_0(R_T^m = s) \quad \forall\, s \in \mathcal{S}.$$

(ii) The agent updates beliefs in response to new information according to *Bayes' rule* using the objective likelihood function:

$$\pi_t(R_T^m = s) = \frac{\pi_{t-1}(R_T^m = s)\, \mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R_T^m = s)}{\mathbb{P}(\theta_t \mid \mathcal{F}_{t-1})},$$

with $\mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}) = \sum_{s' \in \mathcal{S}} \pi_{t-1}(R_T^m = s')\, \mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R_T^m = s')$. ‖

[15] This expositional assumption may seem restrictive, but note that even with multiple possible marginal traders with access to some private information, the logic of the no-trade theorem of Milgrom and Stokey (1982) implies that prices reveal this information, and we treat all signals as public and priced by a single agent for purposes of exposition. We briefly discuss cases in which a no-trade theorem fails to hold — e.g., when agents with heterogeneous beliefs agree to disagree or neglect disagreement — in Section 7.

These two conditions together are equivalent to the agent's beliefs coinciding with the objective probabilities period by period, but we find it useful to be able to consider belief updating separately from the prior in the analysis that follows, which yields the natural definition given above.[16]

Given the maintained assumption of RE, we can define the *risk-neutral belief distribution* without explicitly restricting the agent's utility function or constraint set by applying the same change of measure as defined in (11), using the general SDF M_T / M_t, to her subjective beliefs. This yields a risk-neutral belief distribution $\Pi^*_{t,T} = \{\pi^*_t(R^m_T = s)\}_{s \in \mathcal{S}}$ such that

$$\pi^*_t(R^m_T = s) = \frac{\mathbb{E}_t[M_T / M_t \mid R^m_T = s]}{\mathbb{E}_t[M_T / M_t]} \, \pi_t(R^m_T = s), \tag{15}$$

as in (13). Thus (14) tells us that option prices reveal the agent's risk-neutral beliefs as given here.

As in the example in Section 2, we will state our general results in terms of admissible variation in *conditional* risk-neutral beliefs over pairs of return states. That is, rather than directly restricting the intertemporal behavior of the full distribution $\Pi^*_{t,T}$, we instead consider restrictions on the behavior of the individual entries in the set $\{\widetilde{\pi}^*_{t,j}\}_{j=1,\ldots,J-1}$ defined by

$$\widetilde{\pi}^*_{t,j} \equiv \pi^*_t(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\}) = \frac{\pi^*_t(R^m_T = s_j)}{\pi^*_t(R^m_T = s_j) + \pi^*_t(R^m_T = s_{j+1})}, \tag{16}$$

for $\pi^*_t(R^m_T = s_j) + \pi^*_t(R^m_T = s_{j+1}) > 0$. In words, $\widetilde{\pi}^*_{t,j}$ describes the time-t risk-neutral belief that return state s_j will be realized for the market index at date T, conditional on either return state s_j or s_{j+1} being realized. As in the example in the previous section, we have that the conditional probability of interest as defined in (16) is the "bad-state" (or low-return-state) probability. And by analogy to the notation in (16), we define the expectation under the conditional measure as

$$\widetilde{\mathbb{E}}_t[\,\cdot\,] \equiv \mathbb{E}_t[\,\cdot \mid R^m_T \in \{s_j, s_{j+1}\}]. \tag{17}$$

We consider conditional probabilities as in (16) for purposes of theoretical traction. It will turn out that the space of signal-generating processes over the realization of uncertainty over two states, as is considered when transforming to conditional probabilities, is sufficiently "small" and well-behaved to enable a simple analytic characterization of the admissible variation in risk-neutral beliefs under RE. This characterization of course implies conditions for admissibility for the full distribution of risk-neutral beliefs, but this untransformed distribution proves unwieldy enough that obtaining sharp results is difficult.[17] Further, as will be seen in Section 3.2 below, considering conditional probabilities allows for a weaker restriction on the SDF to allow for identification than would be the case without such a transformation.

[16] We note further, however, that the fact that agents update using the objective likelihood function means that the conditions given in Definition 1 can be thought of as jointly specifying that a suitably enlarged product prior over $\mathcal{S} \times \Theta$ is correctly specified. In other words, while updating behavior may be separable from the prior in an economic sense, one might not think the two are completely distinguishable in a mathematical sense.

[17] We do have additional results for the full distribution available upon request, but have not yet conducted empirical estimation for these results.

Belief variation. To formalize notions of risk-neutral belief volatility and uncertainty, we define the following objects, which are analogous to those defined in Section 2.

DEFINITION 2 (*Movement*). Define *risk-neutral belief movement* for the conditional risk-neutral belief process $\{\widetilde{\pi}^*_{t,j}\}_{0\leqslant t\leqslant T}$ from time t_1 to time $t_2 > t_1$ as

$$m^*_{t_1,t_2,j} \equiv \sum_{t=t_1+1}^{t_2} (\widetilde{\pi}^*_{t,j} - \widetilde{\pi}^*_{t-1,j})^2,$$

and denote movement for the full path by $m^*_j \equiv m^*_{0,T,j}$. ‖

DEFINITION 3 (*Uncertainty*).

(i) Define *risk-neutral uncertainty* for the conditional risk-neutral belief $\widetilde{\pi}^*_{t,j}$ as

$$u^*_{t,j} \equiv (1 - \widetilde{\pi}^*_{t,j})\widetilde{\pi}^*_{t,j}.$$

(ii) Define *risk-neutral uncertainty resolution* for the conditional risk-neutral belief process from time t_1 to time t_2 as

$$r^*_{t_1,t_2,j} \equiv u^*_{t_1,j} - u^*_{t_2,j}.$$

For the full path, $r^*_j \equiv r^*_{0,T,j}$. ‖

See Section 2.1 for discussions of these definitions. We can then measure empirical counterparts for these objects given result (14).[18]

3.2. Restriction on the SDF

We must now confront the joint hypothesis problem, and we attempt to do so in a manner that is sufficiently general and semi-parametric so as to achieve identification in a broad class of models. As above, the absence of arbitrage implies that there is some SDF process that relates the observed risk-neutral beliefs to the objective probabilities describing the true data-generating process as in (11). Thus without any additional restrictions, there is always some sequence $\{M_t\}$ that can in theory be used to transform the observed data to the correct objective probabilities,[19] *even* under the alternative that the subjective probabilities being used by agents to price assets are incorrect (in which case the actual SDF would include these belief distortions). We must accordingly restrict the form that $\{M_t\}$ can take under the maintained null of RE in some way.

[18] We aim to do so without excessive market-microstructure contamination, and the end-of-day sampling we use to do so motivates our discrete-time framework. We note also that these objects are invariant to the addition of nearest-neighbor-interpolated or -extrapolated data points $\widetilde{\pi}^*_{t,j}$, and therefore even if the panel $\{\widetilde{\pi}^*_{t,j}\}_{t,j}$ is unbalanced, it can be made balanced in this way without affecting m^*_j or r^*_j. This is an advantage of this approach relative to, e.g., variance-ratio tests, for which adding additional observations with no change in beliefs reduces the power of the tests.

[19] This is no longer the case in an economy that admits arbitrage opportunities, but we are interested the RE assumption rather than the possibility of arbitrage, so we maintain the no-arbitrage assumption throughout.

We begin by transforming the conditional risk-neutral belief $\widetilde{\pi}^*_{t,j}$ into an odds ratio for state-s_j versus state-s_{j+1} beliefs as follows:

$$\frac{\widetilde{\pi}^*_{t,j}}{1-\widetilde{\pi}^*_{t,j}} = \frac{\mathbb{E}_t[M_T/M_t \mid R^m_T = s_j]}{\mathbb{E}_t[M_T/M_t \mid R^m_T = s_{j+1}]} \frac{\widetilde{\pi}_{t,j}}{1-\widetilde{\pi}_{t,j}}, \tag{18}$$

since the risk-neutral belief for state s_{j+1} conditional on either s_j or s_{j+1} being realized is $1-\widetilde{\pi}^*_{t,j}$, and where $\widetilde{\pi}_{t,j} \equiv \pi_t(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\})$ is the conditional subjective belief. We define the first term on the right side of this equation as $\phi_{t,j}$:

$$\phi_{t,j} \equiv \frac{\mathbb{E}_t[M_T/M_t \mid R^m_T = s_j]}{\mathbb{E}_t[M_T/M_t \mid R^m_T = s_{j+1}]}. \tag{19}$$

This value encodes the slope of the stochastic discount factor across the two adjacent return states s_j and s_{j+1}, and it is a generalization of the structural object of interest in the example in Section 2 as defined in equation (3). In the case of the representative-agent economy specified in (10), this value becomes $\phi_{t,j} = \mathbb{E}_t[U'(C_T) \mid R^m_T = s_j]/\mathbb{E}_t[U'(C_T) \mid R^m_T = s_{j+1}]$. It can accordingly be thought of as the marginal rate of substitution across the two return states, or more generally the riskiness of the bad state s_j relative to the good state s_{j+1}, as encoded in asset prices. We accordingly assume that $\phi_{t,j} \geqslant 1$. This is without loss of generality in theory, as we can relabel the states such that this is true. For empirical implementation, we again use the ordering $s_1 < s_2 < \cdots < s_J$.

Our substantive restriction on the SDF is then as follows. We first state the assumption formally, and then discuss its economic content by way of several remarks and examples.

DEFINITION 4 (*CTI*). The SDF satisfies *conditional transition independence (CTI)* for the return-state pair (s_j, s_{j+1}) and option expiration date T if $\phi_{t,j}$ defined in (19) is constant for all $0 \leqslant t < T$ almost surely, and we denote this constant by ϕ_j. ‖

REMARKS:

1. Stated intuitively, CTI requires that when we observe a change in the risk-neutral odds ratio (18), this is due to a change in the subjective conditional probability $\widetilde{\pi}_{t,j}$ rather than the expected relative severity of the adjacent return states s_j and s_{j+1}. This definition is thus analogous to the assumption discussed in Section 2.3 (Section 2.3). If, as in that case, a representative agent's utility depends only on the maturity value of the market index, then there is a perfect mapping between the terminal state and M_T (with M_T proportional to marginal utility of terminal wealth), which guarantees that $\phi_{t,j}$ is constant.[20] The remainder of this subsection discusses conditions under which this logic can be extended to more general asset-pricing frameworks.
2. Note that we have assumed only that the ratio of the conditional expectations for the SDF is constant over time locally across adjacent states s_j and s_{j+1}. Of course j and $j+1$ are arbitrary,

[20] This holds more generally as long as there is some agent whose indirect utility can be written as a function only of the terminal index value (e.g., an investor retiring at date T with savings fully invested in the market).

but we can importantly assume that CTI holds only for some desired subset of observed return states. This accordingly does *not* require that all changes in the underlying risk-neutral belief distribution in (15) arise from changes in the subjective beliefs term $\pi_t(R_T^m = s_j)$: there may be simultaneous time variation in the values in the numerator and denominator in (19), as in multiple cases discussed below, and the values $\mathbb{E}_t[M_T/M_t]$ and $\mathbb{E}_t[M_T]$ need not be constant.

3. Definition 4 corresponds to a notion of transition or path independence because it implies that the realization of M_T/M_t in return state s_j depend in expectation only on s_j and not on the path of any variables realized between t and T (though the return state itself *can* depend on such a path). This intuition is formalized in Lemma A.1 in Appendix A. We refer to the assumption as *conditional* transition independence to underscore that it requires only constancy of the ratio of conditional expectations of the state-contingent SDF realizations, rather than deterministic state-contingent SDF realizations. The CTI assumption is accordingly less restrictive than the transition-independence assumption used by Martin and Ross (2013) and Ross (2015).[21]

We now turn to a set of examples to illustrate the CTI restriction concretely, all of which work under the maintained hypothesis of rational expectations.

Example 1. Assume an economy with a one-dimensional state variable $A_t\colon \Omega \to \mathbb{R}$ (e.g., productivity, consumption, volatility), with $dV_t^m/dA_t > 0$. This process may be non-stationary but is assumed to satisfy the Markov property, $\mathbb{P}(A_{t+\tau} = a \mid A_t, A_{t-1}, \ldots) = \mathbb{P}(A_{t+\tau} = a \mid A_t)$ for all $\tau \geqslant 0$ and $a \in \mathbb{R}$. Assume further that there exists a representative agent with time-separable utility over the consumption process $\{C_t(A_t)\}$ and that the market index pays dividends according to $\{D_t(A_t)\}$, where these processes are arbitrary but yield a stationary price-dividend ratio. Then CTI holds for any two adjacent return states.

If, in addition, consumption or consumption growth is i.i.d. over t, then CTI holds as well if the agent instead has Epstein–Zin (1989) recursive utility. ‖

While the assumption of a scalar Markov forcing process in this example is restrictive, it nonetheless encompasses some leading cases. For example, with the equilibrium value of log consumption as the state variable, $A_t = c_t \equiv \log(C_t)$, we could have its evolution governed by

$$c_t = g + \rho c_{t-1} + h(c_{t-1})\varepsilon_t,$$

where ε_t is i.i.d. with arbitrary distribution (so it could incorporate the possibility of disasters) and $h(\cdot)$ is a state-dependent volatility function. In any such case, the result in the above example gives that CTI is met.

We can also consider a similar example without assuming the existence of a representative agent, but placing more stringent conditions on utility and stochastic processes, as follows.

[21] Specifically, Ross (2015) writes M_T/M_t as a deterministic function of the current and terminal price states (denoted as V_j^m, V_k^m, respectively) as follows: $M_T/M_t = \delta\, g(V_k^m)/g(V_j^m)$ for some function g and constant δ, for *all* states. See also Heston (2004), Jensen, Lando, and Pedersen (2018), and Walden (2017) for uses of this assumption.

EXAMPLE 2. Assume the existence of some agent with constant-relative-risk-aversion (CRRA) period utility $U(C_t) = (C_t^{1-\gamma} - 1)/(1-\gamma)$ who is almost surely unconstrained in every period. (We need not assume that all of the agent's wealth is invested in the market index nor restrict the source of the agent's income.) Assume further that either of the following conditions holds:

(i) The joint process for the agent's consumption and the market-index dividend is i.i.d. over t, with arbitrary joint distribution over the draws (C_t, D_t).

(ii) The joint process for the agent's consumption growth and the market-index dividend growth is i.i.d. over t, with jointly log-normal draws $(C_t/C_{t-1}, D_t/D_{t-1})$ with arbitrary covariance.

Then CTI holds for any two adjacent return states. ‖

The two preceding examples clarify that the CTI restriction allows for time variation in discount rates and risk premia; see Appendix B.1 for a discussion. Relatedly, both temporary and permanent shocks to consumption and marginal utility (and thus the SDF M_t) are in principle admissible under CTI. For example, with CRRA utility and i.i.d. consumption-growth shocks there are only permanent shocks to $M_t \propto C_t^{-\gamma}$, which does not change $\mathbb{E}_t[M_T/M_t]$ since M_{t+1}/M_t is i.i.d.; meanwhile, an economy with i.i.d. consumption has only transitory shocks to M_t and a fixed value $\mathbb{E}_t[M_T \mid R_T^m = s_j]$ for all s_j.

This flexibility is desirable on both empirical and theoretical grounds. Empirically, Alvarez and Jermann (2005) argue that permanent shocks to the SDF are important for quantitatively matching the observed moments of returns. Theoretically, Borovička, Hansen, and Scheinkman (2016) show that the assumptions used in the empirical application of Ross (2015) to estimate the physical distribution of return states — namely, a finite Markov state space for prices and a transition-independence assumption stronger than ours (see Remark 3) — do not allow for any permanent shocks to the SDF, while many asset-pricing models do feature such shocks prominently.[22]

The next example considers a more fully specified structural macro-finance model that also features time-varying risk premia and permanent SDF shocks (see Bakshi and Chabi-Yo, 2012), and which has been advanced as a rationalization of the excess-volatility puzzle.

EXAMPLE 3. Consider the variable-rare-disasters model of Gabaix (2012), described fully in Appendix B.2. Under the assumptions stated there, given any market-index option horizon T and any (small) positive value δ, there exists a return state $\underline{s}$ such that for all $s_j \geqslant \underline{s}$, the conditional probability of having realized at least one disaster over the life of the option is negligible:

$$\mathbb{P}_0\left(\sum_{t=1}^{T} \mathbb{1}\{\text{disaster}_t\} > 0 \;\middle|\; R_T^m \geqslant \underline{s}\right) < \delta.$$

For all $s_j \geqslant \underline{s}$, CTI holds for any two such adjacent states up to a negligible error, as $\phi_{t,j} = \phi_j + \eta_t$ with $\mathbb{E}[\eta_t] = \mathcal{O}(\delta)$ as $\delta \to 0$. ‖

[22] As can be seen more generally outside the context of our specific examples, models with *only* permanent shocks to the SDF are compatible with the CTI assumption; the assumption might accordingly be considered reasonable to the extent that such models provide accurate approximations to the data.

This result implies that for an economy described by this model, we need only focus attention on conditional probabilities across adjacent return states for which there is little to no probability of having realized a disaster conditional on reaching that state.

Finally, we consider a model that is instructive for the types of environments in which CTI does not hold.

EXAMPLE 4. Consider the external-habit-formation model of Campbell and Cochrane (1999), described fully in Appendix B.3. Under the assumptions stated there, CTI in general fails to hold. ‖

Given this model's specification of habit formation, the path of consumption always matters in a manner not fully accounted for by conditioning on the return state. We note, however, that nothing about our theoretical framework requires considering beliefs over *return* states: all the results below would apply if we were to consider beliefs over the elementary states in Ω, or, in the Campbell–Cochrane case, over the joint realization of the terminal consumption and *surplus* consumption values (see Appendix B.3). But while this allows us to sidestep the issue of CTI in theory, empirical implementation is infeasible: we observe options and risk-neutral beliefs over traded prices rather than over, e.g., surplus-consumption values.[23] This motivates the theoretical exposition we use for beliefs over return states specifically: the exposition in this case is less general than it could be, but it allows us to map directly between our theory and empirics.

We later consider the effects of the violation of CTI implied by the Campbell–Cochrane model in the context of a calibrated simulation study.

4. Theoretical Results

We move now to our main theoretical results. Our bounds consider the relationship between risk-neutral belief movement and risk-neutral uncertainty resolution in Definitions 2–3, and they are accordingly the analogues of the results in Section 2 for the general framework introduced in the previous section. The first subsection below presents and discusses our main belief-volatility bounds; the second subsection discusses how they can be implemented empirically; and the third provides extensions and additional results useful for interpretation of the main results.

4.1. Main Bounds

PROPOSITION 1. *For any return-state pair (s_j, s_{j+1}) meeting CTI, the following bound must hold under rational expectations for arbitrary option expiration date T:*

$$\widetilde{\mathbb{E}}_0[m_j^* - r_j^*] \leqslant \widetilde{\pi}_{0,j}^{*\,2}\left(1 - \frac{1}{\widetilde{\pi}_{0,j}^* + \phi_j(1 - \widetilde{\pi}_{0,j}^*)}\right).$$

[23] The same applies, for example, to economies with recursive preferences and independent volatility shocks (e.g., "Case II" of the long-run-risks model of Bansal and Yaron, 2004), since then marginal utility at T can depend on expected future volatility in addition to terminal consumption (or returns). See Walden (2017) for related discussion.

This result relates the unobserved structural parameter ϕ_j, which corresponds to the slope of the SDF across the two adjacent return states, to a set of observable values. (We discuss the observability of these values in the next subsection, but for now we take as given that they are observable.) Under risk neutrality ($\phi_j = 1$), this upper bound becomes zero: belief movement for $\widetilde{\pi}^*_{t,j}$ must not exceed uncertainty resolution on average conditional on state s_j or s_{j+1} being realized, following Lemma 1.[24] But this bound is otherwise positive, and the *admissible excess movement* in risk-neutral beliefs given by the right side of the inequality increases monotonically in ϕ_j. Movement in risk-neutral beliefs must still correspond on average to the agent learning something about the true terminal state, but in this more general case, the bias in risk-neutral beliefs relative to subjective beliefs induced by risk aversion allows for positive excess movement in those observed beliefs under RE. This result thus formalizes a more general notion of the "correct" amount of belief volatility under rationality, this time as an increasing function of the market's effective risk aversion between the low- and high-return states.

The bound in Proposition 1 is conservative in that it holds over all possible signal-generating processes. The question we ask is effectively, for given values of $\widetilde{\pi}^*_{0,j}$ and ϕ_j, what is the upper bound on $\widetilde{\mathbb{E}}_0[m^*_j - r^*_j]$ that a "malevolent" outside agent could attain given a choice over signal-generating processes? While the proof of the proposition above does not require fully characterizing the maximizing signal-generating process, directly considering the set of possible processes leads to the following additional result.

PROPOSITION 2. *The bound in Proposition 1 is approximately tight: the stated bound holds with strict inequality for any fixed $T < \infty$ as long as $\phi_j > 1$, but there exists a signal-generating process under which*

$$\lim_{T\to\infty} \widetilde{\mathbb{E}}_0[m^*_j - r^*_j] = \widetilde{\pi}^{*2}_{0,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0,j} + \phi_j(1 - \widetilde{\pi}^*_{0,j})}\right).$$

We characterize the maximizing process fully in the proof of the proposition in Appendix A, but it can be intuitively thought of as a "rare-bonanzas" process: with probability $1 - \epsilon$ the agent's belief that the bad-state return will be realized increases slightly, and with probability ϵ the agent receives news that the good state will be realized with certainty. The bad-state risk-neutral probability is upwardly biased relative to the agent's bad-state subjective belief, so maximizing risk-neutral belief movement requires maximizing the size of possible downward revisions to the bad-state belief (see Figure 1). This process achieves this maximization, and in the limit as $T \to \infty$, *all* of the expected belief movement comes from these downward revisions: the upward revisions are infinitesimally small, so the squared change in beliefs given an upward revision disappears given that it is second-order. The conservatism of the bound then depends on the extent to which such a process is unrealistic relative to whatever true process governs agents' beliefs.

While the bound in Proposition 1 in general maps between observed values and the slope of

[24] An associated lower bound of zero can also be shown to hold in this case. We note also that the expectation in the bound is conditional on $R^m_T \in \{s_j, s_{j+1}\}$ in this more general case because conditional physical beliefs are martingales only under the conditional measure; see Lemma A.1 in Appendix A.

the SDF required to rationalize those values, the fact that risk-neutral beliefs are bounded between 0 and 1 by construction implies that the bound is well-defined even for infinitely large risk aversion: there is only so far that risk-neutral beliefs can be distorted relative to subjective beliefs. Thus taking $\phi_j \to \infty$ in that bound yields the following full-identification corollary.

COROLLARY 1. *If* $\widetilde{\mathbb{E}}_0[m_j^* - r_j^*] > \widetilde{\pi}_{0,j}^{*2}$, *then no SDF process meeting CTI can rationalize the variation in risk-neutral beliefs for the given return-state pair.*

In the case that $\widetilde{\mathbb{E}}_0[m_j^* - r_j^*] > \widetilde{\pi}_{0,j}^{*2}$, the conclusion to draw would not necessarily be that there are violations of the no-arbitrage condition; instead, there would be no *rational-expectations* SDF process capable of generating the observed excess movement in risk-neutral beliefs. Thus the *actual* SDF process translating between objective probabilities and risk-neutral beliefs would in this case include a belief distortion that induces excess volatility. Such a finding would be close in spirit to a violation of the "good-deal bounds" of Cochrane and Saá-Requejo (2000): even if the no-arbitrage condition holds, there would be an investment strategy with a large Sharpe ratio under the objective measure that is nonetheless not traded away by investors, because their subjective measure is distorted and thus does not perceive this large Sharpe ratio.[25]

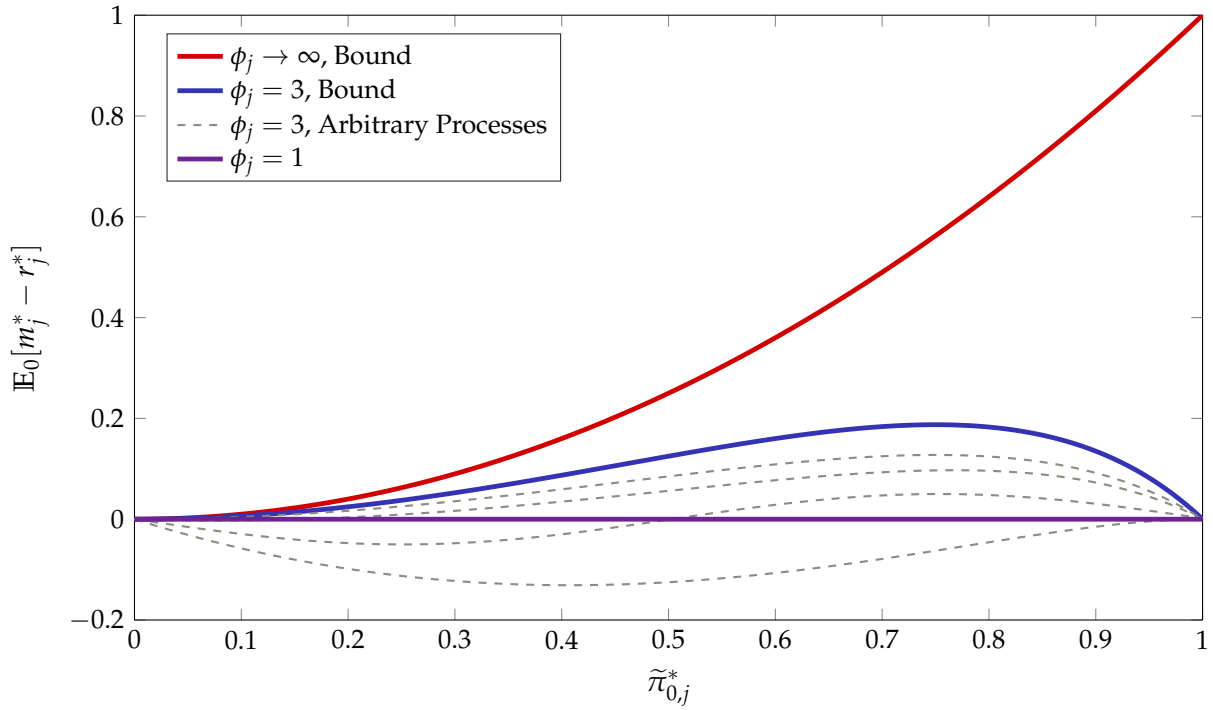
Taken together, Proposition 1 and Corollary 1 characterize the admissible excess movement in risk-neutral beliefs as a function of ϕ_j for any risk-neutral prior. We illustrate these bounds graphically in Figure 2. Starting from the bottom of the chart, the thick purple line corresponds to the bound for the risk-neutral case of $\phi_j = 1$: in this case, excess movement must be zero in expectation regardless of the prior or signal-generating process, from Lemma 1. The thin dashed gray lines correspond to arbitrarily selected signal-generating processes in the case of some risk aversion, $\phi_j = 3$. While there can be positive excess belief movement, this is not necessarily the case for all possible signal-generating processes. Taking the envelope over all of these processes for $\phi_j = 3$ yields the bound shown in the thick blue line. While the admissible excess movement is non-monotonic in the risk-neutral prior in this case, values greater than 0.5 tend to yield greater admissible movement, following the logic of Figure 1. Finally, the thick red line shows the bound for the limiting case $\phi_j \to \infty$, which is equal to the squared risk-neutral prior from Corollary 1.

The non-monotonicity of the bound in the $\phi_j = 3$ case is a general feature of the bounds for $1 < \phi_j < \infty$,[26] and it arises due to the interplay of two competing forces. A greater risk-neutral prior yields more "room" for downward movement of the belief, which under the maximizing signal-generating process increases the expected excess movement. On the other hand, a greater risk-neutral prior corresponds to a greater subjective prior for any given value of ϕ_j, which decreases the likelihood that such a downward movement will be realized. In the limit as $\phi_j \to \infty$, the first force dominates the second, as the underlying subjective prior is pushed arbitrarily close to zero for any given risk-neutral prior given very large values of risk aversion.

[25] This strategy would take the form of betting on mean reversion in risk-neutral beliefs. See also Hansen and Jagannathan (1997) and Hansen (2014) on the possibility of apparent mispricing under no arbitrage.

[26] In particular, one can show that the formula on the right side of Proposition 1 yields an interior maximum for admissible excess movement at $\widetilde{\pi}_{0,j}^* = (4\phi_j - \sqrt{8\phi_j + 1} - 1)/(4(\phi_j - 1)) \in (0, 1)$.

Figure 2: Excess Belief Movement vs. Prior by ϕ_j Under RE



Note: Bounds are obtained from the formulas in Proposition 1 and Corollary 1.

One can also characterize the form of departure from rationality required for this limiting bound to be violated, as in the following result.

PROPOSITION 3. *Assuming CTI holds for the return-state pair (s_j, s_{j+1}), the effects of an incorrect physical prior, $\widetilde{\pi}_{0,j} \neq \mathbb{P}_0(R_T^m = s_j \mid R_T^m \in \{s_j, s_{j+1}\})$, are limited as follows:*

(i) If $\widetilde{\pi}_{0,j}^* < \mathbb{P}_0(R_T^m = s_j \mid R_T^m \in \{s_j, s_{j+1}\})$, then under Bayesian updating, it must be the case that $\widetilde{\mathbb{E}}_0[m_j^* - r_j^*] \leqslant \max\{\widetilde{\pi}_{0,j}^{*2}, (1 - \widetilde{\pi}_{0,j}^*)^2\}$.

