

A New Test of Excess Movement in Asset Prices*

Ned Augenblick
UC Berkeley Haas

Eben Lazarus
MIT Sloan

AUGUST 2022

Abstract

We derive new bounds on the rational variation in asset prices over time. The resulting test requires no proxy for fundamental value, and it allows significantly more flexibility in preferences and discount rates than in standard volatility tests. We gain traction by focusing specifically on risk-neutral beliefs implied by option prices. The core insight is that movement in physical beliefs must be bounded, and although risk preferences distort prices away from beliefs, it is still possible to place bounds on risk-neutral belief movement. Aside from rational expectations, our main assumption is that the slope of the stochastic discount factor is locally constant within an option contract. We show that this assumption allows for an informative null, and it is satisfied in a range of standard frameworks. Implementing our test empirically using index options, we find that there is so much movement in risk-neutral beliefs that the bounds are routinely violated.

KEYWORDS: Volatility bounds, asset prices, beliefs, rational expectations, risk aversion

JEL CODES: D84, G13, G14, G41

*Contact: ned@haas.berkeley.edu and elazarus@mit.edu. Previous versions of this paper circulated as “Restrictions on Asset-Price Movements Under Rational Expectations: Theory and Evidence.” We are grateful to Francesca Bastianello, Laura Blattner, Jaroslav Borovička, John Campbell, Gabriel Chodorow-Reich, Emmanuel Farhi, Xavier Gabaix, Niels Gormsen, Sam Hanson, Bryan Kelly, David Laibson, Ian Martin, Jonathan Parker, Matthew Rabin, Andrei Shleifer, Juan Sotes-Paladino, David Sraer, Jim Stock, Adrien Verdelhan, Jessica Wachter, and seminar and conference participants at Harvard, MIT, UC Berkeley, Chicago Booth, Yale SOM, Stanford GSB, Northwestern Kellogg, Duke Fuqua, LSE, University of Sydney, Bocconi University, City University of Hong Kong, the AEA Annual Meeting, SITE, FIRN, the Chicago Fed Rookie Conference, the NBER Behavioral Finance Meeting, and the NBER Summer Institute Asset Pricing Meeting for advice, comments, and helpful discussions.

1. Introduction

Asset prices are volatile. Whether they are *too* volatile to reflect rational variation in expected future cash flows, though, is by itself an unresolvable question: the true distribution of future cash-flow streams is unobservable, as are aggregate risk and time preferences. The relevant question is instead what *joint* set of assumptions on (i) beliefs and (ii) risk and time preferences — the two elements at the heart of almost all modern theories of asset prices — can be rejected in the data.

In this paper, we make progress on this longstanding question. We show that even with significantly more flexibility in preferences and discount rates than in standard volatility tests, it is possible to place informative bounds on asset-price movements under rational expectations (RE). These new volatility bounds are of theoretical interest in their own right, and they show how the dynamic behavior of certain asset prices can be restricted without knowledge of their true fundamental value. Implementing our bounds empirically, we find that there is so much price variation that the bounds are routinely violated.

To set the stage for our analysis, consider the seminal volatility test provided by [Shiller \(1981\)](#). He finds evidence of excess volatility in equity-index prices relative to a proxy for fundamental value, but this proxy is constructed under the assumption that discount rates over future cash flows are constant over time. Shiller’s test therefore provides a joint rejection of (i) rational pricing and (ii) constant discount rates. But given the ample evidence that discount rates vary over time, it is unclear what to make of this joint rejection. As [Fama \(1991\)](#) argues, “volatility tests are [a] useful way to show that expected returns vary through time,” but they “give no help on the central issue of whether the variation in expected returns is rational.”

Our goal is to understand whether further statements can be made under less restrictive assumptions. As in Shiller’s case, we focus our analysis on expectations over the future value of an equity index. But a key feature distinguishing our analysis is that we consider the behavior of *options* written on the future index value, rather than the behavior of the underlying index itself. This apparently minor change in focus allows us to gain significant theoretical traction, as it reduces the set of possible unobservables driving asset-price variation and therefore the assumptions required to derive volatility bounds. Specifically, as there are multiple options written on the same index with the same expiration date, comparing their *relative* prices over time allows us to discard price variation arising from time discounting (since the risk-free rate affects the prices of all options with the same maturity symmetrically) and common unobservable shocks (e.g., to risk aversion or the quantity of risk) that have the same effect on the prices of options with different strikes.

The comparison of options’ relative prices is captured by the standard transformation of prices into so-called *risk-neutral (RN) beliefs* over the underlying index’s future price. These are the main object of our analysis. The relative value of RN beliefs across index-value states is then a function of (i) the relative likelihood that each state occurs and (ii) the relative value of a marginal dollar in each state (which is driven by risk preferences). While the use of RN beliefs limits the potential drivers of volatility, the main identification challenge remains: we do not observe any direct proxy for the

appropriate physical beliefs or the true level of risk aversion period by period. Therefore, without strong assumptions about beliefs or risk preferences, it would seem difficult to make statements about appropriate prices at a given time, let alone their movement over time. We show that this intuition is incomplete. While it is not possible to pinpoint the correct price at any one time, it is also not necessary to do so: option prices over time are connected because they all concern the terminal value of an index at the same terminal date, and this connection provides restrictions on how they can vary under RE.

To demonstrate the logic of our test, it is useful first to consider the simple case in which we can directly observe a person’s subjective belief $\pi_t(\theta = 1)$ over some binary outcome $\theta \in \{0, 1\}$. Suppose we observe that π_t continually moves from 0.10 to 0.90 as t progresses. While it is of course possible to rationalize this movement ex post by constructing a particular set of signal realizations from some data-generating process (DGP), this amount of movement appears intuitively “rare” for someone with RE. Formalizing this idea, [Augenblick and Rabin \(2021\)](#) note that, when uncertainty is resolved by some period T , the expectation of the sum of squared changes in beliefs across all periods (*belief movement*, $\mathbb{E} \sum_{t=0}^{T-1} (\pi_{t+1} - \pi_t)^2$) must equal *initial uncertainty* ($\pi_0(1 - \pi_0)$) under RE, *regardless* of the DGP. Equivalently, expected *excess movement* — belief movement minus initial uncertainty, which we denote by $\mathbb{E}[X]$ — must always be zero. Intuitively, changes in beliefs must on average correspond to the resolution of uncertainty (from its initial level $\pi_0(1 - \pi_0)$ to 0); rational belief movements must mean the person is learning something about the outcome in question on average. If beliefs are instead continually shifting dramatically relative to initial uncertainty, the null of rationality can be statistically rejected at some confidence level.

The main theoretical contribution of this paper is to show how this logic — that movement in beliefs must correspond on average to reduction in uncertainty — can be used to restrict excess movement in *risk-neutral* beliefs, $\mathbb{E}[X^*]$, in the general case in which only risk-neutral rather than physical beliefs are observable. Our task becomes considerably more difficult in this case: RN beliefs need not follow a martingale given their distortion relative to physical beliefs, and this means that $\mathbb{E}[X^*]$ can be non-zero.¹ But given that the distortion between RN and physical beliefs is indexed by risk aversion over the terminal states, we show in this case that the admissible $\mathbb{E}[X^*]$ can be bounded as a simple function of this risk-aversion value (or, more generally, the slope of the stochastic discount factor across states). The broad intuition is the same: if a person’s valuation for a binary option that pays \$1 in a given state at time T oscillates between \$0.10 and \$0.90, her unobservable risk aversion renders her precise physical beliefs unidentified, but these extreme valuation changes imply belief movements that must be rare regardless of her exact risk aversion.

Our upper bound for excess RN belief movement $\mathbb{E}[X^*]$ is tight in the sense that it is possible to construct a (somewhat perverse) DGP that produces movement that is arbitrarily close to the bound. But as this bound is very conservative by construction, we also provide a tighter bound that holds under somewhat more restrictive assumptions. In both cases, the bound depends on

¹Equivalently, some positive average amount of apparent excess movement in risk-neutral beliefs can be rationalized as long as the person (or market) exhibits some risk aversion over the time- T outcome.

a risk-aversion parameter ϕ that corresponds to the relative marginal value of a dollar across time- T states. For example, if the person is forming beliefs over two possible terminal consumption states $C_T \in \{C_{low}, C_{high}\}$, then $\phi = \frac{U'(C_{low})}{U'(C_{high})}$. In the empirically relevant case where we observe RN beliefs over the index return, then $\phi = \frac{\mathbb{E}[M_{t,T} | R_T=R_{low}]}{\mathbb{E}[M_{t,T} | R_T=R_{high}]}$, where $M_{t,T}$ is the stochastic discount factor (SDF) and R_{low}, R_{high} are arbitrary index return values. Since our bound for $\mathbb{E}[X^*]$ depends on this parameter ϕ , any observed average excess movement in the data is informative as to the minimal ϕ required in order for the bound to be satisfied. We are effectively inverting the logic of the joint hypothesis problem: under our maintained assumptions, the data is informative as to the minimal ϕ needed to rationalize the observed amount of excess movement in RN beliefs.

Our main assumption imposed alongside RE is that ϕ is constant over time within a given belief stream, which generally lasts weeks to months in our empirical exercise. We call this assumption *conditional transition independence (CTI)*. Intuitively, CTI imposes that the marginal value of \$1 in some time- T state does not change *relative* to the marginal value in some other nearby state. In asset-pricing terms, it allows for permanent shocks to the SDF, which may alter $\mathbb{E}_t[M_{t,T}]$ for all states, but which do not affect the *relative* expectations across states. It is thus less restrictive than the assumptions imposed in past work, as it allows for significant variation in discount rates. It is also met either exactly or approximately under many standard modeling frameworks. This should *not*, however, be taken as evidence *in favor* of CTI in fact holding exactly: it is a knife-edge restriction that is unlikely to hold exactly in reality. Instead, it allows for an informative joint null, as a violation of our bounds provides useful information on the changes in modeling assumptions — either RE or CTI — needed to match the price variation observed in the data.

To provide further guidance on this issue, we also consider how our bounds change given variation in ϕ (so that CTI is violated). We argue that such variation is in fact unlikely to meaningfully affect our conclusions. This may seem counterintuitive: one might imagine that if ϕ oscillates between 2.9 and 3.1, for example, then RN belief movement can be unbounded even with no movement in physical beliefs. This logic, however, ignores the fact that the person must also have rational expectations over ϕ , making such persistent oscillations rare. To demonstrate this principle, we first prove theoretically that if ϕ evolves as a martingale or a supermartingale, our bounds continue to hold. We then conduct numerical simulations over thousands of DGPs in which both the underlying state and ϕ are uncertain, so that ϕ is time-varying. We find across all DGPs that even very large initial uncertainty in ϕ has little impact on the average and maximum $\mathbb{E}[X^*]$ statistics. Finally, we simulate RN beliefs in the habit-formation model of [Campbell and Cochrane \(1999\)](#), which features time-varying ϕ . We find that our bounds continue to hold and continue to be conservative. We thus strongly suspect that reasonable changes in ϕ are not driving our findings, though we do not intend for this set of results to be the last word on this question.

After providing our theoretical results and discussing their interpretation and robustness, we then take our bounds to the data. We obtain S&P 500 index option prices from OptionMetrics, and we use standard methods to infer the risk-neutral distribution over index returns for each option expiration date in the sample. In order to map to our two-state theoretical setting, we then translate

each full distribution into a set of binary RN beliefs $\pi_i^*(R_T = \theta_j | R_T \in \{\theta_j, \theta_{j+1}\})$; these correspond to the RN probability that the index return will be equal to (or in a range close to) θ_j , conditional on being either θ_j or θ_{j+1} . (We set our return states to correspond to five-percentage-point bins for the S&P return.) We then implement our theoretical bounds, which allow us to infer the minimal risk-aversion value ϕ (at each point in the return distribution) needed to rationalize the observed variation in RN beliefs over the index return.

We find that extremely high risk aversion is needed to rationalize the observed excess movement under our maintained assumptions. In many cases, there is in fact no value of ϕ under which the tight version of the bound is met, and the conservative version of the bound generally implies implausibly large values for ϕ . This suggests that many leading rational frameworks capable of explaining medium-to-low-frequency variation in asset prices have difficulty rationalizing the large degree of observed medium-to-high-frequency variation in RN beliefs.

Given that we conduct our estimation using variation in index options prices, we must also account for the effect of possible market microstructure noise. In order to do so, we derive an additional theoretical result that describes how microstructure noise biases our estimates through its effect on observed excess movement X^* . This bias depends on the variance of the noise component of observed RN beliefs. We accordingly proceed to estimate this noise variance in the data, by turning to a set of high-frequency index option prices (which we obtain directly from the CBOE) for this purpose. Using these intraday data, we implement the microstructure noise variance estimator proposed by [Li and Linton \(2022\)](#), which is particularly well suited for our purposes. We can then construct an empirical noise correction, removing the effect of noise from X^* before we conduct our estimation. All of our results are noise-corrected in this way, which help to ensure that our findings do not depend on idiosyncrasies specific to the options market.

We then briefly explore possible explanations for our findings, with the goal of providing some positive directions for alternative models. Conducting regressions of RN excess movement on a range of potentially relevant financial variables, we find that X^* has a strong positive relationship with measures related to aggregate equity valuations and volatility, while it is unrelated to option bid-ask spreads or trading volume. This provides further evidence that our findings are not simply reflecting option-market frictions or segmentation; instead, option-price variation contains relevant information about return expectations and risk pricing.

