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

**Ned Augenblick**
*University of California, Berkeley*
*Haas School of Business*

**Eben Lazarus**
*Harvard University*
*Department of Economics*

August 2017

PRELIMINARY DRAFT
*(Comments welcome)*

**Abstract**

We examine restrictions on the variation in asset prices — specifically, variation in risk-neutral beliefs expressed in option prices — that must hold in a broad class of rational-expectations equilibria. We derive an upper bound for intertemporal movement in risk-neutral beliefs under rational expectations, and we show that this implies a lower bound on the curvature of utility (or, more generally, the slope of the stochastic discount factor) required to rationalize the marginal investor's beliefs. Implementing these bounds empirically using risk-neutral distributions measured from S&P 500 index option prices, 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. Under the theory we develop, this can be understood as providing 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.

*Emails: ned@haas.berkeley.edu and elazarus@fas.harvard.edu. Eben Lazarus thanks John Campbell, Gabriel Chodorow-Reich, Emmanuel Farhi, Matthew Rabin, and Jim Stock for their outstanding guidance and support. We are also grateful to Joshua Abel, Arielle Bernhardt, Laura Blattner, William Diamond, Benjamin Enke, Xavier Gabaix, Mark Gertler, Benjamin Golub, Gita Gopinath, Niels Gormsen, Sam Hanson, David Laibson, Matteo Maggiori, José Luis Montiel Olea, Mikkel Plagborg-Møller, Simon Potter, Gianluca Rinaldi, Andrei Shleifer, Emil Siriwardane, Robert Solow, Jann Spiess, David Sraer, Stefanie Stantcheva, Adi Sunderam, Neil Thakral, Johan Walden, Chenzi Xu, and seminar participants at Harvard and UC Berkeley–Haas for advice, comments, and helpful discussions.

# 1. Introduction

The assumption that individuals have rational expectations (RE) over future outcomes is central to many models in macroeconomics and finance. A large body of previous literature has examined this assumption empirically, often in the context of asset markets. Asset prices in part reflect investor expectations over future outcomes, so these markets would seem to provide ideal testing grounds to discipline our understanding of such expectations.

In much of this past literature, researchers have argued that asset prices appear excessively volatile relative to measures of fundamental value under the RE benchmark, and therefore that expectations are overly volatile. But these empirical exercises face a challenge: when observing volatile asset-price changes, how can we disentangle whether these arise due to changes in expectations rather than preferences? Equivalently, when prices decrease, have expectations of future dividends decreased, or has the marginal investor simply become more impatient or risk-averse? This past literature has accordingly appealed to restrictions on risk preferences or the data-generating process for asset returns in making the claim that their findings in fact reflect excess volatility in expectations themselves.[1]

We make progress on these questions by considering a general theoretical framework in which we show that certain restrictions on asset-price movements must hold in a broad class of rational-expectations equilibria. These bounds have empirical counterparts, and they accordingly give us 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 bounds give a precise and conservative value for the curvature of utility — or, in asset-pricing terms, the slope of the stochastic discount factor across states — 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, allowing us to disentangle risk aversion from excess volatility in beliefs in these cases.

We focus our analysis on option prices, and we gain traction theoretically by translating option prices into so-called *risk-neutral beliefs* over the underlying asset's future price. Risk-neutral beliefs represent a probability distribution that would be equal to a hypothetical risk-neutral agent's true belief distribution about the future asset price, but these risk-neutral beliefs are in general distorted relative to the marginal investor's true (or *physical*) beliefs in the case that the investor is risk-averse.[2] 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 the future asset price — so we treat risk-neutral beliefs as observable values, as is standard. These risk-neutral beliefs are then our main object of inter-

[1] As noted in LeRoy and LaCivita (1981), for example, the benchmark tests in Shiller (1979, 1981), and later in De Bondt and Thaler (1985), Campbell and Shiller (1987), and others, can generally be interpreted as joint tests of stationarity, constant discount rates, and rational expectations. We can accordingly interpret their null rejections as joint rejections of these three assumptions. See the end of this section for a full literature review.

[2] Intuitively, with risk aversion, the probability distribution we observe using option prices will overweight states in which the marginal investor has low wealth (e.g., when the underlying asset has a low return), since the investor will be willing to pay more than the actuarially fair value to insure against these high-marginal-utility states.

est, and our bounds show theoretically how their intertemporal behavior is restricted under the maintained assumption of rational expectations.

Economically, these theoretical results reflect the following intuition about belief updating under rational expectations: when individuals' beliefs about a future outcome change, they must on average be getting closer to certainty about the ultimate realization of the event. If not, then they are, loosely, overreacting to new information relative to the rational (Bayesian) benchmark, or underweighting their prior relative to this new information.

For example, consider a hypothetical agent's beliefs over whether a single outcome will occur at some future date $T$. For simplicity, assume for now that this agent is risk-neutral over the outcome's realization. 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 on each successive day, forcing her to reverse her beliefs dramatically on a day-to-day 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, as can be shown formally. 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 in terms of their distance from 0% or 100% — despite their large day-to-day changes.

The intuition in this example is present as well in Augenblick and Rabin (2017), who use a framework similar to ours to prove restrictions on beliefs in the case that we can directly observe those beliefs (as in the above example). When considering beliefs over future asset prices as expressed in options, this is again equivalent to the case in which the marginal investor is risk-neutral. Our task becomes considerably more difficult when the assumption of risk neutrality is dropped. In this case, we observe only risk-neutral beliefs rather than true beliefs, and these risk-neutral beliefs need not follow a martingale given their distortion relative to the true beliefs (which do follow a martingale under RE). The bulk of our theoretical work is thus in extending the logic from the preceding example to this more general case, and showing how risk-neutral belief movement can be similarly restricted under RE.

The key insight we exploit to obtain our main results in this general case is that unlike the underlying asset price or expected returns, beliefs are *bounded*: even with large risk aversion, risk-neutral beliefs can only be distorted a finite distance from the true underlying beliefs, since risk-neutral beliefs lie between 0% and 100% by definition. Using this insight, we obtain our formal notion of excess volatility by showing that intertemporal movement in risk-neutral beliefs is bounded under rational expectations for any given level of risk aversion. This bound thus provides our mapping between movement in asset prices (or risk-neutral beliefs) and the degree of risk aversion required under RE. This mapping arises naturally in this asset-pricing context: we

prove our results in terms of risk-neutral beliefs (because they are observable), and the deviation of these risk-neutral beliefs from the true underlying beliefs depends on the degree of risk aversion of the marginal investor.

Aside from the maintained assumptions of rational expectations and no arbitrage in asset markets, our main theoretical results require only one general assumption on the stochastic discount factor (SDF), the random variable that determine's an asset's ex-ante price by discounting the asset's random future cash flows. In particular, we assume that the realization of the SDF 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. We later extend our theoretical results to incorporate cases in which this assumption is violated.

We then take our bounds to the data. We construct test statistics implied by our theoretical results using daily data on S&P 500 index option prices obtained from OptionMetrics. Our empirical estimation thus gives us information on the degree of risk aversion (or the slope of the SDF) required to rationalize the observed volatility in the risk-neutral beliefs implied by options over the value of the market index.

Conducting this estimation, we find that very high risk aversion is needed to rationalize the data, and in many cases our point estimates indicate that *no* amount of risk aversion is capable of rationalizing the data. Intuitively, the marginal investor's beliefs over the future price of the market index are overly volatile relative to the RE benchmark; equivalently, these beliefs (and the associated risk-neutral beliefs) have a tendency to mean-revert. These results suggest 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 without very high risk aversion.

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.[3] 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 mid prices from daily data to avoid intraday

[3]In this case, one might still consider our findings to be indicative of some apparent inefficiency, but specific to the options market rather than with respect to macroeconomic beliefs more generally. There is indeed evidence of idiosyncracies in this market; for example, various put-selling and related strategies have high measured returns (Coval and Shumway, 2001; Bondarenko, 2003), and Jackwerth (2000) argues that a pricing kernel backed out from options data is nonmonotonic in the index return (though Barone-Adesi, Engle, and Mancini, 2008, argue otherwise). Santa-Clara and Saretto (2009) suggest that margin requirements account for much of the apparent mispricing, while Bollen and Whaley (2004) and Gârleanu, Pedersen, and Poteshman (2009) argue for the importance of demand pressures. However, Bates (2008), Broadie, Chernov, and Johannes (2009), Santa-Clara and Yan (2010), and Constantinides, Jackwerth, and Savov (2013) show evidence suggesting that the apparent mispricings are insignificant once peso-problem-type sampling uncertainty and disaster- or jump-risk premia are accounted for, and are therefore not specific to option markets: "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). There is also some evidence (e.g., Buraschi and Jiltsov, 2006; Ziegler, 2007) for belief-based distortions, but this is similarly consistent with an interpretation based on factors not entirely specific to the option market. We return to this issue in Sections 4–6.

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, which weaken our results slightly but still indicate that very high risk aversion is needed to rationalize the data.

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 finite and reasonable values for risk aversion, 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 reduced-form regressions of our excess volatility measure on a range of macroeconomic statistics, we find that it has a strong positive relationship with measures of macroeconomic uncertainty and no relationship with measures of limits to arbitrage in asset markets.[4]

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 violations of conditional transition independence might affect our conclusions under the maintained null of RE. We solve and simulate the model numerically,[5] and two findings emerge from these simulations: (a) even given the violation of CTI, the simulated risk-neutral beliefs still exhibit substantially less variation than observed in empirical data; (b) perhaps more importantly, when we naïvely apply our theoretical bounds to these simulated beliefs to estimate the risk aversion required to rationalize the simulated data, the bounds are still conservative as estimates of the true underlying risk aversion inherent in the model. This indicates that our bounds still apply approximately under mild violations of CTI, and we formalize this in an additional theoretical result. We consider these results to be suggestive evidence that violations of this sort have difficulty accounting for the excess volatility in beliefs we observe in the data.

**Relation to previous literature and interpretation of results.** Our paper relates to and builds on multiple separate but related strands of previous literature. We follow Shiller (1979, 1981), and later, among others, LeRoy and Porter (1981), De Bondt and Thaler (1985), Campbell and Shiller (1987), and Stein (1989) in conducting empirical tests for excess volatility in asset prices relative

[4]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 admittedly only suggestive and reduced-form.

[5]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.

to the RE benchmark.[6] This literature proposes multiple different tests of excess volatility; taking Shiller (1981) as a benchmark case, he considers the variance of an aggregate equity index price relative to the variance of the present value of the ex-post stream of future dividends discounted at a constant rate, and finds that the former is much greater than the latter. But as noted by LeRoy and LaCivita (1981), this test, and those that followed, are in fact joint tests of (a) stationarity of the relevant data-generating process (for prices, in Shiller's case); (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 these results; meanwhile, much of the modern asset-pricing literature rationalizes the observed price volatility by appealing to time variation in discount rates (see Campbell, 2003, Cochrane, 2011, and Ludvigson, 2013, 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 this additional generality.

First, we must consider derivative prices (specifically, risk-neutral beliefs over index returns) rather than the behavior of the underlying index, as in much of this previous literature.[7] In this way, our work is complementary to recent work by Giglio and Kelly (2017), who document excess volatility for claims on equity and currency volatility, inflation swaps, commodity futures, and credit default swaps. They do so using a framework that, like ours, accounts in principle for discount-rate variation: they exploit the (approximately) linear risk-neutral factor structure of cash flows for those assets in many standard models (as in Chamberlain and Rothschild, 1983), which restricts the price variation in long-maturity claims relative to that of short-maturity claims, and they show using variance-ratio tests that this restriction is violated empirically. Aside from considering different assets — their "affine-$\mathbb{Q}$" framework lends itself to term-structure-like claims (i.e., claims on objects for which all uncertainty is realized by the claim-maturity date), so they cannot consider options on the level of an equity index — we differ further in that our estimation procedure and our focus on the risk-neutral belief distribution allow us to consider the specific behavior of beliefs that leads to the excess volatility documented in both papers, which is less straightforward in the context of variance-ratio tests for prices.[8,9] But despite these differences, we consider the two approaches to be complementary, and both yield similar evidence.

[6] A related literature — e.g., Fama and French (1988a,b), Poterba and Summers (1988), and later Campbell and Thompson (2008) — presents evidence on the predictability of returns. As noted by Campbell and Shiller (1988a,b), excess price volatility implies return predictability, so these papers' findings are related to the findings in the excess-volatility literature.

[7] We have in fact obtained additional results that restrict the variation in the underlying price, but these are preliminary and not yet included in this draft.

[8] There are further technical differences between the Giglio and Kelly (2017) paper and ours. For example, their estimation procedures require the assumptions of either (a) stationarity in the term structure of prices *and* no measurement error in their factor model for short horizons, *or* (b) homoskedastic Gaussian forcing processes for the factor model, while we do not require such restrictions. Giglio and Kelly do, however, consider possible violations of those restrictions and argue that their findings are robust to such violations.

[9] Additional literature examining benchmark models using variance ratios, from which we differ, includes Lo and MacKinlay (1988), Cutler, Poterba, and Summers (1989), Campbell (1991), and Cochrane (1992).

Second, rather than considering a setting allowing for fully binary (rejection vs. non-rejection) empirical tests of benchmark models, our general framework instead allows for a mapping between the observed asset-price variation and the risk aversion required to rationalize the data. (We do, however, show that there are cases in which even infinitely large risk aversion cannot rationalize this observed variation; we postpone a discussion of the interpretation of this result to Section 4.) 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.[10] 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; Martin, 2013; Wachter, 2013) can generate sufficient SDF volatility to rationalize the observed equity premium without requiring high risk aversion. But this is importantly not the case for our results. 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, so that our finding of large belief variation across the entire distribution of beliefs *cannot* be rationalized by features of the data-generating process for consumption or returns; instead, conditional on our identification assumptions, this finding is attributable either to very high risk aversion or a departure from the RE benchmark we consider.[11] Put more simply, 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.

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 and estimates an upper bound on risk aversion using labor-supply behavior, and Chetty (2009) provides a longer survey of the literature in settings different from ours. Within the asset-pricing literature more specifically, see, for example, Alvarez and Jermann (2005), Backus, Chernov, and Zin (2014), 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 rather than asset prices, as, for example, in Frankel and Froot (1987), Han (2008), Greenwood

[10] We 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.

[11] Our framework admits an interpretation of such a rejection as arising due to overreaction to information, and we frequently refer to our findings in this manner. But we should note as well that the set of alternatives being tested against the RE benchmark we consider does in principle include models in which agents do not make this particular statistical error. We postpone consideration of the possible theoretical underpinnings of our empirical results to the Conclusion.

and Shleifer (2014), Gennaioli, Ma, and Shleifer (2016), and Manski (2017), as well as the results of Augenblick and Rabin (2017) for settings in which beliefs are directly observable. Another set of related literature endeavors to measure physical beliefs (rather than risk-neutral beliefs) directly using options data, for the purpose of examining either expectations or preferences. Aït-Sahalia and Lo (2000) and Jackwerth (2000) assume that the conditional subjective (and objective) physical distribution over returns is equal to a function of the observed historical distribution (cf. Aït-Sahalia, 1996, Stutzer, 1996, for the unconditional distribution); Aït-Sahalia, Wang, and Yared (2001), Poteshman (2001), Pan (2002), and Rosenberg and Engle (2002) fit parametric models to prices and/or returns that allow their distributions to be time-varying; Bliss and Panigirtzoglou (2004) restrict the representative-investor utility function to be either power utility or constant-absolute-risk-aversion utility, and then estimate the single unknown parameter in either case by minimizing forecast errors for the distribution of returns constructed using the risk-neutral densities and the assumed utility function.[12] Ross (2015) assumes a Markov process for transitions between return states and a transition-independence assumption on the pricing kernel 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 the above 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.[13]

**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 many-state framework in which an arbitrary stochastic discount factor prices assets, 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 discusses possible theoretical channels for our results and concludes. Proofs are contained in Appendix A, and Appendix B contains additional technical detail from Sections 3–6.

[12] Bates (2003) and Garcia, Ghysels, and Renault (2010) provide surveys of related literature, and see also Backus, Chernov, and Martin (2011) for the application and comparison of multiple related methods to estimate beliefs over consumption disasters. And as discussed in Section 5, the implementation of our empirical tests builds as well on this literature and related econometrics literature: we use the results of Breeden and Litzenberger (1978) and the curve-smoothing methods of Bliss and Panigirtzoglou (2004) and Malz (2014) to measure risk-neutral probabilities using options data, and we use a block bootstrap (Carlstein, 1986) and insights from the partial-identification literature (Imbens and Manski, 2004; Tamer, 2010) for estimation and inference.

[13] Similar parametric concerns apply as well to the option-anomalies literature discussed in Footnote 3, and these concerns are in a sense isomorphic to the peso problem that Broadie, Chernov, and Johannes (2009) argue underlies the anomaly results: 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 ex-ante returns to a put-selling strategy and to incorrectly infer the shape of the pricing kernel over return states. Linn, Shive, and Shumway (2015) argue that this is an empirically relevant concern, and see also Weitzman (2007) for broader discussion.

## 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.[14] 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 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: $C_T \in \{C_{\text{low}}, C_{\text{high}}\}$. To make clear the numeraire in which assets are priced each period, we assume that the agent consumes the exogenous stream $C_t = \overline{C}$ deterministically for all $t < T$, and we focus on the realization of uncertainty as of 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 $\pi_t$ the agent's date-$t$ subjective belief that the bad state $C_{\text{low}}$ 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 (i.e., the history of signals about terminal consumption). This requires that the agent have a correctly specified prior $\pi_0$ and that beliefs are updated according to Bayes' rule in response to new information embedded in $\mathcal{F}_t$. (We postpone a full formal discussion of the requirements of RE to Section 3.) The belief $\pi_t$ 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.[15]

Assume for now that the agent is risk-neutral and does not discount future consumption, so ex-ante utility is $\mathbb{E}_0 \sum_{t=1}^{T} C_t$. 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).[16] With risk neutrality and no discounting, optimality implies that these state prices $q_t(C_{\text{low}})$

[14] The example we use for (a) works through a basic version of the results in Augenblick and Rabin (2017), 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).

[15] 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.

[16] These prices can in general be inferred from the prices of options on the terminal value of the wealth portfolio. A

and $q_t(C_{\text{high}})$ are equal to the agent's actual subjective beliefs $\pi_t$ 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 $t_2 > t_1$:

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

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

Second, we define the *uncertainty* of belief $\pi_t$ as

$$u_t \equiv (1-\pi_t)\pi_t, \tag{2}$$

and *uncertainty resolution* as

$$r_{t_1,t_2} \equiv u_{t_1} - u_{t_2}. \tag{3}$$

Uncertainty intuitively measures the belief's distance from certainty, and it is maximized at $\pi_t = 0.5$. It is equivalent to the variance of the Bernoulli random variable whose realization is given by $\mathbb{1}\{C_T = C_{\text{low}}\}$. 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, as in the current example. Resolution over the entire path is accordingly equal to the ex-ante Bernoulli variance as of period 0.

Movement and uncertainty resolution are restricted under rational expectations according to the following lemma, which formalizes a notion of the "correct" amount of belief volatility over any horizon under rationality.

**Lemma 1 (*Augenblick and Rabin, 2017*).** *Under rational expectations, for any data-generating process, belief movement must equal uncertainty resolution in expectation:*

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

***Proof.*** See Appendix A.

The proof follows from the martingale property of beliefs, and as can be seen in the proof, this equality must in fact hold for any martingale, both unconditionally and conditionally.[17]

standard (European) call option with strike price $K$ pays off $\max\{W_T - K, 0\}$, where $W_T$ is the terminal value of wealth; in the current two-state example, setting $K = C_{\text{high}} - 1$ yields an option that is exactly equivalent to the Arrow-Debreu security that pays off one unit in the high-consumption state. Similar logic applies in more general settings as well, as noted by Breeden and Litzenberger (1978); we defer this discussion to Section 3.

[17] We can interpret $r$ as uncertainty resolution given that the sequence $\{\pi_t\}$ is a *bounded* martingale, but the statement holds even for unbounded martingales (see Augenblick and Rabin, 2017, for further discussion). Further, this result is in fact the discrete-time analogue of a well-known result for continuous-time semimartingales, often used in the context

The restriction in Lemma 1 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 ($r > 0$). And this is true not just for the entire path of beliefs but also on a period-by-period basis, so that belief variation must correspond on average to uncertainty resolution for any subsample or sampling frequency. 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.[18] 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 see how an empirical test based on Lemma 1 differs from a test of whether beliefs follow a martingale,[19] we can rewrite the restriction in the lemma as $\mathbb{E}[\sum_{t=1}^{T}(1 - 2\pi_{t-1})(\pi_t - \pi_{t-1})] = 0$ for the full path. Thus revisions in 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 the test would still reject rationality if such unconditional martingale behavior arose due to revisions in the opposite direction of the initial belief. If beliefs begin at 0.1, they must eventually resolve to $\pi_T = 0$ most of the time (specifically, 90 percent) in order to be correctly calibrated, and the test is capable of detecting mean reversion arising from the violation of this restriction.

## 2.2. Complications Arising from Risk Aversion

An econometrician with access to repeated samples of the belief process for the agent in the previous subsection could test for excess belief movement using Lemma 1, 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 measured beliefs need not follow a martingale with risk aversion, and thus the bulk of our work consists in extending the logic of the example above to accommodate this more general case. The distortion from risk aversion may in theory induce the excess-movement statistic in the previous lemma to be either positive or negative on average, but given that we tend to observe what appears to be positive excess movement, we focus on this case.

of high-frequency volatility measurement; see, for example, Barndorff-Nielsen and Shephard (2001) and Andersen, Bollerslev, Diebold, and Ebens (2001), or Andersen, Bollerslev, and Diebold (2010) for a survey.

[18] The converse holds for $\mathbb{E}[m] < \mathbb{E}[r]$, but we focus on the case of excess volatility given our empirical findings. Further, when we refer to overreaction relative to the rational benchmark, this should be taken to encompass the possibility that the agent is in fact underweighting her prior relative to information from newly observed signals, but we use the overreaction terminology for simplicity. See the Conclusion for further discussion.

[19] Some past literature conducts such direct calibration tests (e.g., Aït-Sahalia, Wang, and Yared, 2001), though other literature — for example, using instrumental-variables approaches (e.g., Hansen and Singleton, 1982) or variance-ratio tests (e.g., Lo and MacKinlay, 1988) — does take advantage of conditioning information, but in a slightly different manner from the one used here. Augenblick and Rabin (2017) conduct simulations comparing these frameworks.

Assume now that the agent in the above example has log utility, $\mathbb{E}_0 \sum_{t=1}^{T} \log(C_T)$, 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. Assume further for exposition that $C_{\text{low}} = 1/2$, $C_{\text{high}} = 2$, and that this is known, as well as the deterministic stream $C_t = \overline{C} = 1$ for $t < T$.