(ii) Otherwise, an incorrect prior cannot by itself lead to $\widetilde{\mathbb{E}}_0[m_j^* - r_j^*] > \widetilde{\pi}_{0,j}^{*2}$.

This result tells us that updating behavior must in general play some role in any outright rejection of RE as in Corollary 1. Part (i) of the proposition shows that there is technically one case in which an incorrect physical prior can lead to $\widetilde{\mathbb{E}}_0[m_j^* - r_j^*] > \widetilde{\pi}_{0,j}^{*2}$, but this case is unlikely empirically: it requires that the prior be so downwardly distorted that the *risk-neutral* belief is below the true conditional *physical* probability, which is a strong requirement given that $\widetilde{\pi}_{0,j}^* \geqslant \widetilde{\pi}_{0,j}$ under our labeling of s_j as the "bad" state.[27]

[27] It is also the case in our empirical setting that $\mathbb{E}[\widetilde{\pi}_{0,j}^*] \approx \mathbb{E}[1 - \widetilde{\pi}_{0,j}^*]$ given the use of conditional beliefs; thus the bound in part (i) is approximately equally tight as the bound in part (ii) in practice. And to see why the finding in case (i) arises, consider a prior distorted so low that $\widetilde{\pi}_{0,j}^* \approx 0$ despite $\mathbb{P}_0(R_T^m = s_j \mid R_T^m \in \{s_j, s_{j+1}\}) \gg 0$; given correct updating, there will be positive excess movement in excess of $\widetilde{\pi}_{0,j}^{*2} \approx 0$ but still less than $(1 - \widetilde{\pi}_{0,j}^*)^2$.

An incorrect physical prior has limited effects because it acts as a one-time distortion in beliefs; while moving back to the "correct" belief in this case does require some excess movement in beliefs, this excess movement is generally not sufficient to produce a full violation of the bound in Proposition 1. Incorrect updating behavior must accordingly be present in such a violation, and the restriction imposed by our bound implies that this incorrect updating behavior necessarily entails excessive volatility in beliefs relative to the degree of uncertainty resolution over time. We discuss theoretical models that are capable of generating this form of incorrect updating behavior in Section 7.

4.2. Observability and Empirical Implementation

The discussion to this point has taken as given that all the values in the bound in Proposition 1 are observable aside from ϕ_j, but this has elided one issue: this bound is stated over the date-0 expectation of excess belief movement in terms of the risk-neutral prior $\widetilde{\pi}^*_{0,j}$, but we observe only one draw $m^*_j - r^*_j$ per expiration date rather than the ex-ante expectation of this random variable. (The expectation is also conditional on $R^m_T \in \{s_j, s_{j+1}\}$, but this outcome is observable and can be conditioned on empirically.) We show now how the use of risk-neutral beliefs series over many expiration dates allows for empirical implementation of our bounds.

We must first generalize the environment and notation slightly. We now assume that we can observe prices of options over the value of the market index on some set of N option-expiration dates $\mathcal{T} \equiv \{t : t \in (T_1, T_2, \ldots, T_N)\}$, or $\{T_i\}_{i=1,\ldots,N}$. For arbitrary expiration date T_i, denote by 0_i the first date on which the price of any such option contract (with arbitrary strike K) is observable.[28] All other objects maintain their previous definitions, but the subscript i is now used to denote a value corresponding to maturity date T_i; for example, risk-neutral belief movement is $m^*_{i,j} = \sum_{t=0_i+1}^{T_i} (\widetilde{\pi}^*_{t,i,j} - \widetilde{\pi}^*_{t-1,i,j})^2$.

The CTI assumption does not specify that the value $\phi_{i,j}$ must be constant across all observed expiration dates $\{T_i\}_i$; we thus need some way of aggregating the observed values $\{m^*_{i,j} - r^*_{i,j}\}_i$ across i to identify an average SDF slope for that return-state pair. Due to Jensen's inequality, we cannot simply insert $\widetilde{\mathbb{E}}[\phi_{i,j}]$ in place of $\phi_{i,j}$ and $\widetilde{\mathbb{E}}[\widetilde{\pi}^*_{0_i,i,j}]$ in place of $\widetilde{\pi}^*_{0_i,i,j}$ when taking the expectation of both sides of the bound in Proposition 1: viewing the expression on the right side of that bound as a function of $\phi_{i,j}$ and $\widetilde{\pi}^*_{0_i,i,j}$, this function's Hessian is not in general negative semidefinite given the non-monotonicity in the bound, so naïvely conducting such replacements will not in general uphold the validity of the bound. However, for any given prior $\widetilde{\pi}^*_{0_i,i,j}$, the bound is concave in $\phi_{i,j}$, as the second partial derivative of that expression with respect to $\phi_{i,j}$ is $-2{\widetilde{\pi}^*_{0_i,i,j}}^2(1 - \widetilde{\pi}^*_{0_i,i,j})^2/(\widetilde{\pi}^*_{0_i,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{0_i,i,j}))^3 \leqslant 0$. This implies the following result, which effectively applies Jensen's inequality for one of several variables.

[28] Formally, 0_i is the minimal t for which there exists a $K \in \mathcal{K}$ such that the call-option price $q^m_{t,i,K}$ is $\mathcal{F}_t$-measurable.

PROPOSITION 4. *For any return-state pair (s_j, s_{j+1}) meeting CTI, the following bound must hold under rational expectations over all option expiration dates:*

$$\widetilde{\mathbb{E}}[m^*_{i,j} - r^*_{i,j}] \leqslant \widetilde{\mathbb{E}}\left[\widetilde{\pi}^{*\,2}_{0_i,i,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0_i,i,j} + \overline{\phi}_j(1 - \widetilde{\pi}^*_{0_i,i,j})}\right)\right],$$

where $\overline{\phi}_j \equiv \max_{\widetilde{\pi}_{0_i,i,j}} \widetilde{\mathbb{E}}[\phi_{i,j} \mid \widetilde{\pi}_{0_i,i,j}]$.

This bound is now implementable empirically: we can measure a sample counterpart of the mean excess-movement statistic on the left side, and the minimum $\overline{\phi}_j$ that solves the bound given the observed excess movement and risk-neutral priors is then a conservative estimate of the maximum conditional-mean SDF slope for the return-state pair in question over all dates T_i. Further, if $\widetilde{\mathbb{E}}[\phi_{i,j} \mid \widetilde{\pi}_{0_i,i,j}] = \widetilde{\mathbb{E}}[\phi_{i,j}]$ — as might be expected to hold approximately, given that the prior over the return may not be especially informative for the expected relative SDF realizations — then $\overline{\phi}_j = \widetilde{\mathbb{E}}[\phi_{i,j}]$, and the proposition allows us to identify a lower bound for this value.

While the values $\{\overline{\phi}_j\}$ will in general differ across return states s_j, we can apply the same result as used in Proposition 4 to obtain a single estimate $\overline{\phi}$ of the required mean SDF slope both across dates *and* return states (for all states meeting CTI) if desired. We discuss our empirical estimation further in Section 5.

4.3. Economic Interpretation and Robustness Results

We now turn to a set of additional results that yield a clearer economic interpretation for ϕ_j, allow us to account for the possibility of mismeasurement or market microstructure noise, and finally extend the results to cases in which conditional transition independence is violated mildly. (For notational simplicity, in this subsection we temporarily return to the environment considered before Section 4.2, with a single expiration date T.)

First, while the results above are convenient to express in terms of the SDF slope ϕ_j given that this allows for closed-form solutions that can be applied across a wide range of structural models regardless of the origin of the SDF, the results also admit an interpretation in terms of the approximate required risk-aversion value for a fictitious representative agent with utility over the terminal value of the market index, as in Section 2.

PROPOSITION 5. *Assume additionally that there is a representative agent with (indirect) utility over time-T wealth, with wealth equal to the market index value, and denote $V^m_j \equiv V^m_0 s_j$. Then relative risk aversion $\gamma_j \equiv -V^m_j U''(V^m_j)/U'(V^m_j)$ is given to a first order around return state s_j by*

$$\gamma_j = \frac{\phi_j - 1}{\Delta_j},$$

where $\Delta_j \equiv (s_{j+1} - s_j)/s_j$ is the percent return deviation between adjacent return states s_j and s_{j+1}.

As in the case in Section 2, relative risk aversion is proportional to $\phi_j - 1$, since this gives the percent decrease in marginal utility in moving from low-return state s_j to high-return state s_{j+1}. This change in marginal utility must be normalized by the consumption increase in moving from s_j to s_{j+1} in order to calculate relative risk aversion, which thus requires dividing through by Δ_j. If, for example, $s_j = 1$, $s_{j+1} = 1.01$, then a value $\phi_j = 1.1$ implies $\gamma_j = 10$.

Given that the expression for γ_j in this proposition is affine in ϕ_j, the aggregation result in Proposition 4 can be applied to conduct valid estimation of $\overline{\gamma}_j \equiv \max_{\widetilde{\pi}_{0_i,i,j}} \widetilde{\mathbb{E}}[\gamma_{i,j} \mid \widetilde{\pi}_{0_i,i,j}]$ across different expiration dates T_i, or additionally across different return states s_j in the case that one would like a single estimate of the average local relative risk aversion value across all such states.

We now provide a result that can be applied to the bounds above to correct for possible mismeasurement of risk-neutral beliefs. Our bounds provide a minimum value of the slope of the SDF required to rationalize the observed variation in risk-neutral beliefs; if some of this variation is in fact arising due to, e.g., transient demand pressures, then we may overestimate this required SDF slope. However, a simple correction can be applied to account for this issue, as follows.

PROPOSITION 6. *Assume that the observed conditional risk-neutral belief $\widehat{\pi}^*_{t,j}$ is measured with error with respect to the true value $\widetilde{\pi}^*_{t,j}$:*

$$\widehat{\pi}^*_{t,j} = \widetilde{\pi}^*_{t,j} + \epsilon_{t,j},$$

where $\widetilde{\mathbb{E}}[\epsilon_{t,j}] = 0$, $\widetilde{\mathbb{E}}[\epsilon_{t,j}\,\epsilon_{t+1,j}] = 0$, and $\widetilde{\mathbb{E}}[\epsilon_{t+k,j}\,\widetilde{\pi}^*_{t+k',j}] = 0$ for $k, k' \in \{0, 1\}$. Denoting the observed one-period expected excess movement statistic by $\widetilde{\mathbb{E}}[\widehat{m}^*_{t,t+1,j} - \widehat{r}^*_{t,t+1,j}]$, its relation to the true value $\widetilde{\mathbb{E}}[m^*_{t,t+1,j} - r^*_{t,t+1,j}]$ is then given by

$$\widetilde{\mathbb{E}}[\widehat{m}^*_{t,t+1,j} - \widehat{r}^*_{t,t+1,j}] = \widetilde{\mathbb{E}}[m^*_{t,t+1,j} - r^*_{t,t+1,j}] + 2\mathrm{Var}(\epsilon_{t,j}).$$

We can thus subtract $2\mathrm{Var}(\epsilon_{t,j})$ from each period's observed excess-movement statistic to identify true excess movement, which can then be used in Proposition 4 after summing over the full path. If measurement error is positively correlated over time rather than uncorrelated, this will reduce the upward bias in measured volatility of beliefs. One might instead worry about negatively correlated measurement errors in the case of bid-ask bounce, but our empirical measurement uses only end-of-day mid-price data, and as shown by Jacod, Li, and Zheng (2017) with high-frequency data, the autocorrelation values for such noise have long died out at one-day lags. We discuss estimation of the value $\mathrm{Var}(\epsilon_{t,j})$ in the next section.

Proposition 6 allows for a correction with respect to empirical misspecification; we turn now to a result that speaks to the possibility of theoretical misspecification. While our identifying assumption, conditional transition independence, is met in some commonly used theoretical frameworks, it is a knife-edge restriction that is unlikely to be met exactly in the data. Further, one may wish to consider the implications of our results for theoretical frameworks in which CTI is not met. We accordingly ask under what conditions the bound holds approximately even when CTI is violated mildly, and we obtain a sufficient condition as follows.

PROPOSITION 7. *Recall the definition* $\phi_{t,j} \equiv \mathbb{E}_t[M_T/M_t \mid R_T^m = s_j]/\mathbb{E}_t[M_T/M_t \mid R_T^m = s_{j+1}]$ *in equation (19), and that conditional transition independence assumes that this value is constant for all* $0 \leqslant t < T$. *Relaxing this assumption, we obtain the following two results:*

(i) If the sequence $\{\phi_{t,j}\}_t$ *is a martingale with respect to filtration* $\mathbb{F}$ *conditional on* $R_T^m \in \{s_j, s_{j+1}\}$, *or* $\phi_{t,j} = \widetilde{\mathbb{E}}_t[\phi_{t+1,j}]$, *then the bound in Proposition 1 applies, with the ex-ante value* $\phi_{0,j}$ *replacing* ϕ_j.

(ii) The bound in Proposition 1 applies to an arbitrarily close approximation within a neighborhood of $\{\phi_{t,j}\}_t$ *being a martingale:* $\forall\, \epsilon > 0$, $\exists\, \delta > 0$ *such that if* $|\widetilde{\mathbb{E}}_t[\phi_{t+1,j}] - \phi_{t,j}| < \delta$ *almost surely for* $t = 0, \ldots, T-1$, *then*

$$\left| \widetilde{\mathbb{E}}_0[m_j^* - r_j^*] - \widetilde{\pi}_{0,j}^{*\,2} \left(1 - \frac{1}{\widetilde{\pi}_{0,j}^* + \phi_{0,j}(1 - \widetilde{\pi}_{0,j}^*)} \right) \right| < \epsilon.$$

Part (i) of the result tells us that a martingale process for $\{\phi_{t,j}\}_t$ yields no more excess movement in beliefs for any given ex-ante value $\phi_{0,j}$ than is expected under the movement-maximizing signal-generating process used in Proposition 1 itself.[29] The sequence $\{\phi_{t,j}\}_t$ is unlikely to follow a martingale exactly. This martingale condition is, however, a convenient benchmark around which to obtain an approximate bound as in part (ii) of the proposition, which shows that the original bound is continuous in the limit as $\widetilde{\mathbb{E}}_t[\phi_{t+1,j}] \to \phi_{t,j}$. As we will see in Section 6, $\phi_{t,j}$ is closely approximated by a martingale in a calibrated version of the Campbell and Cochrane (1999) habit-formation model, as we show in numerical simulations that our bound still holds in this framework despite the violation of CTI. We expect that most models do not generate enough variation in relative expected marginal utility across adjacent market-return states to yield large departures from the $\phi_{t,j}$-martingale benchmark; intuitively, variation in $\phi_{t,j}$ will also be restricted under RE as $\phi_{t,j}$ is itself a function of conditional expectations, and it would require dramatic variation in $\phi_{t,j}$ at a daily frequency to generate the empirical results below.

5. Empirical Estimation and Main Results

Our theory leads to bounds on the variation in risk-neutral beliefs over the value of the market index, which we proceed now to measure in the data.

5.1. Data

We obtain S&P 500 index options data from the OptionMetrics database, which provides end-of-day bid and ask prices for European call and put options on the index value for all Chicago Board Options Exchange–traded strike prices and option expiration dates. We calculate each option's

[29] One might wonder how this can be the case given the measurement-error result in Proposition 6. The key distinction arises from the assumption of i.i.d. noise in this previous result, which induces mean reversion in measured risk-neutral beliefs, while a martingale in $\phi_{t,j}$ does not.

end-of-day price as the midpoint between its bid and ask price. The available sample is from January 1996 through August 2015, which yields data for 4,949 trading dates t for options traded over 685 expiration dates T_i. We refer to the collection of call or put option prices traded across strikes K on date t for a given expiration date T_i as an option *cross-section*.

We apply standard filters to remove outliers and options with poor trading liquidity from the raw options-price data: we drop any options with bid prices of zero (or less than zero), with Black–Scholes implied volatility of greater than 100 percent (for which we discuss measurement below), with greater than six months to maturity (e.g., Constantinides, Jackwerth, and Savov, 2013), and all cross-sections with fewer than three listed prices across different strikes. This procedure yields roughly 6.7 million observed option prices, as compared to 7.4 million option prices before cleaning. Finally, after transforming the observed option prices to risk-neutral beliefs as discussed below, we keep only conditional risk-neutral beliefs $\widetilde{\pi}^*_{t,i,j}$ for which the non-conditional return-state beliefs satisfy $\pi^*_t(R^m_{T_i} = s_j) + \pi^*_t(R^m_{T_i} = s_{j+1}) \geqslant 5\%$, and we label $\widetilde{\pi}^*_{t,i,j}$ as missing when this condition is violated. We do so because conditional beliefs $\widetilde{\pi}^*_{t,i,j}$ are likely to be particularly susceptible to mismeasurement when the underlying beliefs are close to zero.

5.2. Mapping Theory to Data: Return Space and Risk-Neutral Beliefs

Defining the empirical return space. For our baseline estimation, we define our return-state space as in (8) according to the following set:

$$\mathcal{S}_{\text{baseline}} = \exp\left(\{(-\infty, -0.11), -0.09, -0.07, \ldots, 0.07, 0.09, (0.11, \infty)\}\right), \tag{20}$$

so that return states are evenly spaced two percentage points apart, bracketed by the two extreme log-return states $(-\infty, -0.11)$ and $(0.11, \infty)$. With slight abuse of notation, we refer to the log-return states as $s_1 = (-\infty, -0.11), s_2 = -0.09, \ldots$, rather than using s_j to refer to the gross-return state as in the previous section. Further, we in fact typically report results in terms of log *excess* returns, so that we say log-return state s_j is realized if

$$\log(R^{m,e}_{T_i}) \equiv \log(V^m_{T_i}/V^m_{0_i}) - \log(R^f_{0_i,T_i}) = s_j,$$

where V^m_t is the value of the S&P 500, and the $(T_i - 0_i)$-period gross risk-free rate $R^f_{0_i,T_i}$ is measured using the zero-coupon risk-free yield curve provided by OptionMetrics.[30] We do so only for convenience of interpretation; the risk-free rate used here is known ex-ante, so the excess-return states map one-for-one to non-excess-return states as defined in (8) for each option expiration date.

In general, we of course do not observe excess returns of, say, exactly 0.05. Rather than attempting to measure the risk-neutral density exactly at the points in (20), we instead discretize the state space as follows: we say that interior return state s_j ($j \neq 1, J$) is realized if the excess return is in

[30] We linearly interpolate the log risk-free rate between the two nearest maturity dates listed in the OptionMetrics yield curve when necessary. OptionMetrics measures risk-free rates using the London Interbank Offered Rate (LIBOR) where available and otherwise using the Eurodollar futures curve, following market convention.

a neighborhood of one percentage point of s_j, or $\log(R^{m,e}_{T_i}) \in [s_j - 0.01, s_j + 0.01)$. Thus one could equivalently consider our return-state specification to be in terms of two-percentage-point *bins* for excess returns: $\mathcal{S}_{\text{baseline}} = \exp(\{(-\infty, -0.10), [-0.10, -0.08), [-0.08, -0.06), \ldots, [0.06, 0.08), [0.08, 0.10), [0.10, \infty)\})$.

We use two-percentage-point spacing in the return space $\mathcal{S}_{\text{baseline}}$ as defined here in an attempt to balance the tradeoff between (i) measurement accuracy for the risk-neutral beliefs and (ii) the plausibility of our assumption of conditional transition independence as in Definition 4. Wider bins lead to greater accuracy of measurement, but conversely make it less likely that there are no changes in the expected realization of the SDF conditional on reaching a given return state (relative to the realization conditional on the relevant adjacent return state).

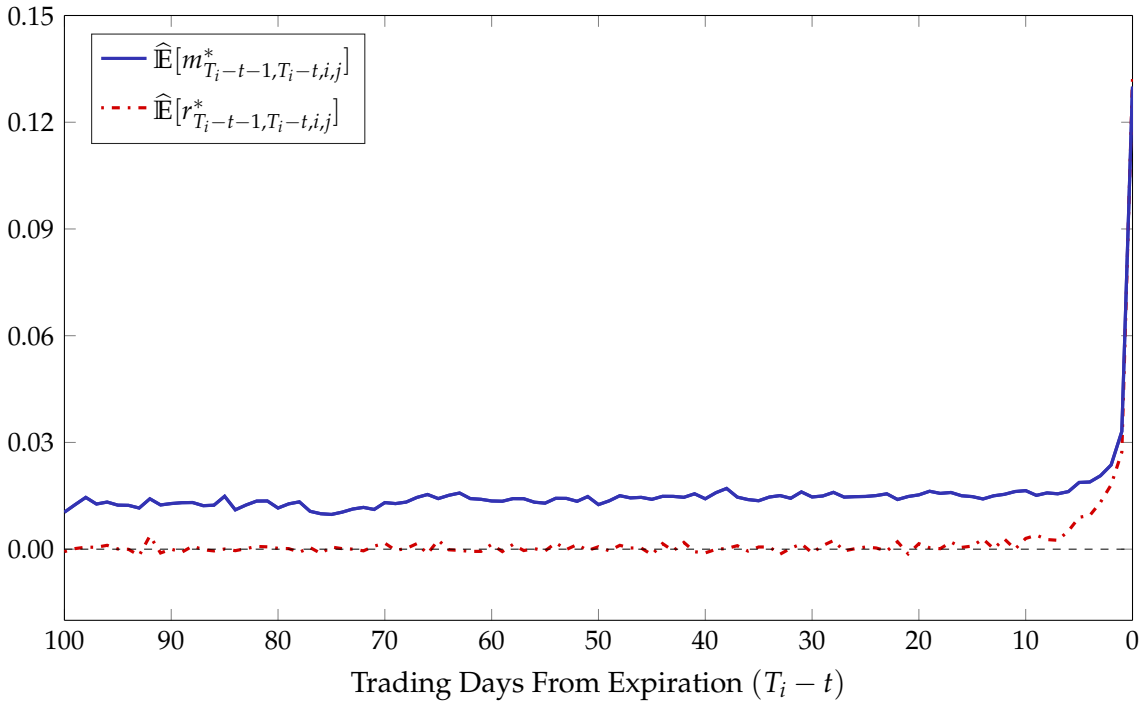
While we report empirical estimates below constructed using conditional risk-neutral beliefs for the extreme state s_1 relative to s_2, and for s_{J-1} relative to s_J, we do so only for completeness, as we do not believe these return-state pairs meet CTI (given the issue discussed just above, or following the logic of an economy with jumps or rare disasters, as in Example 3 in Section 3.2). And when we aggregate our state-by-state estimates of risk aversion required to rationalize the data into a single average risk-aversion value across states, we exclude data from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J) when doing so. This yields an additional de facto filter on the data: this effectively truncates each option cross-section on both sides, using options prices only for strikes with moneyness between 0.9 and 1.1 in conducting our tests. This is slightly more conservative than in most related literature (again, for example, see Constantinides, Jackwerth, and Savov, 2013, who use options with moneyness betwen 0.8 and 1.2).

For some of our additional empirical exercises and robustness tests in Section 6, we consider different specifications of the return-state space. For example, when we re-estimate our empirical tests for observations grouped by time to expiration (i.e., using excess belief movement observed in the last two weeks before expiration, in the third and fourth weeks before expiration, ...), we specify the space as $\mathcal{S}_{\text{alt}} = \exp(\{(-\infty, -0.225), -0.175, -0.125, \ldots, 0.125, 0.175, (0.225, \infty)\})$, to increase the available number of observations for each such time-to-expiration group.[31] We have further experimented with different definitions of the return space in estimating our main results (e.g., defining states in terms of ex-ante option delta, $\Delta_{0_i,i,K} \equiv \partial q^m_{0_i,i,K} / \partial V^m_{0_i}$); in all cases such changes make very little difference for the results reported below. Finally, we use non-annualized returns in our return-space definitions, as this admits comparisons between return states in terms of the percent wealth deviation (rather than the annualized percent wealth deviation) across those states. But we have additionally conducted tests with state spaces instead specified in terms of annualized returns, which again leads to very similar results (all available upon request).

Measuring risk-neutral beliefs. We use techniques similar to those developed in related literature to extract risk-neutral beliefs over the empirical return states from the observed option cross-sections. Our starting point is equation (14), which tells us how to map from option prices to

[31] This increases available observations given our requirement that $\pi^*_t(R^m_{T_i} = s_j) + \pi^*_t(R^m_{T_i} = s_{j+1}) \geqslant 5\%$ and that the distribution of returns tends to be quite concentrated near the mean of the distribution close to the expiration date.

Figure 3: Average One-Day m^* and r^* Across Expiration Dates and States



Notes: Empirical averages $\widehat{\mathbb{E}}[\cdot]$ calculated across all expiration dates T_i and state pairs (s_j, s_{j+1}) in (20), aside from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J).

risk-neutral probabilities. We use this to construct a smooth risk-neutral distribution for returns, largely following the technique proposed by Malz (2014). Appendix B.4 provides a description of this procedure.

5.3. Diagnostic Statistics for Risk-Neutral Beliefs

With the risk-neutral beliefs in hand, we can then calculate the observed conditional risk-neutral belief for state s_j versus s_{j+1} and measure risk-neutral belief movement and uncertainty resolution following Definitions 2–3. We briefly highlight here some high-level diagnostic statistics for belief movement and uncertainty resolution to give a sense of the values for these objects observed in the data.

Figure 3 plots the empirical means, denoted by $\widehat{\mathbb{E}}[\cdot]$, of the one-day risk-neutral belief movement and uncertainty resolution statistics across all expiration dates and state pairs, aside from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J), by trading days from expiration $T_i - t$. The x-axis sorts these days to expiration in decreasing order, so that the left side of the chart corresponds to trading days relatively far from expiration and the right side to days near expiration. Daily movement is plotted in solid blue, and daily uncertainty uncertainty resolution in dash-dotted red.

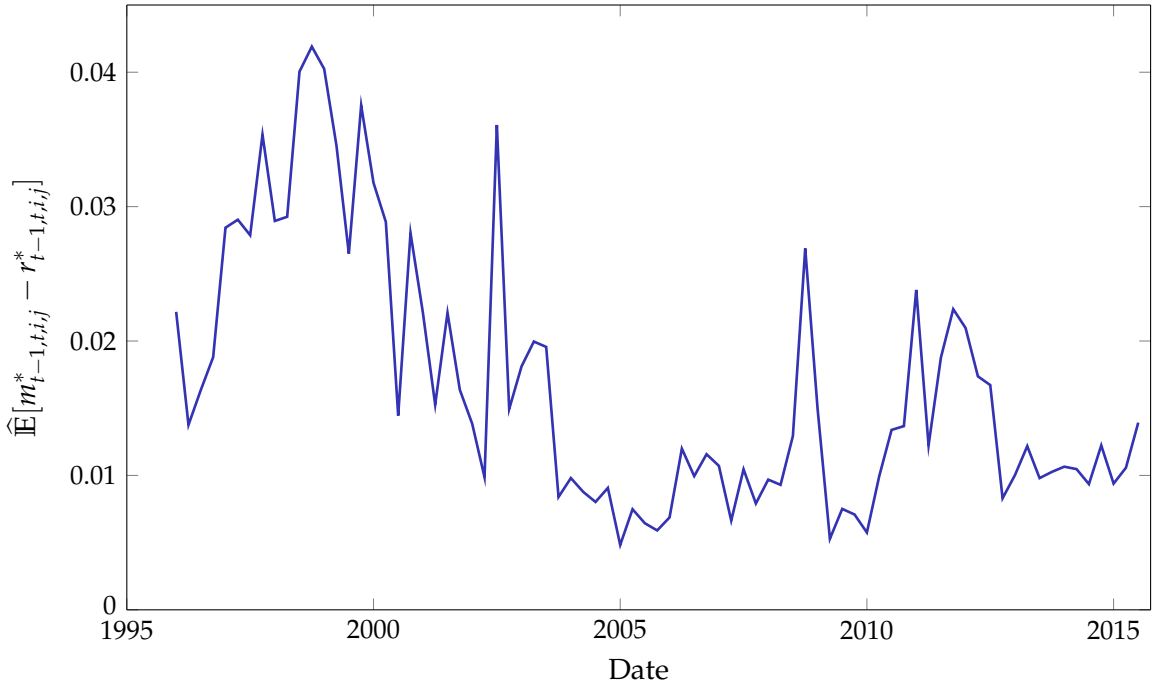
Examining this figure, two facts are immediately apparent. First, for relatively distant horizons from a given option expiration (beyond about ten trading days), there is small but very

consistently positive daily risk-neutral belief movement, while risk-neutral uncertainty resolution is consistently indistinguishable from zero. Second, for the last two weeks before expiration, risk-neutral beliefs move much more strongly, and this movement is in this case matched by strong uncertainty resolution.[32] Thus the distant-horizon beliefs seem more likely to pose challenges for the RE assumption than short-horizon beliefs. Further, the fact that there is zero uncertainty resolution for these distant horizons while movement must be positive by definition indicates that beliefs must essentially be bouncing back and forth without resolving any uncertainty.