Relation to previous literature. In addition to [Shiller \(1981\)](#), we follow, among others, [LeRoy and Porter \(1981\)](#), [De Bondt and Thaler \(1985\)](#), and [Stein \(1989\)](#) in testing for excess volatility in asset prices. [Kleidon \(1986\)](#) emphasizes non-stationarity in accounting for apparent excess volatility; much of the literature since then has emphasized time variation in discount rates ([Cochrane, 2011](#)). We show 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.

In considering derivatives rather than the behavior of the underlying index directly, our work complements [Giglio and Kelly \(2018\)](#), who document excess volatility in long-maturity claims on

volatility, inflation, commodities, and interest rates. They achieve identification by parameterizing the DGP for cash flows on the underlying, whereas we restrict the SDF. Their parameterization — an affine model under the RN measure — applies well to the term-structure-like claims they consider, but not to claims on the equity index itself, to which our framework does apply.²

Our results also complement evidence obtained from survey data, as, for example, in [Greenwood and Shleifer \(2014\)](#) and [De la O and Myers \(2021\)](#), as well as the results of [Augenblick and Rabin \(2021\)](#) for settings with directly observable beliefs. Another set of related literature endeavors to measure physical probabilities or risk preferences indirectly from options data. [Ait-Sahalia, Wang, and Yared \(2001\)](#) test whether option-implied return distributions are well-calibrated; [Ross \(2015\)](#) provides assumptions under which physical beliefs can be recovered from options; [Polkovnichenko and Zhao \(2013\)](#) consider the probability weighting functions consistent with option-price data given an assumption on the form of the weighting function. Our approach differs from this and related work in that we need not measure physical beliefs at all or know the true data-generating process for returns to conduct our tests. This semi-parametric approach ties our theoretical contribution to a line of work including, among many others, [Hansen and Jagannathan \(1991\)](#) and [Kübler and Polemarchakis \(2017\)](#), which consider the identification of structural parameters from different moments of the observable data than the ones we consider here.

Organization. [Section 2](#) introduces our theoretical framework in a simple two-state setting, which allows for clear derivations and intuition for our main results. [Section 3](#) then extends these results to a general asset-pricing setting. [Section 4](#) discusses the main assumption of CTI and the impact of violations on our results. [Section 5](#) provides additional robustness results. We then implement our bounds empirically in [Section 6](#), which describes our data and presents our results. [Section 7](#) concludes. Selected proofs for our main results are provided in [Appendix A](#), and an Online Appendix contains the remaining proofs and additional technical material.

2. Theoretical Framework: Introduction in a Simple Setting

To introduce our framework, we first examine risk-neutral (RN) belief movement in a simple setting with a single individual and two terminal states. This setting helps clarify three issues: (i) the economics underlying restrictions on belief movement under RE when the individual is risk-neutral; (ii) how risk aversion complicates this analysis; and (iii) how we can nonetheless bound RN belief movement. This step-by-step discussion sets the stage for our generalized framework in [Section 3](#).

2.1 Setup and Initial Results

Time is discrete and indexed by $t = 0, 1, 2, \dots, T$. At the beginning of each period, a person observes a signal $s_t \in S$ regarding two mutually exclusive and exhaustive events, which we call *states* $\theta \in \{0, 1\}$. The data-generating process (DGP) is general: signals are drawn from the

²See also [Gandhi, Gormsen, and Lazarus \(2022\)](#) for recent evidence on the term structure of the log equity premium.

discrete signal distribution $DGP(s_t | \theta, H_{t-1})$, where H_t represents the history of signal realizations through t . Define $\mathbb{P}(H_T)$ to be the probability of observing history H_T induced by the DGP, and write $\mathbb{E}[\cdot] \equiv \mathbb{E}^{\mathbb{P}}[\cdot]$ for the expectation under \mathbb{P} . The person's (*physical* or *subjective*) *belief* in state 1 (vs. state 0) at time t given the DGP and history H_t is denoted by $\pi_t(H_t)$. The *belief stream* $\boldsymbol{\pi}(H_t) = [\pi_0, \pi_1(s_1), \pi_2(\{s_1, s_2\}), \dots]$ is the collection of beliefs given history H_t . We often suppress the dependence of these objects on H_t to simplify notation. Given our empirical setting, we focus on *resolving* streams in which the person achieves certainty about the true state by period T with probability 1: $\pi_T(H_T) = \theta \in \{0, 1\}$ for all H_T .

Throughout the paper, we maintain the assumption that the person's beliefs over the terminal state satisfy rational expectations (RE).

ASSUMPTION 1 (RE). Beliefs satisfy $\pi_t(H_t) = \mathbb{E}[\theta | H_t]$ for any H_t .

This assumption states that the agent's beliefs coincide period by period with the true conditional probability of realizing state $\theta = 1$, and the assumption will be satisfied by a person with a correct prior who updates using Bayes' rule according to the true DGP. The assumption is in fact stronger than necessary for our main results: we could instead assume the weaker martingale restriction that $\pi_t = \mathbb{E}[\pi_{t+1} | \pi_t]$, which is implied by [Assumption 1](#). (Under this weaker condition, the person could, for example, ignore some periods' signals.) Assuming RE directly, though, helps streamline our exposition.

Given our asset-pricing setting, we assume that the DGP and the person's physical beliefs cannot be observed directly. Instead, we assume that the econometrician can observe, period by period, the person's willingness to pay for an Arrow-Debreu security that pays \$1 (one unit of the numeraire consumption good) in period T if state θ is realized. Denote this valuation by $q_t(\theta | H_t)$ for each $\theta \in \{0, 1\}$. An object analogous to $q_t(\theta | H_t)$ will be empirically observable using options data for suitably defined states, but we postpone this additional formalism to [Section 3](#).

We start by considering the simple case in which the person values consumption at all periods and in all states equally (i.e., she is risk-neutral and does not discount future consumption). In this risk-neutral case, beliefs are in fact directly inferable from asset values: $q_t(1 | H_t) = \pi(H_t)$ and $q_t(0 | H_t) = 1 - \pi(H_t)$. Testable restrictions on asset values are thus equivalent in this case to restrictions on physical beliefs.

We keep track of the following objects related to the physical belief stream, and we discuss shortly how these objects are restricted under RE. First, total *belief movement* of $\boldsymbol{\pi}$ is defined as the sum of squared changes (or quadratic variation) in beliefs across all periods:

$$m(\boldsymbol{\pi}) \equiv \sum_{t=0}^{T-1} (\pi_{t+1} - \pi_t)^2. \quad (1)$$

Second, *initial uncertainty* of $\boldsymbol{\pi}$ is defined as the variance of the Bernoulli random variable $\mathbb{1}\{\theta = 1\}$ as of $t = 0$:

$$u_0(\boldsymbol{\pi}) \equiv (1 - \pi_0)\pi_0. \quad (2)$$

Given that we focus on resolving belief streams for which $\pi_T \in \{0, 1\}$, final uncertainty is always

zero ($u_T = (1 - \pi_T)\pi_T = 0$). Initial uncertainty u_0 is therefore equal to the total amount of uncertainty reduction for π , $u_0 - u_T$, which is helpful in interpreting some of the following results.

Our main variable of interest will be the difference between movement and initial uncertainty, which — for reasons that will become clear — we call *excess movement*:

$$X(\pi) \equiv m(\pi) - u_0(\pi). \quad (3)$$

Belief movement and initial uncertainty are related under RE according to the following result, which restates a main result in [Augenblick and Rabin \(2021\)](#).

LEMMA 1 ([Augenblick and Rabin, 2021](#)). *Under Assumption 1, for any DGP, expected total belief movement must equal initial uncertainty. Expected excess movement in beliefs must therefore be zero:*


$$\mathbb{E}[X] = 0.$$

[Lemma 1](#) is a straightforward implication of the assumption of martingale beliefs.³ The result formalizes a notion of the “correct” amount of belief volatility under RE, and it motivates referring to X as *excess movement*. Given a set of observed belief streams, one can straightforwardly calculate the sample average of the empirical excess movement statistic and statistically test if it differs from zero. The restriction reflects the intuition that if the person’s beliefs are moving, this movement must on average correspond to learning about the true terminal state (in the sense that uncertainty is resolved from its initial value to 0). Rewriting $\mathbb{E}[X] = 0$ as $\mathbb{E}[\sum_{t=0}^{T-1} (1 - 2\pi_t)(\pi_{t+1} - \pi_t)] = 0$ (see [footnote 3](#)), it is apparent that expected belief movements toward 0.5 (the point of highest uncertainty) lead to a positive $\mathbb{E}[X]$ statistic, and vice versa. So it could be the case that $\mathbb{E}[\pi_{t+1} - \pi_t] = 0$ unconditionally, but a test based on the lemma would still reject the null of RE if, for instance, low values of π_t ($\pi_t < 0.5$) tend to be revised upward ($\pi_{t+1} - \pi_t > 0$), and high values tend to be revised downward.

To clarify our setting and the above lemma, [Figure 1](#) plots belief streams under a two-period DGP. Two fair coins are flipped sequentially at $t = 1$ and 2, generating two possible signals (H or T) at each of these dates. If two heads occur (HH), then state $\theta = 1$ is realized; otherwise, $\theta = 0$ is realized. As the figure shows, the person’s prior π_0 is equal to 0.25 under RE (equal to the probability of two heads). If the first coin flip is tails (shown in the red path), then $\pi_1 = \pi_2 = 0$; if heads then tails are flipped, then $\pi_1 = 0.5$, $\pi_2 = 0$ (light blue); if two heads are flipped, then $\pi_1 = 0.5$, $\pi_2 = 1$ (dark blue). The accompanying table shows belief movement $m = (\pi_1 - \pi_0)^2 + (\pi_2 - \pi_1)^2$ for each possible stream. Weighting the paths by their relative frequencies, expected belief movement is $\mathbb{E}[m] = 0.1875$. This is exactly equal to initial uncertainty $u_0 = (1 - 0.25) \times 0.25 = 0.1875$, illustrating the restriction that $\mathbb{E}[X] = 0$. This applies beyond this two-period example: the lemma implies that *any* DGP for which $\pi_0 = 0.25$ would generate expected movement of 0.1875 under RE.

³To see this, rewrite X as $\sum_{t=0}^{T-1} (2\pi_t - 1)(\pi_t - \pi_{t+1})$. Using the law of iterated expectations on each term in the sum, $\mathbb{E}[(2\pi_t - 1)(\pi_t - \pi_{t+1})] = \mathbb{E}[(2\pi_t - 1)(\pi_t - \mathbb{E}[\pi_{t+1}|\pi_t])]$, which must be zero under the martingale assumption. This result has appeared in other forms in past literature; for one example, see [Barndorff-Nielsen and Shephard \(2001\)](#).

Figure 1: Physical Beliefs and Excess Movement: Two-Period Example



Notes: Signals (s_1, s_2) are generated by sequential fair coin flips. State $\theta = 1$ is realized if two heads occur: $(s_1, s_2) = (H, H)$, labeled *HH*; otherwise $\theta = 0$. Each path in the figure corresponds to a possible belief stream over $\theta = 1$ for a person with RE. In the table, m , u_0 , and X are as defined in (1)–(3), and $\mathbb{E}[m] = 0.25 \times 0.3125 \times 2 + 0.5 \times 0.0625$.

2.2 Risk-Neutral Beliefs: Setup and Identification Challenge

[Lemma 1](#) shows that one can make statements about excess belief movement under RE when beliefs are directly observable or inferrable from asset prices. But the problem of identifying excess movement in asset valuations becomes significantly more complicated when the person is risk averse, as valuations no longer correspond directly to beliefs in this more general case. To see this, assume now that the agent has time-separable utility, with concave period utility function $U(C_t)$, and that she exponentially discounts future consumption with discount factor β . Assume that the state θ determines period- T consumption $C_{T,\theta} \equiv C_T(\theta)$. Valuations for the two Arrow-Debreu securities at time t are now

$$q_t(1|H_t) = \frac{\beta^{T-t} U'(C_{T,1})}{U'(C_t)} \pi_t, \quad q_t(0|H_t) = \frac{\beta^{T-t} U'(C_{T,0})}{U'(C_t)} (1 - \pi_t). \quad (4)$$

Valuations thus no longer directly reveal beliefs. Instead, they are distorted by the relative value of a marginal dollar in state θ in period T versus one in period t . As $q_t(1|H_t)$ and $q_t(0|H_t)$ are similarly distorted by β^{T-t} and $U'(C_t)$, it is useful to focus on their relative valuations. This logic leads to the consideration of *risk-neutral* (RN) beliefs,

$$\pi_t^*(H_t) \equiv \frac{q_t(1|H_t)}{q_t(0|H_t) + q_t(1|H_t)} = \frac{U'(C_{T,1})}{\mathbb{E}_t[U'(C_T)]} \pi_t(H_t) = \frac{\phi \pi_t(H_t)}{1 + (\phi - 1) \pi_t(H_t)}, \quad (5)$$

$$\text{where } \phi \equiv \frac{U'(C_{T,1})}{U'(C_{T,0})}. \quad (6)$$

The definition in (5) follows the usual convention for RN beliefs, and the remaining expressions follow from (4). Like π_t , the RN belief π_t^* corresponds to state $\theta = 1$. The RN belief for $\theta = 0$ can be

similarly defined as $\frac{q_t(0|H_t)}{q_t(0|H_t)+q_t(1|H_t)} = 1 - \pi_t^*$, so the two states' RN beliefs are positive and sum to 1 by construction. We define the RN belief stream $\pi^*(H_t)$, RN belief movement $m^*(\pi^*)$, RN initial uncertainty $u_0^*(\pi^*)$, and RN excess movement $X^*(\pi^*)$ as in [Section 2.1](#), but with RN beliefs π_t^* in the place of physical beliefs π_t .