The econometrician can once again observe Arrow-Debreu state prices over time for the two terminal consumption states, but in this case these prices are not equal to the agent's actual subjective beliefs. Assume the agent's unobservable 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

$$q_t(C_i) = \frac{U'(C_i)}{U'(C_t)} \pi_t \quad \text{for } i \in \{\text{low}, \text{high}\}, \tag{4}$$

so using the parameters given 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 probability* $\pi_t^*$ by dividing the relevant state price through by the sum of state prices (equivalently, the inverse of the gross risk-free rate):

$$\pi_t^* \equiv \frac{q_t(C_{\text{low}})}{q_t(C_{\text{low}}) + q_t(C_{\text{high}})} = \frac{U'(C_{\text{low}})}{\mathbb{E}_t[U'(C_T)]} \pi_t. \tag{5}$$

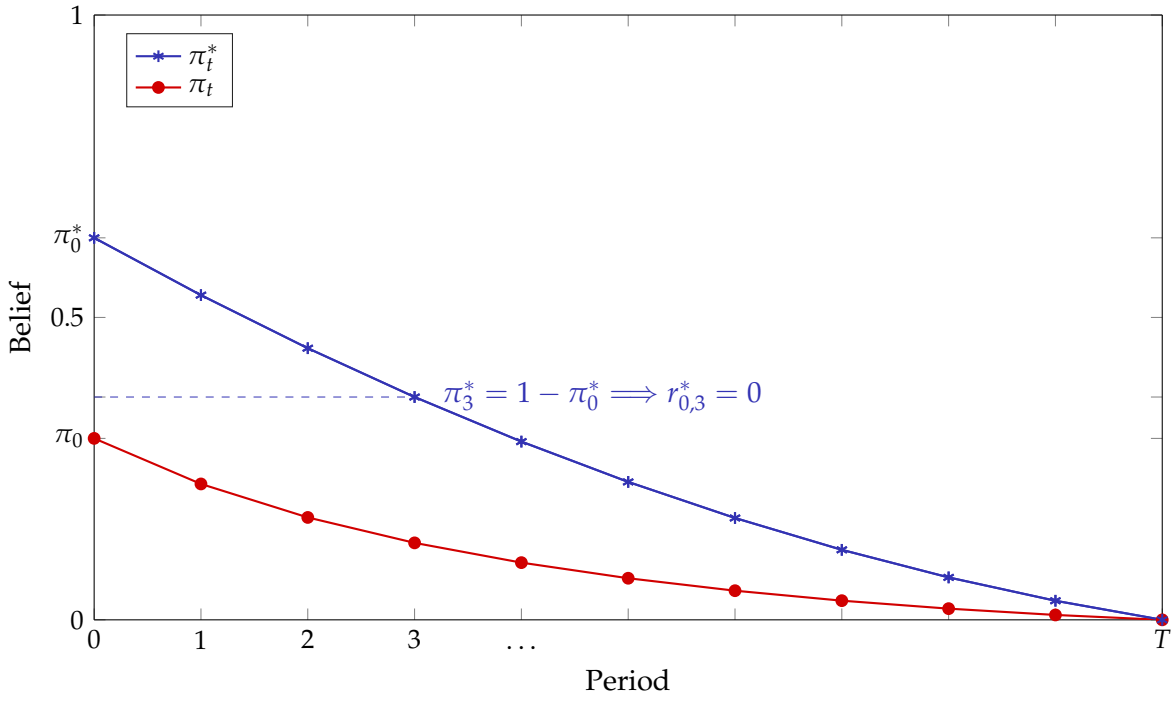
The high-state risk-neutral probability is similarly $\frac{q_t(C_{\text{high}})}{q_t(C_{\text{low}}) + q_t(C_{\text{high}})} = 1 - \pi_t^*$. So the two states' risk-neutral probabilities are both positive and sum to 1 by construction, implying they define a valid probability distribution. As is standard, we refer to them as risk-neutral beliefs because as can be seen in the last expression in (5), the risk-neutral belief and the actual subjective probabilities coincide for a risk-neutral agent, as in the previous subsection. So the risk-neutral beliefs can be interpreted as the actual 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 true subjective belief.

To see how this distortion in observed beliefs could bias inference regarding the volatility of beliefs, 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. $\pi_0$ in this example is equal to 0.3, the value used above. Given this prior for the bad state, a typical realization under rationality will involve beliefs eventually converging to 0 for this state; without yet specifying the information process faced by the agent, the example in the figure simply assumes that the agent receives signals such that her subjective beliefs slowly and monotonically converge to 0 at the terminal date.

The 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 with blue asterisks. Given that $\pi_0^* > 0.5$, the risk-neutral

**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.

belief converging 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 $t_2 = 3$, denoted by $r_{t_1,t_2}^* \equiv (1 - \pi_{t_1}^*)\pi_{t_1}^* - (1 - \pi_{t_2}^*)\pi_{t_2}^*$, 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 $t_2 = 3$.

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. This reflects the logic of the joint hypothesis problem stated by Fama (1970, 1991), or, in the terminology of Borovička, Hansen, and Scheinkman (2016), the "fundamental identification problem."

## 2.3. Identification and Interpretation

We now show how we propose to solve this identification problem in the context of the above example. For simplicity, we focus on beliefs only for periods 0 and 1 (with $T > 1$), and we assume the econometrician observes many draws $(\pi_0^*, \pi_1^*)$ generated by subjective beliefs $(\pi_0, \pi_1)$ 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, as we assume the same utility function and consumption states in terminal period $T$.

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 (4)–(5) 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 $\mathbb{E}[m^* - r^*] > 0$.

However, we can use our knowledge of the data-generating process for beliefs — which is artificially imposed for now, but which is consistent with the structure of the repeated-experiment thought exercise — to achieve identification of the risk-aversion value required for the observed data to be consistent with rationality. In the next section, however, we will no longer assume we have access to the data-generating process underlying the observed belief statistics and take a conservative approach over the entire space of possible data-generating processes.

Denote

$$\phi \equiv \frac{U'(C_{\text{low}})}{U'(C_{\text{high}})}, \tag{6}$$

and 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 $\phi$ and the risk-neutral probability to the associated subjective probability. Equation (5) can be inverted to yield that

$$\pi_t(\phi, \pi_t^*) = \frac{\pi_t^*}{\phi + (1 - \phi)\pi_t^*}. \tag{7}$$

Note that the mapping between $\pi_t$ and $\phi$ is one-to-one for any given observed value $\pi_t^*$, and it is decreasing in $\phi$: as risk aversion increases, the underlying bad-state subjective belief decreases with respect to the observed risk-neutral belief, but $\pi_t$ is of course still bounded below by 0 for any level of risk aversion.

As in the last paragraph of Section 2.1, we can write Lemma 1 in terms of $\pi_t(\phi, \pi_t^*)$ for the current example as $\mathbb{E}[(1 - 2\pi_0(\phi, \pi_0^*))(\pi_1(\phi, \pi_1^*) - \pi_0(\phi, \pi_0^*))] = 0$. Since $\pi_0$ 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, $\mathbb{E}[\pi_1(\phi, \pi_1^*)] = \pi_0(\phi, \pi_0^*)$.[20] We can observe in the data that $\pi_1^*$ takes on two values with equal probability: $\pi_h^* = 0.86$ and $\pi_l^* = 0$. Using this along with (7), the restriction becomes

$$\frac{1}{2}\frac{\pi_h^*}{\phi + (1-\phi)\pi_h^*} = \frac{\pi_0^*}{\phi + (1-\phi)\pi_0^*}, \tag{8}$$

which can be solved to yield

$$\phi = \frac{\pi_0^* \pi_h^*}{\pi_h^*(1+\pi_0^*) - 2\pi_0^*} = 4. \tag{9}$$

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. Note also that $U'(C_{\text{high}}) = U'(C_{\text{low}}) + U''(C_{\text{low}})(C_{\text{high}} - C_{\text{low}}) + \mathcal{O}\big((C_{\text{high}} - C_{\text{low}})^2\big)$ as $C_{\text{high}} \to C_{\text{low}}$ by a Taylor expansion, which can be rearranged to yield

$$\gamma(C_{\text{low}}) \equiv -\frac{C_{\text{low}} U''(C_{\text{low}})}{U'(C_{\text{low}})} = \frac{\phi - 1}{(C_{\text{high}} - C_{\text{low}})/C_{\text{low}}} \tag{10}$$

to a first order, which tells us that risk aversion depends on the ratio of marginal utilities across states relative to the percent consumption (or wealth) 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.

As above, the imposed knowledge of the data-generating process is not central to identification in what follows in Sections 3–4. The true data-generating process is potentially infinite-dimensional — for example, the signal structure at $t+1$ could depend on the value $\pi_t$ in a non-stationary manner — in which case estimation of this process becomes difficult or impossible. We accordingly take a conservative approach and show that there is in fact a single data-generating process that generates the lowest possible $\phi$ required to rationalize the data, and we estimate a lower bound for $\phi$ using $\mathbb{E}[m^* - r^*]$ for this worst-case process.

What *is* central to identification in what follows is an assumption implicit in this example, which is that the value $\phi$ is constant over time. This follows naturally in the above example from the assumption of time-separable utility and fixed state-contingent consumption values. More generally, it requires that the ratio of expected state-contingent realizations of the stochastic discount factor (or marginal utility) be constant from date 0 to $T$. This rules out, for example, models of habit formation as in Campbell and Cochrane (1999); terminal marginal utility will always depend on the path of consumption, and thus the ratio $\frac{\mathbb{E}_t[U'(C_{\text{low}})]}{\mathbb{E}_t[U'(C_{\text{high}})]}$ need not be constant over $t$.[21] We are in effect assuming that when we observe risk-neutral probabilities change, this is because the underlying subjective probabilities are changing rather than the severity of realizing a given state.

[20] As will be seen in the next section, however, using the more powerful restriction $\mathbb{E}[m - r] = 0$ will allow for a simple closed-form bound for $\phi$ given the observed values $m^*$ and $r^*$ given general information structures; the example here simply works through how identification works in a stripped-down case.

[21] It is of course the case that conditioning on a suitably enlarged or redefined state space would yield states across which the ratio of marginal utilities would be constant; for Campbell-Cochrane habit formation, for example, the ratio of marginal utilities across two *surplus-consumption* (i.e., consumption relative to habit) states *is* constant. But this does not address the issue that it is difficult to find empirical counterparts for these states.

One might consider this unreasonable in the context of a rare-disasters model in which bad-state consumption is unknown and depends on the severity of a possible disaster, but we can avoid this issue by conducting empirical tests using conditional probabilities for two terminal states far enough from the left tail of the distribution that this concern is not of first-order relevance. This point is made clearer in the next section.

We note also that this assumption does not rule out time variation in discount rates; instead, it restricts the source of that variation. The price of the terminal consumption claim is given in equilibrium in the above example by $P_t(C_T) = \mathbb{E}_t\left[\beta_t^{T-t}\frac{U'(C_i)}{U'(C_t)}C_i\right]$, where we have relaxed the assumption of no discounting and $\beta_t$ is now the agent's (possibly time-varying) time discount factor, and again $i \in \{\text{low}, \text{high}\}$. 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

$$\begin{aligned}\mathbb{E}_t[R^C_{t,T}] &= \frac{1 - \text{Cov}_t\left(\beta_t^{T-t}\frac{U'(C_i)}{U'(C_t)}, C_i\right)}{\mathbb{E}_t\left[\beta_t^{T-t}\frac{U'(C_i)}{U'(C_t)}\right]} \\ &= \frac{\frac{U'(C_t)}{\beta_t^{T-t}} - \text{Cov}_t(U'(C_i), C_i)}{\mathbb{E}_t[U'(C_i)]},\end{aligned} \tag{11}$$

as usual. For full concreteness, we can write $\mathbb{E}_t[U'(C_i)] = \pi_t U'(C_{\text{low}}) + (1-\pi_t)U'(C_{\text{high}})$, and $\text{Cov}_t(U'(C_i), C_i)$ can be similarly rewritten as a function of $\pi_t$, $C_T$, 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 $\beta_t$.
2. Changes in contemporaneous marginal utility $U'(C_t)$.
3. Changes in the relative probability $\pi_t$.
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 the above 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 will in fact be able to slightly relax these assumptions and allow for independent consumption-growth or marginal-utility shocks for a given return state; the next section more fully discusses the models covered by our assumptions.)

To see one way in which this framework allows for more generality than the standard constant-discount-rate framework of, for example, Shiller (1981), 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 $\pi_t$ is constant at $\pi_t = \pi_0 = 0.3$ for $t < T$, and all other aspects of the example are unchanged. This induces time variation in the price of the consumption claim: again assuming no discounting, we have $P_t(C_T) = 1$ 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 arising from the first term in the numerator of (11). A variance-ratio test under the assumption of constant discount rates would thus spuriously reject the null of RE.

Meanwhile, because the mapping between $\pi_t$ and $\pi_t^*$ is one-to-one for a given $\phi$ in (7), 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 (4)–(5), 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 (11) as

$$\mathbb{E}_t[R^C_{t,T}] = R^f_{t,T} - \frac{\text{Cov}_t(U'(C_i), C_i)}{\mathbb{E}_t[U'(C_i)]}, \tag{12}$$

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, in which $C_t = \overline{C} = 1$ 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 (12) 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 $\text{Var}_t(\beta^{T-t}U'(C_i)/U'(C_t))$; we could rewrite (12) 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 above would nonetheless apply as well with slight modification. We now move to consider a more general framework in which we do not explicitly parameterize the stochastic discount factor or the data-generating process in the manner considered in this section, and show how our analysis works in this more general case.

# 3. General Theoretical Framework

## 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 set of $N$ option expiration dates $\{t : t \in (T_1, T_2, \ldots, T_N)\}$, or $\{T_i\}_{i=1,\ldots,N}$. (The superscript $m$ will generally refer to objects tied to the market index, and the subscript $t$ to $\mathcal{F}_t$-adapted processes.) A European call option on the market index with expiration or maturity date $T_i$ and strike price $K$ has payoff of $X^m_{T_i,K} = \max\{V^m_{T_i} - K, 0\}$, and we denote its time-$t$ price as $q^m_{t,T_i,K}$. We denote by $0_i$ the first date on which the price of any such option contract (with arbitrary strike $K$) is observable; that is, $0_i$ is the minimal $t$ for which there exists a $K \in \mathcal{K} \subseteq \mathbb{R}_+$ such that $q^m_{t,T_i,K}$ is $\mathcal{F}_t$-measurable.[22]

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

$$R^m_{T_i} \equiv \frac{V^m_{T_i}}{V^m_{0_i}} = s, \tag{13}$$

and the set of return states is accordingly a set of discrete values that the market return can take under $\Omega$; 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). Note that these return states in fact correspond to *events* in $\mathcal{F}_{T_i}$ in general, and not necessarily elementary states in $\Omega$.[23] 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 $X_{t+k}$ is given by

$$S_t(X_{t+k}) = \mathbb{E}_t\left[\frac{M_{t+k}}{M_t} X_{t+k}\right], \tag{14}$$

where again $\mathbb{E}_t[\cdot] \equiv \mathbb{E}[\cdot|\mathcal{F}_t]$, and we can initialize $M_0 = 1$.[24]

In the case of a Lucas (1978)–type economy, with a representative agent with consumption

[22] Given put-call parity, the value $0_i$ would be unchanged if we were to additionally consider put options.

[23] We nonetheless refer to these as return states because, following the standard argument, we can think of the random variable $R^m_{T_i}$ as inducing a new probability space with state space $\mathcal{S}$, the image of $\Omega$ under the mapping $R^m_{T_i}$. And 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.

[24] For our purposes, we can define absence of arbitrage following Campbell (2017): we say that there is absence of arbitrage if (i) $S_t(X_{t+k}) \geqslant 0$ for all tradable payoffs $X_{t+k}$ such that $X_{t+k} \geqslant 0$ almost surely, and (ii) $S_t(X_{t+k}) > 0$ for all tradable payoffs $X_{t+k}$ such that $X_{t+k} > 0$ with positive probability. See Duffie (2001), Cochrane (2005), or Campbell (2017) for proofs that absence of arbitrage implies a strictly positive SDF process (and vice versa).

process $\{C_t\}$ and time-separable consumption utility with time discount factor $\beta$ as in the previous section, the SDF evolves according to

$$\frac{M_{t+1}}{M_t} = \beta \frac{U'(C_{t+1})}{U'(C_t)} \tag{15}$$

by the agent's Euler equation, so the SDF can accordingly be interpreted as aggregate marginal utility. But we need not assume the existence of such a representative agent in order for the representation (14) to be valid.

**Risk-neutral measure.** The representation (14) yields that the $k$-period gross risk-free rate is given by

$$R^f_{t,t+k} \equiv \frac{1}{S_t(1_{t+k})} = \frac{1}{\mathbb{E}_t\left[\frac{M_{t+k}}{M_t}\right]}, \tag{16}$$

where $1_{t+k}$ refers to one unit of the numeraire delivered at $t+k$.

We define a new measure $\mathbb{P}^{*T}$ with respect to the objective measure $\mathbb{P}$ according to the Radon-Nikodym derivative

$$\left.\frac{d\mathbb{P}^{*T}}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \frac{M_T/M_t}{\mathbb{E}_t[M_T/M_t]} \tag{17}$$

for arbitrary $T$. Setting $T = t+k$ and using the risk-free rate in (16), we can rewrite this change of measure as $(d\mathbb{P}^{*T}/d\mathbb{P})|_{\mathcal{F}_t} = R^f_{t,T}\frac{M_T}{M_t}$, from which the basic pricing equation (14) can be rewritten as

$$S_t(X_T) = \frac{1}{R^f_{t,T}}\mathbb{E}^{*T}_t[X_T], \tag{18}$$

where $\mathbb{E}^{*T}_t[\cdot]$ is the conditional expectation under $\mathbb{P}^{*T}$, which we write simply as $\mathbb{E}^*_t[\cdot]$ and $\mathbb{P}^*$, respectively, when no confusion will arise. This equation tells us that the price of the random payoff $X_T$ is equal to the expectation of the payoff under $\mathbb{P}^{*T}$, discounted at the $(T-t)$-period risk-free rate, so the change of measure to $\mathbb{P}^{*T}$ incorporates the risk adjustment required to value payoffs at the relevant horizon.

We accordingly refer to $\mathbb{P}^{*T}$ as the *period-T risk-neutral measure*, or just the *risk-neutral measure*.[25] Given that (17) is specified for arbitrary $T$, we can in fact interpret that equation as defining a set of risk-neutral measures $\{\mathbb{P}^{*T}\}_{T\in\mathbb{N}}$.

[25]Some literature instead uses the notation $\mathbb{Q}$ for their definition of the risk-neutral measure or state-price density. Further, $\mathbb{P}^{*T}$ as defined here is also sometimes referred to as the *T-forward measure* rather than the risk-neutral measure; references include Jarrow (1987), Jamshidian (1989), Geman, El Karoui, and Rochet (1995), Hansen and Scheinkman (2016), and Qin and Linetsky (2017). Payoffs are discounted under $\mathbb{P}^{*T}$ 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 asset prices at $T$) are martingales under $\mathbb{P}^{*T}$, e.g., $f^m_{t,T} = \mathbb{E}^*_t[V^m_T]$. It is thus a natural *equivalent martingale measure* (Cox and Ross, 1976; Harrison and Kreps, 1979) for use in considering options over the future market-index value, as in our case.

**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 $T_i$ is

$$\mathbb{P}_t(R^m_{T_i} = s) = \sum_{\omega\,:\,R^m_{T_i}(\omega)=s} \mathbb{P}_t(\omega), \tag{19}$$

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

$$\begin{aligned}
\mathbb{P}^*_t(R^m_{T_i} = s) = \mathbb{P}^{*_{T_i}}_t(R^m_{T_i} = s) &= \sum_{\omega\,:\,R^m_{T_i}(\omega)=s} \mathbb{P}^{*_{T_i}}_t(\omega) \\
&= \sum_{\omega\,:\,R^m_{T_i}(\omega)=s} \frac{M_{T_i}(\omega)/M_t}{\mathbb{E}_t[M_{T_i}/M_t]}\,\mathbb{P}_t(\omega) \\
&= \frac{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s]}{\mathbb{E}_t[M_{T_i}/M_t]}\,\mathbb{P}_t(R^m_{T_i} = s),
\end{aligned} \tag{20}$$

where the last line uses the definition of the conditional expectation and (19).[26] So this gives us a natural generalization of the usual definition of risk-neutral probabilities (e.g., Campbell, 2017), $\mathbb{P}^*_t(\omega) = \frac{M_{T_i}(\omega)/M_t}{\mathbb{E}_t[M_{T_i}/M_t]}\,\mathbb{P}_t(\omega)$, which still holds here for elementary states, to the case we consider in which the return state $s$ can be realized under multiple possible underlying elementary states.

The risk-neutral pricing equation (18) can then be used to show that the date-$t$ schedule of option prices $\{q^m_{t,T_i,K}\}_K$ for a given expiration date $T_i$ reveals the set of risk-neutral probabilities $\{\mathbb{P}^*_t(R^m_{T_i} = s)\}_s$. 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 $s_1 < s_2 < \cdots < s_J$. Given a starting index value $V^m_{0_i}$, the definition of return states (13) gives a one-to-one mapping between return states and date-$T_i$ index values $V^m_{T_i}$, and we accordingly define $V^m_{i,j} \equiv V^m_{0_i} s_j$. Using this notation, the risk-neutral pricing equation (18) can be written for options as

$$\begin{aligned}
q^m_{t,T_i,K} &= \frac{1}{R^f_{t,T_i}} \mathbb{E}^*_t[\max\{V^m_{T_i} - K, 0\}] \\
&= \frac{1}{R^f_{t,T_i}} \left[ \sum_{j\,:\,V^m_{i,j} \geqslant K} (V^m_{i,j} - K)\, \underbrace{\mathbb{P}^*_t(V^m_{T_i} = V^m_{i,j})}_{\mathbb{P}^*_t(R^m_{T_i} = s_j)} \right].
\end{aligned} \tag{21}$$

We assume for notational simplicity that the set $\{V^m_{i,j}\}_j$ coincides with the set of traded option strike prices $\mathcal{K}$; that is, assume that the set of strikes is equal to the set of possible date-$T_i$ index values. (We will see that this can be relaxed.) Thus with $\mathcal{K} = \{K_1, K_2, \ldots, K_J\}$, we have $K_j = V^m_{i,j}$.

[26] Specifically, $\sum (M_{T_i}(\omega)/M_t)\,\mathbb{P}_t(\omega) = \left\{[\sum (M_{T_i}(\omega)/M_t)\,\mathbb{P}_t(\omega)] \,/\, [\sum \mathbb{P}_t(\omega)]\right\} \times \sum \mathbb{P}_t(\omega)$ (with all sums implicitly over $\{\omega : R^m_{T_i}(\omega) = s\}$), and this is equal to $\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s]\,\mathbb{P}_t(R^m_{T_i} = s)$.

Manipulating (21), we find that we can back out the risk-neutral probabilities of interest from option prices as follows:

$$\mathbb{P}_t^*(R_{T_i}^m = s_j) = R_{t,T_i}^f \left[ \frac{q_{t,T_i,K_{j+1}}^m - q_{t,T_i,K_j}^m}{K_{j+1} - K_j} - \frac{q_{t,T_i,K_j}^m - q_{t,T_i,K_{j-1}}^m}{K_j - K_{j-1}} \right]. \tag{22}$$

Appendix A contains a brief derivation of this result. This is the discrete-state equivalent of the classic result of Breeden and Litzenberger (1978), which can be written for continuous spaces as $\partial \mathbb{P}_t^*(R_{T_i}^m \leqslant s_j)/\partial V_{T_i,j}^m = R_{t,T_i}^f(\partial^2 q_{t,T_i,K_j}^m/\partial K_j^2)$. The result in (22) can be seen as the finite-difference approximation of that result for $K_{j+1} - K_j$ and $K_j - K_{j-1}$ small.

To gain intuition for this expression, assume first for simplicity that the gross risk-free rate is $R_{t,T_i}^f = 1$, and that strikes (and terminal price states) are one unit apart, $K_{j+1} = K_j + 1$. Consider a strategy of buying one unit of the strike-$K_j$ call option and selling one unit of the strike-$(K_j + 1)$ option. This strategy costs $q_{t,T_i,K_j}^m - q_{t,T_i,K_{j+1}}^m$ and yields payoff $\mathbb{1}\{V_{T_i}^m \geqslant K_j + 1\}$. *(Include a graph for intuition? Opening this to comments.)* Since the price of an asset or strategy is equal to its expected payoff under the risk-neutral measure by (18), we have $q_{t,T_i,K_j}^m - q_{t,T_i,K_{j+1}}^m = \mathbb{E}_t^*[\mathbb{1}\{V_{T_i}^m \geqslant K_j + 1\}] = \mathbb{P}_t^*(V_{T_i}^m \geqslant K_j + 1)$. In this way we can observe the complement of the cumulative distribution function for the terminal price (and thus for the return state), and differencing that CDF at adjacent points of the distribution allows us to recover the probability of any given return state, as in (22).