To gain a rough understanding for the actual variation in risk-neutral beliefs corresponding to the movement values plotted in Figure 3, recall the definition of one-day movement, $m^*_{t-1,t,i,j} = (\widetilde{\pi}^*_{t,i,j} - \widetilde{\pi}^*_{t-1,i,j})^2$. Thus a value of $m^*_{t-1,t,i,j} = 0.01$ — roughly equal to the values plotted until the last 10 days before expiration — corresponds to a change in beliefs of $\sqrt{0.01} = 0.1$, or 10 percentage points, in either direction. But the beliefs $\widetilde{\pi}^*_{t,i,j}$ are again *conditional* beliefs for state s_j versus s_{j+1}, so these conditional-belief changes are larger than the changes in the full underlying distribution.[33]

Next, Figure 4 plots a time series of average one-day excess risk-neutral belief movement $\widehat{\mathbb{E}}[m^*_{t-1,t,i,j} - r^*_{t-1,t,i,j}]$ across all dates t in a given quarter and across state pairs. This level of time

Figure 4: Average One-Day Excess Belief Movement by Quarter



Notes: Empirical averages $\widehat{\mathbb{E}}[m^*_{t-1,t,i,j} - r^*_{t-1,t,i,j}]$ calculated across all available expiration dates T_i and state pairs (s_j, s_{j+1}) in (20), aside from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J), using all trading dates t within each quarter.

[32] This equivalence is not necessarily mechanical, as ex-ante beliefs (e.g., $\widetilde{\pi}^*_{T_i-1,i,j}$) must be properly calibrated in order for movement to be equal to uncertainty resolution in expectation.

[33] For example, with $\pi^*_{t-1}(s_j) + \pi^*_{t-1}(s_{j+1}) = \pi^*_t(s_j) + \pi^*_t(s_{j+1}) = 0.25$, then $m^*_{t-1,t,i,j} = 0.01$ corresponds to a change $|\pi^*_t(s_j) - \pi^*_{t-1}(s_j)| = \sqrt{0.01} \times 0.25 = 0.025$, or 2.5 percentage points.

aggregation allows us to see medium- to long-term trends and cycles in our measure of the excess volatility of risk-neutral beliefs. There is a mild downward trend in excess volatility over time, with periodic temporary spikes (e.g., in the third quarter of 2008, at the height of the financial crisis). This provides some preliminary evidence that measured excess belief movement is not simply a function of limits to arbitrage in asset markets (rather than changes in underlying beliefs); for example, Du, Tepper, and Verdelhan (2018) document large, persistent post-2008 deviations from covered interest parity in foreign-exchange markets where no such deviations had previously been observed, whereas our measure does not exhibit marked persistence after its 2008 spike. Nonetheless, one might be concerned that the downward trend in $\widehat{\mathbb{E}}[m^*_{t-1,t,i,j} - r^*_{t-1,t,i,j}]$ could reflect improved option-market liquidity over the sample period, in which case the large early-sample excess-movement values could be a function of poor liquidity rather than underlying belief movements. We examine this question further below.

5.4. Estimation and Inference

Estimation. Our theoretical results developed in Section 4 allow us to map directly from empirical excess belief movement to an estimate of the minimal risk aversion, as encoded in the SDF slope across adjacent return states $\overline{\phi}_j$, required to rationalize the data. First, fixing the return-state pair indexed by j (i.e., (s_j, s_{j+1})), define the set of dates $\mathcal{T}_j \equiv \{T_i \in \mathcal{T} : R^m_{T_i} \in \{s_j, s_{j+1}\}\}$, and $N_j \equiv |\mathcal{T}_j|$. A sample estimate of the lower bound for $\overline{\phi}_j$ can then be calculated using Proposition 4 as the value $\widehat{\phi}_j$ solving the following sample moment condition:

$$\frac{1}{N_j} \sum_{i\,:\,T_i \in \mathcal{T}_j} \left[m^*_{i,j} - r^*_{i,j} - \widetilde{\pi}^{*\,2}_{0_i,i,j} \left(1 - \frac{1}{\widetilde{\pi}^*_{0_i,i,j} + \widehat{\phi}_j (1 - \widetilde{\pi}^*_{0_i,i,j})} \right) \right] = 0. \tag{21}$$

We repeat this for each return-state pair, in each case using all observations for which one of the two return states is realized at expiration, to obtain a set of estimates $\{\widehat{\phi}_j\}_j$ for the required local slope of the SDF at each value in the return space $\mathcal{S}_{\text{baseline}}$.

While the estimator in (21) cannot be expressed in closed form given the nonlinearity of the expression in which it appears, there is nonetheless a unique value $\widehat{\phi}_j$ that solves this condition given that the left side of the equation is monotonically increasing in $\widehat{\phi}_j$. The only case in which this is not true is if there is *no* such value such that (21) holds, which, as in the result in Corollary 1, occurs when

$$\frac{1}{N_j} \sum_{i\,:\,T_i \in \mathcal{T}_j} \left[m^*_{i,j} - r^*_{i,j} - \widetilde{\pi}^{*\,2}_{0_i,i,j} \right] > 0. \tag{22}$$

For ease of notation and interpretation of the estimation results reported below, when (22) holds we say that $\widehat{\phi}_j = \infty$, but this should be understood as denoting that even a value of infinity cannot rationalize the observed excess belief movement.

We then define $\overline{\phi}$ to be the average required SDF slope across all option expiration dates *and* return-state pairs for which CTI holds; as discussed in Section 5.2, we assume this is the case

for $j = 2, \ldots, J-2$. Then the estimator $\widehat{\phi}$ for the minimum admissible value for $\overline{\phi}$ solves

$$\frac{1}{J-3}\sum_{j=2}^{J-2}\frac{1}{N_j}\sum_{i\,:\,T_i\in\mathcal{T}_j}\left[m^*_{i,j}-r^*_{i,j}-\widetilde{\pi}^{*\,2}_{0_i,i,j}\left(1-\frac{1}{\widetilde{\pi}^*_{0_i,i,j}+\widehat{\phi}(1-\widetilde{\pi}^*_{0_i,i,j})}\right)\right]=0, \tag{23}$$

and we say that $\widehat{\phi}=\infty$ when

$$\frac{1}{J-3}\sum_{j=2}^{J-2}\frac{1}{N_j}\sum_{i\,:\,T_i\in\mathcal{T}_j}\left[m^*_{i,j}-r^*_{i,j}-\widetilde{\pi}^{*\,2}_{0_i,i,j}\right]>0. \tag{24}$$

Proposition 5 then allows us to map straightforwardly from our estimates $\{\widehat{\phi}_j\}_j$ and $\widehat{\phi}$ to estimates for relative risk aversion for a representative agent: $\widehat{\gamma}_j=(\widehat{\phi}_j-1)/\Delta_j$, and $\widehat{\gamma}=(\widehat{\phi}-1)/\Delta_j$, where $\Delta_j=0.02$ is the percent return deviation between adjacent return states in $\mathcal{S}_{\text{baseline}}$ (or $\Delta_j=0.05$ when using $\mathcal{S}_{\text{alt}}$). We report both SDF-slope estimates ($\{\widehat{\phi}_j\}_j$ and $\widehat{\phi}$) and relative-risk-aversion estimates ($\{\widehat{\gamma}_j\}_j$ and $\widehat{\gamma}$) below; we consider the SDF-slope estimates to be the more directly relevant object of interest given that they do not require the existence of a representative agent with wealth equal to the market index value, but the relative-risk-aversion estimates may be more easily interpretable.

Inference. We must address three issues arising from our framework when constructing confidence intervals for our estimates. First, the moment conditions (21) and (23) are nonlinear in the parameter to be estimated. While this does not in general pose problems for estimation via the generalized method of moments (GMM), in our case it is possible for the estimated values to be arbitrarily large or even nonexistent, as in (22) and (24). Subtracting some multiple of an estimated standard error from a possible point estimate of ∞ is nonsensical (and invalid) in constructing a lower bound for the confidence interval of $\overline{\phi}_j$ or $\overline{\phi}$ with the desired coverage rate, so we cannot pursue this standard GMM approach in our case. Second, there is an unknown dependence structure, both over time and across state pairs, between the members of the set $\{m^*_{i,j}, r^*_{i,j}, \widetilde{\pi}^*_{0_i,i,j}\}_{i,j}$, the observations used to construct our estimates of the parameters of interest. In a standard time-series (or panel-data) setting, we could explicitly account for this dependence structure in our estimation using standard techniques, but in our case the option expiration dates $\{T_i\}$ are not in general evenly spaced. Finally, in constructing our confidence intervals we wish to explicitly take into account that our theoretical results in Section 4 lead to set identification, and more specifically lower bounds for the given parameters of interest, rather than point-identified parameters.

To address these issues, we use a version of a well-known nonparametric bootstrap procedure, the *block bootstrap*, first proposed by Hall (1985) and Carlstein (1986) to account for unknown dependence structures in general settings. See Appendix B.5 for a detailed description of this procedure and a discussion of its asymptotic coverage accuracy. Stated briefly, we resample data from the set $\{m^*_{i,j}, r^*_{i,j}, \widetilde{\pi}^*_{0_i,i,j}\}_{i,j}$ by randomly drawing subsets (or *blocks*) of length D calendar days from the original data and pasting together the resulting observations to obtain a bootstrap dataset.

We then construct a one-sided 95% confidence interval $[CI_{\text{LB},D}, \infty)$ for the parameter of interest — e.g., $\overline{\phi}_j$ — by setting $CI_{\text{LB},D}$ to be the fifth percentile of the distribution of estimates of ϕ_j in the resampled datasets.[34] The set identification implied by our theoretical bounds motivates our use of one-sided rather than two-sided confidence intervals, as the asymptotic coverage rate for our one-sided intervals may in fact be greater than 95% given that we are estimating lower bounds for the parameters of interest; see, for example, Imbens and Manski (2004) and Tamer (2010) for discussions of confidence-set construction in similar partially identified settings.

We experiment with various choices for the block length D in this bootstrap procedure. Our baseline confidence intervals presented in Section 5.6 use block lengths of 45 calendar days, but we also present estimates for block lengths of 90 calendar days, which do not yield materially different estimates. Ultimately, this bandwidth choice reflects a tradeoff between size and power (or size-adjusted power) of the test corresponding to the confidence interval we construct. While there exist results in the relevant literature for the choice of D to minimize the mean squared error of the bootstrap-estimated distribution with respect to the true distribution (e.g., Hall, Horowitz, and Jing, 1995), these results may yield different recommendations than results derived from an optimal-testing standpoint; see Sun, Phillips, and Jin (2008) and Lazarus, Lewis, and Stock (2017) for related discussion. Thus in presenting our results below we simply attempt to give a brief sense of the differences in confidence intervals along the frontier of the size–power tradeoff traced out by different choices of D.

5.5. Accounting for Market Microstructure Noise

As discussed in the context of Proposition 6, we also wish to account for measurement error stemming from possible market microstructure noise in our estimation. That result shows that unlike in the classical errors-in-variables regression model (which leads to attenuation bias), measurement error can in our case lead to an upward bias in the estimated SDF slope required to rationalize the observed variation in risk-neutral beliefs.

Bates (2003) argues that index-option-price measurement error is typically quite small when estimated as part of a parametric factor model. But we nonetheless wish to account for such noise directly for two reasons. First, it is possible that a factor model would attribute some of the observed price variation to a common factor that we would instead prefer to label as idiosyncratic noise for purposes of our testing.[35] Second, even if measurement errors are in fact small in the raw option-price data, these small errors could be magnified when converting to the conditional risk-neutral probabilities we use across adjacent return states.

Rather than parameterizing the source of measurement error — for example, attempting to

[34] This *percentile method* is suboptimal in certain cases, but must be applied in our case given constraints imposed by our data; see Appendix B.5 for further discussion.

[35] For example, changes in net demand pressure in the options market may affect prices, as argued for by Bollen and Whaley (2004) and Gârleanu, Pedersen, and Poteshman (2009). A factor pricing model may interpret movements owing to these changes as stemming from a common factor, but we may wish to include this variation in our notion of noise, as it is not fully correlated with changes in agents' true beliefs over the future index value.

measure it using observed bid-ask spreads — we instead take a nonparametric approach. As in Proposition 6, we postulate that the relation between the observed conditional risk-neutral belief $\widehat{\pi}^*_{t,i,j}$ and its true value $\widetilde{\pi}^*_{t,i,j}$ is $\widehat{\pi}^*_{t,i,j} = \widetilde{\pi}^*_{t,i,j} + \epsilon_{t,i,j}$, where $\epsilon_{t,i,j}$ is uncorrelated with its own lagged values and with risk-neutral beliefs.[36] In order to estimate $\text{Var}(\epsilon_{t,i,j})$, we exploit the restriction that, under the null of RE, the true conditional physical (non-risk-neutral) belief must be an unbiased estimate of the realization of the terminal state; formally, $\widetilde{\pi}_{t,i,j} = \widetilde{\mathbb{E}}_t[\mathbb{1}\{R^m_{T_i} = s_j\}]$. We detail in Appendix B.6 how this can be used to show that the excess risk-neutral belief movement from t to T_i for a fully resolving belief path for which $\widetilde{\pi}^*_{T_i,i,j} \in \{0,1\}$, measured only using beliefs on dates t and T_i (i.e., ignoring the intermediate one-period belief changes), must follow the equality restriction

$$\widetilde{\mathbb{E}}_t[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}] = \frac{\widetilde{\pi}^*_{t,i,j}(1-\widetilde{\pi}^*_{t,i,j})(2\widetilde{\pi}^*_{t,i,j}-1)(\phi_{i,j}-1)}{\widetilde{\pi}^*_{t,i,j}+\phi_{i,j}(1-\widetilde{\pi}^*_{t,i,j})} \tag{25}$$

regardless of the data-generating process, and where $m^*_{t\to T_i,j} \equiv (\widetilde{\pi}^*_{T_i,i,j} - \widetilde{\pi}^*_{t,i,j})^2$ and $r^*_{t\to T_i,j} \equiv (1-\widetilde{\pi}^*_{t,i,j})\widetilde{\pi}^*_{t,i,j} - (1-\widetilde{\pi}^*_{T_i,i,j})\widetilde{\pi}^*_{T_i,i,j}$. By considering only excess risk-neutral belief movement from date t to expiration for a resolving path, $m^*_{t\to T_i,j} - r^*_{t\to T_i,j}$, we avoid the need to consider alternative data-generating processes under which beliefs could follow different paths, since we are using only two points for beliefs rather than the full path. This then yields an exact value for average excess movement as a simple function of $\phi_{i,j}$, the value indexing the bias in risk-neutral beliefs relative to physical beliefs and thus the bias in forecasts of the terminal state using the risk-neutral beliefs.

Our estimate of $\text{Var}(\epsilon_{t,i,j})$ then assumes that any average deviation from the equality restriction in (25) in the observed data arises due to measurement error. We again have that, with measurement error $\epsilon_{t,i,j}$ in risk-neutral beliefs, the observed $t \to T_i$ excess movement is overstated by $2\text{Var}(\epsilon_{t,i,j})$ relative to its true value: $\widetilde{\mathbb{E}}_t[\widehat{m}^*_{t\to T_i,j} - \widehat{r}^*_{t\to T_i,j}] = \widetilde{\mathbb{E}}_t[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}] + 2\text{Var}(\epsilon_{t,i,j})$, as in Proposition 6. Thus, after taking unconditional expectations in (25) and measuring the sample analogues of the objects in that equation for resolving paths, we calculate $\widehat{\text{Var}}(\epsilon_{t,i,j};\overline{\phi}_j)$ as the value equating $\widetilde{\mathbb{E}}[\widehat{m}^*_{t\to T_i,j} - \widehat{r}^*_{t\to T_i,j}] - 2\widehat{\text{Var}}(\epsilon_{t,i,j};\overline{\phi}_j)$ with the right side of (25).[37] We denote the estimate by $\widehat{\text{Var}}(\epsilon_{t,i,j};\overline{\phi}_j)$ because there will be a different noise estimate for each value of $\overline{\phi}_j$ given that the right side of (25) changes with that SDF slope. To distinguish between this set of estimates, we must solve for a fixed point: each possible value $\overline{\phi}_j$ yields a different estimate $\widehat{\text{Var}}(\epsilon_{t,i,j};\overline{\phi}_j)$; following Proposition 6, for each such estimate we subtract $2\widehat{\text{Var}}(\epsilon_{t,i,j};\overline{\phi}_j)$ from each day's observed excess risk-neutral belief movement value $\widehat{m}^*_{t,t+1,i,j} - \widehat{r}^*_{t,t+1,i,j}$ and then re-estimate our period-by-period bound as in (21); this in turn yields an estimate for the SDF slope $\overline{\phi}_j$; and we find the value of $\overline{\phi}_j$ such that this estimate is equal to the value of $\overline{\phi}_j$ used in constructing $\widehat{\text{Var}}(\epsilon_{t,i,j};\overline{\phi}_j)$. This allows us to solve simultaneously for the estimates $\widehat{\phi}_j$ and $\widehat{\text{Var}}(\epsilon_{t,i,j}) = \widehat{\text{Var}}(\epsilon_{t,i,j};\widehat{\phi}_j)$.

The upside of this approach is that, under the null of RE, it yields an unbiased (and thus not

[36] It need not be the case that measurement error in fact exhibits no serial correlation; we need to measure only the component of the measurement error that is uncorrelated with its past values, as any positively correlated component does not increase measured excess risk-neutral belief movement. See after Proposition 6 for related discussion.

[37] There may be heteroskedasticity in these processes, but our main tests use average excess risk-neutral belief movement summed over the entire belief path (see the previous subsection), so for this purpose we need only consider the unconditional variance of the measurement error.

overly conservative) estimate of the measurement-error variance. Further, while the fixed-point estimation procedure appears involved, the intuition behind the approach is straightforward: we are simply finding the measurement-error variance that sets average excess belief movement $\widetilde{\mathbb{E}}[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}]$ equal to the value it would take with correctly calibrated physical beliefs given the estimated SDF slopes. The downside, however, is that under a non-RE alternative, the estimate $\widehat{\text{Var}}(\epsilon_{t,i,j})$ may be biased in either direction: if beliefs are not rational, then they need not be correctly calibrated (so that $\widetilde{\pi}_{t,i,j} \neq \widetilde{\mathbb{E}}_t[\mathbb{1}\{R^m_{T_i} = s_j\}]$), and this would yield incorrect estimates for the noise variance under the procedure above. This of course does not invalidate estimation and inference under the null, but it does mean that the estimated noise-variance values may be uninterpretable from an economic standpoint under a non-RE alternative.

We therefore consider a second approach for estimating the noise variance that is valid regardless of the rationality of expectations, but that is conversely overly conservative in that it will generically overestimate this variance; see Appendix B.6 for details of this alternative approach, which yields economically interpretable estimates regardless of the rationality of expectations. The estimates in this case range from $\widehat{\text{Var}}(\epsilon_{t,i,j}) = 0.005$ to 0.012 for different state pairs indexed by j. The noise values are largest for highly negative return states, and the maximum is attained for conditional beliefs over the pair $(s_j, s_{j+1}) = (-0.07, -0.05)$.

For brevity, we present only noise-corrected estimates using the first approach above, but the second approach yields substantially similar noise-corrected estimates. We note also that the bootstrapped confidence intervals for the noise-corrected estimates incorporate noise-related sampling uncertainty by re-estimating noise as above in each resampled draw of the data.

5.6. Main Results

Table 1 below presents our main results. Panel (a) presents lower-bound estimates for the SDF slopes $\{\overline{\phi}_j\}_j$ (for individual return-state pairs) and $\overline{\phi}$ (overall across all state pairs, excluding the extreme state pairs), constructed using equations (21)–(24). Panel (b) presents the corresponding lower-bound estimates for the relative risk aversion values for a fictitious representative agent, constructed using the result in Proposition 5. All estimates for the individual state pairs reflect variation in risk-neutral probabilities over the excess return being in state s_j versus s_{j+1}, and these estimates ϕ_j and γ_j are presented as corresponding to the midpoint between the two return states, $(s_j + s_{j+1})/2$. For brevity and ease of interpretation, the discussion here considers only the risk-aversion estimates in panel (b).

We begin with the baseline estimates. The individual state-pair estimates in the first row of this panel are all $\widehat{\gamma}_j = \infty$, and the lower bounds of their one-sided 95 percent confidence intervals $[\text{CI}_{\text{LB},D}, \infty)$, where D is the block length in calendar days for the block bootstrap described above, are also $\text{CI}_{\text{LB},D} = \infty$ for $D = 45$ and $D = 90$ in all cases aside from the first state pair, for which those confidence-interval lower bounds are large but finite. The overall estimate at the end of this row is also $\widehat{\gamma} = \infty$, and its confidence intervals allow us to reject a finite value for $\overline{\gamma}$ at the 95 percent level. (It is in fact the case that *none* of the 5,000 bootstrap runs produce a finite overall

Table 1: Main Estimation Results

(a) Lower Bound for SDF Slope

	Individual State Pairs: $\overline{\phi}_j$ by Excess Return											Overall: $\overline{\phi}$
	-10%	-8%	-6%	-4%	-2%	0	2%	4%	6%	8%	10%	
Baseline:	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
($CI_{LB,45}$)	(120)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)
($CI_{LB,90}$)	(90.0)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)
Noise-Adj.:	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
($CI_{LB,45}$)	(86.3)	(∞)	(∞)	(∞)	(∞)	(1.6)	(1.6)	(2.6)	(8.9)	(31.0)	(∞)	(3.0)
($CI_{LB,90}$)	(82.0)	(∞)	(∞)	(∞)	(∞)	(1.6)	(1.6)	(2.6)	(8.7)	(31.5)	(∞)	(2.7)

(b) Lower Bound for Relative Risk Aversion

	Individual State Pairs: $\overline{\gamma}_j$ by Excess Return											Overall: $\overline{\gamma}$
	-10%	-8%	-6%	-4%	-2%	0	2%	4%	6%	8%	10%	
Baseline:	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
($CI_{LB,45}$)	(5,950)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)
($CI_{LB,90}$)	(4,450)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)	(∞)
Noise-Adj.:	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞	∞
($CI_{LB,45}$)	(4,263)	(∞)	(∞)	(∞)	(∞)	(32.0)	(29.0)	(82.0)	(395)	(1,500)	(∞)	(102)
($CI_{LB,90}$)	(4,050)	(∞)	(∞)	(∞)	(∞)	(31.5)	(29.0)	(81.0)	(385)	(1,525)	(∞)	(85.0)

Notes: Estimates are constructed using the bound in Proposition 4 for panel (a), converted to relative risk aversion values using Proposition 5 for panel (b). Baseline estimates use equations (21)–(22) for individual state pairs, where each estimate is presented as corresponding to the excess return at the midpoint of those two states, $(s_j + s_{j+1})/2$, where s_j and s_{j+1} are in the state space $\mathcal{S}_{\text{baseline}} = \exp(\{(-\infty, -0.11), -0.09, -0.07, \ldots, 0.07, 0.09, (0.11, \infty)\})$. Overall estimates use data across all state pairs aside from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J), and the baseline estimates for these values use equations (23)–(24). One-sided 95% confidence intervals using the block bootstrap with 5,000 draws and block length of D calendar days are $[\text{CI}_{\text{LB},D}, \infty)$; we present only the lower bounds of these confidence intervals, $\text{CI}_{\text{LB},D}$, in parentheses below the point estimates, for $D = 45, 90$. Noise-adjusted values estimate $\text{Var}(\epsilon_{t,i,j})$ so that fully resolving paths have implied physical beliefs that are ex-ante correctly calibrated on average (see Appendix B.6), then follow Proposition 6 and subtract two times that estimate from each day's observed excess risk-neutral belief movement value before estimating the respective bounds.

risk-aversion estimate for either $D = 45$ or $D = 90$.) These baseline estimates thus indicate that no amount of risk aversion is capable of rationalizing the observed excess movement in risk-neutral beliefs aggregated across all available return states.

Moving now to the noise-corrected estimates, we obtain identical infinite point estimates, but the lower bounds of the confidence intervals are now finite in many cases. We can reject an overall risk-aversion value below $\overline{\gamma} = 102$ for $D = 45$ in this case, indicating that the risk aversion required to rationalize the data is still very high. We observe further in the noise-corrected esti-

mates that the higher return states are estimated to have lower required risk aversion. This may be a somewhat surprising result in light of the finding that measurement error is greater for the low-return state pairs as discussed in the previous subsection, but excess belief movement is so much larger for those low-return states that this noise differential is rendered irrelevant (and in fact reversed) in the risk-aversion estimates. Finally, the confidence-interval lower bounds are generally quite similar for the 45- versus 90-day block-length specifications, indicating that serial correlation in the excess-belief-movement series does not pose serious issues for inference at least beyond the 45-day horizon.

In light of Proposition 3, we conclude that belief revisions are excessively volatile in all cases for which the data cannot be rationalized with finite risk aversion, as these findings cannot in general be produced solely by miscalibrated priors. Further, the large local risk-aversion bounds at every point of the return distribution and the extremely large overall estimates (which use only beliefs over the excess-return states between -10% and +10%) imply that no feature of the true underlying data-generating process (e.g., volatility in the left tail of the return distribution) can by itself be responsible for these findings of excess belief volatility.

6. Interpretation and Robustness of Empirical Results

We now consider a set of statistical decompositions and additional tests to examine the features of the observed data driving the main results presented in the previous section, as well as the robustness of these results to possible misspecification.

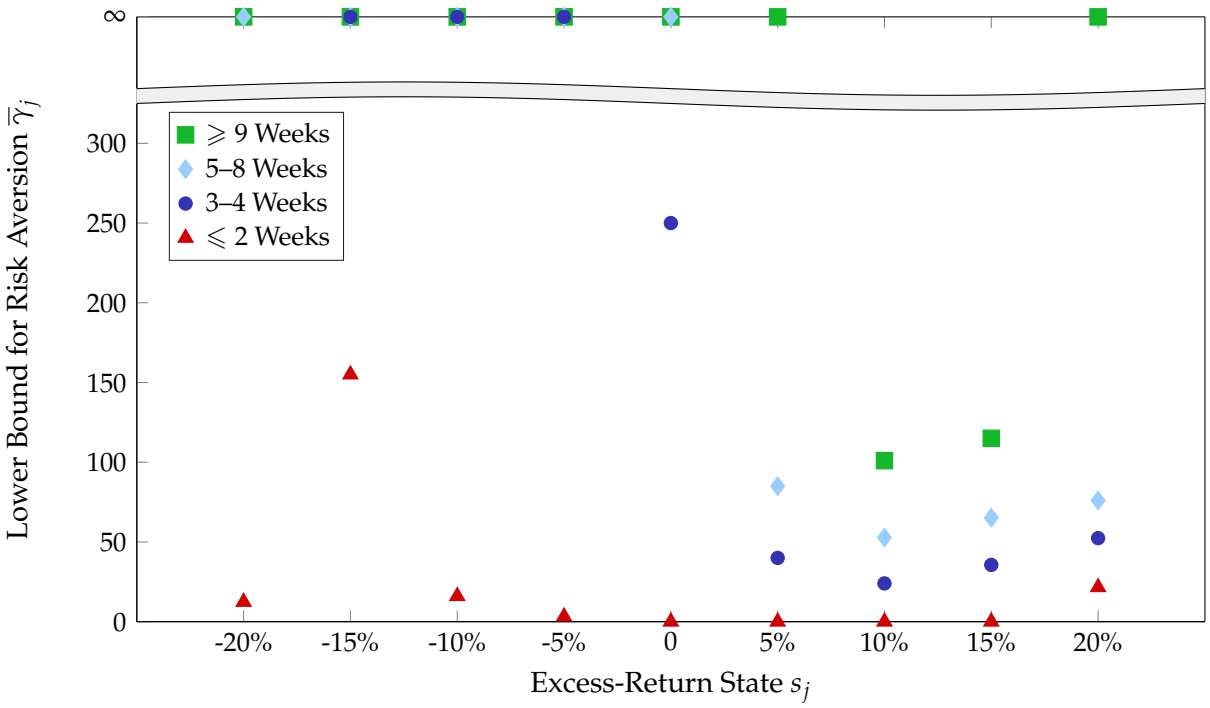
6.1. At What Forecast Horizon Are Beliefs Excessively Volatile?

Figure 3, presented above in Section 5.3, suggests that excess movement in risk-neutral beliefs is concentrated at distant horizons from a given option expiration date. We now consider whether this holds formally when re-estimating our bounds separately for subsets of the risk-neutral belief data split by time to expiration. Figure 5 presents the results from this decomposition. For brevity and interpretation, we consider only estimates of the lower bound of relative risk aversion $\overline{\gamma}_j$ (rather than the SDF slope $\overline{\phi}_j$). We also use the less granular state space $\mathcal{S}_{\text{alt}} = \exp(\{(-\infty, -0.225), -0.175, -0.125, \ldots, 0.125, 0.175, (0.225, \infty)\})$, which increases the number of available observations for the short-horizon subsets. Note that there is a break in the y-axis above the value 300, demarcated with the narrow gray band, to accommodate (or attempt to accommodate) that our longer-horizon estimates are, as above, often $\widehat{\gamma}_j = \infty$.

Beginning at the bottom of the figure, the red triangles present our risk-aversion estimates using only beliefs data for the last two weeks of trading before a given option expiration date T_i. As suggested by Figure 3, the observed variation at these short horizons can be rationalized with reasonable values for risk aversion, and in many cases we cannot rule out risk neutrality.[38] When

[38] This does not necessarily suggest that pricing is in fact risk-neutral, as our estimates are again only lower bounds.