Risk-neutral beliefs are so named because they can be interpreted as the subjective beliefs for a fictitious risk-neutral agent. In general, they represent a pseudo-belief distribution, reflecting a combination of the person's physical beliefs π_t and risk preferences as indexed by ϕ . This object ϕ will be particularly important in our environment. It represents the (constant) marginal rate of substitution across primitive states, as can be seen in (6). Up to a scaling constant, it also serves as an index of relative risk aversion: conducting a Taylor expansion of $U'(C_{T,0})$ around $C_{T,1}$ gives $U'(C_{T,0}) = U'(C_{T,1}) + U''(C_{T,1})(C_{T,0} - C_{T,1}) + \mathcal{O}((C_{T,0} - C_{T,1})^2)$, and thus, to first order,

$$\gamma(C_{T,1}) \equiv -\frac{C_{T,1}U''(C_{T,1})}{U'(C_{T,1})} = \frac{\phi - 1}{(C_{T,0} - C_{T,1})/C_{T,1}}. \quad (7)$$

Relative risk aversion γ thus depends on the ratio of marginal utilities across states ϕ relative to the percent consumption gap across states. Finally, in asset-pricing terms, ϕ is equivalent to the ratio of *stochastic discount factor (SDF)* realizations $M_{t,T}(\theta)$ across the two states:

$$\phi = \frac{M_{t,T}(1)}{M_{t,T}(0)}, \quad \text{where } M_{t,T}(1) \equiv \frac{q_t(1|H_t)}{\pi_t(H_t)}, \quad M_{t,T}(0) \equiv \frac{q_t(0|H_t)}{1 - \pi_t(H_t)}.$$

This SDF-based representation of ϕ will be discussed in detail in the general setting in [Section 3](#).


To complete the setup for the current consumption-based setting, we assume that consumption in state 1 is weakly less than in state 0, $C_{T,1} \leq C_{T,0}$. This is without loss of generality for now, as the states can be relabeled arbitrarily. With concave utility, this labeling of state 1 as the low-consumption state implies that $U'(C_{T,1}) \geq U'(C_{T,0})$ and thus $\phi \geq 1$. Given $\phi \geq 1$, the RN belief π_t^* in general exceeds the subjective belief π_t : the person is willing to pay relatively more for a bad-state consumption claim given her high marginal utility in that state, upwardly biasing the bad-state RN belief relative to π_t .

We wish to make statements similar to [Lemma 1](#), but applicable to observable RN beliefs rather than physical beliefs. Under a risk-neutral expectation $\mathbb{E}^*[\cdot]$ defined such that $\pi_t^*(H_t) = \mathbb{E}^*[\pi_{t+1}(H_{t+1}) | H_t]$, one could in fact apply [Lemma 1](#) directly, as $\mathbb{E}^*[X^*] = 0$. But the frequency of observed RN belief streams is determined by the physical measure rather than the RN measure; that is, we can only observe an empirical counterpart to $\mathbb{E}[X^*]$ rather than $\mathbb{E}^*[X^*]$. And even under the maintained assumption that physical beliefs π_t satisfy RE, the distortion in π_t^* relative to π_t can cause RN movement to differ from RN initial uncertainty on average, so $\mathbb{E}[X^*] \neq 0$. This is the fundamental identification challenge.

For an illustration of this issue, [Figure 2](#) returns to the coin-flip example from [Figure 1](#). In addition to physical beliefs, this figure now shows RN belief streams with $\phi = 3$.⁴ From (5), the

⁴Under risk neutrality, $\phi = 1$. This was implicitly assumed to be the case for [Figure 1](#).

Figure 2: RN Beliefs and Excess Movement: Two-Period Example



Notes: See Figure 1 for description of signals and physical beliefs. Observed RN beliefs are calculated using (5), with the assumption that $\phi = 3$. In the table, m^* , u_0^* , and X^* are calculated as in (1)–(3), with RN beliefs π_t^* in place of π_t .

physical prior $\pi_0 = 0.25$ corresponds to an RN prior of $\pi_0^* = 0.5$. Intuitively, the person perceives state $\theta = 0$ as three times as likely as $\theta = 1$ (*HH*), but values a marginal dollar in $\theta = 1$ three times more than in $\theta = 0$, so $q_0(0) = q_0(1)$ and $\pi_0^* = 0.5$. The same calculations are then applied to obtain π_1^* and π_2^* for each stream shown in the figure. The accompanying table shows RN movement m^* for each stream. Comparing this to the table in Figure 1, it is apparent that $m^* > m$ for streams in which the RN belief has large downward revisions (*HT*, *TT*, and *TH*): π_t^* is biased upward relative to π_t , so these large downward revisions generate more RN movement than physical movement. These streams cause RN excess movement to be positive on average: $\mathbb{E}[X^*] = 0.0625 > 0$.

As this example illustrates, even with rational physical beliefs, one can observe what appears to be excess movement in RN beliefs implied by valuations. So if we naively test for RE using Lemma 1 on observed RN (rather than actual) beliefs, we may spuriously conclude that beliefs are excessively volatile.

2.3 Risk-Neutral Beliefs: Results

The fact that there can be excess movement in RN beliefs even under RE would seem to pose an intractable challenge. But this turns out not to be the end of the story. RN beliefs are not arbitrarily distorted relative to physical beliefs: as in (5), they are linked through the single unobserved parameter ϕ . And regardless of the value of ϕ , RN beliefs must lie between 0 and 1 by definition. These insights will allow us to bound $\mathbb{E}[X^*]$ over all possible DGPs for a given value of ϕ , and over all possible values of ϕ . Because our main results are most straightforwardly stated in the current section's two-state setting, we state them here before discussing how they generalize in Section 3.

Main Results

Before turning to the results, it will be useful to define two additional objects. First, we invert (5) to solve for π_t as a function of π_t^* and ϕ , the solution to which we denote by $\pi_t(\pi_t^*, \phi)$:

$$\pi_t(\pi_t^*, \phi) = \frac{\pi_t^*}{\phi + (1 - \phi)\pi_t^*}. \quad (8)$$

Second, it will be helpful to define the difference in conditional expected X^* across states as

$$\Delta \equiv \mathbb{E}[X^* | \theta = 0] - \mathbb{E}[X^* | \theta = 1]. \quad (9)$$

Given these definitions, we can now provide a number of expressions and bounds for $\mathbb{E}[X^*]$. We assume throughout that [Assumption 1](#) holds.

PROPOSITION 1. *For any DGP,*

$$\begin{aligned} \mathbb{E}[X^*] &= (\pi_0^* - \pi_0)\Delta \\ &= \left(\pi_0^* - \frac{\pi_0^*}{\phi + (1 - \phi)\pi_0^*} \right) (\mathbb{E}[X^* | \theta = 0] - \mathbb{E}[X^* | \theta = 1]). \end{aligned}$$

PROPOSITION 2. *For any DGP and any value for Δ ,*

$$\begin{aligned} \mathbb{E}[X^*] &\leq (\pi_0^* - \pi_0)\pi_0^* \\ &\leq \left(1 - \frac{1}{\phi + (1 - \phi)\pi_0^*} \right) \pi_0^{*2}. \end{aligned} \quad (10)$$

Proofs for this section's main results are provided in [Appendix A](#). [Proposition 1](#) starts from the fact that $\mathbb{E}^*[X^*] = 0$. We then connect $\mathbb{E}^*[X^*]$ to $\mathbb{E}[X^*]$. The key step is to show that *conditional* expectations of X^* under both measures are equal, $\mathbb{E}^*[X^* | \theta] = \mathbb{E}[X^* | \theta]$ for $\theta = 0, 1$, which leads to the stated result. Note that if $\phi = 1$, then $\pi_0^* = \pi_0$ and therefore $\mathbb{E}[X^*] = 0$, as in [Lemma 1](#). As ϕ rises, π_0 drops further below π_0^* , and $\mathbb{E}[X^*]$ departs from 0: the greater is risk aversion, the more one can observe excess RN movement differ from zero on average under RE. The sign and magnitude of this deviation depend on Δ , as explored in more detail later in this section.

Across all DGPs, Δ is bounded above by π_0^* , from which [Proposition 2](#) follows. The version of the bound in (10) is one of our main results. It gives a bound for $\mathbb{E}[X^*]$ as a function of π_0^* and ϕ , regardless of Δ . Under risk neutrality ($\phi = 1$), this upper bound again becomes zero. But the bound is otherwise positive, and the *admissible excess movement* in RN beliefs given by the right side of the inequality increases monotonically in ϕ . This result thus formalizes a more general notion of the admissible amount of belief volatility under rationality, this time as an increasing function of risk aversion across the two states.

Intuitively, movement in RN beliefs must still correspond on average to the agent learning something about the true terminal state, but the bias in RN relative to subjective beliefs induced

by risk aversion allows for positive excess movement in those observed beliefs under RE. This is reflected in the first term in the bound, $(\pi_0^* - \pi_0)$, as this difference increases in ϕ . For the second term (π_0^*) , higher RN priors yield more “room” for downward RN belief movement. This increases admissible excess movement along the lines of the example considered in [Figure 2](#). The bound in [Proposition 2](#) is conservative, as it holds under a “worst-case” DGP with an extreme value for Δ .

Because of the monotonicity of the bound in ϕ , the inequality can equivalently be read as providing a *lower* bound for the unobserved structural parameter ϕ as a function of observables. The observed excess movement in RN belief streams thus provides information on the minimal value of risk aversion necessary for the data to be consistent with RE. But more still can be said: taking $\phi \rightarrow \infty$ in (10) generates the following maximally conservative bound for $\mathbb{E}[X^*]$, which applies for *any* ϕ .

COROLLARY 1. *For any DGP and any values Δ and ϕ ,*

$$\mathbb{E}[X^*] \leq \pi_0^{*2}.$$

[Corollary 1](#) exploits the fact that RN beliefs are bounded between 0 and 1 by construction: there is only so far that RN beliefs can be distorted relative to subjective beliefs, so the bound is well-defined even for arbitrarily high risk aversion.⁵ It states that asset-price movements for which $\mathbb{E}[X^*] > \pi_0^{*2}$ simply cannot be rationalized under RE given constant ϕ . Despite its maximal conservatism, this bound still imposes a meaningful limit on admissible RN excess movement, especially for low π_0^* . For example, with $\pi_0^* = 0.2$, RN excess movement is at most 0.04, regardless of ϕ or the DGP.


Taken together, [Proposition 2](#) and [Corollary 1](#) characterize the maximal admissible excess movement in RN beliefs as a function of ϕ for any RN prior. [Figure 3](#) provides a graphical illustration of these bounds. Starting from the bottom of the chart, the thick purple line corresponds to the bound for $\phi = 1$: in this case, $\mathbb{E}[X^*] = \mathbb{E}[X] = 0$ regardless of the prior or DGP, from [Lemma 1](#). The thin dashed gray lines correspond to arbitrarily selected DGPs in the case of $\phi = 3$. While there can be positive RN excess movement, this is not necessarily the case for all possible DGPs. Taking the envelope over all of these processes for $\phi = 3$ yields the bound from [Proposition 2](#), which is shown in the thick blue line. It is asymmetric around 0.5 as well as non-monotonic in π_0^* .⁶ Finally, the thick red line shows the bound for the limiting case $\phi \rightarrow \infty$, which is equal to the squared RN prior from [Corollary 1](#).

When should we expect to see *negative* RN excess movement — as observed in the lowest dashed gray line in [Figure 3](#) — even with risk aversion? If one is willing to make an assumption on the sign of Δ (discussed shortly), the following stronger bound applies as a corollary of [Proposition 1](#).

⁵In the limit as $\phi \rightarrow \infty$, $\pi_0 \rightarrow 0$ for any π_0^* , so $\pi_0^* - \pi_0 \rightarrow \pi_0^*$.

⁶The second term in the bound (π_0^*) of course increases monotonically, generating asymmetry in the bound around $\pi_0^* = 0.5$ for $\phi > 1$. But the first term in the bound $(\pi_0^* - \pi_0)$ does not increase monotonically for $1 < \phi < \infty$, exerting a countervailing force that generates the observed non-monotonicity.

Figure 3: RN Excess Belief Movement vs. Prior by ϕ Under RE



Note: Theoretical bounds are obtained from the formulas in [Proposition 2](#) and [Corollary 1](#).

COROLLARY 2. If $\mathbb{E}[X^*|\theta = 0] \leq \mathbb{E}[X^*|\theta = 1]$, for any DGP and any value for ϕ ,

$$\mathbb{E}[X^*] \leq 0.$$

When $\Delta < 0$, $\mathbb{E}[X^*]$ is decreasing in ϕ and therefore $\mathbb{E}[X^*] < 0$ for any $\phi > 1$. Consequently, as formalized in [Corollary 2](#), the highest excess movement is $\mathbb{E}[X^*] = 0$. While most of our focus is on the conservative positive upper bounds for $\mathbb{E}[X^*]$, this corollary shows that asset-pricing settings with risk aversion do not necessarily entail positive excess movement in observed beliefs. We now further explore the statistical features of the DGP that are informative about the degree of RN excess movement to be expected under RE.