We see further from (22) that we need not have the set of strikes $\mathcal{K}$ coincide with the full set of possible date-$T_i$ index values to back out the risk-neutral probability $\mathbb{P}_t^*(R_{T_i}^m = s_j)$: we must simply have strikes at $V_{i,j-1}^m$, $V_{i,j}^m$, and $V_{i,j+1}^m$ to back it out exactly, and strikes near those values to obtain an approximation. Thus if options are not traded, for example, for strikes corresponding to the tails of the return distribution, we can still observe the risk-neutral probabilities in the middle of the distribution by using options with strikes corresponding to those values.

The fact that date-$t$ option prices reveal the risk-neutral probabilities for the date-$T_i$ return states without requiring us to account for the value of processes between $t$ and $T_i$ (e.g., one-period risk-free rates or dividends) motivates our use of the particular risk-neutral measure defined in (17).[27]

**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 (14), 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.[28] One might intuitively think of this agent as corresponding to "the market."

[27] For example, Giglio and Kelly (2017) use a one-period-ahead risk-neutral measure and accordingly face this issue with respect to risk-free rates. Separately, there is a recent literature focusing on the converse question of the pricing of so-called dividend strips at short horizons; see van Binsbergen and Koijen (2017) for a summary.

[28] 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.

The agent observes a finite vector of signals $\theta_t \in \Theta$ each period, with $\mathcal{F}_t = \sigma(\theta_\tau, 0 \leqslant \tau \leqslant t)$.[29] The information provided by the date-$t$ signals about date-$T_i$ return states is described by the likelihood function or *signal-generating process* $\mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R^m_{T_i})$, 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 signal history $\theta^{t-1} \equiv (\theta_0, \theta_1, \ldots, \theta_{t-1})$, and the signals are informative about the relative likelihood of arbitrary possible return states $s_a \neq s_b$ as long as $\mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R^m_{T_i} = s_a) \neq \mathbb{P}(\theta_t \mid \mathcal{F}_{t-1}, R^m_{T_i} = s_b)$.

The agent's time-$t$ *subjective belief distribution* over the return states for the market at $T_i$ is denoted by

$$\Pi_{t,T_i} = \{\pi_t(R^m_{T_i} = s)\}_{s \in \mathcal{S}}. \tag{23}$$

The agent brings beliefs $\Pi_{t-1,T_i}$ into the next period, observes the signals $\theta_t$, and then forms new beliefs $\Pi_{t,T_i}$. We assume that the agent has *rational expectations* over the return states for each option maturity date $T_i$, defined as follows in a standard manner (e.g., Muth, 1961, Fama, 1970, and more specifically as first described in Lucas and Prescott, 1971, Lucas, 1972, Green, 1975, Grossman, 1976, and Kreps, 1977).

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

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

$$\pi_{0_i}(R^m_{T_i} = s) = \mathbb{P}_{0_i}(R^m_{T_i} = s) \quad \forall\, s \in \mathcal{S}.$$

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

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

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

These two conditions together are equivalent to the agent's beliefs coinciding with the objective probabilities period by period. Rather than describing RE in this equivalent manner, 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.[30]

(See also Grossman, 1976, 1978.) 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 the Conclusion.

[29] This formulation implies that the history of signals describes the entirety of the agent's $\mathcal{F}_t$-measurable information, and the signal vector should accordingly be thought of as providing all time-$t$-available information about the state of the economy. This assumption is again for expositional simplicity as related to the discussion in Footnote 28.

[30] 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.

Given the maintained assumption of RE as in this definition, we can define the *risk-neutral belief distribution* without explicitly restricting the agent's utility function or constraint set by applying exactly the same change of measure defined in (17), using the general SDF $M_{T_i}/M_t$, to her subjective beliefs. This yields a risk-neutral belief distribution $\Pi^*_{t,T_i} = \{\pi^*_t(R^m_{T_i} = s)\}_{s\in\mathcal{S}}$ such that

$$\pi^*_t(R^m_{T_i} = s) = \frac{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s]}{\mathbb{E}_t[M_{T_i}/M_t]}\,\pi_t(R^m_{T_i} = s), \tag{24}$$

as in (20). Thus the risk-neutral-probability recovery result in (22) 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 $\Pi^*_{t,T_i}$, we instead consider restrictions on the behavior of the individual entries in the set $\{\widetilde{\pi}^*_{t,T_i,j}\}_{j=1,\ldots,J-1}$ defined by

$$\widetilde{\pi}^*_{t,T_i,j} \equiv \pi^*_t(R^m_{T_i} = s_j \mid R^m_{T_i} \in \{s_j, s_{j+1}\}) = \frac{\pi^*_t(R^m_{T_i} = s_j)}{\pi^*_t(R^m_{T_i} = s_j) + \pi^*_t(R^m_{T_i} = s_{j+1})}, \tag{25}$$

for $\pi^*_t(R^m_{T_i} = s_j) + \pi^*_t(R^m_{T_i} = s_{j+1}) > 0$.[31] In words, $\widetilde{\pi}^*_{t,T_i,j}$ describes the time-$t$ risk-neutral belief that return state $s_j$ will be realized for the market index at date $T_i$, 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 (25) is the "bad-state" (or low-return-state) probability given the ordering $s_1 < s_2 < \cdots < s_J$.[32] We collect the $(t, T_i)$ conditional risk-neutral beliefs in the vector $\widetilde{\boldsymbol{\pi}}^*_{t,T_i} \equiv (\widetilde{\pi}^*_{t,T_i,1}, \widetilde{\pi}^*_{t,T_i,2}, \ldots, \widetilde{\pi}^*_{t,T_i,J-1})'$, and we use this vector in our empirical estimation.

Given (24), these conditional risk-neutral beliefs can be written in terms of the underlying subjective beliefs in a manner analogous to the expression derived in (20):

$$\begin{aligned}\widetilde{\pi}^*_{t,T_i,j} &= \frac{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]\,\pi_t(R^m_{T_i} = s_j)}{\sum_{j'\in\{j,j+1\}}\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j'}]\,\pi_t(R^m_{T_i} = s_{j'})} \\ &= \frac{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]}{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} \in \{s_j, s_{j+1}\}]}\,\widetilde{\pi}_{t,T_i,j},\end{aligned} \tag{26}$$

where $\widetilde{\pi}_{t,T_i,j} \equiv \pi_t(R^m_{T_i} = s_j \mid R^m_{T_i} \in \{s_j, s_{j+1}\}) = \pi_t(R^m_{T_i} = s_j)/(\pi_t(R^m_{T_i} = s_j) + \pi_t(R^m_{T_i} = s_{j+1}))$ is the conditional subjective belief. These conditional subjective beliefs describe valid probability measures over the relevant subset of the return-state space, so the requirements of RE in Definition 1

[31] If $\pi^*_t(R^m_{T_i} = s_j) + \pi^*_t(R^m_{T_i} = s_{j+1}) = 0$, then for what follows, we can recursively define $\widetilde{\pi}^*_{t,T_i,j} = \widetilde{\pi}^*_{t-1,T_i,j}$ for any $j$ such that there exists some $\tau < t$ for which $\pi^*_\tau(R^m_{T_i} = s_j) + \pi^*_\tau(R^m_{T_i} = s_{j+1}) > 0$, and similarly for the conditional subjective beliefs after equation (26). Extending the definition in this way upholds the Kolmogorov axioms for both conditional beliefs and further maintains the martingale property of the conditional subjective beliefs under $\mathbb{P}_t(R^m_{T_i})$.

[32] We consider pairs of *adjacent* return states in (25), but we could equally well consider the conditional probability for state $s_j$ relative to $s_{j+k}$ for arbitrary $k$ in deriving our theoretical results. But when conducting our empirical estimation, considering adjacent states allows us to obtain a better estimate of local risk aversion (rather than risk aversion across distant states) required to rationalize the data, which motivates the choice used here.

extend to these beliefs.

We consider conditional probabilities 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.[33] 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.

**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 $\{\widetilde{\pi}^*_{t,T_i,j}\}_{0_i \leqslant t \leqslant T_i}$ from time $t_1$ to time $t_2 > t_1$ as

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

Denote movement for the full path by $m^*_{T_i,j} \equiv m^*_{0_i,T_i,T_i,j}$, and collect the set of movement values over all state pairs in the vector $\boldsymbol{m}^*_{t_1,t_2,T_i} \equiv (m^*_{t_1,t_2,T_i,1}, m^*_{t_1,t_2,T_i,2}, \ldots, m^*_{t_1,t_2,T_i,J-1})'$, and again for the full path $\boldsymbol{m}^*_{T_i} \equiv \boldsymbol{m}^*_{0_i,T_i,T_i}$. ‖

**DEFINITION 3 (*Uncertainty*).**

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

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

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

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

For the full path, $r^*_{T_i,j} \equiv r^*_{0_i,T_i,T_i,j}$, and again define $\boldsymbol{r}^*_{t_1,t_2,T_i} \equiv (r^*_{t_1,t_2,T_i,1}, r^*_{t_1,t_2,T_i,2}, \ldots, r^*_{t_1,t_2,T_i,J-1})'$, and $\boldsymbol{r}^*_{T_i} \equiv \boldsymbol{r}^*_{0_i,T_i,T_i}$. ‖

See Section 2.1 for discussions of these definitions. We can then measure empirical counterparts for these objects given result (22).[34]

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

[34] We aim to do so without excessive contamination from market microstructure noise, as discussed further below, and the infrequent (end-of-day) sampling we use to do so motivates the discrete-time framework used in this section (see also Bertsimas, Kogan, and Lo, 2000). We note further that these objects are invariant to the addition of nearest-

### 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 (17). 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,[35] *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 rational expectations in some way.

We begin by noting that the conditional risk-neutral belief $\widetilde{\pi}^*_{t,T_i,j}$ in (26) can be transformed into a likelihood ratio for state-$s_j$ vs. state-$s_{j+1}$ beliefs as follows:

$$\frac{\widetilde{\pi}^*_{t,T_i,j}}{1-\widetilde{\pi}^*_{t,T_i,j}} = \frac{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]}{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j+1}]} \frac{\widetilde{\pi}_{t,T_i,j}}{1-\widetilde{\pi}_{t,T_i,j}}, \tag{27}$$

since the risk-neutral belief for state $s_{j+1}$ conditional on either $s_j$ or $s_{j+1}$ being realized is $1-\widetilde{\pi}^*_{t,T_i,j}$. We define the first term on the right side of this equation as $\phi_{t,T_i,j}$:

$$\phi_{t,T_i,j} \equiv \frac{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]}{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j+1}]}. \tag{28}$$

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 (6). In the case of the representative-agent economy specified in (15), this value becomes $\phi_{t,T_i,j} = \mathbb{E}_t[U'(C_{T_i}) \mid R^m_{T_i} = s_j]/\mathbb{E}_t[U'(C_{T_i}) \mid R^m_{T_i} = s_{j+1}]$. 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 $\phi_{t,T_i,j} \geqslant 1$. 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 $s_1 < s_2 < \cdots < s_J$.

As can be seen in (28), the SDF slope $\phi_{t,T_i,j}$ in this general case requires the application of the conditional expectation for both states when evaluated as of time $t$, given that the return states do not necessarily correspond to elementary states in $\Omega$. Our more substantive restriction on the SDF is placed on those conditional expectations as follows. We first state the assumption formally, and then discuss its economic content by way of several remarks and examples.

neighbor-interpolated or -extrapolated data points $\widetilde{\pi}^*_{t,T_i,j}$ by definition (as, e.g., in Footnote 31), and therefore even if the panel $\{\widetilde{\pi}^*_{t,T_i}\}_{0_i \leq t \leq T_i}$ is unbalanced, it can be made balanced in this way without affecting $\boldsymbol{m}^*_{T_i}$ or $\boldsymbol{r}^*_{T_i}$. This is an advantage of this approach relative to, e.g., variance-ratio tests (Giglio and Kelly, 2017), or other tests as discussed in Footnote 19, for which adding additional observations with no change in beliefs reduces the power of the tests.

[35] This is no longer the case in an economy that admits arbitrage opportunities, but we are interested in the restrictiveness of the RE assumption (and expectations more generally) rather than the possibility of arbitrage, so we maintain the no-arbitrage assumption throughout. We discuss this issue further in Sections 4 and 6.

**DEFINITION 4 (*CTI*).** The SDF satisfies *conditional transition independence (CTI)* for the return-state pair $(s_j, s_{j+1})$ and option expiration date $T_i$ if $\phi_{t,T_i,j}$ defined in (28) is constant for all $0_i \leqslant t < T_i$ almost surely, and we denote this constant by $\phi_{i,j}$. ‖

**REMARKS:**

1. Stated intuitively, this condition requires that when we observe a change in the risk-neutral likelihood ratio (27), this is due to a change in the subjective conditional probability $\widetilde{\pi}_{t,T_i,j}$ rather than in the expected relative severity of the adjacent return states $s_j$ and $s_{j+1}$ (as encoded in the expected realization of the SDF or marginal utility in the given return states). This definition is thus analogous to the assumption discussed in Section 2.3 (page 14) in the context of the example used there. 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 $M_{T_i}$ (with $M_{T_i}$ proportional to marginal utility of terminal wealth), which guarantees that $\phi_{t,T_i,j}$ is constant.[36] The remainder of this subsection discusses under what conditions this logic can be extended to more recent asset-pricing frameworks.
2. Note that we have assumed only that the ratio of the conditional expectations for the SDF is constant over time across adjacent states $s_j$ and $s_{j+1}$. Of course $j$ and $j+1$ are arbitrary, 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 (24) arise from changes in the subjective beliefs term $\pi_t(R^m_{T_i} = s_j)$: there may be simultaneous time variation in the values in the numerator and denominator in (28), as in multiple cases discussed below, and the values $\mathbb{E}_t[M_{T_i}/M_t]$ and $\mathbb{E}_t[M_{T_i}]$ need not be constant.

   For example, in economies with variable rare disasters, one might worry about the violation of CTI for the left tail of the return distribution. Such an economy generates time variation in $\mathbb{E}_t[M_{T_i}/M_t]$ and thus some variation in raw risk-neutral beliefs arising due to the SDF (or risk aversion) and not the underlying subjective probabilities. But we can restrict attention to conditional probabilities for return states which, conditional on being realized as of time $T_i$, correspond to no disaster having been realized from $0_i$ to $T_i$ (e.g., return states in which returns are positive), and then this concern will no longer apply. It is true in this case that terminal prices or returns even in these high-return states may be affected by the time-$T_i$ value of the forward-looking expected disaster probability or severity, but *prices* need not meet any path-independence restriction; all that is required is that the time-$T_i$ realization of *marginal utility* meet such a restriction for some subset of states. See Example 3.
3. We further assume that CTI holds for those adjacent states only for a given option expiration date $T_i$, which allows for low-frequency changes in the SDF process across different expiration dates, but we will not lean heavily on this generality in what follows. Further, CTI must hold only for $0_i \leqslant t < T_i$ (and not $T_i$) because uncertainty is resolved as of $T_i$.

[36] 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 $T_i$ with savings fully invested in the market).

**4.** Definition 4 corresponds to a notion of transition or path independence because it requires that the realization of $M_{T_i}/M_t$ in return state $s_j$ depend in expectation only on $s_j$ and not on the path of any variables realized between $t$ and $T_i$ (though the return state itself *can* depend on such a path). This intuition is formalized in Lemma A.1 in Appendix A, which we use in proving our main results below.

We refer to the assumption in Definition 4 as *conditional* transition independence, rather than simply transition independence, to underscore that this assumption is less restrictive than the transition-independence assumption used by Ross (2015), among others, and discussed by Martin and Ross (2013). To see this, we note first that a sufficient condition for CTI to hold is for the conditional expectation of SDF growth from $t$ to $T_i$ to admit a factorization

$$\mathbb{E}_t\left[\frac{M_{T_i}}{M_t}\,\middle|\, R^m_{T_i} = s_{j'}\right] = f(s_{j'})\,\psi_t, \tag{29}$$

for $0_i \leqslant t < T_i$ and $j' = j, j+1$, where $f\colon \mathcal{S} \to \mathbb{R}_{++}$ is a function mapping from the induced return space (see Footnote 23) to the positive real line and $\psi_t\colon \Omega \to \mathbb{R}_{++}$ is an $\mathcal{F}_t$-measurable random variable. The transition-independence assumption of, e.g., Ross (2015) specifies instead that such a factorization for the SDF must hold deterministically for all return states rather than just as a conditional expectation for some subset of states.[37] We discuss this distinction further below in the context of our examples.

**5.** The factorization (29) also begins to make clear the conditions that would lead to a violation of CTI. For example, as discussed in Section 2.3 and formalized below, the Campbell and Cochrane (1999) habit-formation model implies that the realization of the SDF at date $T_i$ depends on the path of consumption from $t$ to $T_i$ and not just the terminal return, and therefore a factorization like (29) would require an additional term for that consumption path.

**6.** We note, however, that nothing about our theoretical framework requires considering beliefs over *return* states, and all the results below would apply equally if we were to consider beliefs over the *elementary* states in $\Omega$. In this case, of course, the assumption in Definition 4 would hold trivially, since the elementary state fully determines the realization of the SDF. And even considering beliefs over some more granular partition of the state space than the elementary states themselves, we could in theory define this partition in such a way that this assumption would hold regardless of the setting.

For example, returning to Campbell and Cochrane (1999), the SDF is fully determined by the realization of consumption $C_{T_i}$ and *surplus* consumption $S^c_{T_i}$ (i.e., consumption relative to habit). Therefore since $(M_{T_i}/M_t \mid C_{T_i} = c_j, S^c_{T_i} = s^c_{j'})$ is deterministic for any state $(c_j, s_{j'})$, the SDF would meet CTI for this redefined (consumption) $\times$ (surplus consumption) space. But while this works in theory, empirical implementation is generally infeasible: we observe

[37] Specifically, Ross (2015) writes $M_{T_i}/M_t$ as a deterministic function of the current and terminal price states (which we can denote as $V^m_{i'}, V^m_{j'}$, respectively) as follows: $(M_{T_i}/M_t)(V^m_{i'}, V^m_{j'}) = \delta\, g(V^m_{j'})/g(V^m_{i'})$ for some function $g$ and constant $\delta$, for *all* states. Others making an equivalent assumption include Gerber and Shiu (1994), Heston (2004), Jensen, Lando, and Pedersen (2016), Qin and Linetsky (2016), and Walden (2016).

options and risk-neutral beliefs over traded prices (and functions thereof) rather than over, e.g., surplus-consumption values.[38] 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 the theory and our empirical implementation. And we further use *market* return states (rather than, e.g., individual-equity return states) because the CTI assumption holds in a wider range of contexts for this market return than for individual-asset returns.

We now turn to a set of examples to illustrate the above points more concretely, all of which work under the maintained hypothesis of rational expectations. All proofs for the statements in the examples are contained in Appendix A.

**EXAMPLE 1.** Assume an economy with a one-dimensional state variable $A_t : \Omega \to \mathbb{R}$ (e.g., productivity, consumption, consumption volatility, expected long-run consumption). This process may be non-stationary but is assumed to satisfy the Markov property, $\mathbb{P}(A_{t+\tau} = a \,|\, A_t, A_{t-1}, \ldots) = \mathbb{P}(A_{t+\tau} = a \,|\, A_t)$ for all $\tau \geqslant 0$ and $a \in \mathbb{R}$. Assume further that there exists a representative agent with either time-separable utility or Epstein-Zin (1989) recursive utility over the consumption process $\{C_t(A_t)\}$ and that the market index pays dividends according to $\{D_t(A_t)\}$, where this process is arbitrary but yields a stationary price-dividend ratio. Then conditional transition independence holds for any two adjacent return states. ‖

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

$$c_t = g + \rho c_{t-1} + h(c_{t-1})\varepsilon_t, \tag{30}$$

where $\varepsilon_t$ is i.i.d. with arbitrary distribution (so it could incorporate the possibility of disasters) and $h(\cdot)$ is a state-dependent volatility function. Setting $\rho = 0$, $h(\cdot) = 1$ corresponds to the case of i.i.d. consumption; setting $\rho = 1$, $h(\cdot) = 1$ corresponds to i.i.d. consumption growth (or a random walk with drift $g$); setting $\rho = 0$, $h(c_{t-1}) = (c_{\max} - c_{t-1})^{1/2}$, where $c_{\max}$ is the maximum possible value for log consumption, corresponds to stationary consumption with stochastic volatility; we could in fact modify $g$ to incorporate state-dependent growth or $\rho$ to incorporate state-dependent mean reversion; and so on. We could further model dividends according to $D_t = C_t^\lambda$ with $\lambda \geqslant 0$ (Campbell, 1986; Abel, 1999; Martin, 2013); $\lambda = 1$ corresponds to an unlevered consumption claim,

[38] The same applies, for example, to economies with recursive preferences and independent stochastic-volatility shocks (e.g., "Case II" of the long-run-risks model of Bansal and Yaron, 2004, or more generally models similar to that of Merton, 1973a), since then marginal utility at $T_i$ can depend on expected future volatility in addition to terminal consumption (or returns). One could overcome this with options traded over the joint realization of the index return and volatility to infer the risk-neutral distribution over that two-dimensional space, but such options (sometimes referred to as "Everest options," e.g., Quessette, 2002) are not currently traded outside of non-standardized, over-the-counter trades. See Borovička, Hansen, and Scheinkman (2016) and Walden (2016) for related discussion.

while $\lambda > 1$ represents a levered claim. 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.

**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 conditional transition independence holds for any two adjacent return states. ‖

These two preceding examples clarify that the CTI restriction allows for time variation in discount rates and risk premia; see Section 2.3 for a discussion. Relatedly, these examples illustrate that both temporary and permanent shocks to consumption and marginal utility (and thus the SDF $M_t$) are in principle admissible under CTI, as discussed in Remark 2. 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 $\mathbb{E}_t[M_{T_i}/M_t]$ since $M_{t+1}/M_t$ is i.i.d. (and, as is shown in Appendix A, does not change the relevant ratio of conditional expectations in (28) under the assumptions in the examples); meanwhile, an economy with i.i.d. consumption has only transitory shocks to $M_t$ and a fixed value $\mathbb{E}_t[M_{T_i} \mid R^m_{T_i} = s_j]$ for all $s_j$, so that the factorization (29) is immediate with $\psi_t = M_t^{-1}$.

This flexibility is desirable on both empirical and theoretical grounds. Empirically, Alvarez and Jermann (2005) and Bakshi and Chabi-Yo (2012) argue that permanent shocks to the SDF are quite important for quantitatively matching the observed moments of returns. Theoretically, Borovička, Hansen, and Scheinkman (2016), building on work by Hansen and Scheinkman (2009) and Hansen (2012a), 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 4) — do not allow for any permanent shocks to the SDF, while many asset-pricing models do feature such shocks. Appendix B.1 discusses the formal relation between our conditional transition independence assumption and the decompositions used by Borovička, Hansen, and Scheinkman (2016).

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

**EXAMPLE 3.** Consider the variable-rare-disasters model of Gabaix (2012). Assume a representative agent with CRRA consumption utility and that log consumption $c_t$ and log dividends $d_t$ 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,[39] and $B_{t+1}$ and $F_{t+1}$ are arbitrarily correlated random variables with support $[0, 1]$ that affect consumption and dividend growth respectively in the case of a disaster in period $t+1$, which occurs with probability $p_t$. (Gabaix excludes the i.i.d. shock $\varepsilon^c_{t+1}$ in his main specification, but we include it here for completeness.) Define *resilience* $H_t$ according to $H_t = p_t \mathbb{E}_t[B_{t+1}^{-\gamma} F_{t+1} - 1 \mid \mathbb{1}\{\text{disaster}_{t+1}\}]$, write $H_t = H_* + \widehat{H}_t$, and assume that the variable part follows a *linearity-generating process* (Gabaix, 2009),

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

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

Under the above assumptions, given any market-index option horizon $T_i - 0_i$ and any (small) positive value $\delta$, 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:

$$\mathbb{P}_{0_i}\left( \sum_{t=1}^{T_i} \mathbb{1}\{\text{disaster}_t\} > 0 \;\middle|\; R^m_{T_i} \geqslant \underline{s} \right) < \delta.$$

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

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, as in Remark 5, 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). 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* $S^c_t \equiv (C_t - H_t)/H_t$, assume that $s^c_t \equiv \log(S^c_t)$, $c_t \equiv \log(C_t)$, and log dividends

[39] To be complete with respect to our discrete-state setting above, 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).