Figure 5: Estimates of Relative Risk Aversion: Splits by Time to Expiration



Notes: Estimates are constructed using the bound in Proposition 4, converted to relative risk aversion values using Proposition 5. Within each time-to-expiration subset, each point shows estimate for state pair (s_j, s_{j+1}) plotted at the excess return at the midpoint of those two states, $(s_j + s_{j+1})/2$, where s_j and s_{j+1} are in the state space $\mathcal{S}_{\text{alt}} = \exp(\{(-\infty, -0.225), -0.175, -0.125, \ldots, 0.125, 0.175, (0.225, \infty)\})$. Each estimate in the $\leqslant$ 2 weeks series uses risk-neutral belief movement and uncertainty resolution observations from $t = T_i - 10$ to $t = T_i$, or $m^*_{T_i-10,T_i,i,j}$ and $r^*_{T_i-10,T_i,i,j}$, respectively, and similarly for the remainder of the series. Gray band indicates break in the y-axis. Estimates aggregated across all state pairs (excluding the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J)) for each series: $\widehat{\gamma} = 0$ for $\leqslant$ 2 weeks (95% confidence interval: $[0, \infty)$); $\widehat{\gamma} = 250$ for 3–4 weeks (CI: $[170, \infty)$); $\widehat{\gamma} = \infty$ for 5–8 weeks (CI: $[3980, \infty)$); $\widehat{\gamma} = \infty$ for $\geqslant$ 9 weeks (CI: ∞). CIs calculated using block bootstrap with blocks of 45 days, 5,000 draws.

pooling data across the interior return-state pairs to estimate the average required risk-aversion value across both dates and states (not shown), we obtain an estimate $\widehat{\gamma} = 0$.

As we move up the chart, each successive series of risk-aversion estimates for more-distant horizons from expiration is weakly greater than the preceding series at every point in the return distribution. The estimates using data for trading dates between three and four weeks from expiration (blue circles) are large or infinite for negative excess-return states, and around $\widehat{\gamma}_j = 50$ for positive states; using data from the second-to-last month from expiration (light blue diamonds) yields slightly greater estimates exhibiting a similar pattern; and using data from trading dates more than two months from expiration (green squares) yields $\widehat{\gamma}_j = \infty$ for almost all return states. The pooled estimates are $\widehat{\gamma} = 250$ for the three-to-four-week data (with one-sided 95 percent confidence interval lower bound of 170), and $\widehat{\gamma} = \infty$ for the last two sets of data (with confidence interval lower bounds of 3,980 and ∞, respectively). Each successive pooled estimate is also significantly greater than the previous value at the 95 percent level, with the exception of the last two

estimates (5–8 weeks and $\geqslant$ 9 weeks), which are statistically indistinguishable from one another.

We conclude that beliefs about events in the somewhat-distant future appear to react too strongly to new information, as they are predictably mean-reverting in a manner that induces excess volatility. The same is not true of beliefs over near-term events, which behave as would be expected under RE given reasonable risk-aversion values. The results presented in Section 5 must accordingly be driven largely by belief movement at distant horizons. We note further that the results in Figure 5 do not account for possible measurement error in risk-neutral beliefs. Thus any potential explanation of our baseline results relying on such noise — or underestimation of such noise, in the case of our noise-corrected estimates — must account for the fact that this issue seems not to be present at short belief horizons, when beliefs are particularly volatile.

6.2. At What Frequency Do Beliefs Mean-Revert?

While we use daily data for our baseline estimation, the results presented to this point do not necessarily indicate the presence of strong day-to-day mean-reversion in risk-neutral beliefs; our test is in fact agnostic with respect to the frequency of such reversion in the data, as all that is required to generate positive values of $\widetilde{\mathbb{E}}[m^*_{i,j} - r^*_{i,j}]$ is a tendency for risk-neutral beliefs to revert toward 0.5 on average over the entire series.[39]

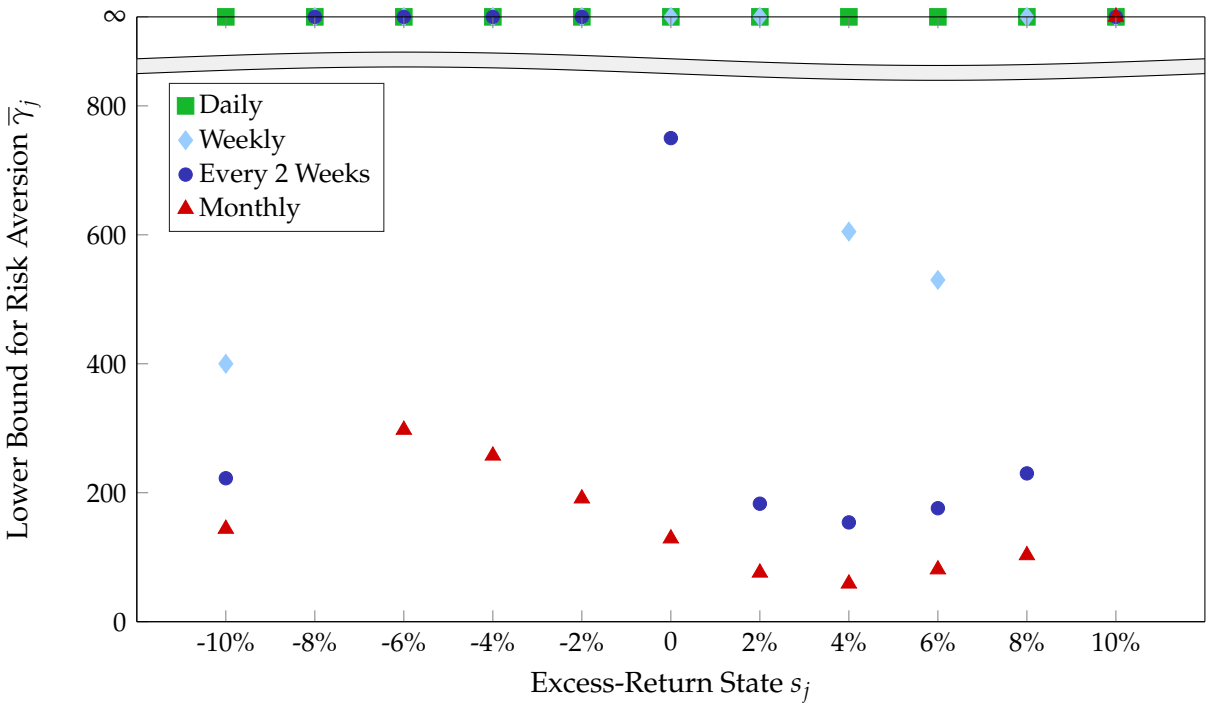
To further examine the belief processes driving our findings, we accordingly re-estimate our bounds using risk-neutral beliefs sampled at different frequencies. In addition to the benchmark daily frequency, we sample beliefs at weekly, biweekly, and monthly frequencies,[40] and then construct associated risk-neutral belief movement and uncertainty resolution statistics in each case; for example, one-week movement realized at trading date t is calculated as $(\widetilde{\pi}^*_{t,i,j} - \widetilde{\pi}^*_{t-5,i,j})^2$. The estimated bounds in all four cases are presented in Figure 6 in a manner similar to that used in Figure 5. (Unlike in that previous figure, however, we once again use the return space $\mathcal{S}_{\text{baseline}} = \exp(\{(-\infty, -0.11), -0.09, -0.07, \ldots, 0.07, 0.09, (0.11, \infty)\})$, as we no longer have any series composed entirely of data for dates close to expiration.) Following usual practice in the related literature, we use Wednesday data for all three non-daily datasets.

Beginning now at the top of the chart and proceeding downward, we again observe a clear monotonic pattern: as we decrease the sampling frequency, the data can be rationalized with weakly decreasing values for relative risk aversion at every point in the return space. This is reflected as well in the pooled estimates across both dates and (interior) return states, which are given by $\widehat{\gamma} = \infty$, ∞, 425, and 123 for daily, weekly, biweekly, and monthly data, respectively,

[39] There is some evidence, for example, that index returns exhibit short-term positive autocorrelation (or *momentum*) but long-term mean-reversion (which may be thought of as a *value* effect in the sense of Fama and French, 1993); see, e.g., Poterba and Summers (1988). But short-term momentum for index returns seems to have declined or disappeared in recent decades; see Froot and Perold (1995), Campbell (2017). Further, it is possible that the underlying index *returns* exhibit momentum at some horizon while changes in the distribution of *beliefs* over future index values in fact exhibit mean-reversion. Separately, the fact that our test is agnostic to the frequency at which beliefs mean-revert distinguishes it from the type of variance-ratio test used by, e.g., Lo and MacKinlay (1988), which requires daily data to test daily mean-reversion, weekly data to test weekly mean-reversion, and so on.

[40] We lack sufficient data to construct quarterly series, as we use only options with less than six months to expiration.

Figure 6: Estimates of Relative Risk Aversion: Splits by Sampling Frequency



Notes: Estimates are constructed using the bound in Proposition 4, converted to relative risk aversion values using Proposition 5. Within each sampling-frequency subset, each point shows estimate for state pair (s_j, s_{j+1}) plotted at the excess return at the midpoint of those two states, $(s_j + s_{j+1})/2$, where s_j and s_{j+1} are in the state space $\mathcal{S}_{\text{baseline}} = \exp(\{(-\infty, -0.11), -0.09, -0.07, \ldots, 0.07, 0.09, (0.11, \infty)\})$. Each estimate in the monthly series uses risk-neutral beliefs $\widetilde{\pi}^*_{t,i,j}$ sampled only on the second Wednesday of each month; the biweekly series samples beliefs on the Wednesdays of the evenly-numbered weeks of the year; the weekly series samples beliefs every Wednesday; and the daily series every trading day. Gray band indicates break in the y-axis. Estimates aggregated across all state pairs (excluding the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J)) for each series: $\widehat{\gamma} = 123$ for monthly (95% confidence interval: $[97, \infty)$); $\widehat{\gamma} = 425$ for biweekly (CI: $[228, \infty)$); $\widehat{\gamma} = \infty$ for weekly (CI: $[4950, \infty)$); $\widehat{\gamma} = \infty$ for daily (CI: ∞). CIs calculated using block bootstrap with blocks of 45 days, 5,000 draws.

which are all mutually statistically different at the 95% level aside from the daily versus weekly estimates. The lower bound of the confidence interval for monthly data is 97. Separately, the return-state-specific monthly estimates also exhibit a roughly decreasing pattern across possible return values.

We thus find that the monthly belief variation can be explained with finite but nonetheless still quite large risk-aversion values. Further, this monthly variation masks additional volatility and required risk aversion at higher sampling frequencies, and so any RE model capable of matching the moments of the risk-neutral belief process at a monthly horizon would seem to possess what might be thought of as incorrect statistical microfoundations.

There are, however, two additional possibilities that may account for the differences between the estimates for the more- versus less-frequently-sampled series. First, it is possible that the greater measured belief volatility at higher sampling frequencies is a result of measurement error, as Figure 6 again does not account for such noise. But in this case, the fact that required

risk aversion is still quite large at a monthly sampling frequency is further indication that such noise does not drive all of our results. Second, we note that the statistical power of our test under non-RE alternatives generally increases with the number of observations in a given belief series $\{\widetilde{\pi}^*_{t,i,j}\}_t$: in the presence of excess belief volatility at the daily frequency, the sum $m^*_{i,j} = \sum_{t=0_i+1}^{T_i}(\widetilde{\pi}^*_{t,i,j} - \widetilde{\pi}^*_{t-1,i,j})^2$ will increase in expectation as we observe more data points between 0_i and T_i, so statistical power for our test will generally be greater for the daily sampling frequency than for the monthly sampling frequency. (The daily data yield an average of 42.7 observations per belief series, while the monthly data give us only 2.9 observations per series across series with at least two observations.) The degree to which this is driving the differences in the estimates in Figure 6 requires knowledge of the data-generating process for beliefs, which may be useful to consider in future work.

6.3. What Macro Statistics Are Correlated with Excess Belief Movement?

We now consider reduced-form evidence on the macroeconomic correlates of excess risk-neutral belief movement. Table 2 presents a set of time-series regressions to this end. The dependent variable in each case is the average one-day value of excess belief movement $m^*_{t-1,1,i,j} - r^*_{t-1,t,i,j}$ by quarter, as plotted above in Figure 4 (so that the average is calculated across all available expiration dates and state pairs, aside from the extreme state pairs, using all trading dates within a quarter).[41] We aggregate to the quarterly level given the frequency of data available for the regressors we consider, and we use quarterly averages of these independent variables when data is available at a higher frequency. Aside from the constant and time trend, all variables (both dependent and independent) are normalized to have unit standard deviation for purposes of interpretation, and we present heteroskedasticity- and autocorrelation-robust t-statistics and p-values using the equal-weighted periodogram estimator of the long-run variance; see Lazarus, Lewis, and Stock (2017) for results on the optimality properties of this estimator.

Moving from left to right across the set of regressions considered, the first column considers the comovement of excess belief movement with commonly used measures of liquidity and limits to arbitrage in asset markets. As a simple proxy for option-market liquidity, we use the quarterly average bid-ask spread across all S&P 500 options in our available OptionMetrics sample, where the average is weighted by the trading volume of each option.[42] For our measure of limits to arbitrage, we follow recent literature in using seasonally adjusted quarterly changes in broker-dealer leverage — measured using the flow of funds accounts published by the Federal Reserve — to proxy for financial constraints faced by arbitrageurs.[43] The coefficients on both regressors are both eco-

[41] In this case (as in Figures 3 and 4), in order to obtain less-noisy estimates of quarterly-average excess belief movement, we use beliefs for all state pairs for each T_i rather than conditioning on $R^m_{T_i} \in \{s_j, s_{j+1}\}$.

[42] This follows, among others, Amihud and Mendelson (1986) and Chordia, Roll, and Subrahmanyam (2008), but we might also wish to consider more direct proxies for the return impact of a transaction, as used, e.g., by Pástor and Stambaugh (2003). Pástor and Stambaugh also discuss why trading volume is inappropriate to use as a measure of liquidity, though in unreported results we find that it is also uncorrelated with excess belief movement.

[43] This measure is proposed and examined by Adrian, Etula, and Muir (2014), Adrian and Shin (2014), and Cho (2018), following theory developed by Shleifer and Vishny (1997) and Gromb and Vayanos (2002).

Table 2: Regressions for Quarterly Average of Excess Belief Movement

	(1)	(2)	(3)	(4)
Liquidity and Limits to Arbitrage				
Bid-Ask Spread	0.2	-0.2	-0.3*	-0.1
	(1.3)	(-0.7)	(-0.4)	(-1.0)
Broker-Dealer Leverage	-0.1	0.1	-0.0	-0.1
	(-0.4)	(0.8)	(-0.7)	(-1.7)
Volatility and Uncertainty				
VIX		0.8**	0.9***	0.6*
		(2.3)	(3.4)	(2.1)
Baker–Bloom–Davis Uncertainty		-0.3	0.1	0.2*
		(-1.2)	(1.5)	(2.2)
Returns and Valuation				
12-Month S&P Return			0.3**	0.3**
			(2.8)	(2.6)
Price to 10-Year Earnings Ratio			0.6***	0.5***
			(4.1)	(4.1)
Time Trend				-0.0*
				(-2.2)
R^2	0.07	0.34	0.72	0.73
N	79	79	79	79

Notes: *** $p < 0.01$; ** $p < 0.05$; * $p < 0.1$. Dependent variable in all regressions is the series in Figure 4, the empirical average $\widehat{\mathbb{E}}[m^*_{t-1,t,i,j} - r^*_{t-1,t,i,j}]$ calculated across all available expiration dates T_i and state pairs (s_j, s_{j+1}) in (20), aside from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J), using all trading dates t within each given quarter. Regressors are correspondingly quarterly averages of each relevant series. All variables (dependent and independent, aside from time trend) are normalized to have unit standard deviation. Constant is included in each regression. Heteroskedasticity- and autocorrelation-robust t-statistics are in parentheses, calculated using the equal-weighted periodogram estimator for the long-run variance with 8 degrees of freedom, as in Lazarus, Lewis, and Stock (2017), and critical values for the p-values are accordingly from the t distribution with 8 degrees of freedom.

nomically and statistically small. This holds true as well in the other regressions with additional regressors. While these results are reduced-form, this nonetheless provides additional suggestive evidence that option-market-specific factors (or related mismeasurement of risk-neutral beliefs) are not the main drivers of our results.

The next column adds measures related to volatility and uncertainty to the regressions. As a proxy for implied volatility, we use the quarterly average of the CBOE volatility index, or VIX.[44] We also consider the Baker, Bloom, and Davis (2016) measure of economic policy uncertainty. The VIX generally has a strong positive relationship with quarterly excess belief movement, while the

[44]While widely used, the VIX in fact measures the risk-neutral *entropy* rather than variance of the distribution of returns; see Martin (2017). Martin shows that a measure of risk-neutral variance, which he terms SVIX, generally tracks the VIX closely, so we use the VIX for simplicity.

uncertainty index has an insignificant relationship aside from column (4). The contemporaneous positive relationship of VIX with belief movement is unsurprising, as this would be expected both under the null of RE and under a non-RE alternative. Under RE, Proposition 1 and Figure 2 show that excess belief movement is maximized in expectation for interior values of the conditional prior $\widetilde{\pi}^*_{0_i,i,j}$, and higher risk-neutral volatility leads to a more-dispersed distribution and therefore conditional risk-neutral beliefs closer to this maximand across the entire belief distribution. Under a non-RE alternative, it is intuitive that greater implied volatility would correspond with greater excess variation in beliefs. Thus these results provide further evidence that our measure of belief movement is in fact reflective of "true" excess volatility.

The third column then considers statistics related to index returns and valuation. We consider both the rolling 12-month S&P 500 return and the index's price to 10-year earnings ratio, also referred to as the cyclically adjusted price-earnings ratio (Campbell and Shiller, 1988; Shiller, 2000); data for the latter is obtained via Robert Shiller's website. We consider 12-month returns to account for possible extrapolation-related excess volatility,[45] and cyclically adjusted price-earnings for valuation-related return predictability.[46] Both measures are significantly positively related to excess movement. But one anomaly is that these positive relationships are with respect to the signed (rather than absolute) values of the regressors, whereas extrapolation-based explanations of excess volatility should in theory be symmetric with respect to over- versus undervaluation. One possibility is that the inclusion of the VIX accounts for such extrapolation on the downside (i.e., when VIX is high and past returns are low), and indeed excluding the VIX renders the coefficients on the absolute valuation measures positive (also unreported).

Finally, column four includes a time trend, which is small but negative, as suggested by Figure 4. The R^2 value for this regression is 0.73, indicating that the statistics we consider account for much of the variation in excess belief movement at the quarterly frequency.

6.4. Simulation Evidence: Are Results Robust to Violations of CTI?

Finally, we conduct a set of numerical simulations to consider the robustness of our empirical results to violations of the assumption of conditional transition independence. As noted in Section 3.2, the habit-formation model of Campbell and Cochrane (1999) violates CTI, and we accordingly consider this model for a first pass at understanding the possible implications of such a violation. The theoretical model is as presented in Appendix B.3, and we adopt the calibration used by Campbell and Cochrane (1999) in the version of their benchmark model with imperfect correlation between consumption and market-dividend growth (where we convert their monthly-frequency parameters to their equivalent daily values). Appendix B.7 contains details of this calibration and our solution procedure and simulations.

Our simulations address two questions. First, does the violation of CTI in the Campbell–Cochrane model generate enough variation in the model-implied risk-neutral beliefs over market

[45] This measure is used, for example, by Greenwood and Shleifer (2014) in related analysis.
[46] Again see, e.g., Campbell and Shiller (1988).

returns to yield excess movement of the magnitude observed in the data? If so, then this would indicate that our empirical results may plausibly be driven by changes in expectations of marginal utility across return states rather than excess variation in the subjective beliefs themselves. Second, when we apply our theoretical bounds to the simulated data, do they provide a valid lower bound on the slope of the SDF across return states despite the model's violation of CTI? This provides evidence on the robustness of the bounds themselves to such a violation.

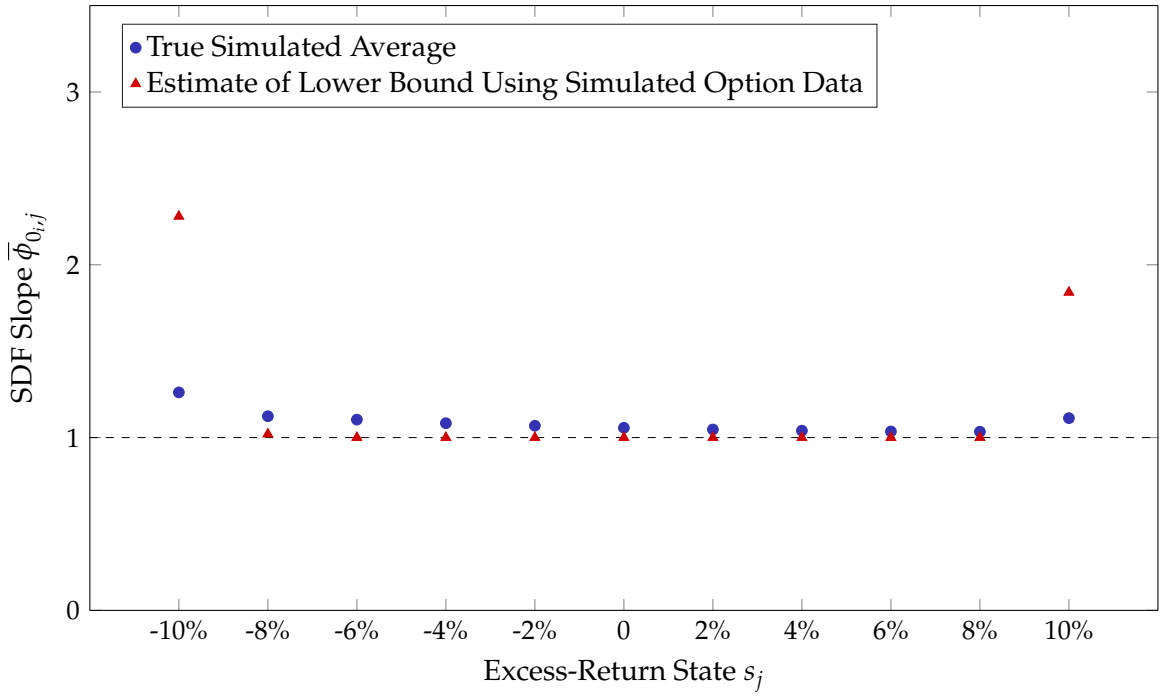
Conducting these simulations requires calculating subjective beliefs for returns over many future horizons as a function of the surplus-consumption state, as well as the set of expected SDF slopes $\{\phi_{t,i,j}\}_{t,i,j}$ in order to translate the subjective beliefs to risk-neutral beliefs. The SDF slope is again given by $\phi_{t,i,j} = \mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]/\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j+1}]$, so calculating this value in turn requires solving for the joint distribution over date-T_i realizations of the SDF and the market return. This is an infinite-dimensional object, so for dimension reduction we iterate backwards using a projection-based approach.[47] We consider 90-day option-expiration horizons (i.e., $T_i - 0_i = 90$), and after solving the model for the price-dividend ratio, we then solve for the joint distribution for returns (from t to T_i) and the SDF at every point in a gridded state space as of $t = T_i - 1$, then $t = T_i - 2$, and so on; see Appendix B.7 for further details. As an example of the output of this solution procedure, Figure B.1 in that appendix shows the joint CDF for the market return from 0_i to T_i and the date-T_i SDF, evaluated as of $t = 0_i$ and with surplus consumption at its steady-state value, $S^c_{0_i} = \overline{S^c}$.

We then conduct 25,000 simulations of the model, where each simulation runs from 0_i to T_i, and for which the initial surplus-consumption state is drawn from its unconditional distribution. For each period in each simulation, we evaluate risk-neutral beliefs over return states at every point in the space $\mathcal{S}_{\text{baseline}}$ as used above and use these to calculate the set of conditional risk-neutral beliefs $\{\widetilde{\pi}^*_{t,i,j}\}_j$. Further, we store the associated set of expected SDF slopes $\{\phi_{t,i,j}\}_j$. We can thus calculate the true average values of these objects of interest, $\overline{\phi}_{0_i,j} \equiv \widehat{\mathbb{E}}[\phi_{0_i,i,j}]$, where $\widehat{\mathbb{E}}[\cdot]$ denotes the expectation over all simulations i and we have fixed the state pair j. And using the risk-neutral beliefs series, we can naïvely apply our theoretical bound in Proposition 4 to obtain lower-bound estimates for those SDF slopes and compare those estimates to the true simulated values. Relative risk aversion for this model's representative agent does not match the definition used in Proposition 5, as this agent's utility does not depend only on terminal wealth (see Campbell and Cochrane, 1999, Section IV.B), so we accordingly present estimates for the SDF slope rather than for relative risk aversion.

Figure 7 presents these simulation results. The blue circles show the true simulated average values of the SDF slopes $\overline{\phi}_{0_i,j}$, while the red triangles show the naïve lower-bound estimates of these values using our theoretical bound on the simulated risk-neutral beliefs data. Considering the first question posed at the outset of this subsection, it is clear in both cases that these SDF-slope values are far below those obtained from our empirical estimates above, so the model does

[47] See Judd (1992), or see Algan, Allais, Den Haan, and Rendahl (2014) for a recent survey, though neither considers solutions for belief distributions.

Figure 7: Estimates of SDF Slope in Habit-Formation Model Simulations



Notes: Estimates are from 25,000 simulations of 90-day periods ($T_i - 0_i = 90$), with initial surplus-consumption ratio $S^c_{0_i}$ drawn from its unconditional distribution. True simulated average value $\overline{\phi}_{0_i,j}$ is equal to average of date-0_i values $\phi_{0_i,i,j} = \mathbb{E}_{0_i}[M_{T_i}/M_{0_i} \mid R^m_{T_i} = s_j]/\mathbb{E}_{0_i}[M_{T_i}/M_{0_i} \mid R^m_{T_i} = s_{j+1}]$ for state pair (s_j, s_{j+1}) across all simulations ($i = 1, 2, \ldots,$ 25,000), where these expectations are evaluated using the solution for the joint CDF of the SDF and the return distribution; see Appendix B.7 for details. The estimate of the lower bound for $\overline{\phi}_{0_i,j}$ is constructed naïvely using the theoretical bound in Proposition 4, using risk-neutral belief movement $m^*_{i,j}$ and uncertainty resolution $r^*_{i,j}$ values across simulations constructed via simulated risk-neutral beliefs $\widetilde{\pi}^*_{t,i,j}$. Each point shows estimate for state pair (s_j, s_{j+1}) plotted at the excess return at the midpoint of those two states, $(s_j + s_{j+1})/2$, where s_j and s_{j+1} are in the state space $\mathcal{S}_{\text{baseline}} = \exp(\{(-\infty, -0.11), -0.09, -0.07, \ldots, 0.07, 0.09, (0.11, \infty)\})$. Aside from the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J), for which the violation of CTI is particularly severe, the naïve estimates from our theoretical bounds are still conservative for the true parameters of interest despite the violation of CTI in this model. Further, the true simulated averages are far below our empirical lower-bound estimates in Section 5.6.

not replicate the observed variation in risk-neutral beliefs even with the violation of CTI.[48]

Also evident in the figure is that, aside from the estimates for the extreme state pairs (s_1, s_2) and (s_{J-1}, s_J), the theoretical bounds yield estimates below the true averages in all cases despite the violation of CTI, and these lower bounds are generally fairly tight. (Following the discussion in Section 5.2, it is unsurprising that the violation of CTI is severe enough for the extreme states that the bounds are no longer valid in these two cases.)

We can understand the validity of the theoretical bound for the interior states by way of Proposition 7, which shows that the bounds hold approximately for violations of CTI for which the $\phi_{t,i,j}$ process is close to a martingale. In our simulations, the values $|\widehat{\mathbb{E}}[\phi_{t+1,i,j} - \phi_{t,i,j}]|$ for different state

[48] It is also the case that these SDF slopes imply lower risk-aversion values than might be expected for the Campbell–Cochrane model, but this arises due to the fact that these are slopes over the index value (and not, e.g., the price-dividend ratio), and that this index does not correspond to the wealth portfolio over the agent's consumption stream.

pairs j range from a minimum of 0.00002 to a maximum of 0.00011, which is not large enough to invalidate the theoretical bounds. (These are figures for the interior state pairs; the figures for the extreme state pairs are more than twice as large.)

Thus the analysis above is robust to the violation of the assumption of CTI considered in these simulations. This is of course only a single illustrative example of a possible violation of CTI, so further work is needed to address other possible violations.[49]

7. Discussion: Theoretical Channels

While our empirics provide evidence on the restrictiveness of a set of benchmark models, we must also consider the possible theoretical underpinnings of our results from a positive standpoint. Violations of our theoretical bounds — or, more generally, findings of large required risk-aversion values for these bounds to be met — arise when beliefs have a tendency to mean-revert. Thus the most straightforward theoretical interpretation of our results points in favor of models in which agents have a tendency to overreact to new information relative to their prior beliefs, especially over events at distant horizons.