How the DGP Determines $\mathbb{E}[X^*]$

[Proposition 1](#) says that the deviation of $\mathbb{E}[X^*]$ from 0 depends on the product of $\pi_0^* - \pi_0$ and $\Delta \equiv \mathbb{E}[X^*|\theta = 0] - \mathbb{E}[X^*|\theta = 1]$. The difference $\pi_0^* - \pi_0$ is always positive and increases in ϕ . But how are the sign and magnitude of Δ related to the DGP? To answer this question, we provide two theoretical results and briefly summarize a set of simulations discussed in detail in [Online Appendix C.1](#). These results will then be useful in interpreting the empirical results to come.

First, given the arbitrary labeling of the two states, there is no reason to expect under RE that Δ should take a particular sign:

PROPOSITION 3. Fixing ϕ , for every RN prior and DGP that leads to a given Δ , there exists a different RN prior and DGP that leads to $-\Delta$.

For any RN prior π_0^* and DGP with some Δ , the RN prior $1 - \pi_0^*$ with the “reversed” DGP will necessarily lead to $-\Delta$. Consequently, there is no reason to assume that $\mathbb{E}[X^*]$ is more likely to be positive than negative given $\phi > 1$.⁷

Next, we summarize numerical simulations for a large set of DGPs as described in Online Appendix C.1. The results there suggest some intuitive comparative statics. First, when the DGP changes such that downward movements in π_t^* become larger, Δ rises and $\mathbb{E}[X^*]$ rises. Second, when upward movements become larger, Δ falls. Third, as π_0^* rises, Δ rises: higher π_0^* allows relatively more downward movement. Thus when the DGP is symmetric (with equal movements up or down), $\Delta < 0$ if $\pi_0^* < .5$, $\Delta = 0$ if $\pi_0^* = .5$, and $\Delta > 0$ if $\pi_0^* > .5$. This suggests, for example, that low π_0^* should lead to a negative Δ unless the DGP is asymmetric. As we show later, the empirical DGPs in our setting appear largely symmetric, so that we estimate $\Delta < 0$ for $\pi_0^* < .5$.

These results also suggest that extreme values of Δ only occur in highly asymmetric DGPs where movements in one direction are large and movements in the other direction are tiny. In fact, our upper bound in (10) is attainable asymptotically given the most asymmetric DGP possible:

PROPOSITION 4. *There exists a sequence of DGPs, indexed by T , for which $\mathbb{E}[X^*]$ approaches the bound in Proposition 2 as $T \rightarrow \infty$. For each DGP in this sequence, downward movements ($\pi_{t+1}^* < \pi_t^*$) are resolving ($\pi_{t+1}^* = 0$) and thus as large as possible, while upward movements are small ($\pi_{t+1}^* - \pi_t^* \rightarrow 0$ as $T \rightarrow \infty$). Meanwhile, the bound holds with strict inequality for any $T < \infty$ as long as $\phi > 1$ and $\pi_0^* \in (0, 1)$.*

One implication of this result is that the bound in Proposition 2 is approximately tight, as one can construct a DGP for which $\mathbb{E}[X^*]$ is close to the bound for large T . Perhaps more important, though, is that it points to the bound’s conservatism: it holds under a somewhat perverse DGP that can be thought of as a “rare bonanzas” process, where with small probability the person receives news that the bad state ($\theta = 1$) will not be realized (so $\pi_{t+1}^* = 0$), and otherwise there is mostly uninformative bad news that increases π_{t+1}^* slightly. More reasonable DGPs, or $T \ll \infty$, will give lower $\mathbb{E}[X^*]$. That said, the conservative bound has the advantage of being very simple and not requiring any estimation of Δ . And as we show below, empirical excess movement is in fact so high that even these conservative bounds are often violated for reasonable values of ϕ .

3. Generalized Theoretical Results for Equilibrium Asset Prices

While the binary-state setting considered in the previous section is useful for exposition, it is also artificial: one cannot obtain a single person’s valuation of Arrow-Debreu claims in observational data alone; there are more than two possible states; and the realized state determines more than just consumption. We thus now consider a general many-state framework for equilibrium asset prices and show how our results extend to this empirically relevant case.

⁷For example, in Figure 2, $\Delta = 0.25$ and $\mathbb{E}[X^*] = 0.0625$. The “reversed” DGP has $\pi_0^* = 0.5$, $\pi_1^* = 0.25$ (with probability p) or 1 ($1 - p$), and $\pi_2^* = 0$ (prob. q if $\pi_1^* = 0.25$) or 1 . With $\phi = 3$, for the physical π_t to be a martingale, one can solve for $p = 5/6$, $q = 9/10$. Further calculation thus gives $\Delta = -0.25$ and $\mathbb{E}[X^*] = -0.0625$, as in the proposition.

3.1 Setup and Notation

Preliminaries: Probability Space, Prices, and Risk-Neutral Probabilities

Time is again indexed by $t \in \{0, 1, 2, \dots\}$, and we consider a discrete probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with filtration $\{H_t\}$.⁸ A realization of the elementary state is denoted by $\omega \in \Omega$. To make our results empirically implementable, we will be concerned with the ex-dividend value of the market index, $V_t^m: \Omega \rightarrow \mathbb{R}_+$, on some option expiration date T . (We will later extend the notation to allow for multiple option expiration dates.) A European call option on the index with strike price K has date- T payoff $(V_T^m - K)^+ = \max(V_T^m - K, 0)$, and its time- t price is $q_{t,K}^m$. These option prices are observable for some set of strike prices $\mathcal{K} \subseteq \mathbb{R}_+$ beginning at date 0. Assuming the absence of arbitrage, there exists a strictly positive *stochastic discount factor* (SDF) $M_{t,T}$ such that option prices satisfy $q_{t,K}^m = \mathbb{E}_t[M_{t,T}(V_T^m - K)^+]$, where $\mathbb{E}_t[\cdot] \equiv \mathbb{E}^\mathbb{P}[\cdot | H_t]$. The SDF can equivalently be written as $M_{t,T} = M_T/M_t$ for strictly positive $\{M_t\}$, with $M_0 = 1$.

Option prices will be of interest for inferring a distribution over the change in value of the market index from 0 to T (rather than consumption, for which options are not directly traded). We say that *return state* $\theta \in \Theta \subset \mathbb{R}_+$ is realized for the market as of date T if $R_T^m \equiv V_T^m/V_0^m = \theta$. The measure $\mathbb{P}: \mathcal{F} \rightarrow [0, 1]$ governs the *objective* or *physical* probabilities of these return states. In this general case, the *risk-neutral (RN) measure* is defined by the change of measure

$$\frac{d\mathbb{P}^*}{d\mathbb{P}} \Big|_{H_t} = \frac{M_{t,T}}{\mathbb{E}_t[M_{t,T}]} = \frac{M_T}{\mathbb{E}_t[M_T]}, \quad (11)$$

and expectations under \mathbb{P}^* are denoted by $\mathbb{E}^*[\cdot]$.⁹ Using this definition of \mathbb{P}^* , the RN probability of return state θ is

$$\mathbb{P}_t^*(R_T^m = \theta) = \frac{\mathbb{E}_t[M_T | R_T^m = \theta]}{\mathbb{E}_t[M_T]} \mathbb{P}_t(R_T^m = \theta). \quad (12)$$

The RN pricing equation $q_{t,K}^m = \mathbb{E}_t^*[(V_T^m - K)^+] / R_{t,T}^f$ can be used to show that the date- t option prices $\{q_{t,K}^m\}_{K \in \mathcal{K}}$ reveal the set of RN probabilities $\{\mathbb{P}_t^*(R_T^m = \theta)\}_{\theta \in \Theta}$. Assume that the set of return states Θ is ordered such that $\theta_1 < \theta_2 < \dots < \theta_J$, and assume for notational simplicity that the set of traded option strikes \mathcal{K} coincides with the set of possible date- T index values (i.e., $\mathcal{K} = \{K_j\}_{j=1}^J$, with $K_j = V_0^m \theta_j$). We can then back out RN probabilities from option prices as follows:

$$\mathbb{P}_t^*(R_T^m = \theta_j) = R_{t,T}^f \left[\frac{q_{t,K_{j+1}}^m - q_{t,K_j}^m}{K_{j+1} - K_j} - \frac{q_{t,K_j}^m - q_{t,K_{j-1}}^m}{K_j - K_{j-1}} \right]. \quad (13)$$

⁸We could consider continuous states with additional technicalities, but do not do so as empirical implementation requires discretization. We note as well that objects analogous to those in Section 2 are given the same denotation here.

⁹As $\text{Price}_t = \mathbb{E}_t[M_{t,T} \text{Payoff}_T]$ for any asset, we have $R_{t,T}^f = \mathbb{E}_t[M_{t,T}]^{-1}$, where $R_{t,T}^f$ is the gross risk-free rate from t to T , and $\text{Price}_t = \mathbb{E}_t^*[\text{Payoff}_T] / R_{t,T}^f$. Thus \mathbb{P}^* incorporates the risk adjustment needed to discount T -payoffs at $R_{t,T}^f$. This \mathbb{P}^* is sometimes referred to as the T -forward measure (e.g., Geman, El Karoui, and Rochet, 1995).

See Online Appendix B.2 for a derivation of this result, which follows from a discrete-state application of the classic result of [Breedon and Litzenberger \(1978\)](#). As in [Section 2.2](#), the above RN probabilities are a more convenient object of analysis than raw option prices.

Beliefs

Aside from assuming no arbitrage, we have not yet taken a stance on the market structure or subjective beliefs underlying prices and RN probabilities. We could in principle pursue a strict mapping from [Section 2](#) to the current case, by assuming a setting in which all individual traders have common beliefs satisfying RE. The assumptions required to generate such an equilibrium are well studied in the literature on information and asset prices following [Radner \(1979\)](#) and [Milgrom and Stokey \(1982\)](#).¹⁰ But rather than maintaining this microeconomic focus on testing the rationality of individual beliefs, we prefer an interpretation in which the “agent” in question is the market as a whole (or alternatively, the marginal trader); this interpretation requires no auxiliary assumptions, and our resultant tests are informative about the efficiency of market valuations.

We therefore assume prices correspond to valuations for an agent (“the market”) who, at the beginning of each period, observes a signal vector $s_t \in \mathcal{S}$ drawn from the distribution $DGP(s_t | \theta, H_{t-1}) = \mathbb{P}_{t-1}(s_t | \theta)$, where θ is the return state realized at T and $H_t = \sigma(s_\tau, 0 \leq \tau \leq t)$ is the Borel σ -algebra representing the history of signal realizations. The agent’s subjective belief distribution over return states is denoted by $\Pi_{t,T} = \{\pi_t(R_T^m = \theta)\}_{\theta \in \Theta}$, where $\pi_t(R_T^m = \theta) \geq 0 \forall \theta \in \Theta$ and $\sum_{\theta \in \Theta} \pi_t(R_T^m = \theta) = 1$. More generally, for any random variable $Y(\omega)$, the agent attaches subjective probability $\pi_t(Y = y)$ to the outcome $Y = y$. We generalize [Assumption 1](#) as:

ASSUMPTION 2 (RE). For any random variable Y , beliefs satisfy $\pi_t(Y = y) = \mathbb{P}_t(Y = y)$ with probability 1 for all t .

This assumption again implies that beliefs satisfy $\pi_t(R_T^m = \theta) = \mathbb{E}[\pi_{t+1}(R_T^m = \theta) | \pi_t(R_T^m = \theta)]$ for all $\theta \in \Theta$. As in [Section 2](#), this martingale condition for beliefs over returns is all that is required for our main results to carry through. The full-RE generalization stated in [Assumption 2](#) is useful for streamlining some of the remaining discussion, as it further implies that *all* conditional expectations — including over the SDF — are martingales with respect to H_t .

Given [Assumption 2](#), we can define the RN belief distribution without explicitly restricting the agent’s utility or constraint set by applying the same change of measure as in [\(11\)](#). This yields the RN belief-distribution $\Pi_{t,T}^* = \{\pi_t^*(R_T^m = \theta)\}_{\theta \in \Theta}$ such that $\pi_t^*(R_T^m = \theta) = \frac{\mathbb{E}_t[M_T | R_T^m = \theta]}{\mathbb{E}_t[M_T]} \pi_t(R_T^m = \theta)$ as in [\(12\)](#), and thus [\(13\)](#) tells us that option prices reveal the agent’s RN beliefs as given here.

¹⁰Complete markets and a common-prior assumption, for example, are sufficient: prices in general reveal information (including private signals) in a rational expectations equilibrium, giving common posteriors. Results under alternative conditions have also been studied extensively (to take one example, see [Blume, Coury, and Easley, 2006](#)).