$d_t \equiv \log(D_t)$ evolve respectively according to

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

where $\varepsilon_{t+1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2_\varepsilon)$ (see Footnote 39), $\eta_{t+1} \overset{\text{i.i.d.}}{\sim} \mathcal{N}(0, \sigma^2_\eta)$, $\text{Corr}(\varepsilon_{t+1}, \eta_{t+1}) = \rho$, and the *sensitivity function* $\lambda(s^c_t)$ is specified as

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

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

Under these assumptions, conditional transition independence in general fails to hold. ‖

Given this model's specification of habit formation, revisions to the expected path of consumption cause revisions to the expected SDF in such a way that it cannot be written in the form (29), $f(s_j)\psi_t$, in the same way across multiple return states $s_j$. Put more simply, the path of consumption always matters. We later consider the effects of such violations of CTI in the context of a calibrated simulation study.

## 4. Theoretical Results

We move now to our main theoretical results. With the exception of the extension considered in Proposition 7, we assume throughout that conditional transition independence holds for some subset of the return-state space. To reduce notational burden, we assume that this CTI-conforming subset of the return space is fixed across all option expiration dates $\{T_i\}_{i=1,\ldots,N}$,[40] and define $\mathcal{S}_{\text{CTI}} \subseteq \{s_j\}_{j=1,\ldots,J-1}$ according to

$$\mathcal{S}_{\text{CTI}} \equiv \left\{s_j : \text{CTI holds for the pair } (s_j, s_{j+1})\right\}. \tag{31}$$

Our bounds consider the relationship between risk-neutral belief movement and risk-neutral uncertainty resolution as defined 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. As usual, all proofs are in Appendix A.

[40] Note also that we have assumed without loss of generality that the return space $\mathcal{S}$ is fixed across expiration dates.

### 4.1. Main Bounds

**PROPOSITION 1.** *For any return-state pair $(s_j, s_{j+1})$ such that $s_j \in \mathcal{S}_{\text{CTI}}$, the following bound must hold under rational expectations for arbitrary option expiration date $T_i$:*

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

This result relates the unobserved structural parameter $\phi_{i,j}$, which corresponds to the slope of the SDF across the two adjacent return states, to a set of observable values.[41] (We discuss the observability of these values by way of an additional result in the next subsection, but for now take as given that they are observable.) Under risk neutrality ($\phi_{i,j} = 1$), this upper bound becomes zero: belief movement must not exceed uncertainty resolution on average, following Lemma 1.[42] 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 $\phi_{i,j}$. 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 for $\phi_{i,j} > 1$ 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.

Unlike in the example considered in Section 2.3, 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 $\widetilde{\pi}^*_{0_i,T_i,j}$ and $\phi_{i,j}$, what signal-generating process would a "malevolent" outside agent who seeks to maximize admissible excess movement choose, and what would the value of admissible excess movement be under that process? While our proof of the proposition above does not require fully characterizing this maximizing signal-generating process, we can in fact show further that such a maximizing process can be characterized in closed form, and that the signal structure under this process stays the same regardless of the current belief $\widetilde{\pi}^*_{t,T_i,j}$. We state this finding in the following proposition.

**PROPOSITION 2.** *There is a unique signal-generating process that maximizes admissible excess movement $\mathbb{E}_{0_i}[m^*_{T_i,j} - r^*_{T_i,j}]$. Further, as $T_i - 0_i \to \infty$, this process yields the bound in Proposition 1.*

The second statement in the above proposition, that the bound in Proposition 1 becomes tight under this particular process as the option-expiration horizon approaches infinity, implies that this

[41] This result can be written equivalently in vector form across all states as

$$\mathbb{E}_{0_i}[\boldsymbol{m}^*_{T_i} - \boldsymbol{r}^*_{T_i}] \leqslant \widetilde{\boldsymbol{\pi}}^{*\prime}_{0_i,T_i} \left\{ I_{J-1} - \text{diag}\left[\widetilde{\boldsymbol{\pi}}^{*\prime}_{0_i,T_i} + \text{diag}(\boldsymbol{\phi}_i')(\iota_{J-1} - \widetilde{\boldsymbol{\pi}}^*_{0_i,T_i})\right]^{-1} \right\} \widetilde{\boldsymbol{\pi}}^*_{0_i,T_i},$$

where $I_{J-1}$ is the $(J-1) \times (J-1)$ identity matrix, $\iota_{J-1}$ is a $(J-1) \times 1$ vector of ones, $\boldsymbol{\phi}_i' = (\phi_{i,1}, \phi_{i,2}, \ldots, \phi_{i,J-1})'$, and the result holds for any entries with indices $j$ such that $s_j \in \mathcal{S}_{\text{CTI}}$.

[42] An associated lower bound of zero can also be shown to hold in this case.

bound is otherwise somewhat more conservative than necessary for finite horizons. 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 $\epsilon$ 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. In the limit as $T_i - 0_i \to \infty$, *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. To the extent that one thinks that such a process is unrealistic relative to whatever true process governs revisions in agents' beliefs, this means that the bound in Proposition 1 is conservative as compared to the admissible excess movement under this true process.

While the bound in Proposition 1 in general maps between observed values and the slope of 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 $\phi_{i,j} \to \infty$ in that bound yields the following full-identification corollary.

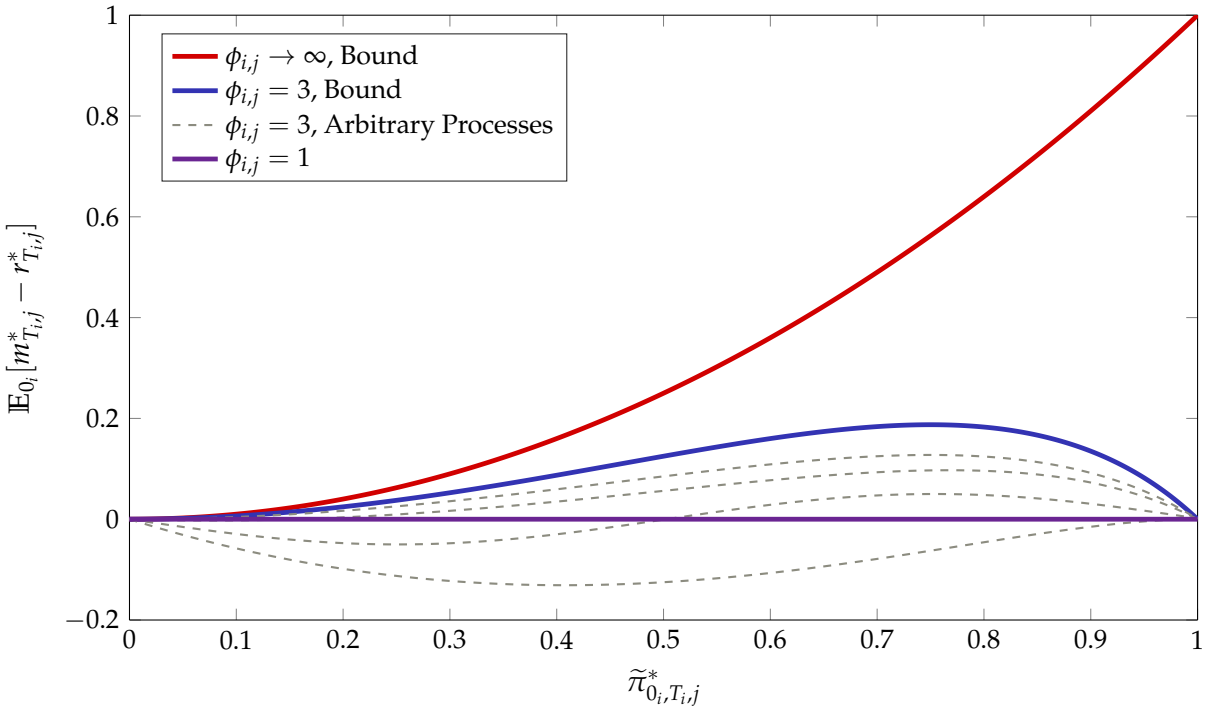
**COROLLARY 1.** *If* $\mathbb{E}_{0_i}[m^*_{T_i,j} - r^*_{T_i,j}] > \widetilde{\pi}^{*\,2}_{0_i,T_i,j}$*, then no SDF process meeting conditional transition independence can rationalize the variation in risk-neutral beliefs for the given return-state pair.*

In the case that the condition in this corollary is found to be met, 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 very 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.[43]

The other possibility in the case of excess volatility is of course that the CTI condition is violated. But such a finding of excess volatility would nonetheless imply that any model with an SDF that does satisfy CTI cannot match the relevant moments of the data without some distortion in beliefs relative to RE.

[43] This strategy would take the form of betting on mean reversion in risk-neutral beliefs. See also Hansen, Heaton, and Luttmer (1995), Hansen and Jagannathan (1997), Hansen (2014), and Ghosh, Julliard, and Taylor (2017) for discussions regarding the possibility of apparent mispricing under the assumption of no arbitrage. The discussion in Hansen (2014) makes clear the distinction of our full-identification result in Corollary 1 from previous literature; in these cases, if "investor beliefs are not in full accord with an underlying model, [this] can result in investor behavior that resembles extreme risk aversion" (p. 948). This intuition is, however, present (and formalized) in our Proposition 1 for finite $\phi_{i,j}$.

**Figure 2: Excess Belief Movement vs. Prior by $\phi_{i,j}$ Under RE**




*Notes:* Bounds are obtained from the formulas in Proposition 1 and Corollary 1, and from Lemma 1 for $\phi_{i,j} = 1$.

Taken together, Proposition 1 and Corollary 1 characterize the admissible excess movement in risk-neutral beliefs as a function of $\phi_{i,j}$ 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 $\phi_{i,j} = 1$: in this case, excess movement must be equal to 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, $\phi_{i,j} = 3$. These lines show that 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 $\phi_{i,j} = 3$ 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 $\phi_{i,j} \to \infty$, which is equal to the squared risk-neutral prior from Corollary 1.

The non-monotonicity of the bound in the $\phi_{i,j} = 3$ case is a general feature of the bounds for $1 < \phi_{i,j} < \infty$,[44] 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

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

risk-neutral prior corresponds to a greater subjective prior for any given value of $\phi_{i,j}$, which decreases the likelihood that such a downward movement will be realized. In the limit as $\phi_{i,j} \to \infty$, 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.

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.** *An incorrect prior, $\pi_{0_i}(R^m_{T_i} = s_{j'}) \neq \mathbb{P}_{0_i}(R^m_{T_i} = s_{j'})$ for $j' \in \{j, j+1\}$, cannot by itself lead to $\mathbb{E}_{0_i}[m^*_{T_i,j} - r^*_{T_i,j}] > \widetilde{\pi}^{*\,2}_{0_i,T_i,j}$.*

This result tells us that updating behavior must play some role in any outright rejection of RE as in Corollary 1. An incorrect prior effectively 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,[45] this excess movement is not enough to generate 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 the Conclusion.

### 4.2. Implementation

The discussion to this point has taken as given that all the values in the bound in Proposition 1 are observable aside from $\phi_{i,j}$, but this has elided one issue: this bound is stated over the date-$0_i$ expectation of excess movement for risk-neutral beliefs over uncertainty resolving at expiration date $T_i$, but we observe only one draw $m^*_{T_i,j} - r^*_{T_i,j}$ per expiration date rather than the ex-ante expectation of this random variable. Thus we need some way of aggregating the observed values $\{m^*_{T_i,j} - r^*_{T_i,j}\}_i$ for a given return state $s_j$ *across* all the observed expiration dates $\{T_i\}_i$, in order to conduct valid inference on the average slope of the SDF for that return-state pair, $\mathbb{E}[\phi_{i,j}]$, over the entire panel of available data.

Due to Jensen's inequality, we cannot simply insert $\mathbb{E}[\phi_{i,j}]$ in place of $\phi_{i,j}$ and $\mathbb{E}[\widetilde{\pi}^*_{0_i,T_i,j}]$ in place of $\widetilde{\pi}^*_{0_i,T_i,j}$ when taking the unconditional 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 $\widetilde{\pi}^*_{0_i,T_i,j}$, this function's Hessian is not in general negative semi-definite given the non-monotonicity in the bound, so naïvely conducting such replacements will not in general uphold the validity of the bound. However, we do have that for any given prior $\widetilde{\pi}^*_{0_i,T_i,j}$, the bound is concave in $\phi_{i,j}$, as the second partial derivative of that expression with respect to $\phi_{i,j}$ is $-2(\phi_{i,j}-1)^2\phi^2_{i,j}/(\widetilde{\pi}^*_{0_i,T_i,j} + \phi_{i,j}(1 - \widetilde{\pi}^*_{0_i,T_i,j})) < 0$. We show in Appendix A that this implies the following result, which effectively applies Jensen's inequality for one of several variables and accordingly places the value $\mathbb{E}[\phi_{i,j}]$ inside the unconditional expectation of the right side of the bound in Proposition 1.

[45] Since the eventual return state $R^m_{T_i}$ is a function of the true underlying elementary state $\omega \in \Omega$, one can think of the "correct" belief from a frequentist standpoint as being either 0 or 1.

**PROPOSITION 4.** *For any return-state pair $(s_j, s_{j+1})$ such that $s_j \in \mathcal{S}_{\text{CTI}}$, the following bound must hold under rational expectations over all option expiration dates:*

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

Given that this bound applies unconditional expectations, it is now implementable empirically: we can measure a sample counterpart of the excess movement statistic on the left side, and the minimum $\mathbb{E}[\phi_{i,j}]$ that solves the bound given the observed excess movement and risk-neutral priors is accordingly our estimate of the average slope of the SDF for the return-state pair in question.[46] And while these estimates will in general differ across return states $s_j$, we can apply this result again to obtain a single estimate of the average 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. Extensions and Further Results

We now turn to a set of additional results that sharpen the bound given above, allow us to account for the possibility of mismeasurement or market microstructure noise, extend the results to cases in which conditional transition independence is violated mildly, and finally yield an interpretation for $\phi_{i,j}$ in terms of a structural risk-aversion value.

First, the bound given above assumes that the conditional belief eventually resolves to certainty as of period $T_i$, or $\widetilde{\pi}^*_{T_i,T_i,j} \in \{0, 1\}$. This is reasonable to assume in a two-state world: in this case, as of period $T_i$, the return state is either $s_j$ or $s_{j+1}$ with certainty. But in the more general multi-state world considered here, we may have that $\pi^*_t(R^m_{T_i} \in \{s_j, s_{j+1}\})$ is quite small (or zero) if it is very unlikely that either of those two return states occurs. In this case, we will treat the associated conditional risk-neutral probability $\widetilde{\pi}^*_{t,T_i,j}$ as missing in our empirical implementation given that these observations would otherwise be particularly susceptible to estimation noise. This is equivalent to fixing $\{\widetilde{\pi}_{t,T_i,j}\}_{\tau<t\leqslant T_i}$ at $\widetilde{\pi}_{\tau,T_i,j}$ for the last date $\tau$ on which $\pi^*_\tau(R^m_{T_i} \in \{s_j, s_{j+1}\})$ is above some threshold (see Footnote 34), implying that it may be the case that $\widetilde{\pi}^*_{T_i,T_i,j} \notin \{0, 1\}$.

The assumption of full resolution of the conditional belief is nonetheless valid in deriving Proposition 1, as shown in the proof of that result, because this assumption in fact biases the bound toward being overly conservative if uncertainty does not necessarily resolve fully. To see this, consider Lemma 1. If expected uncertainty resolution ($\mathbb{E}[r_{0,T}]$ in the lemma) is not equal to the value that would be obtained under full resolution, then admissible belief movement decreases; equivalently, the less beliefs are expected to converge toward certainty, the less they can move under RE. Thus considering the possibility of incomplete resolution of uncertainty in the conditional belief allows for a tighter version of Proposition 1, which we formalize in the following result.

[46] In this context, the risk-neutral prior $\widetilde{\pi}^*_{0_i,T_i,j}$ plays the role of conditioning information in a sense similar to that of Hansen and Richard (1987), and see also Grossman and Shiller (1981) and Hansen (2014).

**PROPOSITION 5.** *For any return-state pair $(s_j, s_{j+1})$ such that $s_j \in \mathcal{S}_{\text{CTI}}$, the following bound must hold under rational expectations for arbitrary option expiration date $T_i$:*

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

The term on the second line accounts for the possible partial (rather than full) resolution of uncertainty: $r^*_{T_i,j} = r^*_{0_i,T_i,T_i,j} = \widetilde{\pi}^*_{0_i,T_i,j}(1 - \widetilde{\pi}^*_{0_i,T_i,j}) - \widetilde{\pi}^*_{T_i,T_i,j}(1 - \widetilde{\pi}^*_{T_i,T_i,j})$ from Definition 3, so the second line in the bound becomes zero, or $\mathbb{E}_{0_i}[r^*_{T_i,j}] = \widetilde{\pi}^*_{0_i,T_i,j}(1 - \widetilde{\pi}^*_{0_i,T_i,j})$, if and only if uncertainty resolves in all cases as of period $T_i$ so that $\widetilde{\pi}^*_{T_i,T_i,j}(1 - \widetilde{\pi}^*_{T_i,T_i,j}) = 0$. Otherwise the term in this second line is positive, as it incorporates the deviation between full uncertainty resolution $\widetilde{\pi}^*_{0_i,T_i,j}(1 - \widetilde{\pi}^*_{0_i,T_i,j})$ and actual expected uncertainty resolution $\mathbb{E}_{0_i}[r^*_{T_i,j}]$. In this way it is in fact a more general version of the bound in Proposition 1. We can then obtain a similar result to Corollary 1 by taking $\phi_{i,j} \to \infty$ in this case, and we can apply unconditional expectations as in Proposition 4 for empirical implementation.[47]

We now provide a result that can be applied to either of the bounds above to correct for possible mismeasurement of risk-neutral beliefs (for example, given issues related to market microstructure). 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, the following proposition shows that a simple correction can be applied to account for this issue.

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

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

*where $\epsilon_{t,T_i,j}$ is i.i.d. over $t$ and $j$, and $\mathbb{E}[\epsilon_{t,T_i,j} \mid \widetilde{\pi}^*_{t,T_i,j}] = 0$. Denoting the observed one-period expected excess movement statistic by $\mathbb{E}_t[\widehat{m}^*_{t,t+1,T_i,j} - \widehat{r}^*_{t,t+1,T_i,j}]$, its relation to the true value $\mathbb{E}_t[m^*_{t,t+1,T_i,j} - r^*_{t,t+1,T_i,j}]$ is then given by*

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

We view the assumption that $\epsilon_{t,T_i,j}$ is i.i.d. to be conservative: to the extent that measurement error is positively correlated over time, this will reduce the upward bias in measured volatility of beliefs.[48] The expectations in this result are time-$t$ conditional expectations, but applying the

[47] In this case we must take a bit more care given that the additional term in Proposition 5 is convex in $\phi_{i,j}$. We avoid this issue by linearizing this second term in $\phi_{i,j}$ in such a way that the associated bound subsumes the bound in Proposition 5 and is thus slightly more conservative than this bound; see Appendix B.2 for details.

[48] One might worry about negatively correlated measurement errors in the case of bid-ask bounce (rather than simply mismeasurement within the bid-ask spread on any given day). 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

law of iterated expectations allows this result to be applied to any of the bounds above. Empirical implementation of this result requires estimation of the value $\text{Var}(\epsilon_{t,T_i,j})$, which we discuss 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* $\phi_{t,T_i,j} \equiv \mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]/\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j+1}]$ *in equation (28), and that conditional transition independence assumes that this value is constant for all* $0_i \leqslant t < T_i$*. Relaxing this assumption, we obtain the following two results:*

*(i) If the sequence* $\{\phi_{t,T_i,j}\}_t$ *is a martingale with respect to the filtration* $\mathbb{F}$*, or* $\phi_{t,T_i,j} = \mathbb{E}_t[\phi_{t+1,T_i,j}]$*, then the bound in Proposition 1 still applies, with the ex-ante value* $\phi_{0_i,T_i,j}$ *replacing the constant* $\phi_{i,j}$*.*

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

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

*These results apply as well to the bounds in Propositions 4 and 5.*

Part (i) of the result tells us that a martingale process for $\{\phi_{t,T_i,j}\}_t$ yields no more excess movement in beliefs for any given ex-ante value $\phi_{0_i,T_i,j}$ than is expected under the movement-maximizing signal-generating process used in Proposition 1 itself.[49] But the sequence $\{\phi_{t,T_i,j}\}_t$ is unlikely to follow a martingale exactly: it is a ratio of two martingales (since conditional expectations are martingales under RE), and such a ratio is not in general itself a martingale, though we provide

for such noise have long died out at one-day lags. The same is true regarding the effects of price discreteness. But demand pressures may lead to positively autocorrelated "noise" (see Section 5.5), in which case the i.i.d. assumption is conservative. It is also the case in what follows below that we need only measure the component of the noise that is uncorrelated with past values, so the true structure of measurement error is in this sense irrelevant: we simply treat $\epsilon_{t,T_i,j}$ as this uncorrelated component and then measure it in Section 5.5.

[49] 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 $\phi_{t,T_i,j}$ does *not* induce mean reversion in measured (or actual) risk-neutral beliefs. The proof of Proposition 7 then shows that the variation arising from this martingale in fact lowers the expected excess belief movement period by period regardless of the variance of the increments of this martingale, because the signal-generating process used in the bound in Proposition 1 can always do slightly better in maximizing excess movement in the absence of this form of variation in $\phi_{t,T_i,j}$.

a necessary and sufficient condition for this to be the case in Appendix B.3. This martingale condition is, however, a convenient benchmark around which to obtain an approximate bound as done in part (ii) of the proposition, which shows that the original bound is continuous in the limit as $\mathbb{E}_t[\phi_{t+1,T_i,j}] \to \phi_{t,T_i,j}$.[50] As we will see in Section 6, $\phi_{t,T_i,j}$ 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. In practice, we expect that most models do not generate enough variation in the relative expected marginal utility values across adjacent market-return states to generate large departures from the $\phi_{t,T_i,j}$-martingale benchmark, but we have yet to conduct an extensive analysis of this conjecture.

Finally, while all the results above are convenient to express in terms of the SDF slope $\phi_{i,j}$ 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, they 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.[51] Following the steps for equation (10) in that section, we obtain the following result.

**PROPOSITION 8.** *In the case that we additionally assume that there is a representative agent with utility (or indirect utility) over time-$T_i$ wealth, with wealth equal to the market index value, relative risk aversion $\gamma_{i,j} \equiv -V_{i,j}^m U''(V_{i,j}^m)/U'(V_{i,j}^m)$ is given to a first order around return state $s_j$ by*

$$\gamma_{i,j} = \frac{\phi_{i,j} - 1}{\Delta_j},$$

*where $\Delta_j \equiv (s_{j+1} - s_j)/s_j \approx \log(s_{j+1}) - \log(s_j)$ 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 — which in this representative-agent framework is equal to the curvature of utility — is proportional to $\phi_{i,j} - 1$, 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 $\Delta_j$. If, for example, $s_j = 1$, $s_{j+1} = 1.01$, then a value $\phi_{i,j} = 1.1$ implies $\gamma_{i,j} = 10$.

Given that the expression for $\gamma_{i,j}$ in this proposition is affine in $\phi_{i,j}$, the aggregation result in Proposition 4 can be applied to conduct valid estimation of $\mathbb{E}[\gamma_{i,j}]$ 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 present empirical results in the following section for both the general SDF slope $\phi_{i,j}$ and the constructed risk-aversion value $\gamma_{i,j}$.