One possible model providing a foundation for such relative overreaction is that of Benjamin, Bodoh-Creed, and Rabin (2017), who formalize the finding of *base-rate neglect* in a body of previous literature (Kahneman and Tversky, 1973; Bar-Hillel, 1980). Their model of base-rate neglect can be described as a departure from the Bayesian-updating requirement of RE in Definition 1 above as follows:

$$\pi_t(R_T^m = s) = \frac{\pi_{t-1}(R_T^m = s)^\alpha \ \mathbb{P}(\theta_t | \ \mathcal{F}_{t-1}, R_T^m = s)}{\sum_{s' \in \mathcal{S}} \pi_{t-1}(R_T^m = s')^\alpha \ \mathbb{P}(\theta_t \,|\, \mathcal{F}_{t-1}, R_T^m = s')}, \tag{26}$$

where $\alpha \in [0, 1)$ and all other notation is as in Section 3. That is, an agent updating under base-rate neglect puts too little weight on her prior beliefs relative to the Bayesian benchmark in which $\alpha = 1$, which induces excess movement in beliefs relative to that benchmark in response to new information. For example, with $\alpha = 0$, the agent acts as if she has a flat prior, $\pi_{t-1}(R_T^m = s_j) = \pi_{t-1}(R_T^m = s_k)$ for all s_j, s_k, and t; in this case, her period-t beliefs will simply reflect the period-t signal received, and her conditional beliefs $\widetilde{\pi}_{t,j} = \pi_t(R_T^m = s_j \mid R_T^m \in \{s_j, s_{j+1}\})$ will oscillate above and below 0.5 as she receives signals favorable to state s_j and s_{j+1}, respectively, inducing excess movement in conditional risk-neutral beliefs as well.

Conversely, correct weighting of the prior but overweighting of the signal likelihood can also provide a foundation for excess belief movement arising from relative overreaction. As in Augenblick and Rabin (2018), one manner in which this can be achieved is by setting $\alpha = 1$ in (26) and then specifying that the second term in both the numerator and the sum in the denominator, $\mathbb{P}(\theta_t | \ \mathcal{F}_{t-1}, R_T^m = s)$, is taken to the power $\beta > 1$. This is a semi-reduced-form way of modeling the "law of small numbers" (Tversky and Kahneman, 1971) for small samples modeled more fully

[49] For example, we are currently considering violations arising from certain forms of stochastic volatility shocks (see Footnote 23), which some literature argues are important for matching features of the option-price data; see, e.g., Chabi-Yo (2012) and Christoffersen, Heston, and Jacobs (2013), though see also Jones (2006) for an alternative perspective.

by Rabin (2002) and Rabin and Vayanos (2010).

Kahneman and Tversky (1972) argue that the law of small numbers arises due to what they term the *representativeness heuristic*, under which agents assess events' probabilities according to the degree to which those events "(i) [are] similar in essential properties to [their] parent population; and (ii) reflect the salient features of the process by which [they are] generated" (p. 431). Some models capable of generating excess volatility accordingly proceed directly from this characterization of people's statistical biases: Barberis, Shleifer, and Vishny (1998), Barberis, Greenwood, Jin, and Shleifer (2015), Hirshleifer, Li, and Yu (2015), and Bordalo, Gennaioli, La Porta, and Shleifer (2017) consider models of investor extrapolation arising from local representativeness and provide empirical support to this end; similarly, Fuster, Laibson, and Mendel (2010) consider a model of *natural expectations*, in which agents form beliefs from simple models fit to the available data.[50] Giglio and Kelly (2018) find that a calibrated version of the Fuster, Laibson, and Mendel model allows them to fit the patterns of excess asset-price volatility documented in their paper.

As emphasized by Benjamin, Rabin, and Raymond (2016), one might expect any or all of the above biases to be operative for a given agent (or agents) depending on features of the environment (e.g., the sample size of observed signals). Similarly, Arrow (1982) argues that *both* base-rate neglect and overweighting of new signals "typif[y] very precisely the excessive reaction to current information which seems to characterize all the securities and futures markets" (p. 5). For our purposes, we wish to know what set of realistic modeling assumptions are capable of parsimoniously explaining the features of our data, and more work is needed to address this question definitively. But one striking feature of our empirical results, as examined in Section 6.1, is that risk-neutral beliefs exhibit less excess movement as information accumulates and expiration becomes nearer, which matches the intuition of Benjamin, Rabin, and Raymond (2016) that the law of small numbers may (as its name might suggest) be less likely to be at play in larger samples.

There are, however, alternative classes of explanations for our empirical findings that do not rely so heavily on individual agents' statistical errors leading to relative overreaction to new information. Models in which agents have heterogeneous beliefs (e.g., Harrison and Kreps, 1978; Scheinkman and Xiong, 2003; Geanakoplos, 2010; Simsek, 2013) imply the possibility of volatile belief changes for the marginal holder of a given asset (or the as-if representative agent) as the identity of that marginal agent changes. These models do, however, require agents either to agree to disagree or to neglect their belief disagreements in order to break the no-trade theorem and generate positive trade in equilibrium (Eyster, Rabin, and Vayanos, 2018); it is accordingly unclear the extent to which these models are "behavioral" versus rational even at the individual level.[51]

[50] These models relate as well to the literature on *experience effects*; see Malmendier, Pouzo, and Vanasco (2017) for a theoretical model as well as a discussion of earlier empirical evidence, and relatedly see Pástor and Veronesi (2009) for a review of *learning* in asset markets. The implications of these and the above models stand in contrast to macroeconomic models of bounded rationality or rational inattention (e.g., Sims, 2003; Woodford, 2009; Gabaix, 2018), which can generate underextrapolation, but it is possible that there are distinct phenomena governing the behavior of the average household versus the average investor.

[51] Similarly, classes of models in which agents' private information can affect prices in general require either the presence of some noise traders (e.g., Shiller, 1984) or widespread departures from rationality (Daniel, Hirshleifer, and Subrahmanyam, 1998); see Eyster, Rabin, and Vayanos (2018) and Campbell (2017) for discussion. Pan and Poteshman

This issue notwithstanding, one may plausibly consider testing the rationality of individual beliefs empirically with access to data on agents' portfolios: if we assume that agents are marginal asset-holders (so that option prices reveal their risk-neutral beliefs) on days on which we observe changes in their portfolio holdings, then we may take advantage of the ability of our empirical tests to be applied at any frequency of data sampling and conduct separate tests of RE for each of these agents.[52]

Separately, models of *ambiguity aversion* (e.g., Gilboa and Schmeidler, 1989; Hansen and Sargent, 2008; Epstein and Schneider, 2010) — which specify that investors have uncertainty not only over future probabilistic outcomes but also over the model specification itself, and respond to this meta-uncertainty by maximizing utility under a worst-case model — are capable of generating belief distortions relative to the correct probability distribution. In this vein, the recent models of Drechsler (2013) and Bidder and Dew-Becker (2016) in particular aim to explain excess equity volatility and, in Drechsler's case, empirical properties of index-option prices.

All alternative models considered here accordingly yield a departure from the RE null in some manner, but assessing their relative capacity to match our observed data with reasonable parameter values will, one hopes, allow us to distinguish between divergent structural explanations for these empirical findings.

8. Conclusion

We consider a general theoretical framework in which we show that the assumption of rational expectations imposes testable restrictions on the time variation in risk-neutral beliefs as expressed in asset prices. Unlike in much of the previous literature, these results do not require any restrictive assumptions on the data-generating process for prices or returns, and they allow for time variation in discount rates. Further, by using asset prices, we do not require direct measures of beliefs over future outcomes, and our bounds exploit intertemporal consistency requirements of beliefs under Bayes' rule without the need for the econometrician to know what agents' beliefs "should" be under RE.[53]

When taken to the data, these bounds give direct evidence on the minimum value of risk aversion required to rationalize the observed behavior of risk-neutral beliefs. Using risk-neutral beliefs over the future value of the S&P 500 index measured from index-option data, we find that very high risk aversion is needed to rationalize the variation in these beliefs, indicating that the RE assumption is quite restrictive; in some cases, no amount of risk aversion is capable of rationalizing this belief movement. These results appear to be driven largely by excessive volatility of beliefs over the index value at distant horizons.

(2006) present empirical evidence suggesting that informed trading is relatively unimportant for index options.

[52] We thank Emmanuel Farhi for this suggestion.

[53] The fact that econometricians testing expectations have no agreed-upon correct model of the world might itself be taken as a priori evidence against the RE hypothesis, but it is of course still possible that agents have correct beliefs on average; we accordingly seek to explicitly measure how restrictive the RE assumption is in the data.

Finally, an additional direction for future work is to consider the application of our theoretical and empirical framework to different asset classes. Long-term interest rates have experienced spells of significant volatility in recent years, as noted by Hanson and Stein (2015) and Farhi and Werning (2017), and the forward premium puzzle in foreign-exchange markets (Hansen and Hodrick, 1980; Fama, 1984) suggests that beliefs in these markets may also be worth examining further. More speculatively, the possibility of excess belief volatility in markets in which risk-neutral beliefs cannot be directly measured has additional implications for real quantities. The unemployment volatility puzzle documented by Shimer (2005) can be recast as a puzzle of excess volatility in the value to an employer of a filled vacancy, an idea pursued by Hall (2017) and Kilic and Wachter (2018); meanwhile, a literature including Iacoviello (2005), He and Krishnamurthy (2013), and Jones, Midrigan, and Philippon (2017) aims to match business-cycle dynamics using models with collateral constraints tied to volatile housing values. The extent to which these swings arise in the data due to excess belief movement, as documented for equity-index prices in this paper, would thus inform our understanding of the causes underlying important macroeconomic dynamics, and we aim to continue to pursue these questions in future work.

Appendix A. Proofs

Proofs for Sections 2 and 3

Proof of Lemma 1. Note that $\mathbb{E}[r_{t_1,t_2}] = \mathbb{E}[u_{t_1} - u_{t_2}]$ is finite under any process since $u_t \in [0, 0.25]$, so the expectation $\mathbb{E}[m_{t_1,t_2}]$ must exist as well under the statement given in the lemma, which we prove now. Consider the conditional expectation of the first term in the movement sum:

$$\begin{aligned}\mathbb{E}_{t_1}[m_{t_1,t_1+1}] = \mathbb{E}_{t_1}[(\pi_{t_1+1} - \pi_{t_1})^2] &= \mathbb{E}_{t_1}[\pi_{t_1+1}^2] - 2\mathbb{E}_{t_1}[\pi_{t_1+1}]\pi_{t_1} + \pi_{t_1}^2 \\ &= \mathbb{E}_{t_1}[\pi_{t_1+1}^2] - 2\pi_{t_1}\pi_{t_1} + \pi_{t_1}^2 \\ &= \mathbb{E}_{t_1}[\pi_{t_1+1}^2] - \pi_{t_1}^2 + \pi_{t_1} - \mathbb{E}_{t_1}[\pi_{t_1+1}] \\ &= \mathbb{E}_{t_1}[(1-\pi_{t_1})\pi_{t_1} - (1-\pi_{t_1+1})\pi_{t_1+1}] = \mathbb{E}_{t_1}[r_{t_1,t_1+1}],\end{aligned}$$

where the second and third lines follow from the martingale property of beliefs under Bayes' rule (see Footnote 7) and the last line rearranges. Repeating and summing across all periods from t_1 to t_2 and applying the law of iterated expectations yields $\mathbb{E}_{t_1}[m_{t_1,t_2}] = \mathbb{E}_{t_1}[r_{t_1,t_2}]$, implying $\mathbb{E}[m_{t_1,t_2}] = \mathbb{E}[r_{t_1,t_2}]$. □

Proof of Equation (14). This follows from a discrete-state application of Breeden and Litzenberger (1978), or see Brown and Ross (1991) for a general version. To review why the stated equation holds, the risk-neutral pricing equation (12) can be written for options as

$$q_{t,K}^m = \frac{1}{R_{t,T}^f}\mathbb{E}_t^*[\max\{V_T^m - K, 0\}] = \frac{1}{R_{t,T}^f}\left[\sum_{j\,:\,K_j \geqslant K}(K_j - K)\underbrace{\mathbb{P}_t^*(V_T^m = K_j)}_{\mathbb{P}_t^*(R_T^m = s_j)}\right].$$

This implies that for two adjacent return states s_{j-1} and s_j,

$$\begin{aligned}q_{t,K_j}^m - q_{t,K_{j-1}}^m &= \frac{1}{R_{t,T}^f}\left[\sum_{j'\geqslant j}(K_{j'} - K_j)\,\mathbb{P}_t^*(V_T^m = K_{j'}) - \sum_{j'\geqslant j-1}(K_{j'} - K_{j-1})\,\mathbb{P}_t^*(V_T^m = K_{j'})\right] \\ &= \frac{1}{R_{t,T}^f}\left[\sum_{j'\geqslant j}(K_{j-1} - K_j)\,\mathbb{P}_t^*(V_T^m = K_{j'})\right] = \frac{1}{R_{t,T}^f}(K_{j-1} - K_j)\left[1 - \mathbb{P}_t^*(V_T^m < K_j)\right].\end{aligned}$$

Rearranging,

$$R_{t,T}^f\frac{q_{t,K_j}^m - q_{t,K_{j-1}}^m}{K_j - K_{j-1}} = \mathbb{P}_t^*(V_T^m < K_j) - 1.$$

Repeating this analysis for the pair s_j and s_{j+1}, we obtain $R_{t,T}^f\frac{q_{t,K_{j+1}}^m - q_{t,K_j}^m}{K_{j+1}-K_j} = \mathbb{P}_t^*(V_T^m < K_{j+1}) - 1$. Subtracting the preceding equation from this equation and using $\mathbb{P}_t^*(R_T^m = s_j) = \mathbb{P}_t^*(V_T^m = K_j)$ yields equation (14). □

Proof of Example 1. We prove the statement separately for the two assumptions on the form of the utility function:

(i) Time-separable utility: Denote $V_j^m \equiv V_0^m s_j$ and $V_{j+1}^m \equiv V_0^m s_{j+1}$, so the event $R_T^m = s_j$ is equivalent to $V_T^m = V_j^m$ given $\mathcal{F}_0$, and similarly for s_{j+1} and V_{j+1}^m. Since $dV_T^m/dA_T > 0$ (and with $\mathbb{P}(V_T^m = V_j^m) > 0$, $\mathbb{P}(V_T^m = V_{j+1}^m) > 0$), there exist unique values a_j and a_{j+1} such that $V_T^m = V_j^m$ if and only if $A_T = a_j$, and $V_T^m = V_{j+1}^m$ if and only if $A_T = a_{j+1}$. Then with $M_T/M_t = \beta^{T-t}U'(C_T)/U'(C_t)$ as in (10) given the assumptions for this example, we have

$$\begin{aligned}\phi_{t,j} \equiv \frac{\mathbb{E}_t[M_T/M_t \mid R_T^m = s_j]}{\mathbb{E}_t[M_T/M_t \mid R_T^m = s_{j+1}]} &= \frac{\mathbb{E}_t[M_T/M_t \mid A_T = a_j]}{\mathbb{E}_t[M_T/M_t \mid A_T = a_{j+1}]} \\ &= \frac{U'(C_T(a_j))}{U'(C_T(a_{j+1}))},\end{aligned}$$

which is almost surely constant, as required for CTI to hold.

(ii) Epstein–Zin (1989) utility: The Epstein–Zin preference recursion is

$$U_t = \left\{(1-\beta)C_t^{1-\frac{1}{\psi}} + \beta\left(\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]\right)^{\frac{1-\frac{1}{\psi}}{1-\gamma}}\right\}^{\frac{1}{1-\frac{1}{\psi}}}. \tag{A.1}$$

It can be shown (e.g., Campbell, 2017, p. 178) that given such preferences the SDF evolves according to

$$\frac{M_{t+1}}{M_t} = \beta\left(\frac{C_{t+1}}{C_t}\right)^{-\frac{1}{\psi}}\left(\frac{U_{t+1}}{\mathbb{E}_t\left[U_{t+1}^{1-\gamma}\right]^{\frac{1}{1-\gamma}}}\right)^{-\left(\gamma-\frac{1}{\psi}\right)},$$

which gives that

$$\frac{M_T}{M_t} = \beta^{T-t}\left(\frac{C_T}{C_t}\right)^{-\frac{1}{\psi}}\prod_{\tau=t}^{T-1}\left(\frac{U_{\tau+1}}{\mathbb{E}_\tau\left[U_{\tau+1}^{1-\gamma}\right]^{\frac{1}{1-\gamma}}}\right)^{-\left(\gamma-\frac{1}{\psi}\right)} \tag{A.2}$$

$$= \beta^{T-t}\left(\frac{C_T}{C_t}\right)^{-\gamma}\prod_{\tau=t}^{T-1}\left(\frac{U_{\tau+1}}{C_{\tau+1}}\right)^{-\left(\gamma-\frac{1}{\psi}\right)}\mathbb{E}_\tau\left[\left(\frac{C_{\tau+1}}{C_\tau}\right)^{1-\gamma}\left(\frac{U_{\tau+1}}{C_{\tau+1}}\right)^{1-\gamma}\right]^{\frac{\gamma-\frac{1}{\psi}}{1-\gamma}}. \tag{A.2′}$$

Denote a_j and a_{j+1} as in part (i). From the first representation of M_T/M_t, equation (A.2), it follows immediately that with i.i.d. consumption (or i.i.d. innovations to an otherwise

predetermined consumption path),

$$\begin{aligned}\phi_{t,j} &= \frac{\mathbb{E}_t[M_T/M_t \mid A_T = a_j]}{\mathbb{E}_t[M_T/M_t \mid A_T = a_{j+1}]} \\ &= \left(\frac{C_T(a_j)}{C_T(a_{j+1})}\right)^{-\frac{1}{\psi}} \left(\frac{U_T(a_j)}{U_T(a_{j+1})}\right)^{-\left(\gamma-\frac{1}{\psi}\right)},\end{aligned}$$

which is almost surely constant given the definition (A.1) and that $E_T[U_{T+1}^{1-\gamma}]$ is constant given the i.i.d. assumption.

When consumption growth C_t/C_{t-1} is i.i.d., note that the scale independence of Epstein–Zin utility in (A.1) allows us to guess and verify that U_t/C_t is constant almost surely. Then from the second representation of M_T/M_t, equation (A.2′), we have in this case that

$$\phi_{t,j} = \left(\frac{C_T(a_j)}{C_T(a_{j+1})}\right)^{-\gamma},$$

completing the proof. □

Proof of Example 2.

(i) Given that the postulated agent is unconstrained, her intertemporal marginal rate of substitution $\beta(C_{t+1}/C_t)^{-\gamma}$ serves as a valid measure of M_{t+1}/M_t by her Euler equation, so

$$\phi_{t,j} = \frac{\mathbb{E}_t[C_T^{-\gamma} \mid V_T^m = V_j^m]}{\mathbb{E}_t[C_T^{-\gamma} \mid V_T^m = V_{j+1}^m]},$$

with V_j^m and V_{j+1}^m as defined in the proof of Example 1. With (C_t, D_t) i.i.d. (and therefore also V_t^m i.i.d.), this value must be constant for $0 \leqslant t < T$.

(ii) Given i.i.d. consumption growth and dividend growth, the market price-to-dividend ratio $V_t^m/D_t = \mathbb{E}_t[\sum_{\tau=1}^{\infty} \beta^\tau (C_\tau/C_t)^{-\gamma}(D_\tau/D_t)]$ must be constant, so the event $V_T^m = V_j^m$ is equivalent to $D_T = D_j$ for some value D_j (and similarly for $j+1$). We thus have

$$\begin{aligned}\phi_{t,j} &= \frac{\mathbb{E}_t[(C_T/C_t)^{-\gamma} \mid D_T/D_t = D_j/D_t]}{\mathbb{E}_t[(C_T/C_t)^{-\gamma} \mid D_T/D_t = D_{j+1}/D_t]} \\ &= \frac{\mathbb{E}_t[(C_T/C_t)^{-\gamma} \mid d_T - d_t = d_j - d_t]}{\mathbb{E}_t[(C_T/C_t)^{-\gamma} \mid d_T - d_t = d_{j+1}/d_t]},\end{aligned}$$

where $d_t \equiv \log(D_t)$ (and similarly for d_j and d_{j+1}). For the numerator,

$$\mathbb{E}_t[(C_T/C_t)^{-\gamma} \mid d_T - d_t = d_j - d_t] = \exp\{\log(\mathbb{E}_t[(C_T/C_t)^{-\gamma} \mid d_T - d_t = d_j - d_t])\}$$

$$= \exp\Big\{ \mathbb{E}_t[-\gamma(c_T - c_t) \mid d_T - d_t = d_j - d_t] + \frac{1}{2}\gamma^2 \operatorname{Var}_t(c_T - c_t \mid d_T - d_t = d_j - d_t) \Big\},$$

where $c_t \equiv \log(C_t)$, and similarly for the denominator with respect to d_{j+1}. By assumption,

$$\begin{pmatrix} \Delta c_{t+1} \\ \Delta d_{t+1} \end{pmatrix} \overset{\text{i.i.d.}}{\sim} \mathcal{N}\left(\begin{pmatrix} \mu_c \\ \mu_d \end{pmatrix}, \begin{pmatrix} \sigma_c^2 & \rho\sigma_c\sigma_d \\ \rho\sigma_c\sigma_d & \sigma_d^2 \end{pmatrix} \right),$$

and thus for $c_T - c_t = \sum_{\tau=t+1}^{T} \Delta c_\tau$, we have

$$\left(\sum_{\tau=t+1}^{T} \Delta c_\tau \,\middle|\, \sum_{\tau=t+1}^{T} \Delta d_\tau = d_j - d_t \right) \sim \mathcal{N}\left((T-t)\mu_c + \rho\frac{\sigma_c}{\sigma_d}\left((d_j - d_t) - (T-t)\mu_d \right), (T-t)\sigma_c^2(1-\rho^2) \right).$$

Thus $\operatorname{Var}_t(c_T - c_t \mid d_T - d_t = d_j - d_t) = \operatorname{Var}_t(c_T - c_t \mid d_T - d_t = d_{j+1} - d_t)$. Further, for $j' = j, j+1$, we have $\mathbb{E}_t[c_T - c_t \mid d_T - d_t = d_{j'} - d_t] = (T-t)\mu_c + \rho\frac{\sigma_c}{\sigma_d}((d_{j'} - d_t) - (T-t)\mu_d)$, and therefore

$$\phi_{t,j} = \exp\left\{ -\gamma\rho\frac{\sigma_c}{\sigma_d}(d_j - d_{j+1}) \right\},$$

which is a constant.[54] □

Proof of Example 3. Gabaix (2012, Theorem 1) shows that

$$V_t^m = \frac{D_t}{1 - e^{-\beta_m}} \left(1 + \frac{e^{-\beta_m - h_*}\widehat{H}_t}{1 - e^{-\beta_m - \phi_H}} \right),$$

where $h_* \equiv \log(1 + H_*)$ and $\beta_m \equiv \beta - g_d - h_*$ (where β is the agent's time discount factor). Thus for any value s and given $\mathcal{F}_0$, there exists some value d_s and function $f(d_s, \widehat{H}_T)$, which is strictly increasing in the first argument and strictly decreasing in the second argument, such that

$$\mathbb{P}_0\left(\sum_{t=1}^{T} \mathbb{1}\{\text{disaster}_t\} > 0 \,\middle|\, R_T^m \geqslant s \right) = \frac{\mathbb{P}_0\left(R_T^m \geqslant s \mid \sum_{t=1}^{T} \mathbb{1}\{\text{disaster}_t\} > 0 \right) \mathbb{P}_0\left(\sum_{t=1}^{T} \mathbb{1}\{\text{disaster}_t\} > 0 \right)}{\mathbb{P}_0(R_T^m \geqslant s)}$$

[54] One may note that given that each cumulant of a sum of i.i.d. variables scales with the number of variates being summed, the above arguments would apply as well to more general (non-log-normal) distributions, but we consider the log-normal case for simplicity.

$$= \frac{\mathbb{P}_0\Big(D_T \geqslant f(d_s, \widehat{H}_T) \,\Big|\, \sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} > 0\Big)\, \mathbb{P}_0\Big(\sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} > 0\Big)}{\mathbb{P}_0\Big(D_T \geqslant f(d_s, \widehat{H}_T)\Big)}.$$

Note now that (i) the innovation to $\widehat{H}_{t+1}$ is independent of the disaster realization; (ii) F_{t+1} (the exponential of the disaster shock to D_t) has support $[0,1]$; and (iii) $\mathbb{P}_t(\varepsilon^d_{t+1} = \epsilon) = \mathcal{O}(e^{-\epsilon^2})$ as $\epsilon \to \infty$. Thus $\mathbb{P}_0(D_T \geqslant f(d_s, \widehat{H}_T) \mid \sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} > 0) = o(\mathbb{P}_0(D_T \geqslant f(d_s, \widehat{H}_T)))$ as $d_s \to \infty$, from which the first statement given in the example follows.

Now define

$$\phi_j \equiv \frac{\mathbb{E}_0[C_T^{-\gamma} \mid V_T^m = V_j^m, \sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} = 0]}{\mathbb{E}_0[C_T^{-\gamma} \mid V_T^m = V_{j+1}^m, \sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} = 0]},$$

where V_j^m and V_{j+1}^m are as defined in the proof of Example 1, and where this value is constant almost surely given that conditional on $\sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} = 0$, the conditions from Example 1 hold. After defining the value η_t such that $\phi_{t,j} = \phi_j + \eta_t$, we then obtain after some algebra that $\eta_t = \mathcal{O}_p(\delta)$ given the first statement shown above, so the second stated result follows. □

Additional Lemmas Used in Proofs for Section 4

Before proceeding to the proofs of our main results, we provide three additional lemmas that are useful in proving those results. As usual, we assume throughout that rational expectations holds.

Lemma A.1. *For some return-state pair (s_j, s_{j+1}), with $\widetilde{\mathbb{P}} \equiv \mathbb{P}(\cdot \mid R_T^m \in \{s_j, s_{j+1}\})$ as per (17), define a new pseudo-risk-neutral measure $\widetilde{\mathbb{P}}^\diamond$ by*

$$\left.\frac{d\widetilde{\mathbb{P}}^\diamond}{d\widetilde{\mathbb{P}}}\right|_{\mathcal{F}_t} = \frac{\widetilde{\pi}^*_{t,j}}{\widetilde{\pi}_{t,j}} \mathbb{1}\{R_T^m = s_j\} + \frac{1 - \widetilde{\pi}^*_{t,j}}{1 - \widetilde{\pi}_{t,j}} \mathbb{1}\{R_T^m = s_{j+1}\}. \tag{A.3}$$

Denote the conditional expectation under $\widetilde{\mathbb{P}}^\diamond$ by $\widetilde{\mathbb{E}}^\diamond_t[\,\cdot\,]$. If conditional transition independence holds for the return-state pair (s_j, s_{j+1}), and $\mathbb{P}_t(R_T^m \in \{s_j, s_{j+1}\}) > 0$, we have that $\widetilde{\mathbb{P}}^\diamond$ serves as a martingale measure for the risk-neutral belief in the sense that

$$\widetilde{\pi}^*_{t,j} = \widetilde{\mathbb{E}}^\diamond_t[\pi^*_{t+1,j}].$$

We conclude from Lemma 1 that for $t_1, t_2 > t_1$,

$$\widetilde{\mathbb{E}}^\diamond_{t_1}[m^*_{t_1,t_2,j}] = \widetilde{\mathbb{E}}^\diamond_{t_1}[r^*_{t_1,t_2,j}],$$

where $m^*_{t_1,t_2,j}$ and $r^*_{t_1,t_2,j}$ are as defined in Definitions 2–3.

Proof of Lemma A.1. From (18), we have after some algebra that

$$\frac{\widetilde{\pi}^*_{t,j}}{\widetilde{\pi}_{t,j}} = \frac{\phi_j}{1 + \widetilde{\pi}_{t,j}(\phi_j - 1)}, \tag{A.4}$$

$$\frac{1-\widetilde{\pi}^*_{t,j}}{1-\widetilde{\pi}_{t,j}} = \frac{1}{1+\widetilde{\pi}_{t,j}(\phi_j-1)}. \tag{A.5}$$

Note therefore that $\widetilde{\mathbb{P}}^\diamond$ is absolutely continuous with respect to $\widetilde{\mathbb{P}}$.

Recall that $\mathcal{F}_t = \sigma(\theta_\tau, 0 \leqslant \tau \leqslant t)$, where $\sigma(\theta_\tau, 0 \leqslant \tau \leqslant t)$ is the σ-algebra generated by the stochastic process $\{\theta_t\}$ and $\theta_t \in \Theta$ is the date-t signal vector. Denote $N_\Theta \equiv |\Theta|$, so that $\theta_t \in \{\theta_1, \theta_2, \ldots, \theta_{N_\Theta}\} = \{\theta_k\}_{k=1,\ldots,N_\Theta}$, and further denote

$$\begin{aligned}
\mathfrak{p}_{t,k} &\equiv \widetilde{\mathbb{P}}_t(\theta_{t+1} = \theta_k), \\
\varrho_{t,k} &\equiv \widetilde{\mathbb{P}}_t(R^m_T = s_j \,|\, \theta_{t+1} = \theta_k), \\
\varrho^*_{t,k} &\equiv \mathbb{P}^*_t(R^m_T = s_j \,|\, \theta_{t+1} = \theta_k, R^m_T \in \{s_j, s_{j+1}\}),
\end{aligned}$$

so that $\widetilde{\pi}_{t+1,j} = \varrho_{t,k}$ if $\theta_{t+1} = \theta_k$, and similarly $\widetilde{\pi}^*_{t+1,j} = \varrho^*_{t,k}$ if $\theta_{t+1} = \theta_k$.