Localization: Conditional Beliefs, SDF Ratio, and Excess Movement

To align with the analysis in [Section 2](#), we consider the behavior of *conditional* RN beliefs over adjacent pairs of return states. That is, rather than directly considering the full distribution $\Pi_{t,T}^*$, we instead consider restrictions on the behavior of the individual entries in $\{\tilde{\pi}_{t,j}^*\}_{j=1}^{J-1}$, defined by

$$\tilde{\pi}_{t,j}^* \equiv \pi_t^*(R_T^m = \theta_j \mid R_T^m \in \{\theta_j, \theta_{j+1}\}) = \frac{\pi_t^*(R_T^m = \theta_j)}{\pi_t^*(R_T^m = \theta_j) + \pi_t^*(R_T^m = \theta_{j+1})},$$

for $\pi_t^*(R_T^m = \theta_j) + \pi_t^*(R_T^m = \theta_{j+1}) > 0$. In words, $\tilde{\pi}_{t,j}^*$ is the RN belief that return state θ_j will be realized, conditional on either θ_j or θ_{j+1} . This binary localization will be useful for two reasons: (i) it will allow us to apply results from the two-state setting of [Section 2](#), and (ii) the main identifying assumption used to derive our tests is less restrictive than it would be without such a transformation (as discussed in [Section 4.1](#) below). Conditional physical beliefs $\tilde{\pi}_{t,j}$ are defined analogously, and the expectation under the conditional physical measure $\tilde{\mathbb{P}}_t$ is $\tilde{\mathbb{E}}_t[\cdot] \equiv \mathbb{E}_t[\cdot \mid R_T^m \in \{\theta_j, \theta_{j+1}\}]$.

In this context, the analogue to ϕ as defined in [Section 2](#) is

$$\phi_{t,j} \equiv \frac{\mathbb{E}_t[M_T \mid R_T^m = \theta_j]}{\mathbb{E}_t[M_T \mid R_T^m = \theta_{j+1}]}, \quad (14)$$

which encodes the slope of the SDF across the adjacent return states θ_j, θ_{j+1} . In a representative-agent economy with consumption C_t , time-separable utility, and discount factor β , the SDF is $M_{t,T} = \beta^{T-t} U'(C_T)/U'(C_t)$, and $\phi_{t,j} = \mathbb{E}_t[U'(C_T) \mid R_T^m = \theta_j]/\mathbb{E}_t[U'(C_T) \mid R_T^m = \theta_{j+1}]$, akin to the marginal rate of substitution in (6). But (14) is general and does not require a representative-agent structure (though we make periodic reference to such an economy for interpretation). Using this definition of $\phi_{t,j}$, RN beliefs satisfy $\frac{\tilde{\pi}_{t,j}^*}{1-\tilde{\pi}_{t,j}^*} = \phi_{t,j} \frac{\tilde{\pi}_{t,j}}{1-\tilde{\pi}_{t,j}}$ and thus $\tilde{\pi}_{t,j}^* = \frac{\phi_{t,j} \tilde{\pi}_{t,j}}{1+(\phi_{t,j}-1)\tilde{\pi}_{t,j}}$, as in (5).

Given a resolving RN belief stream $\pi_j^* = [\tilde{\pi}_{0,j}^*, \dots, \tilde{\pi}_{T,j}^*]$, RN belief movement m_j^* , RN initial uncertainty $u_{0,j}^*$, and RN excess movement X_j^* are as defined in (1)–(3), with $\tilde{\pi}_{t,j}^*$ in place of π_t . We often suppress the dependence on j (writing, e.g., X^*) when considering an arbitrary state pair.

3.2 Identifying Assumptions on ϕ

The analysis in [Section 2](#) considered Arrow-Debreu claims on primitive states (in that case, consumption states). This section's analysis instead considers claims on index-return states, with an eye toward empirical implementation. We must thus confront the joint hypothesis problem. The additional identifying assumptions to be tested jointly with [Assumption 2](#) take the form of two restrictions on $\phi_{t,j}$. We introduce these assumptions here, and we discuss them in much greater detail in the following sections after presenting our main results below.

We first impose an ordering assumption on the states. Return state θ_j here corresponds to state 1 in [Section 2](#) (vs. state 0 for θ_{j+1}). We thus maintain the convention of labeling θ_j as the “bad” state, so that $\phi_{t,j} \geq 1$. While this is an innocuous labeling convention in theory, empirical implementation requires taking a stand on how to distinguish θ_j from θ_{j+1} . We make the intuitive assumption that

the bad state θ_j corresponds to the lower return:

ASSUMPTION 3 (*Positive Risk Aversion in Index Return*). $\phi_{t,j} \geq 1$ with probability 1 for all t, j , where the set of return states Θ is ordered such that $\theta_1 < \theta_2 < \dots < \theta_J$.

We discuss the empirical plausibility of this assumption, and how our results can be modified if it is violated (as might be a concern given the so-called pricing kernel puzzle), in [Section 5](#).

Second, and more substantively, we must impose an assumption on the evolution of the unobserved parameter $\phi_{t,j}$. In [Section 2](#)'s setting, ϕ is naturally constant over t : the terminal states index consumption and thus marginal utility, so the marginal rate of substitution across these primitive states is fixed. For our main analysis, we impose an analogous assumption on $\phi_{t,j}$:

ASSUMPTION 4 (*Constant SDF Ratio, or Conditional Transition Independence*). We say the SDF satisfies *conditional transition independence (CTI)* for return-state pair (θ_j, θ_{j+1}) if $\phi_{t,j}$ is constant with probability 1 for all t . We assume CTI is satisfied for all *interior* state pairs, $j = 2, 3, \dots, J - 2$, and we write $\phi_{t,j} = \phi_j$ for these states.

This assumption imposes that the *relative* expected “severity” of the adjacent return states is constant over a contract, so that the expected SDF (or marginal utility) realization in the low return state θ_j is a constant multiple of that of θ_{j+1} . Unlike in [Section 2](#), this does *not* require that all changes in the underlying RN belief distribution arise from changes in subjective beliefs $\pi_t(R_T^m = \theta_j)$: there may be simultaneous time variation in the values in the numerator and denominator in [\(14\)](#) as long as their ratio is fixed, and the value $\mathbb{E}_t[M_T]$ need not be constant. Note as well that we exclude the extreme state pairs (θ_1, θ_2) and (θ_{J-1}, θ_J) from the constant- ϕ requirement: thinking of θ_1 and θ_J as tail return states, we are allowing for time-varying disaster (or positive jump) severity. The assumption corresponds to a notion of transition or path independence because it implies that the expected relative SDF realizations depend only on the return state pair and not on the path of variables realized between t and T .¹¹

As discussed below in [Section 4](#), this assumption is weaker than the joint assumptions imposed in past volatility tests, and we view it as a reasonable starting point for our analysis. But it is unlikely to hold perfectly in reality, so that section also provides results characterizing RN excess movement when the assumption is violated.

3.3 Main Results in the General Setting

Having completed the formal setup and description of our main assumptions, we turn now to our main results in this more general asset-pricing setting. The bulk of the work in this case is, it turns out, in the setup and notation, as all our main results apply with appropriate relabeling:

PROPOSITION 5. *Under no arbitrage and [Assumptions 2–4](#), for $j = 2, 3, \dots, J - 2$, [Lemma 1](#), [Propositions 1–4](#), and [Corollaries 1–2](#) continue to hold, with $\tilde{\pi}_{t,j}^*$ replacing π_t^* , $\tilde{\pi}_{t,j}$ replacing π_t , X_j^* replacing X^* , ϕ_j replacing ϕ , $\tilde{\mathbb{E}}_0[\cdot]$ replacing $\mathbb{E}[\cdot]$, and with $\Delta_j \equiv \tilde{\mathbb{E}}_0[X_j^* | R_T^m = \theta_{j+1}] - \tilde{\mathbb{E}}_0[X_j^* | R_T^m = \theta_j]$ replacing Δ .*

¹¹This is formalized in [Lemma B.1](#) in the Online Appendix.

The main theoretical complication in applying the results in [Section 2](#) to this setting is in proving that $\tilde{\mathbb{E}}^*[X_j^* | R_T^m = \theta_j] = \tilde{\mathbb{E}}[X_j^* | R_T^m = \theta_j]$.¹² The economic intuition for these results is largely identical to the intuition discussed in [Section 2](#). And while the results above are convenient to express in terms of the SDF slope ϕ_j given that this allows for closed-form solutions that can be applied across a range of structural models regardless of the origin of the SDF, the results also admit an interpretation in terms of the approximate required risk-aversion value for a fictitious representative agent with utility over the terminal value of the market index. Analogous to (7), we have the following result. To avoid repetition, for the remaining results we will continue to assume no arbitrage and that [Assumptions 2–4](#) hold for $j = 2, 3, \dots, J - 2$, unless stated otherwise.

PROPOSITION 6. *Assume additionally that there is a representative agent with (indirect) utility over time- T wealth, with wealth equal to the market index value, and denote $V_j^m \equiv V_0^m \theta_j$. Then local relative risk aversion $\gamma_j \equiv -V_j^m U''(V_j^m) / U'(V_j^m)$ is given to a first order around return state θ_j by*

$$\gamma_j = \frac{\phi_j - 1}{(V_{j+1}^m - V_j^m) / V_j^m}.$$

As in [Section 2](#), γ_j is proportional to $\phi_j - 1$, as this gives the percent decrease in marginal utility in moving from low-return state θ_j to high-return state θ_{j+1} . To calculate relative risk aversion, this change in marginal utility must be normalized by the percent wealth increase $(V_{j+1}^m - V_j^m) / V_j^m$ in moving from θ_j to θ_{j+1} , which is also equal to the percent return deviation $(\theta_{j+1} - \theta_j) / \theta_j$ between the two states. If, for example, $\theta_j = 1$, $\theta_{j+1} = 1.05$, then a value $\phi_j = 1.5$ implies $\gamma_j = 10$.

4. Interpreting and Relaxing the CTI Assumption

The analysis thus far has proceeded under the joint null implied by [Assumptions 2–4](#). We now discuss these assumptions in more detail. What do they entail specifically? If the joint null is rejected in the data, how informative is such a rejection, and how should it be interpreted? We begin by considering CTI, which we consider the most important assumption imposed alongside RE. We first discuss settings in which it does (and doesn't) hold, and how it relates to assumptions imposed in past work. We then provide a set of robustness results when the assumption is relaxed, before turning to the other assumptions in the following section.

4.1 CTI Allows for an Informative Joint Null

We start by illustrating that the constant- ϕ assumption is satisfied across a variety of theoretical settings. We make this point not to assert that CTI must hold in reality, but instead to demonstrate that a violation of our bounds is informative in ruling out a range of well-studied frameworks:¹³

1. Permanent shocks to the SDF that change M_T (and its time- t expectation) in all states are

¹²This is shown in Lemma B.2 for a measure that is observationally equivalent to the RN measure, as is sufficient here.

¹³Derivations for statements 3–6 are in Online Appendix B.3; statements 1–2 are immediate.

admissible under CTI: such shocks affect the numerator and denominator of (14) one-for-one. As Alvarez and Jermann (2005) show, permanent shocks appear to be much more important empirically for SDF variation than transitory shocks. Meanwhile, Borovička, Hansen, and Scheinkman (2016) show that permanent SDF shocks are ruled out in the framework proposed by Ross (2015) for recovering physical probabilities from state prices.

2. If a representative agent’s utility depends only on the maturity value of the market index, then CTI holds. Similarly, if there is *some* agent whose indirect utility can be written as a function only of the terminal index value — for example, an unconstrained investor with horizon T (i.e., $Utility = f(Wealth_T)$) who is fully invested in the market — then CTI holds. This encompasses a setting in which an unconstrained log investor holds the market, which is the leading case considered by Martin (2017), who does not reject it as a benchmark in measuring the equity premium, and Gandhi, Gormsen, and Lazarus (2022), who use this log-utility assumption to derive volatility tests for the term structure of the equity premium.
3. In the variable rare disasters model of Gabaix (2012), CTI holds for all market return-state pairs (θ_j, θ_{j+1}) for which there is negligible probability of having realized a disaster conditional on reaching θ_j .¹⁴ This illustrates the usefulness of the localization provided by considering conditional beliefs: if we are concerned that time-varying disaster risk may affect $\phi_{t,j}$ for states in the tail of the return distribution, we can ignore these states and confine attention to the center of the return distribution, where $\phi_{t,j}$ can be expected to be approximately constant.
4. If a representative agent has Epstein–Zin (1989) recursive utility and holds the market, then CTI holds if any of the following apply: (i) relative risk aversion is $\gamma = 1$, and the elasticity of intertemporal substitution (EIS) ψ and DGP are arbitrary; (ii) γ is arbitrary, the EIS is $\psi = 1$, and log consumption growth follows an AR(1) process; or (iii) γ and ψ are arbitrary, and consumption growth is i.i.d. (a less interesting case). As Martin (2017, Example 4b) notes, case (i) is “considered (and not rejected) by Epstein and Zin (1991) and Hansen and Jagannathan (1991).” Case (ii) provides an approximation to the Bansal and Yaron (2004) long-run risks model, with $\psi = 1$ and with common shocks to current and expected future consumption (as, for example, in Bansal and Yaron, 2000). While assuming $\psi = 1$ often allows for a reasonable approximation to the full model with $\psi \neq 1$ (Dew-Becker and Giglio, 2016), the approximation will degrade given highly persistent growth or volatility processes, and CTI will not hold in this case. Finally, case (iii) is considered by Martin (2013).