[50] We would like to modify this result to relate the value $\epsilon$ to the order of $\delta$, but have not done so yet.

[51] We once again assume in this case that $\phi_{i,j}$ is constant given that CTI holds in such a representative-agent framework, as discussed in Section 2.3.

# 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 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, which we require to measure risk-neutral beliefs as discussed in the next subsection. 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 $\widetilde{\pi}^*_{t,T_i,j} = \pi^*_{0_i}(R^m_{T_i} = s_j)/(\pi^*_{0_i}(R^m_{T_i} = s_j) + \pi^*_{0_i}(R^m_{T_i} = s_{j+1}))$ 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 $\widetilde{\pi}^*_{t,T_i,j}$ as missing when this condition is violated. This procedure is meant to address that conditional beliefs $\widetilde{\pi}^*_{t,T_i,j}$ are likely to be particularly susceptible to mismeasurement when the underlying beliefs are close to zero.

OptionMetrics also 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.[52] 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

[52] 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.

(CBOE) website, which posts these settlement values.

### 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 (13) according to the following set:

$$\mathcal{S}_{\text{baseline}} = \exp\left(\{(-\infty, -0.11), -0.09, -0.07, \ldots, 0.07, 0.09, (0.11, \infty)\}\right), \tag{32}$$

so that return states are evenly spaced two percentage points apart, bracketed by the two extreme log-return states $(-\infty, -0.11)$ and $(0.11, \infty)$. We typically report our estimation results in terms of these log-return states rather than the (approximate) equivalent exponentiated gross-return states $\{(0, 0.89), 0.91, 0.93, \ldots, 1.07, 1.09, (1.11, \infty)\}$; 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

$$\log(R^{m,e}_{T_i}) \equiv \log(V^m_{T_i}/V^m_{T_0}) - \log(R^f_{0_i,T_i}) = s_j, \tag{33}$$

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 as given in (16), measured using the zero-coupon risk-free yield curve provided by OptionMetrics.[53] 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 (13) for each option expiration date.

In general, we of course do not observe excess returns of, say, exactly 0.05, and the risk-neutral probability of realizing this return value is accordingly approximately zero. Rather than attempting to measure the risk-neutral density exactly at the points in (32) (or at additional points), which would yield highly imprecise estimates, we instead discretize the return-state space as follows: we say that interior return state $s_j$ ($j \neq 1, J$) is realized if the excess return is in a neighborhood of one percentage point of $s_j$, or $\log(R^{m,e}_{T_i}) \in [s_j - 0.01, s_j + 0.01)$, and we extend the extreme return states $s_1, s_J$ by 0.01 at the upper and lower ends, respectively, of their return ranges above. Thus one could equivalently consider our return-state specification to be in terms of two-percentage-point *bins* for excess returns: $\mathcal{S}_{\text{baseline}} = \exp(\{(-\infty, -0.10), [-0.10, -0.08), [-0.08, -0.06), \ldots, [0.06, 0.08), [0.08, 0.10), [0.10, \infty)\})$, as we use the measured risk-neutral probabilities for the return states defined by these bins to approximate the probabilities for the points in (32).

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 for the conditional risk-

[53] 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.

neutral beliefs for adjacent return states, 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). To take a degenerate example, it is unlikely that the expected realization of the SDF in the first extreme return state in the binned version of the return-state space, $s_1 = (-\infty, -0.10)$, is constant relative to the expected realization in $s_2 = [-0.10, -0.08)$: if the return being in the range $\log(R^{m,e}_{T_i}) \in [-0.20, -0.15)$ is more likely than in $[-0.15, -0.10)$, then the expected realization of the SDF in $s_1$ will be high; if these relative probabilities change in favor of the range $[-0.15, -0.10)$, then the expected realization of the SDF in $s_1$ will decrease; and these changes can happen without any commensurate changes in the expected SDF in $s_2$. We thus define our return space as in (32) in an attempt to minimize the likelihood of expected-SDF changes arising from changes in the relative sub-probabilities of realizing returns *within* any given interior return state, which could lead to violations of CTI as discussed here.

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 (following the example 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 relative to the discussion in the previous subsection: 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.[54] 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). We construct the state space in this conservative manner in order to avoid possible violations of CTI from jumps or disasters affecting the left tail of the distribution, as in Example 3, as well as to avoid poor measurement of risk-neutral beliefs at the tail ends of the distribution close to option expiration dates.[55]

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.[56] We

[54] In our case, *moneyness* is defined as the ratio of strike price $K$ to the index price $V^m_{T_i}$ corresponding to a log excess return of zero, or $V^m_{T_i} = V^m_{T_0} R^f_{0_i,T_i}$.

[55] Unlike in much of the related literature, we measure risk-neutral beliefs every trading day until expiration where possible; as we will see below, we find interestingly that there seems to be very little excess movement in risk-neutral beliefs for dates very close to expiration, as most of the excess movement appears at more distant horizons.

[56] This widening of the state space increases the number of 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

have further experimented with different definitions of the return space in estimating our main results; for example, we have run specifications using states defined in terms of raw (rather than excess) returns for the S&P 500, as well as states defined in terms of ex-ante option delta, $\Delta_{0_i,T_i,K} \equiv \partial q^m_{0_i,T_i,K}/\partial V^m_{0_i}$, rather than in terms of index returns. 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 for an as-if representative agent; as in Proposition 8, this allows for measurement of the local relative risk aversion of such a fictitious agent. But we have additionally conducted tests with state spaces instead specified in terms of annualized returns, which again leads to very similar results as those reported below.[57]

**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 (22), which tells us how to map from option prices to risk-neutral probabilities.

Consider, for example, an option cross-section for expiration date $T_i$ for which return state $s_3$ in (32) is realized if the S&P 500 value is between 1,000 and 1,020 on date $T_i$. (As above, while we can think of the empirical states as being defined by points in the excess-return space, in practice we measure risk-neutral beliefs over the two-percentage-point neighborhoods surrounding these points.) Equation (22) tells us that the date-$t$ risk-neutral probability of realizing return state $s_3$ can in principle be measured using options with strikes $K_1 = 980$, $K_2 = 1000$, and $K_3 = 1020$, as $\mathbb{P}^*_t(s_3) = R^f_{t,T_i}[(q^m_{t,T_i,K_3} - q^m_{t,T_i,K_2}) - (q^m_{t,T_i,K_2} - q^m_{t,T_i,K_1})]/20$. But this expression assumes that the call-price-schedule slopes $(q^m_{t,T_i,K_3} - q^m_{t,T_i,K_2})/20$ and $(q^m_{t,T_i,K_2} - q^m_{t,T_i,K_1})/20$ are constant over the ranges $K \in [980, 1000]$ and $K \in [1000, 1020]$, respectively; equivalently, it assumes that the state space defined in (32) is a correctly specified and complete description of the true set of possible returns. But this state space is of course only an approximate description of this true set of possible returns, and we must accordingly account for the possibility that the slope of the call-price schedule is not constant over strikes corresponding to the two-percentage-point intervals for returns defined in $\mathcal{S}_{\text{baseline}}$. Further, it may be the case that we lack data on the prices of options with these three strikes $K_1$, $K_2$, and $K_3$ in a given option cross-section, which would require us to obtain approximations for these values even to measure $\mathbb{P}^*_t(s_3)$ even in the case that the return space is correctly specified.

We address these issues in measuring risk-neutral beliefs by constructing a smooth risk-neutral CDF for returns, largely following the technique proposed by Malz (2014), and then differencing this CDF at the endpoints of the intervals defining successive return states. Appendix B.4 provides a detailed description of this procedure, which, stated more simply than in that appendix *(though likely also to be relegated to that appendix at some point)*, requires applying the following steps to the

mean of the distribution close to the expiration date.

[57] These additional results are all available upon request.

raw option-price cross-sections:

1. Transform the collections of call- and put-price cross-sections (for example, for call options on date $t$ for expiration date $T_i$, this set is $\{q^m_{t,T_i,K}\}_{K\in\mathcal{K}}$) 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).

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.[58]

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 $\{\hat{q}^m_{t,T_i,K}\}_{K\in\{200,202,\ldots,4000\}}$.

6. Calculate the risk-neutral CDF for the date-$T_i$ index value at strike price $K$ using $\mathbb{P}^*_t(V^m_{T_i} < K) = 1 + R^f_{t,T_i}(\hat{q}^m_{t,T_i,K} - \hat{q}^m_{t,T_i,K-2})/2$. (See the proof of equation (22) 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 $s_j$,[59] 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

$$\mathbb{P}^*_t(s_j) = \mathbb{P}^*_t(V^m_{T_i} < V^m_{i,j,\max}) - \mathbb{P}^*_t(V^m_{T_i} < V^m_{i,j,\min}), \tag{34}$$

   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 $V^m_{i,j,\min}$, respectively.

[58] 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 $\{q^m_{t,T_i,K}\}_{K\in\mathcal{K}}$ (and similarly for put options), using the numbers given in Section 5.1.

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

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.[60] The particular cubic spline proposed by Malz (2014) and used in step 2 above, a *clamped* cubic spline, is discussed in Appendix B.4; it 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. (The Appendix discusses how we deal with cases in which these features do not hold *inside* the range of observable strikes.)

As noted in step 2, this clamped spline is an *interpolating* spline, as it is restricted to pass through all the observed data points so that the fitted-value set $\{\hat{q}^m_{t,T_i,K}\}$ contains the original values $\{q^m_{t,T_i,K}\}$. 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 accordingly use one such alternative method proposed by Bliss and Panigirtzoglou (2004). See Appendix B.5 for a discussion of this method, as well as empirical results using risk-neutral beliefs calculated in this alternative way; as can be seen there, these results are essentially unchanged as compared to the benchmark results in Section 5.6 below using risk-neutral beliefs measured via the Malz (2014) method above.

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 below; these additional results are available upon request.

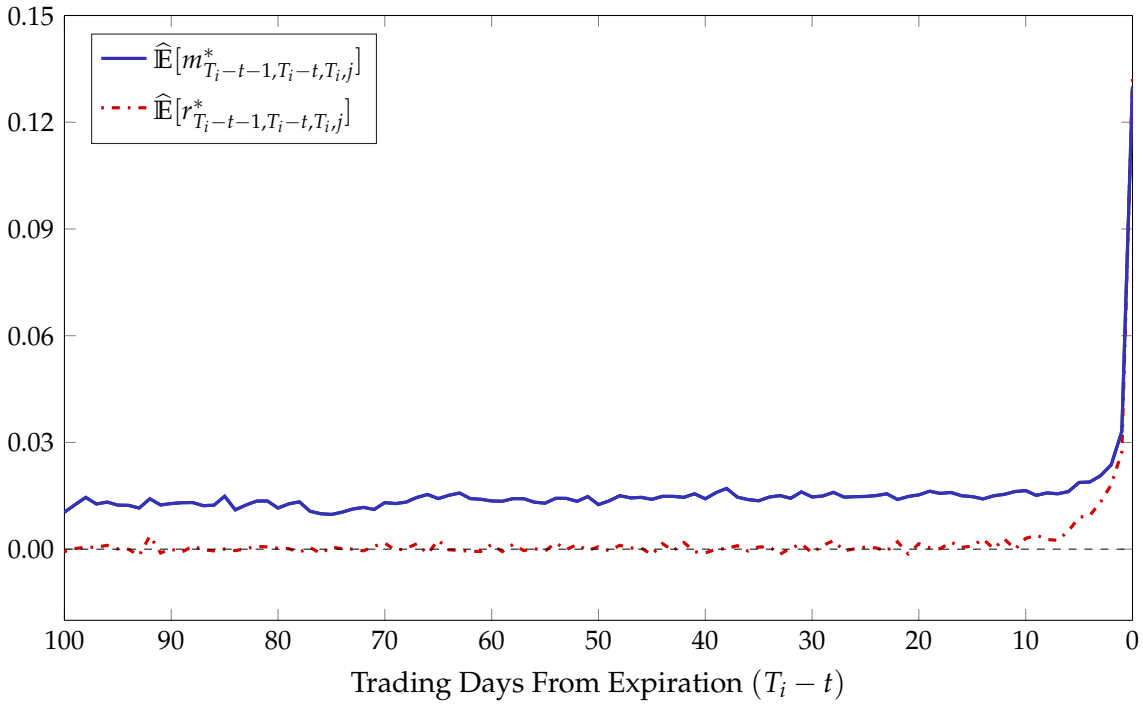
### 5.3. Diagnostic Statistics for Risk-Neutral Beliefs

With the risk-neutral probabilities $\{\mathbb{P}^*_t(s_j)\}_{j,t}$ in in hand, we can then calculate the observed conditional risk-neutral belief for state $s_j$ versus $s_{j+1}$ under the maintained hypothesis of RE as $\widetilde{\pi}^*_{t,T_i,j} = \mathbb{P}^*_t(s_j)/(\mathbb{P}^*_t(s_j) + \mathbb{P}^*_t(s_{j+1}))$ for each state pair and trading date–expiration date pair $(t, T_i)$, following equation (25). We can then measure risk-neutral belief movement and uncertainty resolution following Definitions 2–3, $m^*_{t_1,t_2,T_i,j} = \sum_{\tau=t_1+1}^{t_2}(\widetilde{\pi}^*_{\tau,T_i,j} - \widetilde{\pi}^*_{\tau-1,T_i,j})^2$ and $r^*_{t_1,t_2,T_i,j} = (1-\widetilde{\pi}^*_{t_1,T_i,j})\widetilde{\pi}^*_{t_1,T_i,j} - (1-\widetilde{\pi}^*_{t_2,T_i,j})\widetilde{\pi}^*_{t_2,T_i,j}$, respectively, for arbitrary span $t_1$ to $t_2$ for the relevant state pair and expiration date, and for the full path, $m^*_{T_i,j} = m^*_{0_i,T_i,T_i,j}$, $r^*_{T_i,j} = r^*_{0_i,T_i,T_i,j}$.

Appendix B.6 provides a set of summary statistics for the raw option-price data, the risk-neutral beliefs, and their associated movement and uncertainty resolution. 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; we discuss our construction of these diagnostic statistics as well in Appendix B.6.

[60] 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.

**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 (32), aside from the extreme state pairs $(s_1, s_2)$ and $(s_{J-1}, s_J)$.

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 $T_i$ and state pairs $(s_j, s_{j+1})$ in (32), 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 consistent positive daily risk-neutral belief movement, while risk-neutral uncertainty resolution is very 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.[61] Thus the long-horizon beliefs seem more likely to pose challenges for the RE assumption than short-horizon beliefs. But it remains to be tested formally below whether

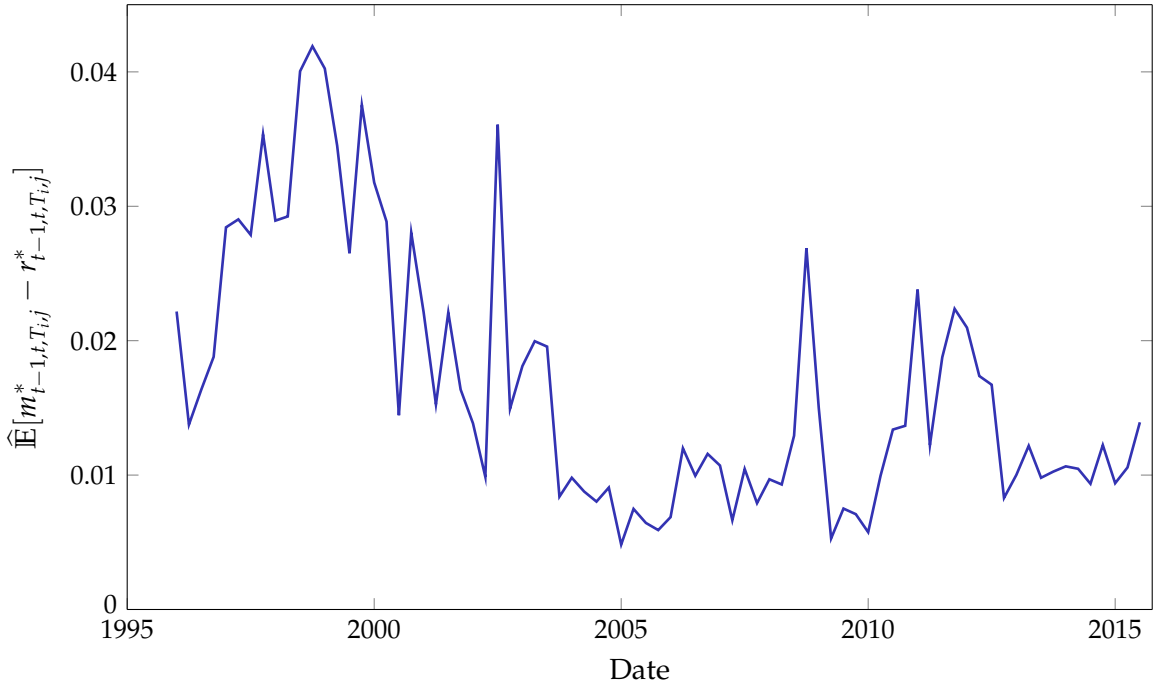
[61] This equivalence is not necessarily mechanical; even in the case that beliefs resolve to $\widetilde{\pi}^*_{T_i,T_i,j} = 0$ or 1, the ex-ante beliefs (e.g., $\widetilde{\pi}^*_{T_i-1,T_i,j}$) must nonetheless be properly calibrated in order for movement to be equal to uncertainty resolution in expectation, as discussed via the example in Section 2. Further, given that there are multiple possible states, it is not necessarily the case that beliefs resolve fully to 0 or 1 even on date $T_i$, in the case that neither state $j$ nor $j+1$ is realized. Thus the rough equivalence of movement and uncertainty resolution for these horizons seems to be some evidence in favor of the RE assumption for events in the near future, as examined further in Section 6.

the gaps between movement and uncertainty resolution until the last two weeks before expiration are statistically large enough to require large values of risk aversion to be rationalized, particularly when taking into account the possibility of measurement error. Regardless, however, the fact that there is zero uncertainty resolution for these 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, $m^*_{t-1,t,T_i,j} = (\widetilde{\pi}^*_{t,T_i,j} - \widetilde{\pi}^*_{t-1,T_i,j})^2$. Thus a value of $m^*_{t-1,t,T_i,j} = 0.01$ — 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 $\widetilde{\pi}^*_{t,T_i,j}$ are again *conditional* beliefs for state $s_j$ versus $s_{j+1}$, so it is possible (and in fact necessary) that these conditional-belief changes are larger than the changes in the full underlying distribution $\{\pi^*_t(s)\}_s$.[62]

Next, Figure 4 plots a time series of average one-day excess risk-neutral belief movement $\widehat{\mathbb{E}}[m^*_{t-1,t,T_i,j} - r^*_{t-1,t,T_i,j}]$ across all dates $t$ in a given quarter and state pairs $(s_j, s_{j+1})$ again aside from the two extreme state pairs. This level of time aggregation allows us to see medium- to long-

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




*Notes:* Empirical averages $\widehat{\mathbb{E}}[m^*_{t-1,t,T_i,j} - r^*_{t-1,t,T_i,j}]$ calculated across all available expiration dates $T_i$ and state pairs $(s_j, s_{j+1})$ in (32), 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.

[62] 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 $m^*_{t-1,t,T_i,j} = 0.01$ 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.

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 recent 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 (2017) document large, persistent post-2008 deviations from covered interest parity in foreign-exchange markets where no such deviations had previously been observed (see also Liao, 2016), whereas our measure is lower during the financial crisis than during much of the pre-2005 period and does not exhibit market persistence after its late-2008 spike. Nonetheless, one might be concerned that the downward trend in $\widehat{\mathbb{E}}[m^*_{t-1,t,T_i,j} - r^*_{t-1,t,T_i,j}]$ 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. But our set of benchmark results in Section 5.6 again includes results allowing for the possibility of measurement error from market microstructure noise, which in theory account for any liquidity-related movement in observed risk-neutral beliefs. And we further examine this question, as well as the relation between excess movement and possible limits to arbitrage, in Section 6, in which we argue again that our findings are not in fact driven by such liquidity or limits-to-arbitrage issues.

### 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 value of risk aversion, as encoded in the average SDF slope across adjacent return states $\mathbb{E}[\phi_{i,j}]$, required to rationalize the data. First, fixing the return-state pair indexed by $j$ (i.e., $(s_j, s_{j+1})$), define $\overline{\phi}_j \equiv \mathbb{E}[\phi_{i,j}]$ to be the average value of $\phi_{i,j}$ for this state pair across all option expiration dates $T_i$ for $i = 1, \ldots, N$. A sample estimate of the lower bound for this value can then be calculated using Proposition 4 as the value $\widehat{\phi}_j$ solving the following sample moment condition:

$$\frac{1}{N}\sum_{i=1}^{N}\left[m^*_{T_i,j} - r^*_{T_i,j} - \widetilde{\pi}^{*\,2}_{0_i,T_i,j}\left(1 - \frac{1}{\widetilde{\pi}^*_{0_i,T_i,j} + \widehat{\phi}_j(1 - \widetilde{\pi}^*_{0_i,T_i,j})}\right)\right] = 0. \tag{35}$$

We repeat this for each return-state pair to obtain a set of estimates $\{\widehat{\phi}_j\}_j$ for the minimum local slope of the SDF at each value in the return space $\mathcal{S}_{\text{baseline}}$.

While the estimator in (35) 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 (35) holds, which, as in the result in Corollary 1, occurs when

$$\frac{1}{N}\sum_{i=1}^{N}\left[m^*_{T_i,j} - r^*_{T_i,j} - \widetilde{\pi}^{*\,2}_{0_i,T_i,j}\right] > 0. \tag{36}$$

For ease of notation and interpretation of the estimation results reported below, when (36) holds we say that $\widehat{\phi}_j = \infty$, but this should in fact be understood as denoting that even a value of infinity cannot rationalize the observed excess belief movement. (It would in this case perhaps be more accurate to say $\widehat{\phi}_j = \varnothing$, but we instead use $\infty$ to make clear that this estimate corresponds to greater excess risk-neutral belief movement than would yield a finite value for $\widehat{\phi}_j$.)

We then define $\overline{\phi} \equiv \mathbb{E}\left[\overline{\phi}_j\right]$ to be the average value of $\phi_{i,j}$ 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 $j = 2, \ldots, J-2$, or equivalently for pairs $(s_2, s_3), \ldots, (s_{J-2}, s_{J-1})$. Then the estimator $\widehat{\phi}$ for the minimum admissible value for $\overline{\phi}$ solves

$$\frac{1}{N(J-3)} \sum_{i=1}^{N} \sum_{j=2}^{J-2} \left[ m^*_{T_i,j} - r^*_{T_i,j} - \widetilde{\pi}^{*\,2}_{0_i,T_i,j} \left( 1 - \frac{1}{\widetilde{\pi}^*_{0_i,T_i,j} + \widehat{\phi}(1 - \widetilde{\pi}^*_{0_i,T_i,j})} \right) \right] = 0, \tag{37}$$

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

$$\frac{1}{N(J-3)} \sum_{i=1}^{N} \sum_{j=2}^{J-2} \left[ m^*_{T_i,j} - r^*_{T_i,j} - \widetilde{\pi}^{*\,2}_{0_i,T_i,j} \right] > 0. \tag{38}$$

In order to obtain sharper results when estimating the lower bounds for $\left\{\overline{\phi}_j\right\}_j$ and $\overline{\phi}$, we in fact use the result in Proposition 5 to allow for only partial resolution of uncertainty for any given state pair as of expiration date $T_i$, as occurs when the index return is not in $s_j$ or $s_{j+1}$. (See Appendix B.2 for more detail.) But as can be seen from that proposition, the expressions for these estimators are substantially similar to but somewhat more complicated than those given in (35)–(38), so we find it more intuitive to consider the estimators arising from Proposition 4 in our exposition above.