Combining (A.3), (A.4), (A.5), and these definitions, we have

$$\begin{aligned}
\widetilde{\mathbb{E}}^\diamond_t[\widetilde{\pi}^*_{t+1,j}] &= \frac{\widetilde{\pi}^*_{t,j}}{\widetilde{\pi}_{t,j}} \sum_{k=1}^{N_\Theta} \mathfrak{p}_{t,k}\, \varrho^*_{t,k}\, \widetilde{\mathbb{E}}_t\big[\mathbb{1}\{R^m_T = s_j\} \,|\, \theta_{t+1} = \theta_k\big] \\
&\quad + \frac{1-\widetilde{\pi}^*_{t,j}}{1-\widetilde{\pi}_{t,j}} \sum_{k=1}^{N_\Theta} \mathfrak{p}_{t,k}\, \varrho^*_{t,k}\, \widetilde{\mathbb{E}}_t\big[\mathbb{1}\{R^m_T = s_{j+1}\} \,|\, \theta_{t+1} = \theta_k\big] \\
&= \frac{\phi_j}{1+\widetilde{\pi}_{t,j}(\phi_j-1)} \sum_{k=1}^{N_\Theta} \mathfrak{p}_{t,k}\, \frac{\phi_j \varrho_{t,k}}{1+\varrho_{t,k}(\phi_j-1)}\, \varrho_{t,k} \\
&\quad + \frac{1}{1+\widetilde{\pi}_{t,j}(\phi_j-1)} \sum_{k=1}^{N_\Theta} \mathfrak{p}_{t,k}\, \frac{\phi_j \varrho_{t,k}}{1+\varrho_{t,k}(\phi_j-1)}\, (1-\varrho_{t,k}) \\
&= \frac{\phi_j}{1+\widetilde{\pi}_{t,j}(\phi_j-1)} \sum_{k=1}^{N_\Theta} \mathfrak{p}_{t,k}\, \frac{\varrho_{t,k}\big(1+\varrho_{t,k}(\phi_j-1)\big)}{1+\varrho_{t,k}(\phi_j-1)} \\
&= \frac{\phi_j}{1+\widetilde{\pi}_{t,j}(\phi_j-1)} \sum_{k=1}^{N_\Theta} \mathfrak{p}_{t,k}\, \varrho_{t,k} \\
&= \frac{\phi_j \widetilde{\pi}_{t,j}}{1+\widetilde{\pi}_{t,j}(\phi_j-1)} \\
&= \widetilde{\pi}^*_{t,j},
\end{aligned}$$

where the second-to-last equality uses that $\widetilde{\pi}_{t,j} = \widetilde{\mathbb{E}}_t[\widetilde{\pi}_{t+1,j}]$, as can be seen from the law of iterated expectations given that $\widetilde{\pi}_{t,j} = \mathbb{E}_t[\mathbb{1}\{R^m_T = s_j\} \,|\, R^m_T \in \{s_j, s_{j+1}\}] = \widetilde{\mathbb{E}}_t[\mathbb{1}\{R^m_T = s_j\}] = \widetilde{\mathbb{E}}_t[\widetilde{\mathbb{E}}_{t+1}[\mathbb{1}\{R^m_T = s_j\}]] = \widetilde{\mathbb{E}}_t[\widetilde{\pi}_{t+1,j}]$, and the last equality above again uses (A.4).

The fact that $\widetilde{\mathbb{E}}^\diamond_{t_1}[m^*_{t_1,t_2,j}] = \widetilde{\mathbb{E}}^\diamond_{t_1}[r^*_{t_1,t_2,j}]$ for arbitrary t_1 and $t_2 > t_1$ then follows immediately from the proof of Lemma 1. □

LEMMA A.2. *For any return-state pair (s_j, s_{j+1}) meeting CTI, risk-neutral belief movement must satisfy the following for $j' = j, j+1$:*

$$\widetilde{\mathbb{E}}_0^{\diamond}[m_j^* \mid R_T^m = s_{j'}] = \widetilde{\mathbb{E}}_0[m_j^* \mid R_T^m = s_{j'}].$$

Proof of Lemma A.2. Define the path of risk-neutral beliefs by $\mathcal{B}_j \equiv (\widetilde{\pi}_{0,j}^*, \widetilde{\pi}_{1,j}^*, \ldots, \widetilde{\pi}_{T,j}^*)$, and some arbitrary realization for that path by $\mathfrak{b}_j$. The realization of m_j^* depends on the path of risk-neutral beliefs, so denote $m_j^* = m_j^*(\mathcal{B}_j) = \sum_{t=1}^T (\widetilde{\pi}_{t,j}^* - \widetilde{\pi}_{t-1,j}^*)^2$.

For any $\mathfrak{b}_j$ such that $\widetilde{\pi}_{T,j}^* = 1$ (i.e., $R_T^m = s_j$), the definition of $\widetilde{\mathbb{P}}^{\diamond}$ in (A.3) gives that

$$\widetilde{\mathbb{P}}_0^{\diamond}(\mathcal{B}_j = \mathfrak{b}_j) = \frac{\widetilde{\pi}_{0,j}^*}{\widetilde{\pi}_{0,j}} \widetilde{\mathbb{P}}(\mathcal{B}_j = \mathfrak{b}_j), \tag{A.6}$$

and further $\widetilde{\mathbb{P}}_0^{\diamond}(R_T^m = s_j) = (\widetilde{\pi}_{0,j}^* / \widetilde{\pi}_{0,j}) \widetilde{\mathbb{P}}_0(R_T^m = s_j)$ trivially. Combining these two equations yields $\widetilde{\mathbb{P}}_0^{\diamond}(\mathcal{B}_j = \mathfrak{b}_j \mid R_T^m = s_j) = \widetilde{\mathbb{P}}_0(\mathcal{B}_j = \mathfrak{b}_j \mid R_T^m = s_j)$. (Intuitively, all paths ending in $\widetilde{\pi}_{T,j}^* = 1$ receive the same change of measure under $\widetilde{\mathbb{P}}^{\diamond}$ relative to $\widetilde{\mathbb{P}}$, so probabilities conditional on $R_T^m = s_j$ are preserved, and similarly for $R_T^m = s_{j+1}$.) Thus

$$\begin{aligned}
\widetilde{\mathbb{E}}_0^{\diamond}[m_j^* \mid R_T^m = s_j] &= \sum_{\mathfrak{b}_j \,:\, \widetilde{\pi}_{T,j}^* = 1} m_j^*(\mathfrak{b}_j) \, \widetilde{\mathbb{P}}_0^{\diamond}\big(\mathcal{B}_j = \mathfrak{b}_j \,\big|\, R_T^m = s_j\big) \\
&= \sum_{\mathfrak{b}_j \,:\, \widetilde{\pi}_{T,j}^* = 1} m_j^*(\mathfrak{b}_j) \, \widetilde{\mathbb{P}}_0\big(\mathcal{B}_j = \mathfrak{b}_j \,\big|\, R_T^m = s_j\big) \\
&= \widetilde{\mathbb{E}}_0[m_j^* \mid R_T^m = s_j].
\end{aligned}$$

The same steps apply for $R_T^m = s_{j+1}$: in this case, for any $\mathfrak{b}_j$ such that $\widetilde{\pi}_{T,j}^* = 0$, (A.6) becomes $\widetilde{\mathbb{P}}_0^{\diamond}(\mathcal{B}_j = \mathfrak{b}_j) = (1 - \widetilde{\pi}_{0,j}^*)/(1 - \widetilde{\pi}_{0,j}) \, \widetilde{\mathbb{P}}(\mathcal{B}_j = \mathfrak{b}_j)$, and $\widetilde{\mathbb{P}}_0^{\diamond}(R_T^m = s_{j+1}) = (1 - \widetilde{\pi}_{0,j}^*)/(1 - \widetilde{\pi}_{0,j}) \, \widetilde{\mathbb{P}}_0(R_T^m = s_{j+1})$, so that again $\widetilde{\mathbb{P}}_0^{\diamond}(\mathcal{B}_j = \mathfrak{b}_j \mid R_T^m = s_{j+1}) = \widetilde{\mathbb{P}}_0(\mathcal{B}_j = \mathfrak{b}_j \mid R_T^m = s_{j+1})$, and thus $\widetilde{\mathbb{E}}_0^{\diamond}[m_j^* \mid R_T^m = s_{j+1}] = \widetilde{\mathbb{E}}_0[m_j^* \mid R_T^m = s_{j+1}]$. □

LEMMA A.3. *Consider a return-state pair (s_j, s_{j+1}) meeting CTI, and an arbitrary date τ such that $0 \leqslant \tau < T$. Denote $\widetilde{\varepsilon} \equiv (1 - \widetilde{\pi}_{\tau,j}^*)/(T - \tau)$. Postulate a signal-generating process for $\tau < t \leqslant T$ such that if $\widetilde{\pi}_{t-1,j}^* = 0$ then $\widetilde{\pi}_{t,j}^* = 0$ with probability 1, and otherwise*

$$\widetilde{\pi}_{t,j}^* = \begin{cases} 0, & \text{with probability } x_{t-1} \\ \widetilde{\pi}_{t-1,j}^* + \widetilde{\varepsilon}, & \text{with probability } 1 - x_{t-1}, \end{cases}$$

$$\text{with} \quad x_{t-1} = 1 - \left[\frac{\widetilde{\pi}_{t-1,j}^*}{\phi_j + (1 - \phi_j)(\widetilde{\pi}_{t-1,j}^*)} \right] \Bigg/ \left[\frac{\widetilde{\pi}_{t-1,j}^* + \widetilde{\varepsilon}}{\phi_j + (1 - \phi_j)(\widetilde{\pi}_{t-1,j}^* + \widetilde{\varepsilon})} \right].$$

This "rare-bonanzas" process is consistent with rational expectations. Further, under this process,

$$\widetilde{\mathbb{E}}_\tau[m^*_{\tau,T,j} - r^*_{\tau,T,j}] = \frac{\widetilde{\pi}^*_{\tau,j}(1-\widetilde{\pi}^*_{\tau,j})(\phi_j - 1)(\widetilde{\pi}^*_{\tau,j} - \widetilde{\varepsilon})}{\widetilde{\pi}^*_{\tau,j} + \phi_j(1-\widetilde{\pi}^*_{\tau,j})}.$$

Proof. The first bracketed term on the right side of the definition of x_{t-1} is equal to $\widetilde{\pi}_{t-1,j}$ from equation (18), and the second term is equal to $\widetilde{\pi}_{t,j}$ in the case that $\widetilde{\pi}^*_{t,j} = \widetilde{\pi}^*_{t-1,j} + \widetilde{\varepsilon}$; denote this latter value as $\widetilde{\pi}^H_{t,j} \equiv (\widetilde{\pi}^*_{t-1,j} + \widetilde{\varepsilon})/(\phi_j + (1-\phi_j)(\widetilde{\pi}^*_{t-1,j} + \widetilde{\varepsilon}))$. Thus by construction $\widetilde{\pi}_{t-1,j} = x_{t-1} \times 0 + (1 - x_{t-1})\widetilde{\pi}^H_{t,j}$, so the physical belief process is consistent with RE, as stated. (It is trivial from Definition 1 to construct a signal-generating process $\{\mathbb{P}(\theta_t | \mathcal{F}_{t-1}, R^m_T = s)\}$ that, conditional on the prior $\Pi_{\tau,T}$, delivers the postulated belief process.) Note also that given the definition of $\widetilde{\varepsilon}$, this belief process generates $\widetilde{\pi}^*_{T,j} \in \{0,1\}$ almost surely so that $R^m_T \in \{s_j, s_{j+1}\}$.

For the second statement in the lemma, note that for this constructed process, we have

$$\begin{aligned}\mathbb{E}_\tau[m^*_{\tau,\tau+1,j}] &= x_\tau \widetilde{\pi}^{*2}_{\tau,j} + (1-x_\tau)\widetilde{\varepsilon}^2 \\ &= \frac{\phi_j + (1-\phi_j)\widetilde{\varepsilon}}{\phi_j + (1-\phi_j)\widetilde{\pi}^*_{\tau,j}}\widetilde{\pi}^*_{\tau,j}\widetilde{\varepsilon}.\end{aligned}$$

Conducting this calculation for each subsequent period and summing over periods through T, we therefore obtain

$$\begin{aligned}\mathbb{E}_\tau[m^*_{\tau,T,j}] &= \sum_{k=0}^{T-\tau-1} \Bigg[\overbrace{\frac{\phi_j + (1-\phi_j)\widetilde{\varepsilon}}{\phi_j + (1-\phi_j)(\widetilde{\pi}^*_{\tau,j} + k\widetilde{\varepsilon})}(\widetilde{\pi}^*_{\tau,j} + k\widetilde{\varepsilon})\widetilde{\varepsilon}}^{\mathbb{E}_\tau[m^*_{\tau+k,\tau+k+1,j} \,|\, \widetilde{\pi}^*_{\tau+k,j} > 0]} \\ &\qquad \times \underbrace{\prod_{\ell=0}^{k-1} \frac{\widetilde{\pi}^*_{\tau,j} + \ell\widetilde{\varepsilon}}{\phi_j + (1-\phi_j)(\widetilde{\pi}^*_{\tau,j} + \ell\widetilde{\varepsilon})} \frac{\phi_j + (1-\phi_j)(\widetilde{\pi}^*_{\tau,j} + (\ell+1)\widetilde{\varepsilon})}{\widetilde{\pi}^*_{\tau,j} + (\ell+1)\widetilde{\varepsilon}}}_{\mathbb{P}_\tau(\widetilde{\pi}^*_{\tau+k,j} > 0)} \Bigg] \\ &= \sum_{k=0}^{T-\tau-1} \Bigg[\frac{\phi_j + (1-\phi_j)\widetilde{\varepsilon}}{\phi_j + (1-\phi_j)(\widetilde{\pi}^*_{\tau,j} + k\widetilde{\varepsilon})}(\widetilde{\pi}^*_{\tau,j} + k\widetilde{\varepsilon})\widetilde{\varepsilon} \\ &\qquad \times \frac{\widetilde{\pi}^*_{\tau,j}}{\phi_j + (1-\phi_j)\widetilde{\pi}^*_{\tau,j}} \frac{\phi_j + (1-\phi_j)(\widetilde{\pi}^*_{\tau,j} + k\widetilde{\varepsilon})}{\widetilde{\pi}^*_{\tau,j} + k\widetilde{\varepsilon}} \Bigg] \\ &= \frac{\widetilde{\pi}^*_{\tau,j}(1-\widetilde{\pi}^*_{\tau,j})}{\phi_j + (1-\phi_j)\widetilde{\pi}^*_{\tau,j}}[\phi_j + (1-\phi_j)\widetilde{\varepsilon}],\end{aligned}$$

so

$$\mathbb{E}_\tau[m^*_{\tau,T,j} - r^*_{\tau,T,j}] = \frac{\widetilde{\pi}^*_{\tau,j}(1-\widetilde{\pi}^*_{\tau,j})}{\phi_j + (1-\phi_j)\widetilde{\pi}^*_{\tau,j}}[\phi_j + (1-\phi_j)\widetilde{\varepsilon}] - \widetilde{\pi}^*_{\tau,j}(1-\widetilde{\pi}^*_{\tau,j})$$

$$= \frac{\widetilde{\pi}^*_{\tau,j}(1-\widetilde{\pi}^*_{\tau,j})(\phi_j - 1)(\widetilde{\pi}^*_{\tau,j} - \widetilde{\varepsilon})}{\widetilde{\pi}^*_{\tau,j} + \phi_j(1-\widetilde{\pi}^*_{\tau,j})}.$$

Then using that $\widetilde{\mathbb{E}}_\tau[\cdot] = \mathbb{E}_\tau[\cdot]$ under the stated process, as $R^m_T \in \{s_j, s_{j+1}\}$ almost surely as above, yields the stated result. □

Proofs for Section 4

Proof of Proposition 1. Denote $\Delta_{t,j} \equiv \widetilde{\pi}^*_{t,j} - \widetilde{\pi}_{t,j}$. Then

$$\begin{aligned}
\mathbb{E}_0[m^*_j \mid R^m_T \in \{s_j, s_{j+1}\}] &= \widetilde{\pi}_{0,j}\,\mathbb{E}_0[m^*_j \mid R^m_T = s_j] + (1-\widetilde{\pi}_{0,j})\,\mathbb{E}_0[m^*_j \mid R^m_T = s_{j+1}] \\
&= \widetilde{\pi}^*_{0,j}\,\mathbb{E}_0[m^*_j \mid R^m_T = s_j] + (1-\widetilde{\pi}^*_{0,j})\,\mathbb{E}_0[m^*_j \mid R^m_T = s_{j+1}] \\
&\quad + \Delta_{0,j}\Big\{\mathbb{E}_0[m^*_j \mid R^m_T = s_{j+1}] - \mathbb{E}_0[m^*_j \mid R^m_T = s_j]\Big\} \\
&= \widetilde{\pi}^*_{0,j}\,\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] + (1-\widetilde{\pi}^*_{0,j})\,\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}] \\
&\quad + \Delta_{0,j}\Big\{\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}] - \mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j]\Big\} \\
&= \mathbb{E}^\diamond_0[r^*_j \mid R^m_T \in \{s_j, s_{j+1}\}] + \Delta_{0,j}\Big\{\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}] - \mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j]\Big\}, \qquad \text{(A.7)}
\end{aligned}$$

where the third equality uses Lemma A.2 and the fourth uses Lemma A.1.

Lemma A.1 also implies that

$$\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}] = \widetilde{\pi}^*_{0,j} - \frac{\widetilde{\pi}^*_{0,j}}{1-\widetilde{\pi}^*_{0,j}}\,\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j], \qquad \text{(A.8)}$$

so using this in the term in braces in (A.7), we obtain

$$\begin{aligned}
\mathbb{E}_0[m^*_j - r^*_j \mid R^m_T \in \{s_j, s_{j+1}\}] &= \Delta_{0,j}\widetilde{\pi}^*_{0,j} - \Delta_{0,j}\left(1 + \frac{\widetilde{\pi}^*_{0,j}}{1-\widetilde{\pi}^*_{0,j}}\right)\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] \\
&\leqslant \Delta_{0,j}\widetilde{\pi}^*_{0,j}, \qquad \text{(A.9)}
\end{aligned}$$

since $\Delta_{0,j}\left(1 + \frac{\widetilde{\pi}^*_{0,j}}{1-\widetilde{\pi}^*_{0,j}}\right)\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] \geqslant 0$.

Now

$$\begin{aligned}
\Delta_{0,j} &= \widetilde{\pi}^*_{0,j} - \widetilde{\pi}_{0,j} \\
&= \widetilde{\pi}^*_{0,j} - \frac{\widetilde{\pi}^*_{0,j}}{\widetilde{\pi}^*_{0,j} + \phi_j(1-\widetilde{\pi}^*_{0,j})}
\end{aligned}$$

$$= \widetilde{\pi}^*_{0,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0,j} + \phi_j(1-\widetilde{\pi}^*_{0,j})}\right), \tag{A.10}$$

where the second line uses (18). Substituting the above into the preceding inequality,

$$\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] = \mathbb{E}_0[m^*_j - r^*_j \mid R^m_T \in \{s_j, s_{j+1}\}] \leqslant \widetilde{\pi}^{*2}_{0,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0,j} + \phi_j(1-\widetilde{\pi}^*_{0,j})}\right). \qquad \square$$

Proof of Proposition 2. First consider the statement for fixed $T < \infty$. If $\phi_j > 1$, then $\Delta_{0,j} > 0$ from (A.10). (This requires $\widetilde{\pi}^*_{0,j} \in (0,1)$, but ϕ_j is undefined if this is not the case given its definition in (19), so assuming $\phi_j > 1$ implies this holds as well.) The only inequality applied in the proof of Proposition 1 is that $\Delta_{0,j}(1 + \widetilde{\pi}^*_{0,j}/(1-\widetilde{\pi}^*_{0,j}))\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] \geqslant 0$, as used in (A.9). In the current case with $\phi_j > 1$, we have $\Delta_{0,j}(1 + \widetilde{\pi}^*_{0,j}/(1-\widetilde{\pi}^*_{0,j})) > 0$, so it remains to be shown that $\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] > 0$. Assume toward a contradiction that $\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] = 0$, which requires that $\widetilde{\pi}^*_{t,j} = \widetilde{\pi}^*_{t-1,j}$ almost surely (conditional on $R^m_T = s_j$) since $m^*_j = \sum_{t=1}^T (\widetilde{\pi}^*_{t,j} - \widetilde{\pi}^*_{t-1,j})^2 \geqslant 0$. But this implies that $\widetilde{\pi}^*_{T,j} = \widetilde{\pi}^*_{0,j}$ almost surely, while it must be the case that $\widetilde{\pi}^*_{T,j} = 1$ for any path conditional on $R^m_T = s_j$. Thus we have a contradiction, and $\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] > 0$. We conclude that for finite T, the bound in Proposition 1 must hold with strict inequality as long as $\phi_j > 1$.

For the second part of the statement, we proceed constructively using the signal-generating process postulated in Lemma A.3. Setting $\tau = 0$ and $T \to \infty$ yields $\widetilde{\varepsilon} \to 0$ and therefore, under this process, the second result in that lemma yields

$$\begin{aligned}\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] &\xrightarrow{T\to\infty} \frac{\widetilde{\pi}^*_{0,j}(1-\widetilde{\pi}^*_{0,j})(\phi_j - 1)\widetilde{\pi}^*_{0,j}}{\widetilde{\pi}^*_{0,j} + \phi_j(1-\widetilde{\pi}^*_{0,j})} \\ &= \widetilde{\pi}^{*2}_{0,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0,j} + \phi_j(1-\widetilde{\pi}^*_{0,j})}\right),\end{aligned}$$

as in Proposition 1. $\square$

Proof of Proposition 3. We show first that the effect of an incorrect prior on $\widetilde{\pi}^*_{0,j}$ is isomorphic to a change in ϕ_j under RE. Note from (18) that

$$\frac{\widetilde{\pi}^*_{0,j}}{1-\widetilde{\pi}^*_{0,j}} = \phi_j \frac{\widetilde{\pi}_{0,j}}{1-\widetilde{\pi}_{0,j}}. \tag{A.11}$$

Under RE, $\widetilde{\pi}_{0,j} = \mathbb{P}_0(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\})$, and Proposition 1 follows for $\{\widetilde{\pi}^*_{t,j}\}$ given any value ϕ_j. But if the prior is incorrect, define

$$\xi \equiv \frac{\mathbb{P}_0(R^m_T = s_{j+1} \mid R^m_T \in \{s_j, s_{j+1}\})}{\mathbb{P}_0(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\})}\frac{\widetilde{\pi}_{0,j}}{1-\widetilde{\pi}_{0,j}} \in (0, \infty).$$

Then defining $\check{\phi}_j \equiv \xi\phi_j$, (A.11) yields that by construction,

$$\frac{\widetilde{\pi}^*_{0,j}}{1-\widetilde{\pi}^*_{0,j}} = \check{\phi}_j \frac{\mathbb{P}_0(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\})}{\mathbb{P}_0(R^m_T = s_{j+1} \mid R^m_T \in \{s_j, s_{j+1}\})},$$

so that under Bayesian updating for $t = 1, \ldots, T$, the risk-neutral probabilities $\{\widetilde{\pi}^*_{0,j}\}$ follow a stream that would be consistent with RE but with the transformation $\check{\phi}_j$ with respect to the physical probabilities (rather than ϕ_j, as would be the case with a correct prior).

We now have two cases we must consider. We begin with case (ii), in which $\check{\phi}_j \geqslant 1$ so that $\widetilde{\pi}^*_{0,j} \geqslant \mathbb{P}_0(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\})$. In this case, the proof of Proposition 1 applies for $\{\widetilde{\pi}^*_{t,j}\}$ with respect to $\check{\phi}_j$, so that

$$\begin{aligned}\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] &\leqslant \widetilde{\pi}^{*2}_{0,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0,j} + \check{\phi}_j(1-\widetilde{\pi}^*_{0,j})}\right) \\ &\leqslant \widetilde{\pi}^{*2}_{0,j},\end{aligned}$$

so the stated claim holds in this case.

Now, for case (i), if the prior distortion is such that $\check{\phi}_j \in (0,1)$, then the proof of Proposition 1 no longer applies, since it used that $\Delta_{0,j} \equiv \widetilde{\pi}^*_{0,j} - \mathbb{P}_0(R^m_T = s_j \mid R^m_T \in \{s_j, s_{j+1}\}) \geqslant 0$ from $\phi_j \geqslant 1$ whereas now we have this $\Delta_{0,j} < 0$. But we note that we can use the following rearrangement of (A.8) to substitute into (A.7) rather than using (A.8) as in that proof:

$$\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_j] = (1-\widetilde{\pi}^*_{0,j}) - \frac{1-\widetilde{\pi}^*_{0,j}}{\widetilde{\pi}^*_{0,j}}\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}],$$

which, used in (A.7), yields

$$\begin{aligned}\mathbb{E}_0[m^*_j - r^*_j \mid R^m_T \in \{s_j, s_{j+1}\}] &= -\Delta_{0,j}(1-\widetilde{\pi}^*_{0,j}) + \Delta_{0,j}\left(1 + \frac{1-\widetilde{\pi}^*_{0,j}}{\widetilde{\pi}^*_{0,j}}\right)\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}] \\ &\leqslant -\Delta_{0,j}(1-\widetilde{\pi}^*_{0,j}),\end{aligned}$$

since $\Delta_{0,j}\left(1 + \frac{1-\widetilde{\pi}^*_{0,j}}{\widetilde{\pi}^*_{0,j}}\right)\mathbb{E}^\diamond_0[m^*_j \mid R^m_T = s_{j+1}] \leqslant 0$. Then using (A.10), we obtain

$$\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] \leqslant \widetilde{\pi}^*_{0,j}\left(\frac{1}{\widetilde{\pi}^*_{0,j} + \check{\phi}_j(1-\widetilde{\pi}^*_{0,j})} - 1\right)(1-\widetilde{\pi}^*_{0,j}). \tag{A.12}$$

We now have two subcases to consider. If $\widetilde{\pi}^*_{0,j} < 1/2$, then

$$\widetilde{\pi}^*_{0,j}\left(\frac{1}{\widetilde{\pi}^*_{0,j} + \check{\phi}_j(1-\widetilde{\pi}^*_{0,j})} - 1\right) = (1-\widetilde{\pi}^*_{0,j})\frac{\widetilde{\pi}^*_{0,j} + \check{\phi}_j\widetilde{\pi}^*_{0,j}}{\widetilde{\pi}^*_{0,j} + \check{\phi}_j(1-\widetilde{\pi}^*_{0,j})}$$

$$< (1 - \widetilde{\pi}^*_{0,j}),$$

and therefore, substituting into (A.12),

$$\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] \leqslant (1 - \widetilde{\pi}^*_{0,j})^2.$$

Meanwhile, if $\widetilde{\pi}^*_{0,j} \geqslant 1/2$, then

$$\begin{aligned}\widetilde{\pi}^*_{0,j}\left(\frac{1}{\widetilde{\pi}^*_{0,j} + \check{\phi}_j(1 - \widetilde{\pi}^*_{0,j})} - 1\right) &= \widetilde{\pi}^*_{0,j}\frac{(1 - \widetilde{\pi}^*_{0,j}) + \check{\phi}_j(1 - \widetilde{\pi}^*_{0,j})}{\widetilde{\pi}^*_{0,j} + \check{\phi}_j(1 - \widetilde{\pi}^*_{0,j})} \\ &\leqslant \widetilde{\pi}^*_{0,j},\end{aligned}$$

and therefore in this case

$$\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] \leqslant \widetilde{\pi}^*_{0,j}(1 - \widetilde{\pi}^*_{0,j}) \leqslant \widetilde{\pi}^{*2}_{0,j}.$$

Combining these two subcases yields that for case (i),

$$\widetilde{\mathbb{E}}_0[m^*_j - r^*_j] \leqslant \max\{\widetilde{\pi}^{*2}_{0,j}, (1 - \widetilde{\pi}^*_{0,j})^2\},$$

as stated, completing the proof. □

Proof of Proposition 4. Denote the upper bound for admissible excess movement in Proposition 1 (using the notation from Section 4.2) by $\mathcal{M}\colon [1,\infty) \times [0,1] \to [0,1]$, with

$$\mathcal{M}(\phi_{i,j}, \widetilde{\pi}^*_{0_i,i,j}) \equiv \widetilde{\pi}^{*2}_{0_i,i,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0_i,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{0_i,i,j})}\right).$$

As in the text,

$$\frac{\partial^2 \mathcal{M}}{\partial \phi^2_{i,j}} = -\frac{2\widetilde{\pi}^{*2}_{0_i,i,j}(1 - \widetilde{\pi}^*_{0_i,i,j})^2}{\left(\widetilde{\pi}^*_{0_i,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{0_i,i,j})\right)^3} \leqslant 0.$$

For any arbitrary realization of the prior, $\widetilde{\pi}^*_{0_i,i,j} = \varrho$, Jensen's inequality accordingly yields that

$$\widetilde{\mathbb{E}}\left[\mathcal{M}(\phi_{i,j}, \widetilde{\pi}^*_{0_i,i,j}) \,\middle|\, \widetilde{\pi}^*_{0_i,i,j} = \varrho\right] \leqslant \mathcal{M}\left(\widetilde{\mathbb{E}}[\,\phi_{i,j} \mid \widetilde{\pi}^*_{0_i,i,j} = \varrho\,],\, \varrho\right).$$