As the above examples illustrate, the assumption of CTI is significantly weaker than the assumption of constant discount rates.¹⁵ This should again not be taken as evidence *in favor* of CTI in fact holding; instead, it allows for an informative joint test whose null includes a range of models that have been advanced as rationalizations of the excess-volatility puzzle.

¹⁴Formally, for any δ , there exists a $\underline{\theta}$ such that $\forall \theta_j \geq \underline{\theta}, \mathbb{P}_0(\sum_{t=1}^T \mathbb{1}\{\text{disaster}_t\} > 0 \mid R_T^m \geq \underline{\theta}) < \delta$, so the conditional probability of having realized a disaster before T is negligible. For all $\theta_j \geq \underline{\theta}$, CTI holds for the state pair (θ_j, θ_{j+1}) up to a negligible error, as $\phi_{t,j} = \phi_j + \eta_t$ for ϕ_j constant and $\eta_t = o_p(1)$ as $\delta \rightarrow 0$. Again see Online Appendix B.3 for detail.

¹⁵Online Appendix C.2 discusses the relationship between RN beliefs and discount rates in greater detail.

That said, there are also well-known models under which CTI does *not* hold. Such models are also useful for illustrating the content of the assumption:

5. In the habit-formation model of [Campbell and Cochrane \(1999\)](#), CTI fails to hold: the path of consumption matters in a manner not fully accounted for by conditioning on the return state.
6. In the [Basak \(2000\)](#) model with heterogeneous beliefs and extraneous (non-fundamental) risk, CTI fails to hold as long as extraneous risk is priced.

We view the fact that CTI rules out habit-like models to be a downside of the current framework, which we address below by means of robustness results and simulations. The fact that we rule out models of belief volatility induced by dynamic belief heterogeneity, though, is by design: if ϕ is time-varying due to changing weights being assigned to different individual agents — and therefore the as-if representative agent has excessively volatile beliefs — then this represents a meaningful alternative to our null.

4.2 Relaxing CTI Theoretically

While CTI allows for an informative null in our main analysis, it is also a knife-edge restriction that is unlikely to hold exactly. We thus now investigate how such a violation affects RN excess movement. One might worry that even small fluctuations in ϕ could generate dramatic violations of our bounds. For example, suppose that π_t is constant at 0.5 and ϕ_t changes back and forth from 1 to 1.5 repeatedly. Without any movement in physical beliefs, π_t^* will vary repeatedly between 0.5 and 0.6, leading to unbounded movement as $T \rightarrow \infty$. But this argument overlooks a core insight: under RE, ϕ_t cannot undergo such repeated oscillation if its variance is bounded, because it is itself a function of martingale conditional expectations. We pursue this logic more formally here.

Dropping the constant- ϕ_t assumption comes at the cost of much of the parsimony in our previous analysis: SDF expectations non-linearly map to ϕ_t , which then non-linearly combines with π_t to determine π_t^* , which is then non-linearly mapped into X . Perhaps surprisingly, though, we can still make theoretical statements for a meaningful subset of DGPs. Consider ϕ_t as defined in (14) (continuing to suppress j subscripts). The agent now also receives signals $s_{\phi_t} \in \mathcal{S}_{\phi_t}$ to learn about ϕ_t over time. For simplicity, assume that uncertainty over ϕ_t evolves on a binomial tree ($|\mathcal{S}_{\phi_t}| = 2$), though this is not necessary for the following statements. We also assume that π_t and ϕ_t do not change (relative to their $t - 1$ values) together in the same period.¹⁶ Then the following holds.

PROPOSITION 7. *If ϕ_t evolves as a martingale or supermartingale ($\mathbb{E}_t[\phi_{t+1}] \leq \phi_t$) and $\text{Var}(\phi_t) < \infty$, then the bounds in [Proposition 2](#) and [Corollary 1](#) (and their counterparts in [Proposition 5](#)) continue to apply, with ϕ_0 replacing ϕ .*

The proof of [Proposition 7](#) demonstrates that the variation in π_0^* arising when ϕ_t is a non-degenerate (super)martingale always *strictly lowers* excess RN belief movement, rendering the

¹⁶In a previous working paper version of this manuscript, we prove the result for arbitrary values of $|\mathcal{S}_{\phi_t}|$, at greatly increased complexity and little added insight. As for the second assumption, we can, for example, split each period into two sub-periods, with ϕ_t changing in the first half of the period and π_t in the second. The general problem quickly becomes intractable without this assumption, though we provide additional simulation evidence relaxing it below.

main bound in (10) even more conservative. Changes in ϕ_t do cause π_t^* to change, which adds movement. But when ϕ_t is a supermartingale, this movement works *against* the upward bias in RN beliefs themselves and decreasing $\pi_t^* - \pi_t$ in expectation. This reduces potential *future* excess movement more than enough to offset the increase in period t .

As for the interpretation of the supermartingale restriction, from the definition of covariance, $\phi_{t,j}$ is a supermartingale if and only if $\text{Cov}_t(\phi_{t+1,j}, \mathbb{E}_t[M_T | R_T^m = \theta_{j+1}]) \geq 0$. Interpreting the SDF M_T as proportional to $U'(C_T)$, this requires that risk aversion (encoded in $\phi_{t+1,j}$) be positively correlated with expected marginal utility (MU) in the high-consumption state. This is an intuitively reasonable restriction, as it implies bad news (higher expected MU) generally arrives at the same time for both states, with the expected low-consumption MU increasing more than its high-consumption counterpart.¹⁷ The converse ($\text{Cov}_t(\phi_{t+1,j}, \mathbb{E}_t[M_{t,T} | R_T^m = \theta_{j+1}]) < 0$), by contrast, requires risk aversion to increase in general in response to *good* news about MU in the good state.

4.3 Relaxing CTI in Numerical Simulations


Given the complexity of the setting, it is difficult to make analytical statements when ϕ_t does not satisfy the assumptions of Proposition 7. Instead, we numerically simulate DGPs in which ϕ_t can vary with more freedom. The agent learns about the terminal state, and the SDF realization in both the j and $j + 1$ states, over time, under binary DGPs with different combinations of signal strengths and uncertainty for all three objects. We describe the setting more fully in Online Appendix C.3, and we outline the results of these simulations here.

Figure 4 plots distributions for the estimated $\mathbb{E}[m^*]$ (rather than $\mathbb{E}[X^*]$, as $\mathbb{E}[m^*]$ is what changes with the DGP here) across these simulations. Each DGP is simulated repeatedly to obtain an estimated $\mathbb{E}[m^*]$ for that DGP. Each line represents a different $\mathbb{E}[m^*]$ distribution given variation in the signal strengths for θ , with the different lines showing different signal strengths for learning about the conditional values of M_T (and thus ϕ). In all cases, $\pi_0^* = 0.5$ and $\phi_0 = 3$.

The black line (“No ϕ Uncertainty”) shows a baseline with $\phi_t = \phi = 3$ for all t . As in Section 2.3, when signals are symmetric, $\mathbb{E}[m^*] = u_0^* = 0.25$, and very asymmetric DGPs produce the tails. Up to smoothing noise, $\mathbb{E}[m^*]$ never crosses the theoretical upper bound of 0.375 from (10). In the dark gray lines (“Low ϕ Uncertainty”), $\phi_0 = 3$ as before, but ϕ_t now varies such that the ex ante standard deviation of ϕ_T is $\sigma_\phi \equiv \text{SD}_0(\phi_T) = 0.36$. Using Proposition 6, if return states θ_j and θ_{j+1} differ by 5% (as in our empirical setting), this corresponds to relative risk aversion of $\gamma_0 = 40$ and standard deviation for γ_T of $\sigma_\gamma = 7.2$. Changing ϕ has virtually no effect regardless of the signal structure: average $\mathbb{E}[m^*]$ rises by 0.0012, and the number of DGPs for which $\mathbb{E}[m^*]$ exceeds the bound rises by just 0.00007 percentage points (pp). In the medium gray lines (“Medium ϕ Uncertainty”), $\sigma_\phi = 1.62$ (or $\sigma_\gamma = 32.4$). Even with such sizable variation, average $\mathbb{E}[m^*]$ rises by 0.006, and DGPs above the bound by 0.0003 pp. In the light gray lines (“High ϕ Uncertainty”), $\sigma_\phi = 3.0$ ($\sigma_\gamma = 60$). Average $\mathbb{E}[m^*]$ still only increases by 0.015, and DGPs above the bound by 0.0012 pp. In all cases, the bound in Corollary 1 for $\phi \rightarrow \infty$ holds for 100 percent of the simulations.

¹⁷For further discussion of such a restriction in a different context, see Lazarus (2019).

Figure 4: RN Belief Movement Distributions with Time-Varying ϕ_t



Notes: This figure shows the results of simulations studying the impact of time-varying ϕ_t on the distribution of $\mathbb{E}[m^*]$ given different DGPs with possibly asymmetric signal strengths about θ and ϕ_T over time. The dark black line (“No ϕ Uncertainty”) shows the distribution when there is no uncertainty about ϕ . Each line in the slightly lighter dark gray set (“Low ϕ Uncertainty”) represents the equivalent distribution for a DGP that also contains uncertainty about the SDF in both states. In this case, $\phi_0 = 3$, but ϕ_T can vary from 2.14 to 4.2. The set of gray lines (“Medium ϕ Uncertainty”) allow ϕ_T to vary from 1.5 to 6, and the set of light-colored lines (“High ϕ Uncertainty”) allow ϕ_T to vary from 1 to 9.

We conclude, somewhat surprisingly, that even significant uncertainty in ϕ_T has limited impact on our bounds under these binary DGPs. When the agent updates her SDF expectations, these updates must still respect RE. Thus even large values for σ_ϕ do not allow for arbitrary oscillations of ϕ_t , as information about ϕ_T is revealed gradually over time. While there might exist DGPs in which time-varying ϕ_t generates a meaningful effect in an otherwise standard rational setting, we suspect given our results that they would require rather perverse time-varying structures.

To test this supposition in a less abstract environment in which CTI is violated, we proceed to solve and simulate the [Campbell and Cochrane \(1999\)](#) habit formation model using that paper’s baseline calibration. We find in this case that for the interior state pairs, ϕ_t is closely approximated by a martingale and therefore does not generate additional $\mathbb{E}[X^*]$. Our bounds thus continue to hold and continue to be conservative. See Online Appendix C.4 for details.¹⁸

4.4 Aggregating Over Belief Streams

The previous subsection considers the effect of time variation in ϕ_j *within* a given RN belief stream (i.e., within a single option contract, lasting weeks or months). But even if the effects of such higher-frequency variation are small, it may be less palatable to assume that ϕ_j is constant *across*

¹⁸One might reasonably argue that the habit model is the wrong place to start when looking for plausible alternatives, as it is designed to match low-frequency variation. This is in fact the point of the exercise: it shows that to generate bounds violations, alternative models are needed to produce the high- and medium-frequency variation that we measure.

belief streams, given the possibility of lower-frequency variation over the span of years. We now ask how such variation affects our bounds when considering multiple belief streams.

Answering this question is also important for empirical implementation. The bounds in [Proposition 5](#) are stated as date-0 expectations conditional on the RN prior, but we observe only one draw X_j^* per expiration date rather than the expectation of this statistic for a given $\tilde{\pi}_{0,j}^*$. Estimation thus requires aggregating over multiple streams with different $\tilde{\pi}_{0,j}^*$ and, as above, potentially different ϕ_j . Thus, generalizing slightly, assume now that we can observe index options for N expiration dates $\mathcal{T} \equiv \{T_i\}_{i=1}^N$, so i indexes belief streams (and their DGPs). We use $\phi_{i,j}$ for the SDF ratio for expiration date T_i and state pair (θ_j, θ_{j+1}) ; RN beliefs are $\tilde{\pi}_{t,i,j}^*$; and RN excess movement is $X_{i,j}^*$. We again often suppress j for an arbitrary state pair. Due to Jensen's inequality, we cannot simply use $\mathbb{E}[\phi_i]$ in place of ϕ_i or $\mathbb{E}[\pi_{0,i}^*]$ in place of $\pi_{0,i}^*$ when taking the expectation of both sides of the results in [Propositions 1, 2](#) and [5](#) over all i . However, the following generalizations do hold:

PROPOSITION 8. Define $\bar{\phi} \equiv \max_{\pi_{0,i}^*} \mathbb{E}[\phi_i | \pi_{0,i}^*]$. We have:

- (i) GENERALIZATION OF [PROPOSITION 1](#): If $\text{Cov}(\pi_{0,i}, \Delta_i) = 0$, and $\pi_{0,i}^*$ is constant across i (i.e., fixing a given $\pi_{0,i}^*$), then over all streams,

$$\mathbb{E}[X_i^*] \leq \max \left\{ 0, \left(\pi_{0,i}^* - \frac{\pi_{0,i}^*}{\mathbb{E}[\phi_i] + (1 - \mathbb{E}[\phi_i])\pi_{0,i}^*} \right) \mathbb{E}[\Delta_i] \right\}. \quad (15)$$

- (ii) GENERALIZATION OF [PROPOSITION 2](#): Over all streams, we have without any additional assumptions that

$$\mathbb{E}[X_i^*] \leq \mathbb{E} \left[\left(\pi_{0,i}^* - \frac{\pi_{0,i}^*}{\bar{\phi} + (1 - \bar{\phi})\pi_{0,i}^*} \right) \pi_{0,i}^* \right], \quad (16)$$

or, fixing a given $\pi_{0,i}^*$, $\mathbb{E}[X_i^*] \leq \left(\pi_{0,i}^* - \frac{\pi_{0,i}^*}{\mathbb{E}[\phi_i] + (1 - \mathbb{E}[\phi_i])\pi_{0,i}^*} \right) \pi_{0,i}^*$.