Proposition 8 then allows us to map straightforwardly from our estimates $\left\{\widehat{\phi}_j\right\}_j$ 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 $\mathcal{S}_{\text{alt}}$). These give lower-bound estimates for local relative risk aversion around return state $s_j$ (i.e., $\overline{\gamma}_j = \mathbb{E}[\gamma_{i,j}]$ for a given $j$) and average relative risk aversion across all return states meeting CTI (i.e., $\overline{\gamma} = \mathbb{E}[\overline{\gamma}_j]$), respectively. We report both SDF-slope estimates ($\left\{\widehat{\phi}_j\right\}_j$ and $\widehat{\phi}$) and relative-risk-aversion estimates ($\left\{\widehat{\gamma}_j\right\}_j$ 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 from an economic standpoint.

**Inference.** We must address three issues arising from our framework when constructing confidence intervals for our estimates. First, the moment conditions (35) and (37) 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 (36) and (38). This non-compactness of the parameter

space violates the usual conditions for the consistency and asymptotic normality of the GMM estimator.[63] More simply, subtracting some multiple of an estimated standard error from a possible point estimate of $\infty$ (or, as above, $\varnothing$) is both nonsensical and also invalid in constructing a lower bound for the confidence interval of $\overline{\phi}_j$ or $\overline{\phi}$ with the desired coverage rate. Second, there is an unknown dependence structure, both over time and across state pairs, between the members of the set $\{m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,T_i,j}\}_{i,j}$, 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.[64] 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. We proceed as follows *(guessing this will all go in an appendix as well, so for those reading the preliminary draft, again feel free to skip)*:

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.[65] 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 $\{(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,T_i,j})_j\}_i$, into these constructed blocks according to the block in which expiration date $T_i$ appears for each set $(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,T_i,j})_j$; that is, $b_\ell = \{(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,T_i,j})_j : T_i \in \mathcal{T}_\ell\}$ 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 $(b^\dagger_1, b^\dagger_2, \ldots, b^\dagger_B)$.

[63] See, for example, Hansen (1982) and Newey and McFadden (1994). The latter paper considers some conditions under which consistency can be met without compactness, but these conditions are violated in our case and this further fails to address the issue of inference; Hansen (2012b) shows that compactness is unnecessary when the moment conditions can be written using the expectation of a random data matrix multiplied by a linear function of the parameters to be estimated, but this is not the case here.

[64] For example, if the set of time values between $0_i$ and $T_i$ overlaps partially with values between $0_{i+1}$ and $T_{i+1}$, this likely introduces correlation between the excess risk-neutral belief movement statistics for $T_i$ and $T_{i+1}$ in a manner that is difficult to model.

[65] 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.

4. For each resampled dataset, re-estimate the lower bounds for $\{\overline{\phi}_j\}_j$ and $\overline{\phi}$ as discussed above (see equations (35)–(38)).

5. Construct a lower bound for a $100\alpha\%$ confidence interval for each value in $\{\overline{\phi}_j\}_j$ and for $\overline{\phi}$ as the $1-\alpha$ 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 $\text{CI}_j = [\mathcal{Q}^\dagger_{j,1-\alpha}, \infty)$, where $\mathcal{Q}^\dagger_{j,1-\alpha}$ is the $1-\alpha$ quantile of the bootstrap distribution of the estimates of the lower bound for $\overline{\phi}_j$. We refer to $\mathcal{Q}^\dagger_{j,1-\alpha}$ 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 $\alpha = 0.95$ 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, Petersen (2009) or 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, $\{(m^*_{T_i,j}, r^*_{T_i,j}, \widetilde{\pi}^*_{0_i,T_i,j})_j\}_i$ — 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.[66] This coverage rate may 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; see, for example, Imbens and Manski (2004), Chernozhukov, Hong, and Tamer (2007), and Tamer (2010) for discussions of confidence-set construction in similar partially identified settings.

[66] 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 nonoverlapping block bootstrap (NBB); others (Künsch, 1989; Liu and Singh, 1992; Bühlmann, 1994) have proposed a *moving* block bootstrap (MBB) using overlapping blocks, among other alternatives (e.g., Politis and Romano, 1994). 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; Hansen, 2017; 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 $\widehat{\phi}_j - \mathcal{Q}^{\dagger\dagger}_{j,\alpha}$, where $\mathcal{Q}^{\dagger\dagger}_{j,\alpha}$ is the $\alpha$ 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 $\alpha$ quantile is equal to $\infty$, 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, though it does not yield any asymptotic refinement (as should be expected given that it uses a non-pivotal statistic; see Horowitz, 2001).

We experiment with various choices for the number of blocks $B$ in this bootstrap procedure. Our baseline confidence intervals presented in Section 5.6 use a value of $B$ chosen so that the block length $D = \lceil (T_N + 1)/B \rceil$ is equivalent to 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 $B$ 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 $B$.

### 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 (2000, 2003) argues that index-option-price measurement error is typically quite small when estimated as part of a parametric factor model; see also Vijh (1990) and Andersen, Fusari, and Todorov (2015). 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.[67] 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 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 $\widehat{\pi}^*_{t,T_i,j}$ and its true value $\widetilde{\pi}^*_{t,T_i,j}$ is $\widehat{\pi}^*_{t,T_i,j} = \widetilde{\pi}^*_{t,T_i,j} + \epsilon_{t,T_i,j}$, where $\epsilon_{t,T_i,j}$ is i.i.d.[68] We consider two distinct approaches for estimating $\text{Var}(\epsilon_{t,T_i,j})$.

[67] For example, changes in net demand pressure in the options market may affect prices, as argued for by Bollen and Whaley (2004), Gârleanu, Pedersen, and Poteshman (2009), and Muravyev (2016). Given that such pressure is naturally correlated over options and trading dates, a factor pricing model may interpret movements owing to these changes as stemming from a common factor. (See Stock and Watson, 2016, for an introduction to such models.) But while there would be nothing incorrect about such inference, for our purposes we may wish to include this variation in our notion of noise: this variation is not fully correlated with changes in agents' true beliefs over the future index value, so to the extent that we wish to measure the rationalizability of such macroeconomic expectations (rather than, e.g., the rationalizability of the option prices themselves; see Footnote 3), we would want to include this variation as part of our microstructure noise measurement.

[68] For our purposes, it need not be the case that the true measurement error is in fact i.i.d.; we need to measure only

The first approach exploits 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, $\widetilde{\pi}_{t,T_i,j} = \mathbb{E}_t[\mathbb{1}\{R^m_{T_i} = s_j\} \mid R^m_{T_i} \in \{s_j, s_{j+1}\}]$. We detail in Appendix B.7 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 $\widetilde{\pi}^*_{T_i,T_i,j} \in \{0,1\}$, measured only using beliefs on dates $t$ and $T_i$ (i.e., ignoring the intermediate one-period belief changes), must follow the equality restriction

$$\mathbb{E}_t[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}] = \frac{\widetilde{\pi}^*_{t,T_i,j}(1-\widetilde{\pi}^*_{t,T_i,j})(2\widetilde{\pi}^*_{t,T_i,j}-1)(\phi_{i,j}-1)}{\widetilde{\pi}^*_{t,T_i,j} + \phi_{i,j}(1-\widetilde{\pi}^*_{t,T_i,j})} \tag{39}$$

regardless of the data-generating process, and where $m^*_{t\to T_i,j} \equiv (\widetilde{\pi}^*_{T_i,T_i,j} - \widetilde{\pi}^*_{t,T_i,j})^2$ and $r^*_{t\to T_i,j} \equiv (1-\widetilde{\pi}^*_{t,T_i,j})\widetilde{\pi}^*_{t,T_i,j} - (1-\widetilde{\pi}^*_{T_i,T_i,j})\widetilde{\pi}^*_{T_i,T_i,j} = (1-\widetilde{\pi}^*_{t,T_i,j})\widetilde{\pi}^*_{t,T_i,j}$ (where the last equality again uses that we are conditioning on full resolution, $\widetilde{\pi}^*_{T_i,T_i,j} \in \{0,1\}$). 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 first estimate of $\mathrm{Var}(\epsilon_{t,T_i,j})$ then assumes that any average deviation from the equality restriction in (39) in the observed data arises due to measurement error. We again have that, with measurement error $\epsilon_{t,T_i,j}$ in risk-neutral beliefs, the observed $t \to T_i$ excess movement is overstated by $2\mathrm{Var}(\epsilon_{t,T_i,j})$ relative to its true value: $\mathbb{E}_t[\widehat{m}^*_{t\to T_i,j} - \widehat{r}^*_{t\to T_i,j}] = \mathbb{E}_t[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}] + 2\mathrm{Var}(\epsilon_{t,T_i,j})$, as in Proposition 6. Thus, after taking unconditional expectations in (39) and measuring the sample analogues of the objects in that equation for resolving paths, we calculate $\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\overline{\phi}_j)$ as the value equating $\mathbb{E}[\widehat{m}^*_{t\to T_i,j} - \widehat{r}^*_{t\to T_i,j}] - 2\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\overline{\phi}_j)$ with the right side of (39).[69] We denote the estimate by $\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\overline{\phi}_j)$ because there will be a different noise estimate for each value of $\overline{\phi}_j$ given that the right side of (39) 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 $\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\overline{\phi}_j)$; following Proposition 6, for each such estimate we subtract $2\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\overline{\phi}_j)$ from each day's observed excess risk-neutral belief movement value $\widehat{m}^*_{t,t+1,T_i,j} - \widehat{r}^*_{t,t+1,T_i,j}$ and then re-estimate our period-by-period bound as in (35); 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 $\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\overline{\phi}_j)$. This allows us to solve simultaneously for the estimates $\widehat{\phi}_j$ and $\widehat{\mathrm{Var}}(\epsilon_{t,T_i,j}) = \widehat{\mathrm{Var}}(\epsilon_{t,T_i,j};\widehat{\phi}_j)$. See Appendix B.7 for further details of this procedure.

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

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.

[69] There may be conditional 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 $\mathbb{E}[m^*_{t\to T_i,j} - r^*_{t\to T_i,j}]$ 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 $\widehat{\text{Var}}(\epsilon_{t,T_i,j})$ may be biased in either direction: if beliefs are not rational, then they need not be correctly calibrated (so that $\widetilde{\pi}_{t,T_i,j} \neq \mathbb{E}_t[\mathbb{1}\{R^m_{T_i} = s_j\} \mid R^m_{T_i} \in \{s_j, s_{j+1}\}]$), 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, which is our main concern, 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. We note first that given $\widehat{\pi}^*_{t,T_i,j} = \widetilde{\pi}^*_{t,T_i,j} + \epsilon_{t,T_i,j}$, the unconditional variance (or zeroth autocovariance) and the first autocovariance of the observed beliefs are given respectively by

$$\text{Var}(\widehat{\pi}^*_{t,T_i,j}) = \text{Var}(\widetilde{\pi}^*_{t,T_i,j}) + \text{Var}(\epsilon_{t,T_i,j}), \tag{40}$$

$$\text{Cov}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j}) = \text{Cov}(\widetilde{\pi}^*_{t,T_i,j}, \widetilde{\pi}^*_{t-1,T_i,j}). \tag{41}$$

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

$$\begin{aligned}\text{Var}(\widehat{\pi}^*_{t,T_i,j}) - \text{Cov}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j}) &= \text{Var}(\widetilde{\pi}^*_{t,T_i,j}) - \text{Cov}(\widetilde{\pi}^*_{t,T_i,j}, \widetilde{\pi}^*_{t-1,T_i,j}) + \text{Var}(\epsilon_{t,T_i,j}) \\ &\geqslant \text{Var}(\epsilon_{t,T_i,j}).\end{aligned} \tag{42}$$

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

$$\widehat{\text{Var}}(\epsilon_{t,T_i,j}) = \widehat{\text{Var}}(\widehat{\pi}^*_{t,T_i,j}) - \widehat{\text{Cov}}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j}), \tag{43}$$

which requires measuring only the variance and first autocovariance of the observed conditional risk-neutral beliefs process (again see Footnote 69). 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.

The estimates obtained from this procedure are then directly economically interpretable regardless of the rationality of expectations. The estimates in this case range from $\widehat{\text{Var}}(\epsilon_{t,T_i,j}) = 0.005$ 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 $(s_j, s_{j+1}) = (-0.07, -0.05)$. See Appendix B.7 for additional discussion and comparison with the first method.

As discussed above, the noise estimate (43) is likely to be overly conservative: if the *true* process $\{\widetilde{\pi}^*_{t,T_i,j}\}$ has a nonzero white-noise component, then this component will be lost when making the conservative assumption that $\text{Var}(\widetilde{\pi}^*_{t,T_i,j}) = \text{Cov}(\widetilde{\pi}^*_{t,T_i,j}, \widetilde{\pi}^*_{t-1,T_i,j})$ as in (43). Further, the conservatism in (43) is layered on top of the conservatism embodied in our theoretical bounds, as these bounds are constructed under the worst-case (or excess-movement-maximizing) data-generating process, as discussed in Proposition 2. Thus when presenting noise-corrected estimates using this second approach, we assume that $\text{Var}(\epsilon_{t,T_i,j}) = 0.5\,[\text{Var}(\widehat{\pi}^*_{t,T_i,j}) - \text{Cov}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j})]$, which is (roughly) equivalent to the assumption that 50 percent of the white-noise (or non-serially-correlated) component of the observed belief process is due to noise and 50 percent is "real."[70] The choice of 50 percent is of course arbitrary, but it yields noise-corrected estimates that are on average close to the estimates from the first approach described above.[71] So the first approach provides a valid average noise correction, while the second approach provides more readily interpretable noise estimates and further insight into the extent to which noise affects the individual-state-pair estimates differentially across states (which may be lost in the first approach).

We have also recently obtained a sample of daily option-price data directly from the CBOE, which reports transaction prices in addition to bid and ask prices; we will accordingly be able to re-estimate our bounds on the transaction-price data with somewhat less concern about noise. Further, we aim to implement additional theoretical results for the full histogram of risk-neutral beliefs (as mentioned briefly in Footnote 33), for which measurement error is less of a concern than for conditional risk-neutral beliefs as used here.

### 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 (35)–(38). 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 8. 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 $\phi_j$ and $\gamma_j$ 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

[70] This statement holds only roughly because the serial correlation in the non-white-noise component of $\widetilde{\pi}^*_{t,T_i,j}$ will generally be less than 1; for example, consider the model $\widetilde{\pi}^*_{t,T_i,j} = c + \rho(\widetilde{\pi}^*_{t-1,T_i,j} - \eta_{t-1}) + \eta_t$ with $\eta_t$ i.i.d. and AR(1) coefficient $0 < \rho < 1$, which yields a first-order autocorrelation value of $\rho < 1$ and thus a positive autocovariance deviation $\text{Var}(\widehat{\pi}^*_{t,T_i,j}) - \text{Cov}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j})$. We are thus attributing slightly more of the variation in observed beliefs to noise than necessary under the assumption that 50 percent of the white-noise component of the observed belief process is measurement error.

[71] We again construct these noise-corrected estimates following Proposition 6, subtracting $2\widehat{\text{Var}}(\epsilon_{t,T_i,j})$ from each observed day's observed excess risk-neutral belief movement value $\widehat{m}^*_{t,t+1,T_i,j} - \widehat{r}^*_{t,t+1,T_i,j}$ before constructing our estimators as in equations (35)–(38).

## Table 1: Main Estimation Results

### (a) Lower Bound for SDF Slope

| | Individual State Pairs: $\overline{\phi}_j$ by Excess Return | | | | | | | | | | | Overall: |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | -10% | -8% | -6% | -4% | -2% | 0 | 2% | 4% | 6% | 8% | 10% | $\overline{\phi}$ |
| Baseline: | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ |
| ($CI_{LB,45}$) | (120) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) |
| ($CI_{LB,90}$) | (90.0) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) |
| Noise-Adjusted Estimates: | | | | | | | | | | | | |
| 1. | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ |
| ($CI_{LB,45}$) | (86.3) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (1.6) | (1.6) | (2.6) | (8.9) | (31.0) | ($\infty$) | (3.0) |
| ($CI_{LB,90}$) | (82.0) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (1.6) | (1.6) | (2.6) | (8.7) | (31.5) | ($\infty$) | (2.7) |
| 2. | 12.2 | $\infty$ | $\infty$ | $\infty$ | $\infty$ | 10.1 | 2.7 | 4.5 | 10.4 | 130 | $\infty$ | 27.5 |
| ($CI_{LB,45}$) | (8.2) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (3.3) | (1.7) | (2.8) | (5.4) | (12.6) | ($\infty$) | (6.0) |
| ($CI_{LB,90}$) | (7.5) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (3.1) | (1.7) | (2.6) | (5.2) | (11.1) | ($\infty$) | (5.4) |

### (b) Lower Bound for Relative Risk Aversion

| | Individual State Pairs: $\overline{\gamma}_j$ by Excess Return | | | | | | | | | | | Overall: |
|---|---|---|---|---|---|---|---|---|---|---|---|---|
| | -10% | -8% | -6% | -4% | -2% | 0 | 2% | 4% | 6% | 8% | 10% | $\overline{\gamma}$ |
| Baseline: | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ |
| ($CI_{LB,45}$) | (5,950) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) |
| ($CI_{LB,90}$) | (4,450) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) |
| Noise-Adjusted Estimates: | | | | | | | | | | | | |
| 1. | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ | $\infty$ |
| ($CI_{LB,45}$) | (4,263) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (32.0) | (29.0) | (82.0) | (395) | (1,500) | ($\infty$) | (102) |
| ($CI_{LB,90}$) | (4,050) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (31.5) | (29.0) | (81.0) | (385) | (1,525) | ($\infty$) | (85.0) |
| 2. | 560 | $\infty$ | $\infty$ | $\infty$ | $\infty$ | 455 | 85.0 | 174 | 470 | 6,450 | $\infty$ | 1,325 |
| ($CI_{LB,45}$) | (360) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (116) | (35.0) | (88.0) | (218) | (580) | ($\infty$) | (248) |
| ($CI_{LB,90}$) | (325) | ($\infty$) | ($\infty$) | ($\infty$) | ($\infty$) | (106) | (33.0) | (78.0) | (208) | (505) | ($\infty$) | (218) |

*Notes:* Lower-bound estimates are constructed using the bound in Proposition 5 for panel (a), converted to relative risk aversion values using Proposition 8 for panel (b). Baseline estimates use equations (35)–(36) 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 (37)–(38). One-sided 95% confidence intervals using the block bootstrap with 5,000 draws and block length of $D$ calendar days are $[\text{CI}_{\text{LB},D}, \infty)$; for simplicity, we present only the lower bounds of these confidence intervals, $\text{CI}_{\text{LB},D}$, in parentheses below the point estimates, and we present CIs for $D = 45, 90$. First method for noise-adjusted estimates uses $\widehat{\text{Var}}(\epsilon_{t,T_i,j})$ estimated so that fully resolving paths have implied physical beliefs that are ex-ante correctly calibrated on average; second method uses $\widehat{\text{Var}}(\epsilon_{t,T_i,j}) = 0.5[\widehat{\text{Var}}(\widehat{\pi}^*_{t,T_i,j}) - \widehat{\text{Cov}}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j})]$. After obtaining $\widehat{\text{Var}}(\epsilon_{t,T_i,j})$, both methods 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.

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 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, the first method (which estimates the noise variance so as to yield correctly calibrated physical beliefs on average) yields 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. Under the second noise-correction method (which estimates the measurement-error variance using the deviation between the variance and the first autocovariance of the risk-neutral beliefs series), roughly half of the point estimates are finite, but the estimates and their confidence-interval lower bounds are again very large.[72] As in the discussion in Section 5.5, we prefer the first set of noise-corrected estimates given that this method is valid under RE regardless of the data-generating process for beliefs.

We observe further in both sets of noise-corrected estimates 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. Additional work is needed to more fully understand this pattern, but it does appear to be a robust finding, as it is again evident in the statistical decompositions considered below in Section 6. We note also that 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 can 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 be generated 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.

[72] When we vary the fraction of $\widehat{\text{Var}}(\widehat{\pi}^*_{t,T_i,j}) - \widehat{\text{Cov}}(\widehat{\pi}^*_{t,T_i,j}, \widehat{\pi}^*_{t-1,T_i,j})$ imputed as measurement-error noise, we find that the estimated risk-aversion values remain very large roughly until we assume that more than 2/3 of that variance-autocovariance gap is due to noise.

# 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?

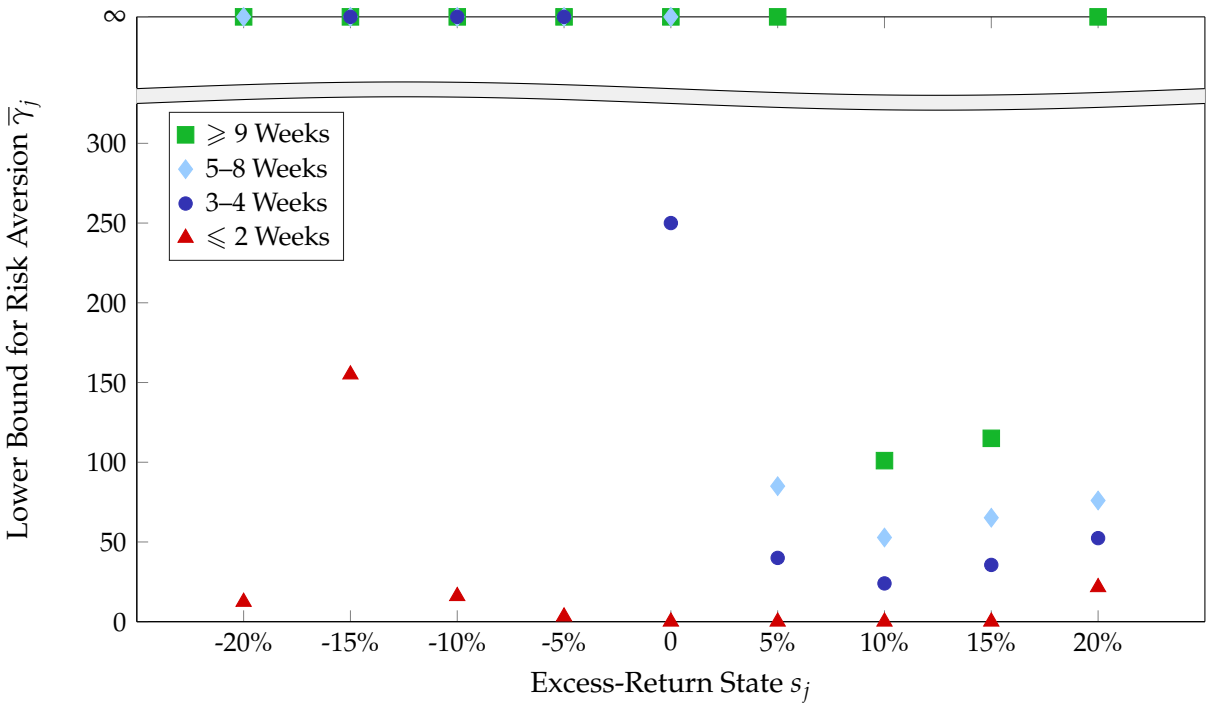
Figure 3, presented above in Section 5.3, suggests that excess movement in risk-neutral beliefs is concentrated at distant horizons from a given option expiration date. We now consider whether this holds formally when re-estimating our bounds separately for subsets of the risk-neutral belief data split by time to expiration. Figure 5 presents the results from this decomposition. For purposes of 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: the distribution of returns is often concentrated near its conditional mean close to the expiration date, and we require $[\pi_t^*(R_{T_i}^m = s_j) + \pi_t^*(R_{T_i}^m = s_{j+1})] \geqslant 5\%$ to avoid using poorly measured conditional beliefs. 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 excess-belief-movement 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 finite and reasonable values for risk aversion, and in many cases we cannot rule out risk-neutrality.[73] And when pooling data across the interior return-state pairs to estimate the average required risk-aversion value across both dates and states, we obtain an estimate $\widehat{\gamma} = 0$ (not shown in the figure).

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 $\infty$, 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

[73] Note that this does not necessarily suggest that pricing is in fact risk-neutral, as our estimates are again only lower bounds for the risk aversion needed to rationalize the data.