Thus with $\widetilde{\mathbb{E}}[\mathcal{M}(\phi_{i,j}, \widetilde{\pi}^*_{0_i,i,j}) \mid \widetilde{\pi}^*_{0_i,i,j}] \leqslant \mathcal{M}(\widetilde{\mathbb{E}}[\phi_{i,j} \mid \widetilde{\pi}^*_{0_i,i,j}],\, \widetilde{\pi}^*_{0_i,i,j})$, we have that

$$\begin{aligned}\widetilde{\mathbb{E}}[m^*_{i,j} - r^*_{i,j}] \leqslant \widetilde{\mathbb{E}}\left[\mathcal{M}(\phi_{i,j}, \widetilde{\pi}^*_{0_i,i,j})\right] &\leqslant \widetilde{\mathbb{E}}\left[\mathcal{M}\left(\widetilde{\mathbb{E}}[\phi_{i,j} \mid \widetilde{\pi}^*_{0_i,i,j}],\, \widetilde{\pi}^*_{0_i,i,j}\right)\right] \\ &\leqslant \widetilde{\mathbb{E}}\left[\mathcal{M}\left(\overline{\phi}_j, \widetilde{\pi}^*_{0_i,i,j}\right)\right],\end{aligned}$$

where $\overline{\phi}_j \equiv \max_{\widetilde{\pi}_{0_i,i,j}} \widetilde{\mathbb{E}}[\phi_{i,j} \mid \widetilde{\pi}_{0_i,i,j}]$, and where the first line uses Proposition 1 and applies the law of iterated expectations and the second uses $\partial \mathcal{M} / \partial \phi_{i,j} \geqslant 0$. Substituting the definition of $\mathcal{M}(\cdot,\cdot)$ into this inequality yields Proposition 4. □

Proof of Proposition 5. Using $(V^m_{j+1} - V^m_j)/V^m_j = (V^m_0 s_{j+1} - V^m_0 s_j)/(V^m_0 s_j) = (s_{j+1} - s_j)/s_j \equiv \Delta_j$, the result then follows immediately from equation (6), with V^m_j and V^m_{j+1} replacing C_{low} and C_{high}, respectively. □

Proof of Proposition 6. Starting with belief movement,

$$\begin{aligned}
\widetilde{\mathbb{E}}[\widehat{m}^*_{t,t+1,j}] &= \widetilde{\mathbb{E}}[(\widehat{\pi}^*_{t+1,j} - \widehat{\pi}^*_{t,j})^2] \\
&= \widetilde{\mathbb{E}}\left[\left((\widetilde{\pi}^*_{t+1,j} - \widetilde{\pi}^*_{t,j})^2 + (\epsilon_{t+1,j} - \epsilon_{t+1,j})\right)^2\right] \\
&= \widetilde{\mathbb{E}}[m^*_{t,t+1,j}] + 2\widetilde{\mathbb{E}}[\widetilde{\pi}^*_{t+1,j}\epsilon_{t+1,j} - \widetilde{\pi}^*_{t,j}\epsilon_{t+1,j} - \widetilde{\pi}^*_{t+1,j}\epsilon_{t,j} + \widetilde{\pi}^*_{t,j}\epsilon_{t,j}] + \widetilde{\mathbb{E}}[(\epsilon_{t+1,j} - \epsilon_{t,j})^2] \\
&= \widetilde{\mathbb{E}}[m^*_{t,t+1,j}] + \widetilde{\mathbb{E}}[\epsilon^2_{t,j} + \epsilon^2_{t+1,j}].
\end{aligned}$$

For uncertainty resolution,

$$\begin{aligned}
\widetilde{\mathbb{E}}[\widehat{r}^*_{t,t+1,j}] &= \widetilde{\mathbb{E}}[(\widetilde{\pi}^*_{t,j} + \epsilon_{t,j})(1 - \widetilde{\pi}^*_{t,j} - \epsilon_{t,j}) - (\widetilde{\pi}^*_{t+1,j} + \epsilon_{t+1,j})(1 - \widetilde{\pi}^*_{t+1,j} - \epsilon_{t+1,j})] \\
&= \widetilde{\mathbb{E}}[\widetilde{r}^*_{t,t+1,j}] + \widetilde{\mathbb{E}}[\epsilon^2_{t+1,j} - \epsilon^2_{t,j}].
\end{aligned}$$

Combining these two, with $\text{Var}(\epsilon_{t,j}) \equiv \widetilde{\mathbb{E}}[(\epsilon_{t,j} - \widetilde{\mathbb{E}}[\epsilon_{t,j}])^2] = \widetilde{\mathbb{E}}[\epsilon^2_{t,j}]$,

$$\widetilde{\mathbb{E}}[\widehat{m}^*_{t,t+1,j} - \widehat{r}^*_{t,t+1,j}] = \widetilde{\mathbb{E}}[m^*_{t,t+1,j} - r^*_{t,t+1,j}] + 2\text{Var}(\epsilon_{t,j}).$$ □

Appendix B. Additional Material

B.1. Risk-Neutral Beliefs and Discount Rates

We again work in the context of the example in Section 2 for simplicity of exposition. The price of the terminal consumption claim is given in equilibrium in by $P_t(C_T) = \mathbb{E}_t\left[\beta_t^{T-t}\frac{U'(C_T)}{U'(C_t)}C_T\right]$, where we have relaxed the assumption of no discounting and β_t is now the agent's (possibly time-varying) time discount factor. Defining the gross return $R^C_{t,T} \equiv \frac{C_T}{P_t(C_T)}$, rearranging this equation

for $P_t(C_T)$ yields

$$\begin{aligned}\mathbb{E}_t[R^C_{t,T}] &= \frac{1 - \text{Cov}_t\left(\beta_t^{T-t}\frac{U'(C_T)}{U'(C_T)}, C_T\right)}{\mathbb{E}_t\left[\beta_t^{T-t}\frac{U'(C_T)}{U'(C_t)}\right]} \\ &= \frac{\frac{U'(C_t)}{\beta_t^{T-t}} - \text{Cov}_t(U'(C_T), C_T)}{\mathbb{E}_t[U'(C_T)]},\end{aligned} \tag{B.1}$$

as usual. For full concreteness, we can write $\mathbb{E}_t[U'(C_T)] = \pi_t U'(C_{\text{low}}) + (1-\pi_t)U'(C_{\text{high}})$ in our two-state example, and $\text{Cov}_t(U'(C_T), C_T)$ can be similarly rewritten as a function of π_t, C_T, and $U'(C_T)$. This decomposition makes clear that intertemporal discount-rate variation can arise from four sources:

1. Changes in the time discount factor β_t.
2. Changes in contemporaneous marginal utility $U'(C_t)$.
3. Changes in the relative probability π_t.
4. Changes in state-contingent terminal consumption C_i and/or state-contingent marginal utility $U'(C_i)$.

Our framework thus allows for discount-rate variation arising from the first three sources, but not the last one. One might not consider this to be particularly restrictive in the context of this example; in theory, we can *define* the states such that the realization of the state fully determines consumption and marginal utility. But when taken to the data, we define states by the return on the market index, in which case this does become more restrictive. (We in fact slightly relax these assumptions and allow for independent consumption-growth or marginal-utility shocks for a given return state; Section 3.2 more fully discusses the models covered by our assumptions.)

Now consider the specification of the example in which the deterministic consumption stream for $t < T$ is given by $(C_0, C_1, C_2, C_3, \ldots, C_{T-1}) = (1, 1/2, 1, 1/2, \ldots)$ but that π_t is constant at $\pi_t = \pi_0 = 0.5$ for $t < T$, as on Section 2.3. Again assume for simplicity that $\beta = 1$. As noted there, because the mapping between π_t and π_t^* is one-to-one for a given ϕ in (4), measured risk-neutral beliefs would be constant for $t < T$ in this case: risk-neutral beliefs are invariant to changes in the risk-free rate arising from proportional changes to Arrow-Debreu state prices across the two states, as can be seen in equations (1)–(2), and all discount-rate changes for the consumption claim are in fact driven by the risk-free rate in this case. The gross $(T-t)$-period risk-free rate with $\beta = 1$ is $R^f_{t,T} = \frac{U'(C_t)}{\mathbb{E}_t[U'(C_i)]}$ in equilibrium; we can thus rewrite (B.1) as

$$\mathbb{E}_t[R^C_{t,T}] = R^f_{t,T} - \frac{\text{Cov}_t(U'(C_T), C_T)}{\mathbb{E}_t[U'(C_T)]}, \tag{B.2}$$

and the second term is constant for $t < T$ under the current assumptions. But we need not restrict ourselves to settings in which all discount-rate variation arises due to changes in the risk-free rate.

The previous specification of the example (on Section 2.3), in which $\pi_0 = 0.3$, $C_t = \overline{C} = 1$ for $t < T$ and $\pi_1 = 0$ or 0.6 with equal probability, has no equity premium at $t = 1$ if $\pi_1 = 0$ since pricing is risk-neutral in this case (given that there is no risk); meanwhile, if $\pi_1 = 0.6$, then $\mathbb{E}_1[R^C_{1,T}] > R^f_{1,T}$ since the second term in (B.2) is positive. So the framework is capable of achieving identification in cases in which both the risk-free rate and risk premia are time-varying.

More generally, this example shows that the framework can handle cases in which an object that can be intuitively thought of as the quantity of aggregate risk is time-varying. As in Hansen and Jagannathan (1991), the conditional risk premium on any asset depends on the conditional volatility of the stochastic discount factor, which in this case is given for the horizon $T - t$ by $\text{Var}_t(\beta^{T-t}U'(C_T)/U'(C_t))$; we could rewrite (B.2) in terms of this value if desired. In the current example, this value is again equal to 0 at $t = 1$ if $\pi_1 = 0$, while it is positive if $\pi_1 = 0.6$. Further, while relative risk aversion (and thus the aggregate "price" of risk) is constant in the current example, nothing about the example restricts utility to take this form; we could, e.g., specify exponential utility and thus obtain time-varying relative risk aversion, and the analysis in Section 2.3 and here would nonetheless apply as well with slight modification.

B.2. Description of Gabaix (2012) Rare-Disasters Model for Example 3

Assume a representative agent with CRRA consumption utility, and assume that log consumption $c_t \equiv \log(C_t)$ and log dividends $d_t \equiv \log(D_t)$ evolve respectively according to

$$c_{t+1} = c_t + g_c + \varepsilon^c_{t+1} + \log(B_{t+1})\mathbb{1}\{\text{disaster}_{t+1}\},$$
$$d_{t+1} = d_t + g_d + \varepsilon^d_{t+1} + \log(F_{t+1})\mathbb{1}\{\text{disaster}_{t+1}\},$$

where $(\varepsilon^c_{t+1}, \varepsilon^d_{t+1})'$ is i.i.d. bivariate normal with mean zero and arbitrary covariance and is independent of all disaster-related variables,[55] and B_{t+1} and F_{t+1} are arbitrarily correlated random variables with support $[0, 1]$ (or some discretization thereof) that affect consumption and dividends respectively in the case of a disaster in period $t + 1$, which occurs with probability p_t. Define *resilience* H_t according to $H_t = p_t\mathbb{E}_t[B^{-\gamma}_{t+1}F_{t+1} - 1 \mid \mathbb{1}\{\text{disaster}_{t+1}\}]$, write $H_t = H_* + \widehat{H}_t$, and assume that the variable part follows

$$\widehat{H}_{t+1} = \frac{1 + H_*}{1 + H_t}e^{-\phi_H}\widehat{H}_t + \varepsilon^H_{t+1},$$

where $\mathbb{E}_t[\varepsilon^H_{t+1}] = 0$ and this shock is independent from all other shocks. Then the statements in Example 3 follow.

[55] To be complete with respect to our discrete-state setting, we can assume $(\varepsilon^c_{t+1}, \varepsilon^d_{t+1})'$ is in fact an appropriately discretized normal distribution (e.g., a shifted binomial distribution).

B.3. Description of Campbell–Cochrane (1999) Habit-Formation Model for Example 4

Assume a representative agent with utility $\mathbb{E}_0\{\sum_{t=0}^{\infty}\beta^t[(C_t - H_t)^{1-\gamma} - 1]/(1-\gamma)\}$, where C_t is consumption and H_t is the level of habit, taken as exogenous by the agent. Defining the *surplus-consumption ratio* $S_t^c \equiv (C_t - H_t)/H_t$, assume that $s_t^c \equiv \log(S_t^c)$, $c_t \equiv \log(C_t)$, and log dividends $d_t \equiv \log(D_t)$ evolve respectively according to

$$\begin{aligned} s_{t+1}^c &= (1-\phi)\bar{s}^c + \phi s_t^c + \lambda(s_t^c)\varepsilon_{t+1}, \\ c_{t+1} &= g + c_t + \varepsilon_{t+1}, \\ d_{t+1} &= g + d_t + \eta_{t+1}, \end{aligned}$$

where $\varepsilon_{t+1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2)$ (see Footnote 55), $\eta_{t+1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma_\eta^2)$, $\text{Corr}(\varepsilon_{t+1}, \eta_{t+1}) = \rho$, and the *sensitivity function* $\lambda(s_t^c)$ is specified as

$$\lambda(s_t^c) = \left[\frac{1}{\overline{S}^c}\sqrt{1 - 2(s_t^c - \bar{s}^c)} - 1\right] \mathbb{1}\{s_t^c \leqslant s_{\max}^c\},$$

where $\overline{S}^c = \exp(\bar{s}^c) = \sigma_\varepsilon\sqrt{\gamma/(1-\phi)}$ is the assumed steady-state surplus-consumption ratio and $s_{\max}^c = \bar{s}^c + (1 - \overline{S}^c)^2/2$. Then the statement in Example 4 follows.

B.4. Measurement of Risk-Neutral Distribution

Before detailing measurement of the risk-neutral distribution, we note that we must collect additional data in order to follow the procedure below. In particular, OptionMetrics reports a risk-free zero-coupon yield curve across multiple maturities, as well as the underlying end-of-day S&P 500 index price. We use the risk-free rate at the relevant horizon as an input in our measurement of risk-neutral beliefs, and we use the index price to observe the ex-post return state for each option expiration date T_i and assign probability 1 to that state on date T_i. But the settlement value for many S&P 500 options in fact reflects the opening (rather than closing) price on the expiration date; for example, the payoff for the traditional monthly S&P 500 option contract expiring on the third Friday of each month depends on the opening S&P index value on that third Friday morning, while the payoff for the more recently introduced end-of-month option contract depends on the closing S&P index value on the last business day of the month.[56] To obtain the ex-post return state for A.M.-settled options, we hand-collect the option settlement values for these expiration dates from the Chicago Board Options Exchange (CBOE) website, which posts these values.

Then, as introduced in Section 5.2, we measure the risk-neutral distribution for returns by applying the following steps to the observed option-price cross-sections, following Malz (2014):

1. Transform the collections of call- and put-price cross-sections (for example, for call options on

[56] See http://www.cboe.com/SPX for further detail. For our dataset, 441 of the 685 option expiration dates correspond to A.M.-settled options.

date t for expiration date T_i, this set is $\{q^m_{t,i,K}\}_{K\in\mathcal{K}}$) into Black–Scholes implied volatilities.

2. Fit a cubic spline to interpolate a smooth function between the points in the resulting implied-volatility schedule for each trading date–expiration date pair (separately for the call- and put-option values). The spline is *clamped*: its boundary conditions are that the slope of the spline at the minimum and maximum values of the knot points $\{q^m_{t,i,K}\}_{K\in\mathcal{K}}$ is equal to 0; further, to extrapolate outside of the range of observed knot points, set the implied volatilities for those unobserved strikes equal to the implied volatility for the closest observed strike (i.e., maintain a slope of 0 for the implied-volatility schedule outside the observed range).

3. Evaluate this spline (separately for calls and puts) at 1,901 strike prices, for S&P index values ranging from 200 to 4,000 (so that the evaluation strike prices are $K = 200, 202, \ldots, 4000$), to obtain a set of implied-volatility values across this fine grid of possible strike prices.[57]

4. Average the separate call- and put-option implied-volatility values from the previous step at each strike for each (t, T_i) pair, to obtain a single implied-volatility schedule across strikes for each such (t, T_i) pair. (Given put-call parity, the implied-volatility values for calls and puts should in theory be equal at a given strike; in practice, they tend to differ slightly given market microstructure issues, so using the mean of the two values is a simple way of averaging out the effects of such idiosyncratic noise. This step is the only point of distinction between our procedure and that of Malz, who assumes access to a single implied-volatility schedule and thus does not consider call and put prices separately.)

5. Invert the single resulting smoothed 1,901-point implied-volatility schedule for each (t, T_i) pair to transform these values back into call prices, and denote this fitted call-price schedule as $\{\hat{q}^m_{t,i,K}\}_{K\in\{200,202,\ldots,4000\}}$.

6. Calculate the risk-neutral CDF for the date-T_i index value at strike price K using $\mathbb{P}^*_t(V^m_{T_i} < K) = 1 + R^f_{t,T_i}(\hat{q}^m_{t,i,K} - \hat{q}^m_{t,i,K-2})/2$. (See the proof of equation (14) in Appendix A for a derivation of this result; the index-value distance between the two adjacent strikes is equal to 2 given that we evaluate the spline at intervals of two index points.)

7. Defining $V^m_{i,j,\max}$ and $V^m_{i,j,\min}$ to be the date-T_i index values corresponding to the upper and lower bounds, respectively, of the bin defining return-state s_j,[58] we then calculate the risk-neutral probability that return state s_j will be realized at date T_i, referred to with slight notational abuse as $\mathbb{P}^*_t(s_j)$, as

$$\mathbb{P}^*_t(s_j) = \mathbb{P}^*_t(V^m_{T_i} < V^m_{i,j,\max}) - \mathbb{P}^*_t(V^m_{T_i} < V^m_{i,j,\min}),$$

where the CDF values are taken from the previous step using linear interpolation between whichever two strike values $K \in \{200, 202, \ldots, 4000\}$ are nearest to $V^m_{i,j,\max}$ and $V^m_{i,j,\min}$, respec-

[57] This set of ∼1,900 strike prices is on average about 20 times larger than the set of strikes for which there are prices in the data, as there is a mean of roughly 94 observed values in a typical set $\{q^m_{t,i,K}\}_{K\in\mathcal{K}}$ (and similarly for put options), using the numbers given in Section 5.1.

[58] That is, formally, $V^m_{i,j,\min} = R^f_{0_i,T_i} V^m_{T_0} \exp(s_j - 0.01)$ and $V^m_{i,j,\max} = R^f_{0_i,T_i} V^m_{0_i} \exp(s_j + 0.01)$. For example, for excess return state s_2, we have $V^m_{i,j,\min} = R^f_{0_i,T_i} V^m_{0_i} \exp(-0.10)$ and $V^m_{i,j,\max} = R^f_{0_i,T_i} V^m_{T_0} \exp(-0.08)$.

tively.

Note that we transform the option prices into Black–Scholes implied volatilities simply for purposes of fitting the cubic spline and then transform these implied volatilities back into call prices before calculating risk-neutral beliefs, so this procedure does *not* require the Black–Scholes model to be correct.[59] The clamped cubic spline proposed by Malz (2014), and used in step 2 above, is chosen to ensure that the call-price schedule obtained in step 5 is decreasing and convex with respect to the strike price outside the range of observable strike prices, as required under the restriction of no arbitrage. Violations of these restrictions *inside* the range of observable strikes, as observed infrequently in the data, generate negative implied risk-neutral probabilities; in any case that this occurs, we set the associated risk-neutral probability to 0.

As noted in step 2, the clamped spline is an *interpolating* spline, as it is restricted to pass through all the observed data points so that the fitted-value set $\{\hat{q}^m_{t,i,K}\}$ contains the original values $\{q^m_{t,i,K}\}$. Some alternative methods for measuring risk-neutral beliefs use smoothing splines that are not constrained to exhibit such interpolating behavior. To check the robustness of our results to the choice of measurement technique, we have accordingly used one such alternative method proposed by Bliss and Panigirtzoglou (2004). Empirical results obtained using risk-neutral beliefs calculated in this alternative manner are essentially unchanged as compared to the benchmark results in Section 5.6, and are available upon request.

We have also conducted robustness tests with respect to the fineness of the grid on which we evaluate the spline in step 4 and calculate the risk-neutral CDF in step 6, with results from these exercises indistinguishable from the benchmark results; these additional results are also available upon request.

B.5. Block Bootstrap Procedure and Asymptotic Validity

Our block-bootstrap method for inference, as introduced in Section 5.4, proceeds as follows:

1. Recalling that we observe data for $T_N + 1$ trading dates, we first define B blocks of $(T_N + 1)/B$ trading dates each: the first block (b_1) contains values for t of $\mathcal{T}_1 = \{0, 1, \ldots, (T_N + 1)/B - 1\}$, the second ($b_2$) contains values $\mathcal{T}_2 = \{(T_N + 1)/B, \ldots, 2(T_N + 1)/B - 1\}$, and so on.[60] In the case that $(T_N + 1)/B$ is non-integer-valued, we set the lengths of the first $B - 1$ blocks to $\lceil (T_N + 1)/B \rceil$, where $\lceil \cdot \rceil$ denotes the ceiling or least-greater-integer function, and the last block is correspondingly smaller and contains the remaining points.

2. We then divide the observations in our sample, which we now write as $\{(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,i,j})_j\}_i$, into these constructed blocks according to the block in which expiration date T_i appears for

[59] We conduct this transformation following Malz (2014), as well as much of the related literature, which argues that these smoothing procedures tend to perform slightly better in implied-volatility space than in the option-price space given the convexity of option-price schedules; see Malz (1997) for a discussion.

[60] We construct blocks using trading dates rather than expiration-date indices i since the expiration dates are unevenly spaced, and we accordingly expect that the dependence structure in the data is more closely related to the date than the index value for the expiration date.

each set $(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,i,j})_j$; that is, $b_\ell = \{(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,i,j})_j : T_i \in \mathcal{T}_\ell\}$ for $\ell = 1, 2, \ldots, B$.

3. We construct $\mathscr{S}$ resampled datasets; we typically set $\mathscr{S} = 5000$. For each resampled dataset, we randomly draw B complete blocks of data with replacement from the set of blocks constructed in the previous steps and paste them together to form our new sample. That is, denoting the first resampled block for a given round $b^\dagger_1$, the second $b^\dagger_2$, and so on, this resampled set of observations is $(b^\dagger_1, b^\dagger_2, \ldots, b^\dagger_B)$.

4. For each resampled dataset, re-estimate the lower bounds for $\{\overline{\phi}_j\}_j$ and $\overline{\phi}$ as discussed in the main text (see equations (21)–(24)).

5. Construct a lower bound for a $100(1-\alpha)\%$ confidence interval for each value in $\{\overline{\phi}_j\}_j$ and for $\overline{\phi}$ as the α quantile of the bootstrap distribution of estimates for the parameter of interest from the resampled datasets above. This yields a confidence interval for, e.g., $\overline{\phi}_j$ given by $\text{CI}_j = [\mathcal{Q}^\dagger_{j,\alpha}, \infty)$, where $\mathcal{Q}^\dagger_{j,\alpha}$ is the α quantile of the bootstrap distribution of the estimates of the lower bound for $\overline{\phi}_j$. We refer to $\mathcal{Q}^\dagger_{j,\alpha}$ as $\text{CI}_{\text{LB},D}$, where $D = \lceil (T_N + 1)/B \rceil$ is the block length and "LB" refers to the lower bound of the associated confidence interval. We typically set $\alpha = 0.05$ to obtain 95 percent confidence intervals.

Note that the groupings of return-state pairs (indexed by $j = 1, \ldots, J-1$) are fully preserved in this resampling procedure for each set of observations indexed by i (corresponding to the option expiration date) within each block, as we split the observations into blocks only by time and not by return states. (This motivates the slight change of notation introduced in part 2 above.) We do so in order to obtain valid inference for the aggregate value $\overline{\phi}$, which uses observations for state pairs $(s_2, s_3), \ldots, (s_{J-2}, s_{J-1})$, in the face of arbitrary dependence for the observations across those state pairs and a fixed number of return states J (whereas we assume $N \to \infty$, and further $B \to \infty$ according to a sequence such that $(T_N + 1)/B \to \infty$). In this way our procedure is in fact a *panel* (or *cluster*) *block bootstrap*; see, for example, Palm, Smeekes, and Urbain (2011).

Lahiri (2003, Theorem 3.2) then provides a weak condition on the strong mixing coefficient of the relevant stochastic process — in our case, $\{(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,i,j})_j\}_i$ — under which the blocks are asymptotically independent and the bootstrap distribution estimator is consistent for the true distribution under the asymptotics above, so that our confidence intervals from step 5 have asymptotic coverage probability of at least 95% for the population parameters of interest in the presence of nearly arbitrary (stationary) autocorrelation and heteroskedasticity.[61] This coverage rate may

[61] There are additional conditions required for the result of Lahiri (2003, Theorem 3.2) to hold, but they will hold trivially in our context under the RE null given the boundedness of the relevant belief statistics. Our block bootstrap is a non-overlapping block bootstrap (NBB); others (Künsch, 1989; Liu and Singh, 1992) have proposed a *moving* block bootstrap (MBB) using overlapping blocks, among other alternatives. While the MBB has efficiency gains relative to the NBB (Hall, Horowitz, and Jing, 1995), these are "likely to be very small in applications" (Horowitz, 2001, p. 3190), so we use the NBB for computational convenience.

We note also that the confidence intervals constructed in step 5 apply the so-called *percentile method* (Efron, 1979), which has well-known issues in the presence of asymmetries in the finite-sample distribution for the relevant statistic (Schenker, 1985; Hall, 1988; cf. the Percentile Interval Lemma of Efron and Tibshirani, 1993, p. 173, for a rejoinder). We would instead ideally construct the lower bound of the interval as $\widehat{\phi}_j - \mathcal{Q}^{\dagger\dagger}_{j,1-\alpha}$, where $\mathcal{Q}^{\dagger\dagger}_{j,1-\alpha}$ is the $1-\alpha$ quantile of the distribution of $\widehat{\phi}^\dagger_j - \widehat{\phi}_j$, where $\widehat{\phi}^\dagger_j$ is the bootstrap estimate. But in many cases in our estimation this $1-\alpha$ quantile is

in fact be greater than 95% given that we are estimating lower bounds for the parameters of interest rather than the parameters themselves, and this motivates our use of one-sided rather than two-sided confidence intervals, as discussed in Section 5.4.

B.6. Details of Estimation of Measurement Error

Main estimation approach. For given $\phi_{i,j}$ and $\widetilde{\pi}^*_{t,i,j}$, we have from (18) that

$$\widetilde{\pi}_{t,i,j} = \frac{\widetilde{\pi}^*_{t,i,j}}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{t,i,j})}.$$

Thus under RE, conditional on $\widetilde{\pi}^*_{T_i,i,j} \in \{0,1\}$, we have as of time t that

$$\widetilde{\pi}^*_{T_i,i,j} = \widetilde{\pi}_{T_i,i,j} = \begin{cases} 1, & \text{with probability } \frac{\widetilde{\pi}^*_{t,i,j}}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1-\widetilde{\pi}^*_{t,i,j})} \\ 0, & \text{with probability } \frac{\phi_{i,j}(1-\widetilde{\pi}^*_{t,i,j})}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1-\widetilde{\pi}^*_{t,i,j})}. \end{cases}$$

This yields expected $t \to T_i$ movement of

$$\begin{aligned} \widetilde{\mathbb{E}}_t[m^*_{t\to T_i,j}] &= \widetilde{\mathbb{E}}_t[(\widetilde{\pi}^*_{T_i,i,j} - \widetilde{\pi}^*_{t,i,j})^2] \\ &= \widetilde{\mathbb{E}}_t[\widetilde{\pi}^{*2}_{T_i,i,j}] - 2\widetilde{\mathbb{E}}_t[\widetilde{\pi}^*_{T_i,i,j}]\widetilde{\pi}^*_{t,i,j} + \widetilde{\pi}^{*2}_{t,i,j} \\ &= \frac{\widetilde{\pi}^*_{t,i,j}}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{t,i,j})} - 2\frac{\widetilde{\pi}^*_{t,i,j}}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{t,i,j})}\widetilde{\pi}^*_{t,i,j} + \widetilde{\pi}^{*2}_{t,i,j} \\ &= \frac{\widetilde{\pi}^*_{t,i,j}(1 - 2\widetilde{\pi}^*_{t,i,j})}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{t,i,j})} + \widetilde{\pi}^{*2}_{t,i,j}. \end{aligned}$$

Further, given $\widetilde{\pi}^*_{T_i,i,j} \in \{0,1\}$, expected resolution is $\widetilde{\mathbb{E}}_t[r^*_{t\to T_i,j}] = (1 - \widetilde{\pi}^*_{t,i,j})\widetilde{\pi}^*_{t,i,j}$. Thus

$$\begin{aligned} \widetilde{\mathbb{E}}_t[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}] &= \frac{\widetilde{\pi}^*_{t,i,j}(1 - 2\widetilde{\pi}^*_{t,i,j})}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{t,i,j})} + \widetilde{\pi}^*_{t,i,j}(2\widetilde{\pi}^*_{t,i,j} - 1) \\ &= \frac{\widetilde{\pi}^*_{t,i,j}(1 - \widetilde{\pi}^*_{t,i,j})(2\widetilde{\pi}^*_{t,i,j} - 1)(\phi_{i,j} - 1)}{\widetilde{\pi}^*_{t,i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{t,i,j})}, \end{aligned}$$

as stated in (25).