- (iii) GENERALIZATION OF [COROLLARY 1](#): Over all streams, without any additional assumptions,

$$\mathbb{E}[X_i^*] \leq \mathbb{E} \left[\pi_{0,i}^{*2} \right].$$

- (iv) GENERALIZATION OF [COROLLARY 2](#): If $\Delta_i \leq 0$ for all i , then over all streams,

$$\mathbb{E}[X_i^*] \leq 0.$$

Only the analogue to [Proposition 1](#) requires an additional assumption. The original formula includes the product of $\pi_0^* - \pi_0$ and Δ , so the covariance between $\pi_{0,i}$ and Δ_i across DGPs affects the generalized bound. For simplicity, we set this covariance to zero, which is equivalent to assuming no relationship between the asymmetry of the DGP and ϕ . This part also holds fixing $\pi_{0,i}^*$; this is sufficient for our purposes, as our empirical results for this less-conservative bound will generally be conditional on a given $\pi_{0,i}^*$.¹⁹ Part (ii), meanwhile, generalizes (10) by applying

¹⁹The constant- $\pi_{0,i}^*$ bounds can also be read as bounds for the conditional expectation $\mathbb{E}[X_i^* | \pi_{0,i}^*]$ given $\mathbb{E}[\phi_i | \pi_{0,i}^*]$.

Jensen's inequality for one of several variables, as the second partial derivative of that bound in $\phi_{i,j}$ is negative. The bound in (16) is thus even more conservative than the original bound for a single stream. The bounds that do not depend on $\phi_{i,j}$ (parts (iii)–(iv)) apply as previously stated.

These bounds are now empirically implementable, and the minimum $\bar{\phi}$ that solves (16) is a conservative estimate of the maximal conditional-mean SDF slope for the return-state pair in question. If no such $\bar{\phi}$ solves the bound, then the bound in part (iii) is violated. Finally, reintroducing dependence on the state pair j , it is likely that the values $\bar{\phi}_j$ vary over j . But the same steps used for Proposition 8 to take expectations over i can also be applied to take expectations over j , thereby obtaining a single estimate $\bar{\phi}$ aggregated over both streams *and* return states (for all states meeting CTI) when desired.

5. Robustness to Additional Assumptions

Having considered Assumption 4 in detail, we turn to the other two assumptions to continue our discussion on how to interpret a rejection of the joint null. We consider Assumption 2 first, as the robustness result in this case will be a useful stepping stone in discussing Assumption 3. Finally, we provide a result accounting for possible mismeasurement or market microstructure noise.

5.1 RE and the Interpretation of Bound Violations

We have assumed throughout that Assumption 2 holds, which in general requires both (i) a correct subjective prior and (ii) rational updating using the true signal distribution as the likelihood. One natural question is whether an incorrect prior by itself can generate violations of the upper bound for X^* , or whether excessive movement in general requires incorrect updating. The following straightforward proposition makes clear that the latter is likely to be necessary for a bound violation. We continue to adopt the notation from Section 2 for clarity, but the following should be understood to apply for conditional beliefs for some state j .²⁰ Assume that Assumptions 3–4 continue to hold.

PROPOSITION 9. *In place of Assumption 2, assume that the agent has an incorrect prior, $\pi_0 \neq \mathbb{P}_0(\theta)$, but updates correctly, in the sense that $\pi_t \propto \pi_{t-1} DGP(s_t | \theta, H_{t-1})$. Define $\check{\phi} \equiv \phi L$, where $L \equiv \frac{\pi_0 / (1 - \pi_0)}{\mathbb{P}_0(\theta) / (1 - \mathbb{P}_0(\theta))}$ indexes the prior belief distortion, with $0 < L < \infty$. Then:*

- (i) *For all H_t , the agent's RN beliefs π_t^* are equivalent to the RN beliefs of a fictitious agent whose physical beliefs $\check{\pi}_t$ satisfy Assumption 2 but who has $\check{\phi}$ in place of ϕ .*
- (ii) *If $\check{\phi} \geq 1$, then all previously stated restrictions on $\mathbb{E}[X^*]$ continue to hold, with $\check{\phi}$ in place of ϕ and $\check{\pi}_0$ in place of π_0 . In particular, one cannot in this case have $\mathbb{E}[X^*] > \pi_0^{*2}$.*
- (iii) *If $\check{\phi} < 1$ so that $\pi_0^* < \mathbb{P}_0(\theta) = \check{\pi}_0$, then the bound expressed in Proposition 2 becomes $\mathbb{E}[X^*] \leq (\check{\pi}_0 - \pi_0^*)(1 - \pi_0^*)$, and Corollary 1 becomes $\mathbb{E}[X^*] \leq (1 - \pi_0^*)^2$. Thus regardless of $\check{\phi}$, it must be the case that $\mathbb{E}[X^*] \leq \max(\pi_0^{*2}, (1 - \pi_0^*)^2)$.*

²⁰For example, the incorrect prior is $\check{\pi}_{0,j} \neq \mathbb{P}_0(R_T^m = \theta_j | R_T^m \in \{\theta_j, \theta_{j+1}\})$, and $L \equiv \frac{\check{\pi}_{0,j} / (1 - \check{\pi}_{0,j})}{\mathbb{P}_0(R_T^m = \theta_j) / (1 - \mathbb{P}_0(R_T^m = \theta_j))}$.

Part (i) formalizes that risk aversion is isomorphic to an incorrect prior, in that both have the same effect on π_t^* relative to the objective $\mathbb{P}_t(\theta)$. Thus with a suitably altered value of ϕ , the bounds generally cover the case of an incorrect prior, as in part (ii). The only case in which this argument requires slight amendment is when the prior is so downwardly distorted that $\pi_0^* < \mathbb{P}_0(\theta)$. Even in this case, though, a slightly altered version of [Corollary 1](#) still applies, as in part (iii). An incorrect prior acts as a one-time belief distortion; while reverting to the correct belief in this case does require some excess movement, this is generally not sufficient for a full violation of the bound in [Proposition 2](#). In general, then, incorrect updating behavior must be present in such a violation.²¹

5.2 Robustness to $\phi < 1$

We now turn to [Assumption 3](#), which imposes that $\phi_j \geq 1$ with return states ordered such that $\theta_j < \theta_{j+1}$. This requires that the expected SDF realization in the low-return state be greater than its expected realization in the high-return state. A line of work beginning with [Jackwerth \(2000\)](#) and [Ait-Sahalia and Lo \(2000\)](#), however, argues that the SDF does not decrease monotonically with the index return in options data. This finding, often referred to as the “pricing kernel puzzle,” would imply a violation of [Assumption 3](#) for at least some subset of the return space. How should our results be interpreted in light of this possibility?

Given such a violation, it turns out that one can still make theoretical statements providing an upper bound for RN excess movement. As shown in [Proposition 9\(i\)](#), assuming a different value for ϕ is formally equivalent to assuming a distortion in the agent’s prior belief. The case with $\phi < 1$ is equivalent to a large downward distortion in the prior, as in part (iii) of that proposition. This equivalence leads to the following corollary, which holds under [Assumptions 2](#) and [4](#).

COROLLARY 3. *If $\phi < 1$ rather than $\phi \geq 1$ in [Assumption 3](#), then the bound from [Proposition 2](#) becomes*

$$\mathbb{E}[X^*] \leq (\pi_0 - \pi_0^*)(1 - \pi_0^*) = \left(1 - \frac{1}{\phi^{-1} + (1 - \phi^{-1})(1 - \pi_0^*)}\right) (1 - \pi_0^*)^2,$$

and [Corollary 1](#) becomes $\mathbb{E}[X^*] \leq (1 - \pi_0^*)^2$. For any ϕ , therefore, $\mathbb{E}[X^*] \leq \max(\pi_0^{*2}, (1 - \pi_0^*)^2)$.

The main bounds thus apply with minor modification, effectively flipping the role of the two states when $\phi < 1$. This entails replacing π_0^* with $1 - \pi_0^*$, and for [Proposition 2](#), replacing ϕ with ϕ^{-1} as well. Thus it is *not* the case that anything goes when [Assumption 3](#) is violated: excess movement is bounded no matter what, and its upper bound is a function of $\max(\phi, \phi^{-1})$, which in both cases indexes risk aversion across the two states. Bounds violations, meanwhile, retain their interpretation regardless of ϕ .

While [Corollary 3](#) shows that one can still make statements about excess movement when $\phi < 1$, the bound in this case would have to be estimated separately from our main bounds derived for the

²¹Note that even if ϕ is constant under the agent’s subjective belief, incorrect updating will induce a probability distortion such that the *actual* SDF (assessed under \mathbb{P}) will feature a time-varying ϕ_t . This illustrates the deep connection between violations of [Assumptions 2](#) and [4](#).

$\phi \geq 1$ case. So for the sake of simplicity, we empirically estimate only the main bounds. We do so in part because we estimate the bounds separately for multiple points in the return space. Thus even if one is concerned about the pricing kernel puzzle affecting our results, one can confine empirical attention to our estimates for return ranges for which the puzzle does *not* emerge. For example, within this literature, it is a robust finding that the estimated SDF declines monotonically for the range of negative index returns; see, for example, [Driessen, Koeter, and Wilms \(2022, Figure 4\)](#) or [Schreindorfer and Sichert \(2022, p. 4\)](#). And as will be seen in [Section 6](#), our empirical estimates for ϕ_j for these negative return states are just as high (in fact somewhat higher) than for positive return states, indicating that the results are unlikely to be driven by violations of [Assumption 3](#).

5.3 Robustness to Measurement Error

The robustness results above speak to the possibility of theoretical misspecification. As the final step in making our bounds implementable, we now consider how to account for possible *empirical* misspecification arising from mismeasurement or microstructure noise in RN beliefs. The bounds provide a minimum value of ϕ required to rationalize the observed variation in RN beliefs; if some of this variation is in fact arising due to noise, then we may overestimate this required ϕ . A simple correction can be applied to our bounds to account for this issue, as shown in the following result. Given that noise arises period-by-period, we first define one-period analogues for our statistics: denote RN movement between t and $t + 1$ by $m_{t,t+1}^* \equiv (\pi_t^* - \pi_{t+1}^*)^2$, RN uncertainty at t by $u_t^* \equiv (1 - \pi_{t+1}^*)\pi_t^*$, and RN excess movement between t and $t + 1$ as $X_{t,t+1}^* \equiv m_{t,t+1}^* - (u_t^* - u_{t+1}^*)$. Similar to [Augenblick and Rabin \(2021, Section II.E\)](#), we then have the following:

PROPOSITION 10. *Assume that the observed $\hat{\pi}_t^*$ is measured with error with respect to the true π_t^* :*

$$\hat{\pi}_t^* = \pi_t^* + \epsilon_t, \tag{17}$$

where $\mathbb{E}[\epsilon_t] = 0$, $\mathbb{E}[\epsilon_t \epsilon_{t+1}] = 0$, and $\mathbb{E}[\epsilon_{t+k} \pi_{t+k'}^*] = 0$ for $k, k' \in \{0, 1\}$. Denoting the observed one-period RN excess movement statistic by $\hat{X}_{t,t+1}^*$, its relation to the true value $X_{t,t+1}^*$ in expectation is

$$\mathbb{E}[\hat{X}_{t,t+1}^* - X_{t,t+1}^*] = 2\text{Var}(\epsilon_t).$$

We can thus subtract $2\text{Var}(\epsilon_{t,j})$ from each period's observed excess-movement statistic to obtain an unbiased true excess movement value, which can then be used in our bounds after summing over the full stream. If measurement error is positively correlated over time rather than uncorrelated, this will reduce the upward bias in measured X^* . We discuss estimation of $\text{Var}(\epsilon_t)$, as well as autocovariance statistics for ϵ_t , in [Section 6.2](#).

6. Empirical Estimation and Results

Our theory leads to bounds on the variation in RN beliefs over the market index return, which we proceed now to measure in the data. We begin by describing how we map from theory to data and how we estimate microstructure noise; we then summarize the data before turning to our main results. We conclude with a discussion of the reduced-form correlates of RN excess movement.

6.1 Data and Risk-Neutral Distribution

Data. Our main source for S&P 500 index options data is the OptionMetrics database, which provides end-of-day prices for European call and put options for all strike prices and option expiration dates traded on the Chicago Board Options Exchange (CBOE). The sample runs from January 1996 through December 2018. We augment this data with intraday price quotes obtained directly from the CBOE for a subset of trading days in our sample, in order to account for market microstructure noise; this additional data is described further in [Section 6.2](#).

We apply standard filters to remove outliers and options with poor trading liquidity from the OptionMetrics data, with details provided in Online Appendix C.5.²² Two aspects of this data cleaning bear mention here. First, while our bounds apply for belief streams of arbitrary length, we follow past literature (e.g., [Christoffersen, Heston, and Jacobs, 2013](#); [Martin, 2017](#)) and consider options with maturity of at most one year. Second, after transforming the filtered prices to RN beliefs (described below), we keep only conditional RN belief observations $\tilde{\pi}_{t,i,j}^*$ for which the non-conditional beliefs satisfy $\pi_t^*(R_{T_i}^m = \theta_j) + \pi_t^*(R_{T_i}^m = \theta_{j+1}) \geq 5\%$, as conditional beliefs are likely to be particularly susceptible to mismeasurement when the underlying beliefs are close to zero.