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




*Notes:* Lower-bound estimates are constructed using the bound in Proposition 5, converted to relative risk aversion values using Proposition 8. 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 $m^*_{T_i-10,T_i,T_i,j}$ and $r^*_{T_i-10,T_i,T_i,j}$, respectively, and similarly for the remainder of the series. Gray band denotes 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: $\infty$). CIs calculated using block bootstrap with blocks of 45 days, 5,000 draws.

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 (though not *excessively* volatile relative to uncertainty resolution, as evident in Figure 3).

### 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; the excess-belief-movement statistic we use to construct our tests 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 $\mathbb{E}[m^*_{T_i,j} - r^*_{T_i,j}]$ is a tendency for risk-neutral beliefs to revert toward 0.5 on average over the entire series.[74]

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,[75] 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 $(\widetilde{\pi}^*_{t,T_i,j} - \widetilde{\pi}^*_{t-5,T_i,j})^2$. 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; see the figure's footnote for further detail.

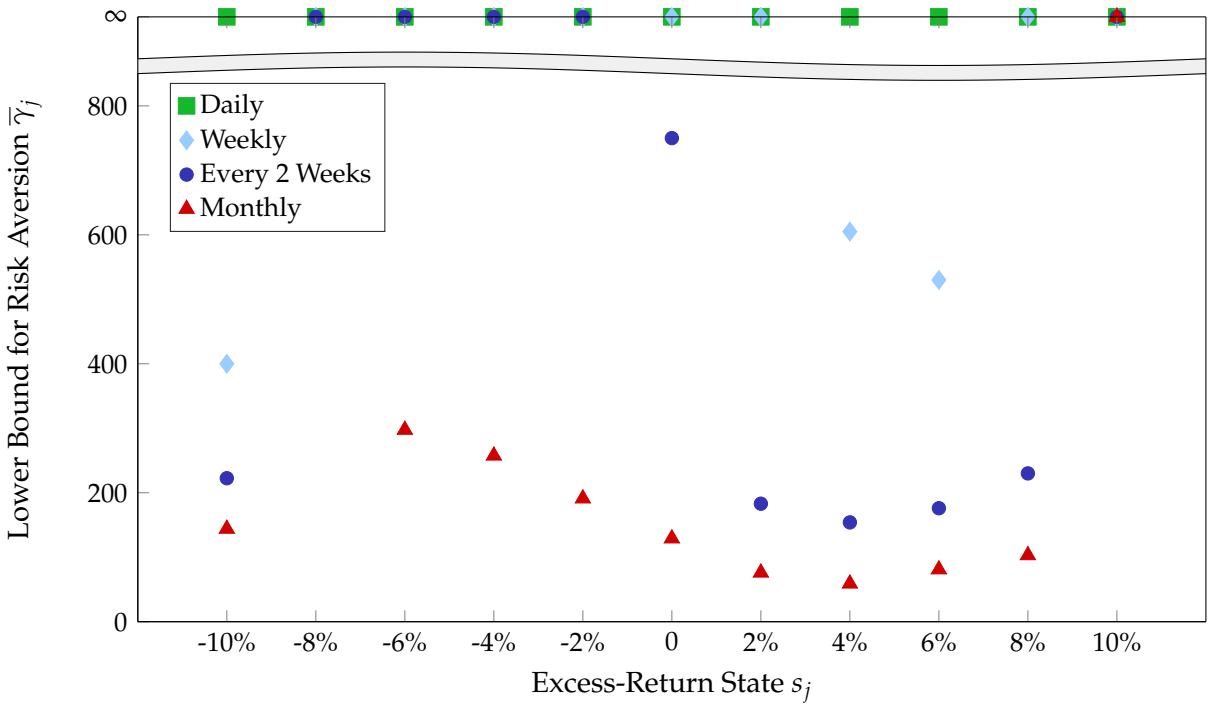
Beginning now at the top of the chart and proceeding downward, we again observe a clear monotonic pattern: as we decrease the sampling frequency, the data can be rationalized with weakly decreasing values for relative risk aversion at every point in the return space. This is reflected as well in the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

[74] There is some evidence, for example, that equity-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, but for the aggregate index); see, e.g., Poterba and Summers (1988), Hirshleifer (2001), and Asness, Moskowitz, and Pedersen (2013). DeLong et al. (1990), Barberis, Shleifer, and Vishny (1998), and Hong and Stein (1999) provide unified theories explaining both facts. Short-term momentum for index returns seems, however, 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 at a given expiration date $T_i$ 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.

[75] 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:* Lower-bound estimates are constructed using the bound in Proposition 5, converted to relative risk aversion values using Proposition 8. 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 constructs risk-neutral belief movement and uncertainty resolution using risk-neutral beliefs $\widetilde{\pi}^*_{t,T_i,j}$ 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 denotes 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: $\infty$). CIs calculated using block bootstrap with blocks of 45 days, 5,000 draws.

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 $\{\widetilde{\pi}^*_{t,T_i,j}\}_t$: in the presence of excess belief volatility at the daily frequency, the sum $m^*_{T_i,j} = \sum_{t=0_i+1}^{T_i} (\widetilde{\pi}^*_{t,T_i,j} - \widetilde{\pi}^*_{t-1,T_i,j})^2$ 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 ob-

servations 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 $m^*_{t-1,1,T_i,j} - r^*_{t-1,t,T_i,j}$ 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). 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. Appendix B.8 contains detailed descriptions of the construction of these series. 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.[76] 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.[77] The coefficients on both regressors are both economically 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

[76] This follows, among others, Amihud and Mendelson (1986) and Chordia, Roll, and Subrahmanyam (2000, 2008), but we might also wish to consider more direct proxies for the return impact of a transaction, as used, e.g., by Amihud (2002), Pástor and Stambaugh (2003), and Acharya and Pedersen (2005). Pástor and Stambaugh (2003) also discuss why trading volume is inappropriate to use as a measure of liquidity, though in unreported results we find that it is uncorrelated with excess belief movement.

[77] This measure is proposed and examined by Adrian, Etula, and Muir (2014), Adrian and Shin (2014), He, Kelly, and Manela (2016), and Cho (2017), among others, following theoretical insights developed by Shleifer and Vishny (1997), Xiong (2001), 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 $\widehat{\mathbb{E}}[m^*_{t-1,t,T_i,j} - r^*_{t-1,t,T_i,j}]$ calculated across all available expiration dates $T_i$ and state pairs $(s_j, s_{j+1})$ in (32), 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; see Appendix B.8 for details of construction of these variables. 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 from the $t$ distribution with 8 degrees of freedom.

proxy for implied volatility, we use the quarterly average of the CBOE volatility index, or VIX.[78] 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 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 $\widetilde{\pi}^*_{0_i,T_i,j}$, and higher risk-neutral volatility leads to a more-dispersed distribution and therefore

[78] 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.

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, 1988a; 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,[79] and cyclically adjusted price-earnings for valuation-related return predictability (see Footnote 6). 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. And when absolute values are used for the valuation measures (unreported in the table), the relationships are insignificant or even negative. 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 exactly as presented in Example 4 in Section 3.2, 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.9 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 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,

[79] This measure is used by, for example, Greenwood and Shleifer (2014) and Gennaioli, Ma, and Shleifer (2016) in related analysis.

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.

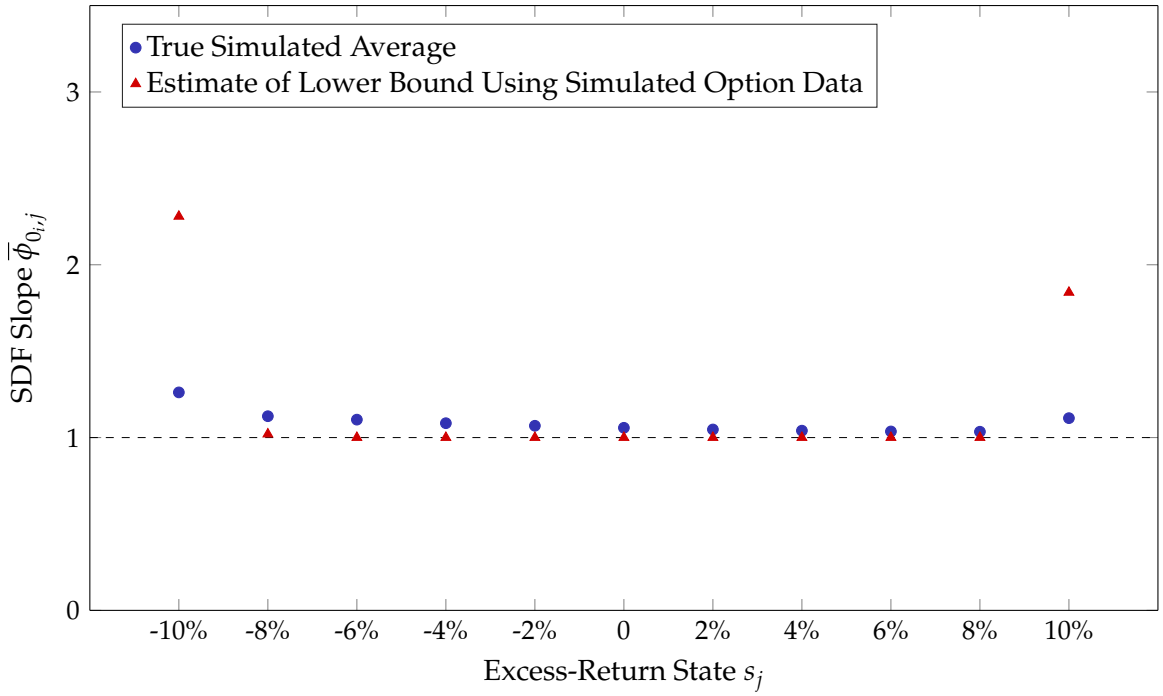
Conducting these simulations requires calculating subjective beliefs for returns over many future horizons as a function of the surplus-consumption state, as well as the set of expected SDF slopes $\{\phi_{t,T_i,j}\}_{t,T_i,j}$ in order to translate the subjective beliefs to risk-neutral beliefs. The SDF slope is again given by $\phi_{t,T_i,j} = \mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_j]/\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j+1}]$, 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 (see Judd, 1992, or see Algan, Allais, Den Haan, and Rendahl, 2014, and Maliar and Maliar, 2014, for recent surveys, though none of these consider solutions for belief distributions). 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.9 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, $S^c_{0_i} = \overline{S^c}$ *(for those reading the preliminary version, for which the appendices are completely incomplete, the footnote of that figure briefly describes the projection-based solution algorithm)*.

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 $\{\widetilde{\pi}^*_{t,T_i,j}\}_j$. Further, we store the associated set of expected SDF slopes $\{\phi_{t,T_i,j}\}_j$. We can thus calculate the true average values of these objects of interest, $\overline{\phi}_{0_i,j} \equiv \widehat{\mathbb{E}}[\phi_{0_i,T_i,j}]$, 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 5 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 8, 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 not replicate the observed variation in risk-neutral beliefs even with the violation of CTI.[80]

[80] It is also the case that these SDF slopes imply lower risk-aversion values than might be expected for the Campbell-

**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 $\phi_{0_i,T_i,j} = \mathbb{E}_{0_i}[M_{T_i}/M_{0_i} \mid R^m_{T_i} = s_j]/\mathbb{E}_{0_i}[M_{T_i}/M_{0_i} \mid R^m_{T_i} = s_{j+1}]$ 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.9 for details. The estimate of the lower bound for $\overline{\phi}_{0_i,j}$ is constructed naïvely using the theoretical bound in Proposition 5, using risk-neutral belief movement $m^*_{T_i,j}$ and uncertainty resolution $r^*_{T_i}$ values across all simulations constructed via simulated risk-neutral beliefs $\widetilde{\pi}^*_{t,T_i,j}$; again see Appendix B.9 for details. 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.

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, and for this reason we again exclude these states when aggregating to obtain single estimates in the empirical analysis above.)

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 $\phi_{t,T_i,j}$ process is close to a martingale. In our simulations, the values $\left|\widehat{\mathbb{E}}[\phi_{t+1,T_i,j} - \phi_{t,T_i,j}]\right|$ for dif-

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 given that the index's dividends are only weakly correlated with consumption and marginal utility.

ferent state 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.[81]

## 7. Discussion and 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.[82]

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 average 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.

While our results 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.

[81] For example, we are currently considering violations arising from certain forms of stochastic volatility shocks (see Footnote 38), which a large literature argues are important for matching various features of the option-price data; see, e.g., Bakshi, Cao, and Chen (1997), Duffie, Pan, and Singleton (2000), Coval and Shumway (2001), Rosenberg and Engle (2002), and Christoffersen, Heston, and Jacobs (2013), though see also Jones (2006) and Cuesdeanu and Jackwerth (2016) for slightly different interpretations of the available evidence. Preliminary results for these additional simulations suggest similar findings to those for the habit-formation simulations presented here, but more work is required. We also have some additional (currently overly complicated) theoretical results considering the implications of more general violations of CTI, which would allow for more general empirical tests.

[82] 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.

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:

$$\pi_t(R^m_{T_i} = s) = \frac{\pi_{t-1}(R^m_{T_i} = s)^\alpha \ \mathbb{P}(\theta_t | \ \mathcal{F}_{t-1}, R^m_{T_i} = s)}{\sum_{s' \in \mathcal{S}} \pi_{t-1}(R^m_{T_i} = s')^\alpha \ \mathbb{P}(\theta_t \,|\, \mathcal{F}_{t-1}, R^m_{T_i} = s')}, \tag{44}$$

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, $\pi_{t-1}(R^m_{T_i} = s_a) = \pi_{t-1}(R^m_{T_i} = s_b)$ for all $s_a$, $s_b$, and $t$; in this case, her period-$t$ beliefs will simply reflect the period-$t$ signal received, and her conditional beliefs $\widetilde{\pi}_{t,T_i,j} = \pi_t(R^m_{T_i} = s_j)/(\pi_t(R^m_{T_i} = s_j) + \pi_t(R^m_{T_i} = s_{j+1}))$ will bounce back and forth above and below 0.5 as she receives signals favorable to state $s_j$ and $s_{j+1}$, respectively, inducing excess movement in 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 (2017), one manner in which this can be achieved is by setting $\alpha = 1$ in (44) and then specifying that the second term in both the numerator and the sum in the denominator, $\mathbb{P}(\theta_t | \ \mathcal{F}_{t-1}, R^m_{T_i} = s)$, 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 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); see also Tversky and Kahneman (1974). 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; similarly, Fuster, Laibson, and Mendel (2010) and Fuster, Hebert, and Laibson (2012) consider a model of *natural expectations*, in which agents form beliefs from simple models fit to the available data.[83] Giglio and Kelly (2017) find that a calibrated version of the Fuster, Laibson, and Mendel (2010) model allows them to fit the patterns of excess asset-price volatility documented in their paper.

[83]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; Angeletos and Lian, 2017; Gabaix, 2017), which can generate underextrapolation, but it is possible that there are distinct phenomena governing the behavior of the average household versus the average investor.

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). And for an earlier example of similar logic, 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 psychologically 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 and face short-sale or borrowing constraints (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, 2015); it is accordingly unclear the extent to which these models are "behavioral" versus rational even at the individual level.[84] 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.[85]

Separately, models of *ambiguity aversion* (e.g., Gilboa and Schmeidler, 1989; Hansen and Sargent, 2001, 2008; Epstein and Schneider, 2010) — which specify that investors have uncertainty not only over model parameters 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, as discussed by Hansen (2014).[86] In this

[84] Similarly, classes of models in which agents' private information can affect prices in general require either the presence of some less-than-fully-rational noise traders (e.g., Shiller, 1984; Campbell, Grossman, and Wang, 1993) or widespread departures from rationality (Daniel, Hirshleifer, and Subrahmanyam, 1998); see Eyster, Rabin, and Vayanos (2015) and Campbell (2017) for further discussion. Pan and Poteshman (2006) present empirical evidence suggesting that informed trading is relatively unimportant for index options.

[85] We thank Emmanuel Farhi for this suggestion.

[86] For assessing certain moments of the data — for example, excess returns — these distortions can be isomorphic to increases in risk aversion under the correct probability distribution. But the agent's belief distribution itself is nonetheless distorted relative to the true measure. See Campbell (2017, Chapter 6.5) for a simple exposition, and as above see Hansen (2014) for further discussion.

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.

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), Wright (2016), 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 (2017); meanwhile, a literature including Iacoviello (2005), He and Krishnamurthy (2013), and Midrigan and Philippon (2016) 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 first 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:

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

where the second and third lines follow from the martingale property of beliefs under Bayes' rule (see Footnote 15) 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 $\mathbb{E}[m_{t_1,t_2}] = \mathbb{E}[r_{t_1,t_2}]$. □

***Proof of Equation (22).*** Equation (21) implies that for two adjacent return states $s_{j-1}$ and $s_j$,

$$\begin{aligned}q^m_{t,T_i,K_j} - q^m_{t,T_i,K_{j-1}} &= \frac{1}{R^f_{t,T_i}}\left[\sum_{j'\geqslant j}(V^m_{i,j'} - V_{i,j})\,\mathbb{P}^*_t(V^m_{T_i} = V^m_{i,j'}) - \sum_{j'\geqslant j-1}(V^m_{i,j'} - V_{i,j-1})\,\mathbb{P}^*_t(V^m_{T_i} = V^m_{i,j'})\right] \\ &= \frac{1}{R^f_{t,T_i}}\left[\sum_{j'\geqslant j}(V^m_{i,j-1} - V^m_{i,j})\,\mathbb{P}^*_t(V^m_{T_i} = V^m_{i,j'})\right] \\ &= \frac{1}{R^f_{t,T_i}}(V^m_{i,j-1} - V^m_{i,j})\left[1 - \mathbb{P}^*_t(V^m_{T_i} < V_{i,j})\right].\end{aligned}$$

Dividing through by $(V^m_{i,j-1} - V^m_{i,j})/R^f_{t,T_i} = (K_{j-1} - K_j)/R^f_{t,T_i}$ and rearranging,

$$R^f_{t,T_i}\frac{q^m_{t,T_i,K_j} - q^m_{t,T_i,K_{j-1}}}{K_j - K_{j-1}} = \mathbb{P}^*_t(V^m_{T_i} < V_{i,j}) - 1.$$

Repeating this analysis for the pair $s_j$ and $s_{j+1}$, we obtain $R^f_{t,T_i}\frac{q^m_{t,T_i,K_{j+1}} - q^m_{t,T_i,K_j}}{K_{j+1}-K_j} = \mathbb{P}^*_t(V^m_{T_i} < V_{i,j+1}) - 1$. Subtracting the preceding equation from this equation and using $\mathbb{P}^*_t(R^m_{T_i} = s_j) = \mathbb{P}^*_t(V^m_{T_i} = V^m_{i,j})$ yields equation (22). □

### Additional Lemmas Used in Proofs for Section 4

Before proceeding to the proofs of our main results, we provide two additional lemmas that are useful in proving those results.

**Lemma A.1.** *For some return-state pair $(s_j, s_{j+1})$, denote the path of conditional risk-neutral beliefs from date $t+1$ through $T_i$ by $\mathcal{B}_{t,T_i,j} \equiv (\widetilde{\pi}^*_{t+1,T_i,j}, \widetilde{\pi}^*_{t+2,T_i,j}, \ldots, \widetilde{\pi}^*_{T_i,T_i,j})$, and some arbitrary realization for that path by $\mathfrak{b}_{t,T_i,j}$. If conditional transition independence holds for this return-state pair, then the following two statements must hold for $0_i \leqslant t < T_i$:*

*(i) The time-$t$ conditional risk-neutral probability of any possible path ending in $\widetilde{\pi}^*_{T_i,T_i,j} = 1$ (or equivalently any path such that $R^m_{T_i} = s_j$), with such a path denoted by $\mathfrak{b}^1_{t,T_i,j}$, is given by*

$$\mathbb{P}^*_t\left(\mathcal{B}_{t,T_i,j} = \mathfrak{b}^1_{t,T_i,j} \;\middle|\; R^m_{T_i} \in \{s_j, s_{j+1}\}\right) = \frac{\widetilde{\pi}^*_{t,T_i,j}}{\widetilde{\pi}_{t,T_i,j}} \mathbb{P}_t\left(\mathcal{B}_{t,T_i,j} = \mathfrak{b}^1_{t,T_i,j} \;\middle|\; R^m_{T_i} \in \{s_j, s_{j+1}\}\right).$$

*(ii) The time-$t$ conditional risk-neutral probability of any possible path ending in $\widetilde{\pi}^*_{T_i,T_i,j} = 0$ (or equivalently any path such that $R^m_{T_i} = s_{j+1}$), with such a path denoted by $\mathfrak{b}^0_{t,T_i,j}$, is given by*

$$\mathbb{P}^*_t\left(\mathcal{B}_{t,T_i,j} = \mathfrak{b}^0_{t,T_i,j} \;\middle|\; R^m_{T_i} \in \{s_j, s_{j+1}\}\right) = \frac{1-\widetilde{\pi}^*_{t,T_i,j}}{1-\widetilde{\pi}_{t,T_i,j}} \mathbb{P}_t\left(\mathcal{B}_{t,T_i,j} = \mathfrak{b}^0_{t,T_i,j} \;\middle|\; R^m_{T_i} \in \{s_j, s_{j+1}\}\right).$$

*We conclude that when evaluating the conditional probability of the realization of any such path, the change of measure (or risk adjustment) from $\mathbb{P}_t$ to $\mathbb{P}^*_t$ does not depend on the path itself aside from its terminal value $\widetilde{\pi}^*_{T_i,T_i,j}$. This result further implies that the conditional expectation of the realization of the SDF in state $s_{j'}$, for $j' = j, j+1$, is path-independent for $0_i \leqslant t < T_i$:*

$$\mathbb{E}_t\left[\frac{M_{T_i}}{M_t} \;\middle|\; \mathcal{B}_{t,T_i,j} = \mathfrak{b}_{t,T_i,j}\,, R^m_{T_i} = s_{j'}\right] = \mathbb{E}_t\left[\frac{M_{T_i}}{M_t} \;\middle|\; R^m_{T_i} = s_{j'}\right]$$

*for any path $\mathfrak{b}_{t,T_i,j}$ such that $R^m_{T_i} = s_{j'}$.*

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

$$\mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j'}] = \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j'}].$$

***Proof of Lemma A.2.*** The realization of $m^*_{T_i,j}$ depends on the path of risk-neutral beliefs, so denote $m^*_{T_i,j} = m^*_{T_i,j}(\mathcal{B}_{t,T_i,j})$. Then using that $\mathbb{E}_t[M_{T_i}/M_t \mid \mathcal{B}_{t,T_i,j} = \mathfrak{b}_{t,T_i,j}\,, R^m_{T_i} = s_{j'}] = \mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j'}]$ from Lemma A.1, we have that

$$\mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j'}] = \sum_{\mathfrak{b}_{t,T_i,j}\,:\,R^m_{T_i} = s_{j'}} m^*_{T_i,j}(\mathfrak{b}_{t,T_i,j})\, \mathbb{P}^*_t\left(\mathcal{B}_{t,T_i,j} = \mathfrak{b}_{t,T_i,j} \;\middle|\; R^m_{T_i} = s_{j'}\right)$$

$$= \sum_{\mathfrak{b}_{t,T_i,j} \,:\, R^m_{T_i} = s_{j'}} m^*_{T_i,j}(\mathfrak{b}_{t,T_i,j}) \, \frac{\mathbb{E}_t[M_{T_i}/M_t \mid \mathcal{B}_{t,T_i,j} = \mathfrak{b}_{t,T_i,j}, R^m_{T_i} = s_{j'}]}{\mathbb{E}_t[M_{T_i}/M_t \mid R^m_{T_i} = s_{j'}]}$$

$$\times \mathbb{P}_t\big(\mathcal{B}_{t,T_i,j} = \mathfrak{b}_{t,T_i,j} \;\big|\; R^m_{T_i} = s_{j'}\big)$$

$$= \sum_{\mathfrak{b}_{t,T_i,j} \,:\, R^m_{T_i} = s_{j'}} m^*_{T_i,j}(\mathfrak{b}_{t,T_i,j}) \, \mathbb{P}_t\big(\mathcal{B}_{t,T_i,j} = \mathfrak{b}_{t,T_i,j} \;\big|\; R^m_{T_i} = s_{j'}\big)$$

$$= \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j'}]. \qquad \square$$

## Proofs for Section 4

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

$$\begin{aligned}
\mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} \in \{s_j, s_{j+1}\}] &= \widetilde{\pi}_{0_i,T_i,j}\, \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j] + (1 - \widetilde{\pi}_{0_i,T_i,j})\, \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] \\
&= \widetilde{\pi}^*_{0_i,T_i,j}\, \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j] + (1 - \widetilde{\pi}^*_{0_i,T_i,j})\, \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] \\
&\quad + \Delta_{0_i,j}\Big\{\mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] - \mathbb{E}_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j]\Big\} \\
&= \widetilde{\pi}^*_{0_i,T_i,j}\, \mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j] + (1 - \widetilde{\pi}^*_{0_i,T_i,j})\, \mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] \\
&\quad + \Delta_{0_i,j}\Big\{\mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] - \mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j]\Big\} \\
&= \mathbb{E}^*_{0_i}[r^*_{T_i,j} \mid R^m_{T_i} \in \{s_j, s_{j+1}\}] + \Delta_{0_i,j}\Big\{\mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] - \mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j]\Big\},
\end{aligned}$$

where the third equality uses Lemma A.2 and the fourth uses Lemma 1 (which applies immediately in this case).