Alternative estimation approach. We note first that given $\widehat{\pi}^*_{t,i,j} = \widetilde{\pi}^*_{t,i,j} + \epsilon_{t,i,j}$, the unconditional variance (or zeroth autocovariance) and the first autocovariance of the observed beliefs are given

equal to ∞, so we instead apply the method in step 5. This method is nonetheless first-order correct as per the result of Lahiri (2003) given a normal (and thus symmetric) limiting distribution around $\phi_j < \infty$, as would be obtained under the RE null, though it does not yield any asymptotic refinement (as should be expected given that it uses a non-pivotal statistic; see Horowitz, 2001).

respectively by

$$\begin{aligned}\mathrm{Var}(\widehat{\pi}^*_{t,i,j}) &= \mathrm{Var}(\widetilde{\pi}^*_{t,i,j}) + \mathrm{Var}(\epsilon_{t,i,j}),\\ \mathrm{Cov}(\widehat{\pi}^*_{t,i,j}, \widehat{\pi}^*_{t-1,i,j}) &= \mathrm{Cov}(\widetilde{\pi}^*_{t,i,j}, \widetilde{\pi}^*_{t-1,i,j}).\end{aligned}$$

Given $\mathrm{Cov}(\widetilde{\pi}^*_{t,i,j}, \widetilde{\pi}^*_{t-1,i,j}) \leqslant \mathrm{Var}(\widetilde{\pi}^*_{t,i,j})$, we can thus obtain a bound for $\mathrm{Var}(\epsilon_{t,i,j})$ using only the observed data as follows:

$$\begin{aligned}\mathrm{Var}(\widehat{\pi}^*_{t,i,j}) - \mathrm{Cov}(\widehat{\pi}^*_{t,i,j}, \widehat{\pi}^*_{t-1,i,j}) &= \mathrm{Var}(\widetilde{\pi}^*_{t,i,j}) - \mathrm{Cov}(\widetilde{\pi}^*_{t,i,j}, \widetilde{\pi}^*_{t-1,i,j}) + \mathrm{Var}(\epsilon_{t,i,j})\\ &\geqslant \mathrm{Var}(\epsilon_{t,i,j}).\end{aligned}$$

We accordingly construct our conservative estimate of the measurement-error variance as

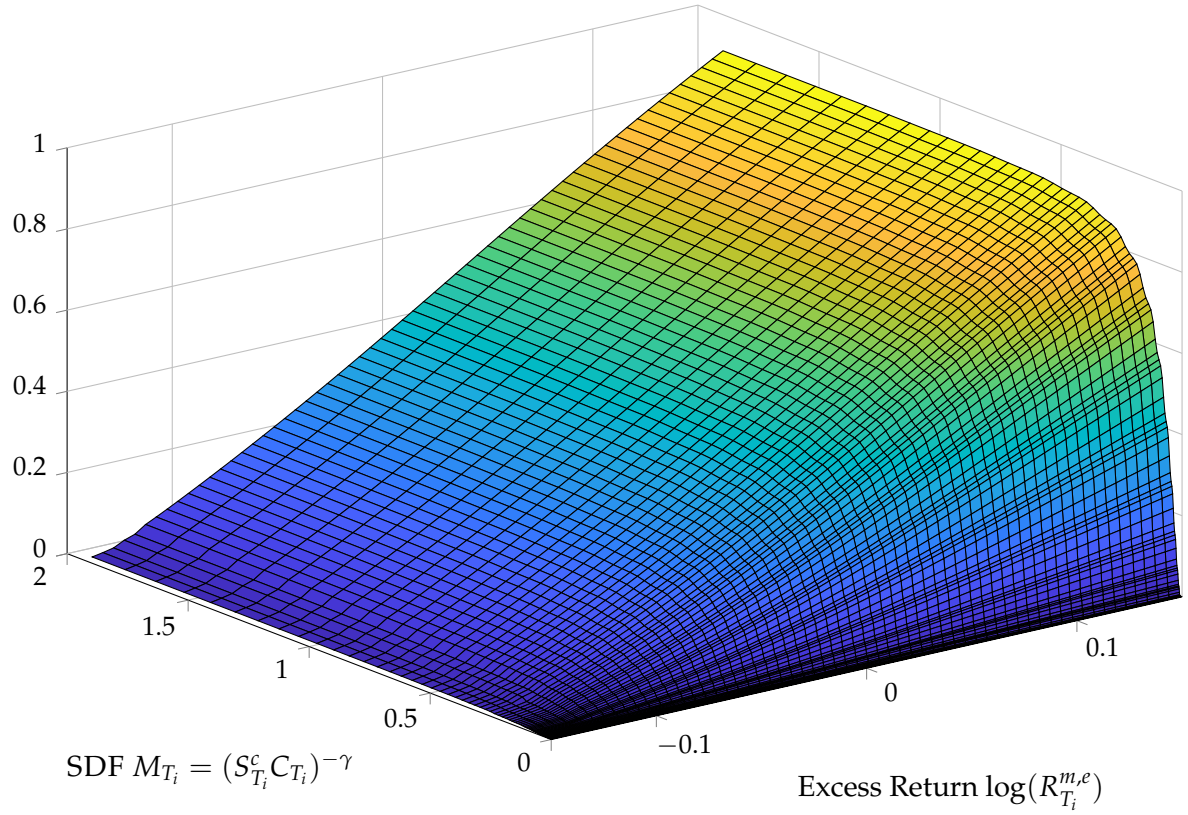
$$\widehat{\mathrm{Var}}(\epsilon_{t,i,j}) = \widehat{\mathrm{Var}}(\widehat{\pi}^*_{t,i,j}) - \widehat{\mathrm{Cov}}(\widehat{\pi}^*_{t,i,j}, \widehat{\pi}^*_{t-1,i,j}),$$

which requires measuring only the variance and first autocovariance of the observed conditional risk-neutral beliefs process (again see Footnote 37). We construct this estimate separately for risk-neutral beliefs corresponding to each return-state pair (s_j, s_{j+1}), as it may be the case that risk-neutral beliefs are measured more accurately at different points of the distribution.

B.7. Details of Solution Method and Simulations for Model in Section 6.4

See Appendix B.3 for a description of the model, and the calibrated parameters are identical to those used by Campbell and Cochrane (1999, Table 1), converted to daily values, for the version of their model with imperfectly correlated consumption and dividends. The initial market index value is normalized to $V^m_{0_i} = 1$, and the joint CDF for the SDF realization and the return as a function of the current surplus-consumption state is then solved by iterating backwards from T_i: after solving the model for the price-dividend ratio as a function of the surplus-consumption value, we then calculate the $T_i - 1$ CDF for any possible surplus-consumption value by integrating over the distributions of shocks to consumption (and thus surplus consumption) and dividends at T_i; we then project this CDF onto an interpolating cubic spline over the three dimensions $(S^c_{T_i-1}, M_{T_i}, \log(R^{m,e}_{T_i}))$; we then calculate the $T_i - 2$ CDF by integrating over the distribution of shocks at $T_i - 1$ *and* the projection solutions for the conditional distribution functions for $(T_i - 1) \rightarrow T_i$ obtained in the previous step; and so on. These CDFs are then used for the model simulation results presented in Section 6.4, and see Figure B.1 for an example CDF arising from our calibration and solution procedure.

Figure B.1: Solution for Joint CDF for SDF and Return: Example at $S_t^c = \overline{S^c}$, $t = 0_i$



Notes: Vertical axis shows joint physical CDF for date-T_i realization of marginal utility (which is proportional to SDF) and excess return for market index value from 0_i to T_i, with CDF evaluated at date $t = 0_i$ at steady-state surplus-consumption ratio $S^c_{0_i} = \overline{S^c}$, or equivalently $\mathbb{P}_{0_i}(M_{T_i} \leqslant m, \log(R^{m,e}_{T_i}) \leqslant r \mid S^c_{0_i} = \overline{S^c})$ across values m and r. Model is described in Appendix B.3, and calibrated parameters are identical to those used by Campbell and Cochrane (1999, Table 1), converted to daily values, for the version of their model with imperfectly correlated consumption and dividends.

References

ADRIAN, T., E. ETULA, AND T. MUIR (2014): "Financial Intermediaries and the Cross-Section of Asset Returns," *Journal of Finance*, 69, 2557–2596.

ADRIAN, T. AND H. S. SHIN (2014): "Procyclical Leverage and Value-at-Risk," *Review of Financial Studies*, 27, 373–403.

ALGAN, Y., O. ALLAIS, W. J. DEN HAAN, AND P. RENDAHL (2014): "Solving and Simulating Models with Heterogeneous Agents and Aggregate Uncertainty," in *Handbook of Computational Economics*, ed. by K. Schmedders and K. L. Judd, Amsterdam: Elsevier, vol. 3, chap. 6, 277–324.

ALVAREZ, F. AND U. J. JERMANN (2005): "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth," *Econometrica*, 73, 1977–2016.

AMIHUD, Y. AND H. MENDELSON (1986): "Asset Pricing and the Bid-Ask Spread," *Journal of Financial Economics*, 17, 223–249.

ANDERSEN, T. G., T. BOLLERSLEV, AND F. X. DIEBOLD (2010): "Parametric and Nonparametric Volatility Measurement," in *Handbook of Financial Econometrics*, ed. by Y. Aït-Sahalia and L. P. Hansen, Amsterdam: Elsevier, vol. 1, chap. 2, 67–137.

ARROW, K. J. (1982): "Risk Perception in Psychology and Economics," *Economic Inquiry*, 20, 1–9.

AUGENBLICK, N. AND M. RABIN (2018): "Belief Movement, Uncertainty Reduction, and Rational Updating," *UC Berkeley-Haas and Harvard University Mimeo*.

BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): "Measuring Economic Policy Uncertainty," *Quarterly Journal of Economics*, 131, 1593–1636.

BAKSHI, G. AND F. CHABI-YO (2012): "Variance Bounds on the Permanent and Transitory Components of Stochastic Discount Factors," *Journal of Financial Economics*, 105, 191–208.

BANSAL, R. AND A. YARON (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.

BAR-HILLEL, M. (1980): "The Base-Rate Fallacy in Probabilistic Judgments," *Acta Psychologica*, 44, 211–233.

BARBERIS, N., R. GREENWOOD, L. JIN, AND A. SHLEIFER (2015): "X-CAPM: An Extrapolative Capital Asset Pricing Model," *Journal of Financial Economics*, 115, 1–24.

BARBERIS, N., A. SHLEIFER, AND R. VISHNY (1998): "A Model of Investor Sentiment," *Journal of Financial Economics*, 49, 307–343.

BARNDORFF-NIELSEN, O. E. AND N. SHEPHARD (2001): "Non-Gaussian Ornstein-Uhlenbeck-Based Models and Some of Their Uses in Financial Economics," *Journal of the Royal Statistical Society, Series B*, 63, 167–241.

BARRO, R. J. (2006): "Rare Disasters and Asset Markets in the Twentieth Century," *Quarterly Journal of Economics*, 121, 823–866.

BATES, D. S. (2003): "Empirical Option Pricing: A Retrospection," *Journal of Econometrics*, 116, 387–404.

BENJAMIN, D. J., A. BODOH-CREED, AND M. RABIN (2017): "The Dynamics of Base-Rate Neglect," *Unpublished Manuscript*.

BENJAMIN, D. J., M. RABIN, AND C. RAYMOND (2016): "A Model of Nonbelief in the Law of Large Numbers," *Journal of the European Economic Association*, 14, 515–544.

BIDDER, R. AND I. DEW-BECKER (2016): "Long-Run Risk Is the Worst-Case Scenario," *American Economic Review*, 106, 2494–2527.

BLACK, F. AND M. SCHOLES (1973): "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637.

BLISS, R. R. AND N. PANIGIRTZOGLOU (2004): "Option-Implied Risk Aversion Estimates," *Journal of Finance*, 59, 407–446.

BOLLEN, N. P. AND R. E. WHALEY (2004): "Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?" *Journal of Finance*, 59, 711–753.

BORDALO, P., N. GENNAIOLI, R. LA PORTA, AND A. SHLEIFER (2017): "Diagnostic Expectations and Stock Returns," *NBER Working Paper No. 23863*.

BOROVIČKA, J., L. P. HANSEN, AND J. A. SCHEINKMAN (2016): "Misspecified Recovery," *Journal of Finance*, 71, 2493–2544.

BREEDEN, D. T. AND R. H. LITZENBERGER (1978): "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621–651.

BROADIE, M., M. CHERNOV, AND M. JOHANNES (2009): "Understanding Index Option Returns," *Review of Financial Studies*, 22, 4493–4529.

BROWN, D. J. AND S. A. ROSS (1991): "Spanning, Valuation and Options," *Economic Theory*, 1, 3–12.

CAMPBELL, J. Y. (2003): "Consumption-Based Asset Pricing," in *Handbook of the Economics of Finance*, ed. by G. M. Constantinides, M. Harris, and R. M. Stulz, Amsterdam: Elsevier, vol. 1B, chap. 13, 801–885.

——— (2017): *Financial Decisions and Markets: A Course in Asset Pricing*, Princeton, NJ: Princeton University Press.

CAMPBELL, J. Y. AND J. H. COCHRANE (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107, 205–251.

CAMPBELL, J. Y. AND R. J. SHILLER (1987): "Cointegration and Tests of Present Value Models," *Journal of Political Economy*, 95, 1062–1088.

——— (1988): "Stock Prices, Earnings, and Expected Dividends," *Journal of Finance*, 43, 661–676.

CARLSTEIN, E. (1986): "The Use of Subseries Values for Estimating the Variance of a General Statistic from a Stationary Sequence," *Annals of Statistics*, 14, 1171–1179.

CHABI-YO, F. (2012): "Pricing Kernels with Stochastic Skewness and Volatility Risk," *Management Science*, 58, 624–640.

CHAMBERLAIN, G. AND M. ROTHSCHILD (1983): "Arbitrage, Factor Structure, and Mean-Variance Analysis on Large Asset Markets," *Econometrica*, 51, 1281–1304.

CHETTY, R. (2006): "A New Method of Estimating Risk Aversion," *American Economic Review*, 96, 1821–1834.

——— (2009): "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods," *Annual Review of Economics*, 1, 451–488.

CHO, T. (2018): "Turning Alphas into Betas: Arbitrage and Endogenous Risk," *London School of Economics Mimeo*.

CHORDIA, T., R. ROLL, AND A. SUBRAHMANYAM (2008): "Liquidity and Market Efficiency," *Journal of Financial Economics*, 87, 249–268.

CHRISTOFFERSEN, P., S. HESTON, AND K. JACOBS (2013): "Capturing Option Anomalies with a Variance-Dependent Pricing Kernel," *Review of Financial Studies*, 26, 1962–2006.

COCHRANE, J. H. (2011): "Presidential Address: Discount Rates," *Journal of Finance*, 66, 1047–1108.

COCHRANE, J. H. AND J. SAÁ-REQUEJO (2000): "Beyond Arbitrage: Good-Deal Asset Price Bounds in Incomplete Markets," *Journal of Political Economy*, 108, 79–119.

CONSTANTINIDES, G. M., J. C. JACKWERTH, AND A. SAVOV (2013): "The Puzzle of Index Option Returns," *Review of Asset Pricing Studies*, 3, 229–257.

COVAL, J. D. AND T. SHUMWAY (2001): "Expected Option Returns," *Journal of Finance*, 56, 983–1009.

DANIEL, K., D. HIRSHLEIFER, AND A. SUBRAHMANYAM (1998): "Investor Psychology and Security Market Under- and Overreactions," *Journal of Finance*, 53, 1839–1886.

DE BONDT, W. F. M. AND R. THALER (1985): "Does the Stock Market Overreact?" *Journal of Finance*, 40, 793–805.

DRECHSLER, I. (2013): "Uncertainty, Time-Varying Fear, and Asset Prices," *Journal of Finance*, 68, 1843–1889.

DU, W., A. TEPPER, AND A. VERDELHAN (2018): "Deviations from Covered Interest Rate Parity," *Forthcoming, Journal of Finance*.

EFRON, B. (1979): "Bootstrap Methods: Another Look at the Jackknife," *The Annals of Statistics*, 7, 1–26.

EFRON, B. AND R. J. TIBSHIRANI (1993): *An Introduction to the Bootstrap*, New York: Chapman & Hall.

EPSTEIN, L. G. AND M. SCHNEIDER (2010): "Ambiguity and Asset Markets," *Annual Review of Financial Economics*, 2, 315–346.

EPSTEIN, L. G. AND S. E. ZIN (1989): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

EYSTER, E., M. RABIN, AND D. VAYANOS (2018): "Financial Markets where Traders Neglect the Informational Content of Prices," *Forthcoming, Journal of Finance*.

FAMA, E. F. (1984): "Forward and Spot Exchange Rates," *Journal of Monetary Economics*, 14, 319–338.

FAMA, E. F. AND K. R. FRENCH (1993): "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

FARHI, E. AND I. WERNING (2017): "Monetary Policy, Bounded Rationality, and Incomplete Markets," *NBER Working Paper No. 23281*.

FROOT, K. A. AND A. F. PEROLD (1995): "New Trading Practices and Short-Run Market Efficiency," *Journal of Futures Markets*, 15, 731–765.

FUSTER, A., D. LAIBSON, AND B. MENDEL (2010): "Natural Expectations and Macroeconomic Fluctuations," *Journal of Economic Perspectives*, 24, 67–84.

GABAIX, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, 127, 645–700.

——— (2018): "A Behavioral New Keynesian Model," *Harvard University Mimeo*.

GARCIA, R., E. GHYSELS, AND É. RENAULT (2010): "The Econometrics of Option Pricing," in *Handbook of Financial Econometrics*, ed. by Y. Aït-Sahalia and L. P. Hansen, Amsterdam: Elsevier, vol. 1, chap. 9, 479–552.

GÂRLEANU, N., L. H. PEDERSEN, AND A. M. POTESHMAN (2009): "Demand-Based Option Pricing," *Review of Financial Studies*, 22, 4259–4299.

GEANAKOPLOS, J. (2010): "The Leverage Cycle," in *NBER Macroeconomics Annual 2009, Volume 24*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, Chicago, IL: University of Chicago Press, chap. 1, 1–65.

GEMAN, H., N. EL KAROUI, AND J.-C. ROCHET (1995): "Changes of Numéraire, Changes of Probability Measure and Option Pricing," *Journal of Applied Probability*, 32, 443–458.

GENNAIOLI, N., Y. MA, AND A. SHLEIFER (2016): "Expectations and Investment," in *NBER Macroeconomics Annual 2015, Volume 30*, ed. by M. Eichenbaum and J. Parker, Chicago, IL: University of Chicago Press, chap. 5, 379–431.

GIGLIO, S. AND B. KELLY (2018): "Excess Volatility: Beyond Discount Rates," *Quarterly Journal of Economics*, 133, 71–127.

GILBOA, I. AND D. SCHMEIDLER (1989): "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

GREEN, J. R. (1975): "Information, Efficiency, and Equilibrium," *Harvard Institute of Economic Research Discussion Paper No. 284*.

GREENWOOD, R. AND A. SHLEIFER (2014): "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 27, 714–746.

GROMB, D. AND D. VAYANOS (2002): "Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs," *Journal of Financial Economics*, 66, 361–407.

HALL, P. (1985): "Resampling a Coverage Pattern," *Stochastic Processes and Their Applications*, 20, 231–246.

——— (1988): "Theoretical Comparison of Bootstrap Confidence Intervals," *Annals of Statistics*, 16, 927–953.

HALL, P., J. L. HOROWITZ, AND B.-Y. JING (1995): "On Blocking Rules for the Bootstrap with Dependent Data," *Biometrika*, 82, 561–574.

HALL, R. E. (2017): "High Discounts and High Unemployment," *American Economic Review*, 107, 305–330.

HANSEN, L. P. (2014): "Nobel Lecture: Uncertainty Outside and Inside Economic Models," *Journal of Political Economy*, 122, 945–987.

HANSEN, L. P. AND R. J. HODRICK (1980): "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis," *Journal of Political Economy*, 88, 829–853.

HANSEN, L. P. AND R. JAGANNATHAN (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.

——— (1997): "Specification Errors in Stochastic Discount Factor Models," *Journal of Finance*, 52, 557–590.

HANSEN, L. P. AND T. J. SARGENT (2008): *Robustness*, Princeton, NJ: Princeton University Press.

HANSEN, L. P. AND J. A. SCHEINKMAN (2017): "Stochastic Compounding and Uncertain Valuation," in *After the Flood: How the Great Recession Changed Economic Thought*, ed. by E. L. Glaeser, T. Santos, and E. G. Weyl, Chicago, IL: University of Chicago Press, chap. 2, 21–50.

HANSON, S. G. AND J. C. STEIN (2015): "Monetary Policy and Long-Term Real Rates," *Journal of Financial Economics*, 115, 429–448.

HARRISON, J. M. AND D. M. KREPS (1978): "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations," *Quarterly Journal of Economics*, 92, 323–336.

HE, Z. AND A. KRISHNAMURTHY (2013): "Intermediary Asset Pricing," *American Economic Review*, 103, 732–770.

HESTON, S. L. (2004): "Option Valuation with Infinitely Divisible Distributions," *Quantitative Finance*, 4, 515 – 524.

HIRSHLEIFER, D., J. LI, AND J. YU (2015): "Asset Pricing in Production Economies with Extrapolative Expectations," *Journal of Monetary Economics*, 76, 87–106.

HOROWITZ, J. L. (2001): "The Bootstrap," in *Handbook of Econometrics*, ed. by J. J. Heckman and E. Leamer, Amsterdam: Elsevier, vol. 5, chap. 52, 3159–3228.

IACOVIELLO, M. (2005): "House Prices, Borrowing Constraints, and Monetary Policy in the Business Cycle," *American Economic Review*, 95, 739–764.

IMBENS, G. W. AND C. F. MANSKI (2004): "Confidence Intervals for Partially Identified Parameters," *Econometrica*, 72, 1845–1857.

JACKWERTH, J. C. (2000): "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.

JACOD, J., Y. LI, AND X. ZHENG (2017): "Statistical Properties of Microstructure Noise," *Econometrica*, 85, 1133–1174.

JAMSHIDIAN, F. (1989): "An Exact Bond Option Formula," *Journal of Finance*, 44, 205–209.

JENSEN, C. S., D. LANDO, AND L. H. PEDERSEN (2018): "Generalized Recovery," *Forthcoming, Journal of Financial Economics*.

JONES, C., V. MIDRIGAN, AND T. PHILIPPON (2017): "Household Leverage and the Recession," *IMF and NYU Mimeo*.

JONES, C. S. (2006): "A Nonlinear Factor Analysis of S&P 500 Index Option Returns," *Journal of Finance*, 61, 2325–2363.

JUDD, K. L. (1992): "Projection Methods for Solving Aggregate Growth Models," *Journal of Economic Theory*, 58, 410–452.

KAHNEMAN, D. AND A. TVERSKY (1972): "Subjective Probability: A Judgment of Representativeness," *Cognitive Psychology*, 3, 430–454.

——— (1973): "On the Psychology of Prediction," *Psychological Review*, 80, 237–251.

KILIC, M. AND J. A. WACHTER (2018): "Risk, Unemployment, and the Stock Market: A Rare-Event-Based Explanation of Labor Market Volatility," *Forthcoming, Review of Financial Studies*.

KLEIDON, A. W. (1986): "Variance Bounds Tests and Stock Price Valuation Models," *Journal of Political Economy*, 94, 953–1001.

KÜNSCH, H. R. (1989): "The Jackknife and the Bootstrap for General Stationary Observations," *Annals of Statistics*, 17, 1217–1241.

LAHIRI, S. N. (2003): *Resampling Methods for Dependent Data*, New York: Springer.

LAZARUS, E., D. J. LEWIS, AND J. H. STOCK (2017): "The Size-Power Tradeoff in HAR Inference," *Harvard University Mimeo*.

LEROY, S. F. AND C. J. LACIVITA (1981): "Risk Aversion and the Dispersion of Asset Prices," *Journal of Business*, 54, 535–547.

LEROY, S. F. AND R. D. PORTER (1981): "The Present-Value Relation: Tests Based on Implied Variance Bounds," *Econometrica*, 49, 555–574.

LINN, M., S. SHIVE, AND T. SHUMWAY (2018): "Pricing Kernel Monotonicity and Conditional Information," *Review of Financial Studies*, 31, 493–531.

LIU, R. AND K. SINGH (1992): "Moving Blocks Jackknife and Bootstrap Capture Weak Dependence," in *Exploring the Limits of the Bootstrap*, ed. by R. LePage and L. Billard, New York: Wiley, chap. 11, 225–248.

LO, A. W. AND A. C. MACKINLAY (1988): "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test," *Review of Financial Studies*, 1, 41–66.

LUCAS, R. E. (1972): "Expectations and the Neutrality of Money," *Journal of Economic Theory*, 4, 103–124.

——— (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46, 1429–1445.

LUCAS, R. E. AND E. C. PRESCOTT (1971): "Investment Under Uncertainty," *Econometrica*, 39, 659–681.

MALMENDIER, U., D. POUZO, AND V. VANASCO (2017): "A Theory of Experience Effects," *UC Berkeley and Stanford University Mimeo*.

MALZ, A. M. (1997): "Option-Implied Probability Distributions and Currency Excess Returns," *Federal Reserve Bank of New York Staff Report No. 32*.

——— (2014): "A Simple and Reliable Way to Compute Option-Based Risk-Neutral Distributions," *Federal Reserve Bank of New York Staff Report No. 677*.

MANSKI, C. F. (2017): "Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise," *NBER Working Paper No. 23418*.

MARSH, T. A. AND R. C. MERTON (1986): "Dividend Variability and Variance Bounds Tests for the Rationality of Stock Market Prices," *American Economic Review*, 76, 483–498.

MARTIN, I. (2017): "What Is the Expected Return on the Market?" *Quarterly Journal of Economics*, 132, 367–433.

MARTIN, I. AND S. ROSS (2013): "The Long Bond," *London School of Economics and MIT Sloan Mimeo*.

MEHRA, R. AND E. C. PRESCOTT (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15, 145–161.

MILGROM, P. AND N. STOKEY (1982): "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, 26, 17–27.

MUTH, J. F. (1961): "Rational Expectations and the Theory of Price Movements," *Econometrica*, 29, 315.

PALM, F. C., S. SMEEKES, AND J.-P. URBAIN (2011): "Cross-Sectional Dependence Robust Block Bootstrap Panel Unit Root Tests," *Journal of Econometrics*, 163, 85–104.

PAN, J. AND A. M. POTESHMAN (2006): "The Information in Option Volume for Future Stock Prices," *Review of Financial Studies*, 19, 871–908.

PÁSTOR, L. AND R. F. STAMBAUGH (2003): "Liquidity Risk and Expected Stock Returns," *Journal of Political Economy*, 111, 642–685.

PÁSTOR, L. AND P. VERONESI (2009): "Learning in Financial Markets," *Annual Review of Financial Economics*, 1, 361–383.

POTERBA, J. M. AND L. H. SUMMERS (1988): "Mean Reversion in Stock Prices: Evidence and Implications," *Journal of Financial Economics*, 22, 27–59.

RABIN, M. (2002): "Inference By Believers in the Law of Small Numbers," *Quarterly Journal of Economics*, 117, 775–816.

RABIN, M. AND D. VAYANOS (2010): "The Gambler's and Hot-Hand Fallacies: Theory and Applications," *Review of Economic Studies*, 77, 730–778.

RIETZ, T. A. (1988): "The Equity Risk Premium: A Solution," *Journal of Monetary Economics*, 22, 117–131.

ROSS, S. (2015): "The Recovery Theorem," *Journal of Finance*, 70, 615–648.

SANTA-CLARA, P. AND S. YAN (2010): "Crashes, Volatility, and the Equity Premium: Lessons from S&P 500 Options," *Review of Economics and Statistics*, 92, 435–451.

SCHEINKMAN, J. AND W. XIONG (2003): "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111, 1183–1220.

SCHENKER, N. (1985): "Qualms About Bootstrap Confidence Intervals," *Journal of the American Statistical Association*, 80, 360–361.

SHILLER, R. J. (1979): "The Volatility of Long-Term Interest Rates and Expectations Models of the Term Structure," *Journal of Political Economy*, 87, 1190–1219.

——— (1981): "Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?" *American Economic Review*, 71, 421–436.

——— (1984): "Stock Prices and Social Dynamics," *Brookings Papers on Economic Activity*, 2, 457–498.

——— (2000): *Irrational Exuberance*, Princeton, NJ: Princeton University Press.

SHIMER, R. (2005): "The Cyclical Behavior of Equilibrium and Vacancies Unemployment," *American Economic Review*, 95, 25–49.

SHLEIFER, A. AND R. W. VISHNY (1997): "The Limits of Arbitrage," *Journal of Finance*, 52, 35–55.

SIMS, C. A. (2003): "Implications of Rational Inattention," *Journal of Monetary Economics*, 50, 665–690.

SIMSEK, A. (2013): "Belief Disagreements and Collateral Constraints," *Econometrica*, 81, 1–53.

SUN, Y., P. C. B. PHILLIPS, AND S. JIN (2008): "Optimal Bandwidth Selection in Heteroskedasticity–Autocorrelation Robust Testing," *Econometrica*, 76, 175–194.

TAMER, E. (2010): "Partial Identification in Econometrics," *Annual Review of Economics*, 2, 167–195.

TVERSKY, A. AND D. KAHNEMAN (1971): "Belief in the Law of Small Numbers," *Psychological Bulletin*, 76, 105–110.

WACHTER, J. A. (2013): "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *Journal of Finance*, 68, 987–1035.

WALDEN, J. (2017): "Recovery with Unbounded Diffusion Processes," *Review of Finance*, 21, 1403–1444.

WOODFORD, M. (2009): "Information-Constrained State-Dependent Pricing," *Journal of Monetary Economics*, 56, S100–S124.