Lemma 1 also implies that

$$\mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_{j+1}] = \widetilde{\pi}^*_{0_i,T_i,j} - \frac{\widetilde{\pi}^*_{0_i,T_i,j}}{1 - \widetilde{\pi}^*_{0_i,T_i,j}}\, \mathbb{E}^*_{0_i}[m^*_{T_i,j} \mid R^m_{T_i} = s_j],$$

so using this in the term in braces above, we obtain

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

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

Evaluating $\Delta_{0_i,j}$ yields the stated expression for the conditional expectation. *(Need to finish typing this up.)* $\square$

# Appendix B. Additional Material

**B.1. CTI and the Decompositions of Borovička, Hansen, and Scheinkman (2016)**

**B.2. Linearization of the Second Term in Proposition 5 for Empirical Implementation**

**B.3. Condition Under Which $\phi_{t,T_i,j}$ Is a Martingale**

**B.4. Details of Measurement of Risk-Neutral Distribution**

**B.5. Alternative Measurement of Risk-Neutral Distribution**

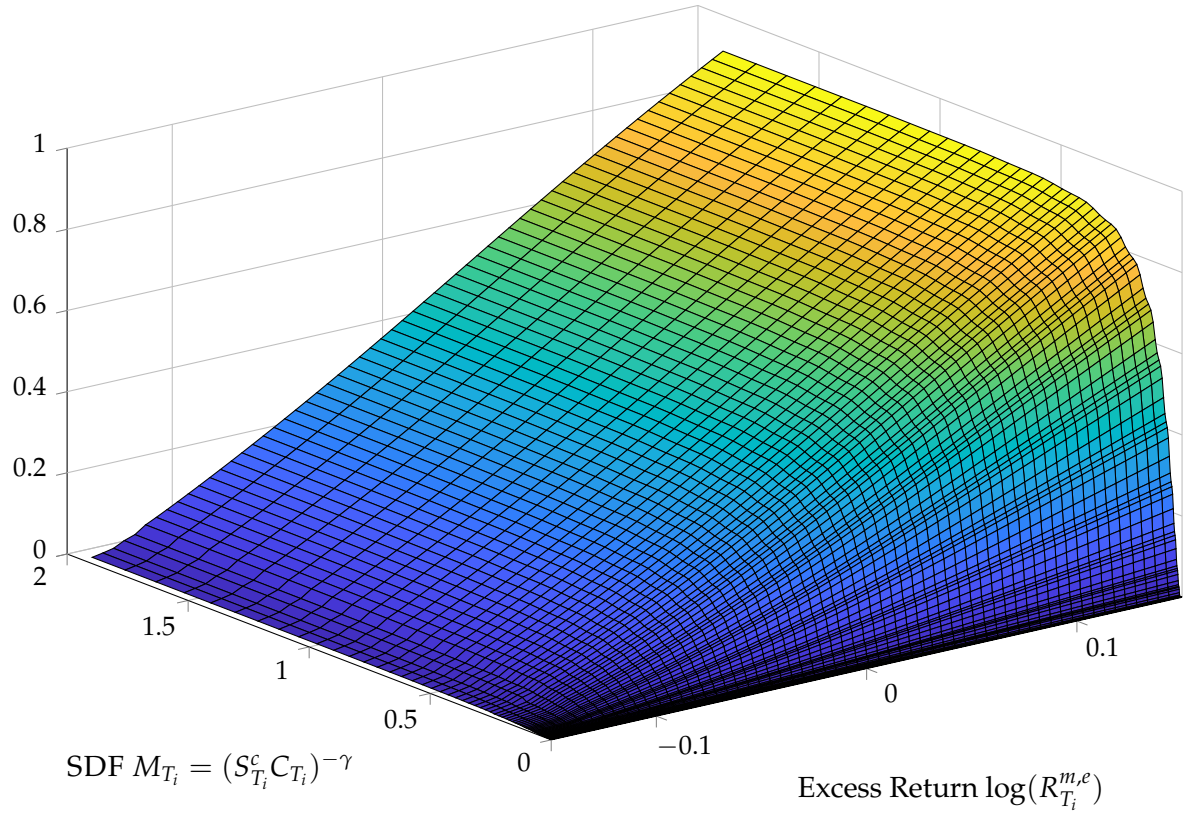
**B.6. Summary Statistics for Option-Price Data and Risk-Neutral Beliefs**

**B.7. Details of Estimation of Measurement Error**

**B.8. Details of Variable Construction for Table 2**

**B.9. Details of Solution Method and Simulations for Model in Section 6.4**

**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 Example 4, and calibrated parameters are identical to those used by Campbell and Cochrane (1999), converted to daily values, for the version of their model with imperfectly correlated consumption and dividends. Initial market index value is normalized to $V^m_{0_i} = 1$. This joint CDF is 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 $(T_i - 1) \to T_i$ obtained in the previous step; and so on. These CDFs are then used for the model simulation results presented in Section 6.4, but the CDFs themselves (e.g., the one presented here) are outputs of our model *solution* rather than outputs from the simulations.

# References


ABEL, A. B. (1999): "Risk Premia and Term Premia in General Equilibrium," *Journal of Monetary Economics*, 43, 3–33.

ACHARYA, V. V. AND L. H. PEDERSEN (2005): "Asset Pricing with Liquidity Risk," *Journal of Financial Economics*, 77, 375–410.

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.

AÏT-SAHALIA, Y. (1996): "Nonparametric Pricing of Interest Rate Derivative Securities," *Econometrica*, 64, 527–560.

AÏT-SAHALIA, Y. AND A. W. LO (1998): "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," *Journal of Finance*, 53, 499–547.

——— (2000): "Nonparametric Risk Management and Implied Risk Aversion," *Journal of Econometrics*, 94, 9–51.

AÏT-SAHALIA, Y., Y. WANG, AND F. YARED (2001): "Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?" *Journal of Econometrics*, 102, 67–110.

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. (2002): "Illiquidity and Stock Returns: Cross-Section and Time-Series Effects," *Journal of Financial Markets*, 5, 31–56.

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.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND H. EBENS (2001): "The Distribution of Realized Stock Return Volatility," *Journal of Financial Economics*, 61, 43–76.

ANDERSEN, T. G., N. FUSARI, AND V. TODOROV (2015): "Parametric Inference and Dynamic State Recovery From Option Panels," *Econometrica*, 83, 1081–1145.

ANGELETOS, G.-M. AND C. LIAN (2017): "Dampening General Equilibrium: From Micro Elasticities to Macro Effects," *MIT Mimeo*.

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

ASNESS, C. S., T. J. MOSKOWITZ, AND L. H. PEDERSEN (2013): "Value and Momentum Everywhere," *Journal of Finance*, 68, 929–985.

AUGENBLICK, N. AND M. RABIN (2017): "Belief Movement, Uncertainty Reduction, and Rational Updating," *UC Berkeley-Haas and Harvard University Mimeo*.

BACKUS, D., M. CHERNOV, AND I. MARTIN (2011): "Disasters Implied by Equity Index Options," *Journal of Finance*, 66, 1969–2012.

BACKUS, D., M. CHERNOV, AND S. ZIN (2014): "Sources of Entropy in Representative Agent Models," *Journal of Finance*, 69, 51–99.

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

BAKSHI, G., C. CAO, AND Z. CHEN (1997): "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52, 2003–2049.

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.

BARONE-ADESI, G., R. F. ENGLE, AND L. MANCINI (2008): "A GARCH Option Pricing Model with Filtered Historical Simulation," *Review of Financial Studies*, 21, 1223–1258.

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

BATES, D. S. (2000): "Post-'87 Crash Fears in the S&P Futures Option Market," *Journal of Econometrics*, 94, 181–238.

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

——— (2008): "The Market for Crash Risk," *Journal of Economic Dynamics and Control*, 32, 2291–2321.

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.

BERTSIMAS, D., L. KOGAN, AND A. W. LO (2000): "When Is Time Continuous?" *Journal of Financial Economics*, 55, 173–204.

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.

BONDARENKO, O. (2003): "Statistical Arbitrage and Securities Prices," *Review of Financial Studies*, 16, 875–919.

BORDALO, P., N. GENNAIOLI, R. LA PORTA, AND A. SHLEIFER (2017): "Diagnostic Expectations and Stock Returns," *University of Oxford, Universita Bocconi, Brown University, and Harvard University Mimeo*.

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.

BÜHLMANN, P. (1994): "Blockwise Bootstrapped Empirical Processes for Stationary Sequences," *Annals of Statistics*, 22, 995–1012.

BURASCHI, A. AND A. JILTSOV (2006): "Model Uncertainty and Option Markets with Heterogeneous Beliefs," *Journal of Finance*, 61, 2841–2897.

CAMPBELL, J. Y. (1986): "Bond and Stock Returns in a Simple Exchange Model," *Quarterly Journal of Economics*, 101, 785–803.

——— (1991): "A Variance Decomposition for Stock Returns," *Economic Journal*, 101, 157–179.

——— (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, Forthcoming.

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., S. J. GROSSMAN, AND J. WANG (1993): "Trading Volume and Serial Correlation in Stock Returns," *Quarterly Journal of Economicsa*, 108, 905–939.

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

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

——— (1988b): "The Dividend-Price Ratio and Expectations of Future Dividends and Discount Factors," *Review of Financial Studies*, 1, 195–228.

CAMPBELL, J. Y. AND S. B. THOMPSON (2008): "Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average?" *Review of Financial Studies*, 21, 1509–1531.

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.

CARR, P. AND J. YU (2012): "Risk, Return, and Ross Recovery," *Journal of Derivatives*, 20, 38–59.

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

CHERNOZHUKOV, V., H. HONG, AND E. TAMER (2007): "Estimation and Confidence Regions for Parameter Sets in Econometric Models," *Econometrica*, 75, 1243–1284.

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. (2017): "Turning Alphas into Betas: Arbitrage and Endogenous Risk," *Harvard University Mimeo*.

CHORDIA, T., R. ROLL, AND A. SUBRAHMANYAM (2000): "Market Liquidity and Trading Activity," *Journal of Finance*, 56, 501–530.

——— (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. (1992): "Explaining the Variance of Price-Dividend Ratios," *Review of Financial Studies*, 5, 243–280.

——— (2005): *Asset Pricing*, Princeton, NJ: Princeton University Press, revised ed.

——— (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 S. PERRAKIS (2009): "Mispricing of S&P 500 Index Options," *Review of Financial Studies*, 22, 1247–1277.

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.

COX, J. C. AND S. A. ROSS (1976): "The Valuation of Options for Alternative Stochastic Processes," *Journal of Financial Economics*, 3, 145–166.

CUESDEANU, H. AND J. C. JACKWERTH (2016): "The Pricing Kernel Puzzle: Survey and Outlook," *University of Konstanz Mimeo*.

CUTLER, D. M., J. M. POTERBA, AND L. H. SUMMERS (1989): "What Moves Stock Prices?" *Journal of Portfolio Management*, 15, 4–12.

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.

DELONG, J. B., A. SHLEIFER, L. H. SUMMERS, AND R. WALDMANN (1990): "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance*, 45, 379–395.

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

DU, W., A. TEPPER, AND A. VERDELHAN (2017): "Deviations from Covered Interest Rate Parity," *Federal Reserve Board, Columbia University, and MIT Sloan Mimeo*.

DUFFIE, D. (2001): *Dynamic Asset Pricing Theory*, Princeton, NJ: Princeton University Press, third ed.

DUFFIE, D., J. PAN, AND K. J. SINGLETON (2000): "Transform Analysis and Asset Pricing for Affine Jump-Diffusions," *Econometrica*, 68, 1343–1376.

DYBVIG, P. H. AND L. C. G. ROGERS (1997): "Recovery of Preferences from Observed Wealth in a Single Realization," *Review of Financial Studies*, 10, 151–174.

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 (2015): "Financial Markets where Traders Neglect the Informational Content of Prices," *NBER Working Paper No. 21224*.

FAMA, E. F. (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work," *Journal of Finance*, 25, 383–417.

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

——— (1991): "Efficient Capital Markets: II," *Journal of Finance*, 46, 1575–1617.

FAMA, E. F. AND K. R. FRENCH (1988a): "Dividend Yields and Expected Stock Returns," *Journal of Financial Economics*, 22, 3–25.

——— (1988b): "Permanent and Temporary Components of Stock Prices," *Journal of Political Economy*, 96, 246–273.

——— (1993): "Common Risk Factors in the Returns on Stocks and Bonds," *Journal of Financial Economics*, 33, 3–56.

FARHI, E., S. P. FRAIBERGER, X. GABAIX, R. RANCIERE, AND A. VERDELHAN (2015): "Crash Risk in Currency Markets," *Harvard University, NYU, IMF, and MIT Mimeo*.

FARHI, E. AND I. WERNING (2017): "Monetary Policy, Bounded Rationality, and Incomplete Markets," *Harvard University and MIT Mimeo*.

FRANKEL, J. A. AND K. A. FROOT (1987): "Using Survey Data to Test Standard Propositions Regarding Exchange Rate Expectations," *American Economic Review*, 77, 133–153.

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., B. HEBERT, AND D. LAIBSON (2012): "Natural Expectations, Macroeconomic Dynamics, and Asset Pricing," in *NBER Macroeconomics Annual 2011*, ed. by D. Acemoglu and M. Woodford, Chicago: University of Chicago Press, vol. 26, chap. 1, 1–48.

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

GABAIX, X. (2009): "Linearity-Generating Processes: A Modeling Tool Yielding Closed Forms for Asset Prices," *NYU Mimeo*.

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

——— (2017): "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: University of Chicago Press, chap. 5, 379–431.

GERBER, H. U. AND E. S. W. SHIU (1994): "Option Pricing by Esscher Transforms," *Transactions of the Society of Actuaries*, 46, 99–140.

GHOSH, A., C. JULLIARD, AND A. P. TAYLOR (2017): "What Is the Consumption-CAPM Missing? An Information-Theoretic Framework for the Analysis of Asset Pricing Models," *Review of Financial Studies*, 30, 442–504.

GIGLIO, S. AND B. T. KELLY (2017): "Excess Volatility: Beyond Discount Rates," *University of Chicago Mimeo*.

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.

GROSSMAN, S. (1976): "On the Efficiency of Competive Stock Markets Where Trades Have Diverse Information," *Journal of Finance*, 31, 573–585.

——— (1978): "Further Results on the Informational Efficiency of Competitive Stock Markets," *Journal of Economic Theory*, 18, 81–101.

GROSSMAN, S. J. AND R. J. SHILLER (1981): "The Determinants of the Variability of Stock Market Prices," *American Economic Review: Papers and Proceedings*, 71, 222–227.

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.

HAN, B. (2008): "Investor Sentiment and Option Prices," *Review of Financial Studies*, 21, 387–414.

HANSEN, B. E. (2017): "Econometrics," *Unpublished Manuscript, University of Wisconsin*.

HANSEN, L. P. (1982): "Large Sample Properties of Generalised Method of Moments Estimators," .

——— (2012a): "Dynamic Valuation Decomposition Within Stochastic Economies," *Econometrica*, 80, 911–967.

——— (2012b): "Proofs for Large Sample Properties of Generalized Method of Moments Estimators," *Journal of Econometrics*, 170, 325–330.

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

HANSEN, L. P., J. HEATON, AND E. G. LUTTMER (1995): "Econometric Evaluation of Asset Pricing Models," *Review of Financial Studies*, 8, 237–274.

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 S. F. RICHARD (1987): "The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models," *Econometrica*, 55, 587–613.

HANSEN, L. P. AND T. J. SARGENT (2001): "Robust Control and Model Uncertainty," *American Economic Review: Papers and Proceedings*, 91, 60–66.

——— (2008): *Robustness*, Princeton, NJ: Princeton University Press.

HANSEN, L. P. AND J. A. SCHEINKMAN (2009): "Long-Term Risk: An Operator Approach," *Econometrica*, 77, 177–234.

——— (2016): "Stochastic Compounding and Uncertain Valuation," *University of Chicago and Columbia University Mimeo*.

HANSEN, L. P. AND K. J. SINGLETON (1982): "Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica*, 50, 1269–1286.

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.

——— (1979): "Martingales and Arbitrage in Multiperiod Securities Markets," *Journal of Economic Theory*, 20, 381–408.

HE, Z., B. KELLY, AND A. MANELA (2016): "Intermediary Asset Pricing: New Evidence From Many Asset Classes," *Forthcoming, Journal of Financial Economics*.

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. (2001): "Investor Psychology and Asset Pricing," *Journal of Finance*, 56, 1533–1597.

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

HONG, H. AND J. C. STEIN (1999): "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *Journal of Finance*, 54, 2143–2184.

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.

HUANG, C.-F. AND H. PAGÈS (1992): "Optimal Consumption and Portfolio Policies with an Infinite Horizon: Existence and Convergence," *Annals of Applied Probability*, 2, 36–64.

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.

JARROW, R. A. (1987): "The Pricing of Commodity Options with Stochastic Interest Rates," *Advances in Futures and Options Research*, 2, 19–45.

JENSEN, C. S., D. LANDO, AND L. H. PEDERSEN (2016): "Generalized Recovery," *Copenhagen Business School 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 (2017): "Risk, Unemployment, and the Stock Market: A Rare-Event-Based Explanation of Labor Market Volatility," *University of Pennsylvania Mimeo*.

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

KRAUS, A. AND G. A. SICK (1980): "Distinguishing Beliefs and Preferences in Equilibrium Prices," *Journal of Finance*, 35, 335–344.

KREPS, D. M. (1977): "A Note on 'Fulfilled Expectations' Equilibria," *Journal of Economic Theory*, 14, 32–43.

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.

LIAO, G. Y. (2016): "Credit Migration and Covered Interest Rate Parity," *Project on Behavioral Finance and Financial Stability Working Paper No. 2016-07*.

LINN, M., S. SHIVE, AND T. SHUMWAY (2015): "Pricing Kernel Monotonicity and Conditional Information," *University of Michigan and University of Notre Dame Mimeo*.

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.

LUDVIGSON, S. C. (2013): "Advances in Consumption-Based Asset Pricing: Empirical Tests," in *Handbook of the Economics of Finance*, ed. by G. M. Constantinides, M. Harris, and R. M. Stulz, Amsterdam: Elsevier, vol. 2B, chap. 12, 799–906.

MALIAR, L. AND S. MALIAR (2014): "Numerical Methods for Large-Scale Dynamic Economic Models," in *Handbook of Computational Economics*, ed. by K. Schmedders and K. L. Judd, Amsterdam: Elsevier, vol. 3, chap. 7, 325–477.

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. (2013): "Consumption-Based Asset Pricing with Higher Cumulants," *Review of Economic Studies*, 80, 745–773.

——— (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.

MERTON, R. C. (1973a): "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–887.

——— (1973b): "Theory of Rational Option Pricing," *Bell Journal of Economics and Management Science*, 4, 141–183.

MIDRIGAN, V. AND T. PHILIPPON (2016): "Household Leverage and the Recession," *NYU Mimeo*.

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

MURAVYEV, D. (2016): "Order Flow and Expected Option Returns," *Journal of Finance*, 71, 673–708.

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

NEWEY, W. K. AND D. MCFADDEN (1994): "Large Sample Estimation and Hypothesis Testing," in *Handbook of Econometrics*, ed. by R. F. Engle and D. McFadden, Amsterdam: Elsevier, vol. 4, chap. 36, 2111–2245.

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. (2002): "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study," *Journal of Financial Economics*, 63, 3–50.

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.

PETERSEN, M. A. (2009): "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches," *Review of Financial Studies*, 22, 435–480.

POLITIS, D. N. AND J. P. ROMANO (1994): "Large Sample Confidence Regions Based on Subsamples Under Minimal Assumptions," *Annals of Statistics*, 22, 2031–2050.

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

POTESHMAN, A. M. (2001): "Underreaction, Overreaction, and Increasing Misreaction to Information in the Options Market," *Journal of Finance*, 56, 851–876.

QIN, L. AND V. LINETSKY (2016): "Positive Eigenfunctions of Markovian Pricing Operators: Hansen-Scheinkman Factorization, Ross Recovery, and Long-Term Pricing," *Operations Research*, 64, 99–117.

——— (2017): "Long-Term Risk: A Martingale Approach," *Econometrica*, 85, 299–312.

QUESSETTE, R. (2002): "New Products, New Risks," *Risk*, 15, 97–100.

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.

ROSENBERG, J. V. AND R. F. ENGLE (2002): "Empirical Pricing Kernels," *Journal of Financial Economics*, 64, 341–372.

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

SANTA-CLARA, P. AND A. SARETTO (2009): "Option Strategies: Good Deals and Margin Calls," *Journal of Financial Markets*, 12, 391–417.

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.

SIRIWARDANE, E. (2015): "The Probability of Rare Disasters: Estimation and Implications," *Harvard Business School Working Paper 16-061*.

STEIN, J. (1989): "Overreactions in the Options Market," *Journal of Finance*, 44, 1011–1023.

STOCK, J. H. AND M. W. WATSON (2016): "Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and H. Uhlig, Amsterdam: Elsevier, vol. 2A, chap. 8, 415–525.

STUTZER, M. (1996): "A Simple Nonparametric Approach to Derivative Security Valuation," *Journal of Finance*, 51, 1633–1652.

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.

——— (1974): "Judgment Under Uncertainty: Heuristics and Biases," *Science*, 185, 1124–1131.

VAN BINSBERGEN, J. H. AND R. S. KOIJEN (2017): "The Term Structure of Returns: Facts and Theory," *Journal of Financial Economics*, 124, 1–21.

VIJH, A. M. (1990): "Liquidity of the CBOE Equity Options," *Journal of Finance*, 45, 1157–1179.

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

WALDEN, J. (2016): "Recovery with Unbounded Diffusion Processes," *UC Berkeley-Haas Mimeo*.

WEITZMAN, M. L. (2007): "Subjective Expectations and Asset-Return Puzzles," *American Economic Review*, 97, 1102–1130.

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

WRIGHT, J. H. (2016): "Options-Implied Probability Density Functions for Real Interest Rates," *International Journal of Central Banking*, 12, 129–149.

XIONG, W. (2001): "Convergence Trading with Wealth Effects: An Amplification Mechanism in Financial Markets," *Journal of Financial Economics*, 62, 247–292.

ZIEGLER, A. (2007): "Why Does Implied Risk Aversion Smile?" *Review of Financial Studies*, 20, 859–904.