Empirical return space. For our baseline estimation, we define the return state space Θ in terms of log excess return intervals:

$$\Theta = R_{0,T_i}^f \exp\{(-\infty, -0.2], (-0.2, -0.15], (-0.15, -0.1], \dots, (0.1, 0.15], (0.15, 0.2], (0.2, \infty)\},$$

where R_{0,T_i}^f is the gross risk-free rate from 0_i to T_i .²³ In words, return state 1 is realized if the log excess S&P return from 0_i to T_i is less than $-0.2 \approx -20\%$; state 2 is realized if the excess return is in the five-percentage-point bin between -0.2 and -0.15 ; and so on. Abusing notation slightly, we often refer to states by the right end of their associated excess-return bin: $\theta_1 = -0.2$, $\theta_2 = -0.15$, \dots , $\theta_9 = 0.2$, $\theta_{10} = \infty$. The binary conditional beliefs to be used in our tests are $\tilde{\pi}_{t,i,j}^* \equiv \pi_t^*(R_{T_i}^m = \theta_j | R_{T_i}^m \in \{\theta_j, \theta_{j+1}\})$, so $\tilde{\pi}_{t,i,j}^*$ again corresponds to the probability that the low state j (e.g., $\theta_2 = -15\%$) will be realized, conditional on j or $j+1$ (in this case, conditional on

²²The raw dataset contains 12.4 million option prices; the filtered data, which are then used to measure the RN distribution, contain 4.3 million option prices corresponding to 5,537 trading dates and 991 option expiration dates. The majority of the difference is attributable to our use of only out-of-the-money call and put strikes.

²³We use excess returns for convenience of interpretation. Following [van Binsbergen, Diamond, and Grotteria \(2022\)](#), we measure R_{0,T_i}^f directly from the options prices by applying the put-call parity relationship; again see Appendix C.5.

an excess return between -20% and -10%). We again use the right end of the return bin for j in referencing statistics for pair (θ_j, θ_{j+1}) , corresponding to the midpoint of the two return bins.

This five-percentage-point partition of the return space reflects a desire to balance (i) measurement accuracy for the RN beliefs and (ii) plausibility of our assumption of constant ϕ_j (CTI). Wider bins lead to greater measurement accuracy, but make it less likely that there are no changes in the expected SDF realization conditional on a given return state θ_j relative to θ_{j+1} . We report empirical estimates below for all adjacent state pairs for completeness, but as discussed after [Assumption 4](#), it is unlikely that CTI is met for the extreme state pairs (θ_1 relative to θ_2 , and $\theta_9 = \theta_{J-1}$ vs. $\theta_{10} = \theta_J$). Our focus is thus on the interior state pairs with low-return states $\theta_2, \dots, \theta_8$; in particular, when we aggregate our state-by-state estimates of ϕ_j required to rationalize the data into a single average value $\bar{\phi}$ across states, we use only these interior states.²⁴

Risk-neutral beliefs. To extract a risk-neutral distribution over the return states in Θ from the observed option cross-sections, we use standard tools from the option-pricing literature. Our starting point is equation (13), which tells us how to map from option prices to RN beliefs. We use this to construct a smooth RN distribution for returns, largely following the technique proposed by [Malz \(2014\)](#); Appendix C.5 provides a detailed description. With the RN beliefs $\pi_{t,i,j}^*$ in hand, we can then calculate conditional beliefs straightforwardly as $\tilde{\pi}_{t,i,j}^* = \pi_{t,i,j}^* / (\pi_{t,i,j}^* + \pi_{t,i,j+1}^*)$. We then use the resulting conditional RN belief streams to calculate the excess movement statistics $X_{i,j}^*$ needed to implement our bounds. Our general results in [Section 3](#) restrict the expectation of $X_{i,j}^*$ conditional on state θ_j or θ_{j+1} being realized — in particular, [Propositions 5](#) and [8](#) give bounds for $\tilde{\mathbb{E}}[X_{i,j}^*]$ — and we accordingly keep only observations for which $\tilde{\pi}_{T,i,j}^* = 0$ or 1 ex post; for example, if the total excess return on the market over the life of option contract i is -14%, then we keep only $X_{i,2}^*$ (θ_2 ranges from -20% to -15% return, so $\tilde{\pi}_{T,i,2}^* = 0$) and $X_{i,3}^*$ ($\tilde{\pi}_{T,i,3}^* = 1$).

Simplifying notation. Having clarified how the relevant empirical objects ($\tilde{\pi}_{t,i,j}^*$, $X_{i,j}^*$) depend on the contract i and state pair j , in what follows we generally drop the cumbersome use of i , j , and $\tilde{\cdot}$, and again write π_t^* for $\tilde{\pi}_{t,i,j}^*$, X^* for $X_{i,j}^*$, and so on. Similarly, we often drop the “conditional” qualifier when referring to conditional RN belief statistics.

6.2 Noise Estimation

As in [Proposition 10](#), we also wish to account for measurement error stemming from possible microstructure noise in our estimation. That result shows that unlike in the classical errors-in-variables regression model (which leads to attenuation), measurement error in our case can increase the observed variation X^* and thereby lead to an upward bias in the estimated SDF slope needed to rationalize the data. With noise described by $\hat{\pi}_t^* = \pi_t^* + \epsilon_t$ as in (17), [Proposition 10](#) tells us that we must estimate $\text{Var}(\epsilon_t)$ in order to eliminate this bias. We turn to a sample of high-frequency option prices to estimate this noise variance in our RN beliefs data.

²⁴This yields an additional de facto data filter, as we are effectively considering only option strikes with moneyness between 0.8 and 1.2 (following, e.g., [Constantinides, Jackwerth, and Savov, 2013](#)).

Specifically, we obtain minute-by-minute price quotes on S&P index options for a subset of trading days directly from the CBOE. For each available option expiration date on each such trading day, we recalculate the RN belief distribution at the end of each minute using exactly the same procedure as described in [Section 6.1](#). As this requires calculating 390 sets of RN beliefs for each trading day (9:30 AM–4:00 PM), this procedure would be computationally infeasible if applied to our entire sample of 5,537 trading days (each of which has an average of 11 available option expiration dates, generating 60,543 (t, T) combinations). We accordingly select 30 trading days at random from within our available sample period, and use the minute-by-minute quotes to calculate intraday RN distributions for these days.²⁵

We then use tools from the literature on microstructure noise to estimate $\text{Var}(\epsilon_t)$ using these intraday data. The intuition for this strategy — as described, for example, by [Zhang, Mykland, and Ait-Sahalia \(2005\)](#) — is best understood by assuming temporarily that the noise ϵ_t in (17) is i.i.d., while the true π_t^* changes smoothly over time. Given high-frequency option data, imagine calculating movement using the observed beliefs, $(\hat{\pi}_{t+h}^* - \hat{\pi}_t^*)^2$, with less and less time h between consecutive observations. As one decreases h to 0, the noise swamps the true variation: since $(\pi_{t+h}^* - \pi_t^*)^2 \rightarrow 0$, we have $\mathbb{E}[(\hat{\pi}_{t+h}^* - \hat{\pi}_t^*)^2] \rightarrow 2\text{Var}(\epsilon_t)$. Thus in this simple example, $\text{Var}(\epsilon_t)$ can be estimated by calculating the quadratic variation in RN beliefs sampled at a high frequency.

In practice, one would expect the data to contain both non-i.i.d. noise ϵ_t and jumps in the true process π_t^* , and it is desirable to use a noise estimation method that is robust to these features. One such estimator for $\text{Var}(\epsilon_t)$ is the ReMeDI (“Realized moMents of Disjoint Increments”) estimator proposed by [Li and Linton \(2022\)](#). This estimator takes the average product of *disjoint increments* of the observed process, $(\hat{\pi}_t^* - \hat{\pi}_{t-h}^*)(\hat{\pi}_t^* - \hat{\pi}_{t+h}^*)$.²⁶ The idea is that even if the true process features jumps so that $\mathbb{E}[(\pi_{t+h}^* - \pi_t^*)^2] > 0$, its increments over non-overlapping windows are still approximately uncorrelated. [Li and Linton \(2022, Theorem 4.1\)](#) show that this estimator is consistent for $\text{Var}(\epsilon_t)$ for quite general dependent noise processes and for π_t^* in a general class of semimartingales. It also performs well in simulations and empirical applications.

Using this ReMeDI estimator on our minute-by-minute data, we estimate $\text{Var}(\epsilon_t) = \text{Var}(\epsilon_{t,i,j})$ separately for each combination of trading day t , expiration date T_i , and state pair j in our intraday sample. We then match the noise estimates (which are obtained for a subsample of days) to the observed excess movement observations in our original data. Technical details on our estimation and matching procedure are provided in [Online Appendix C.6](#). Finally, we subtract $2\widehat{\text{Var}}(\epsilon_{t,i,j})$ from $\widehat{X}_{t,t+1,i,j}^*$ to obtain a *noise-adjusted* estimate of one-day excess movement following [Proposition 10](#), and we sum these noise-adjusted one-day values over the full stream to obtain noise-adjusted $X_{i,j}^*$.

We discuss the magnitude of the noise estimates in the next subsection. The ReMeDI procedure also allows for estimation of the intraday autocovariances of the noise ϵ_t . These autocovariances are estimated to be positive for small lag values, but they die out quickly and are precisely

²⁵This yields an intraday data set roughly twice as large as the original one, as $30 \times 390 \times 11 \approx 130,000$.

²⁶More formally, the estimator is $\widehat{\text{Var}}(\epsilon_t) = \frac{1}{N_{\epsilon,n}} \sum_{i=2k_n}^{N_{\epsilon,n}-k_n} (\hat{\pi}_{t_i}^* - \hat{\pi}_{t_i-2k_n}^*)(\hat{\pi}_{t_i}^* - \hat{\pi}_{t_i+k_n}^*)$, where $N_{\epsilon,n}$ is the number of observations over a fixed span (in our case, one trading day) and k_n is a tuning parameter.

estimated near zero for noise observations more than an hour apart. This justifies the assumption in [Proposition 10](#) that end-of-day noise observations are uncorrelated, $\mathbb{E}[\epsilon_t \epsilon_{t+1}] = 0$, as ultimately we care about noise only to the extent that it affects our excess movement statistics at a daily frequency.

Our main results in [Section 6.4](#) use the noise-adjusted excess movement data. All standard errors and confidence intervals are based on a bootstrap procedure (detailed in [Section 6.4](#)) that accounts for the sampling uncertainty in the above noise estimation and averaging procedure.

6.3 Excess Movement: Descriptive Statistics and Figures

[Table 1](#) summarizes the average RN excess movement \bar{X}^* overall (across all interior state pairs and expiration dates) and by subsample. Excess movement is difficult to interpret without some normalization. The first two columns thus divide \bar{X}^* by average initial uncertainty \bar{u}_0 . As in [Augenblick and Rabin \(2021\)](#), this normalized statistic can be interpreted as the percent by which movement exceeds initial uncertainty and thus uncertainty resolution. These values are quite high in our data: for the noise-adjusted statistics, there is on average 123% more movement than initial uncertainty. These values decrease for return states in the middle of the distribution. The early sample has high but noisy X^* statistics, but these averages remain high until the most recent subsample. And higher priors π_0^* correspond with greater X^* , in line with our bounds given $\phi \geq 1$.

The next two columns of [Table 1](#) instead normalize \bar{X}^* by the average contract length \bar{T} , so the resulting statistics can be interpreted as excess movement per day.²⁷ Under this normalization, there is now no clear pattern for average excess movement across bins, as the more-extreme return states also tend to have longer contract lengths.²⁸ This pattern is in fact quite consistent in our data: longer contract lengths tend to coincide with more excess movement, as RN beliefs bounce up and down over the length of a contract. This is inconsistent with RE, under which excess movement in subjective beliefs should not depend at all on the horizon at which uncertainty is resolved. For the splits by date and by prior, the basic patterns discussed above are still present here.

Comparing the raw and noise-adjusted values makes clear that despite the substantial excess movement in the noise-adjusted statistics, noise does represent a meaningful portion (about 1/3) of the raw X^* data. The raw and noise-adjusted mean for \bar{X}^*/\bar{T} differ by about 0.002, so $\widehat{\text{Var}}(\epsilon_t) \approx 0.002/2 = 0.001$. The standard deviation of ϵ_t is thus roughly 0.03 per day. For the return-state splits, noise tends to be lowest for returns near the center of the distribution, as is intuitive.

Next, [Figure 5](#) provides a visual summary of the X^* statistics relative to the bounds. The blue curves describe the raw and noise-adjusted local-average X^* statistics as one varies the RN prior π_0^* ; these curves are the same in both panels. As one would expect, there is very little excess movement for RN priors near 0 or 1, but excess movement is positive for intermediate π_0^* values for which there is greater initial uncertainty. We compare these values to the theoretical bounds

²⁷For a rough idea of the actual variation in RN beliefs given these values, consider a pair $(t, t+1)$ for which $\pi_{t+1}^* = 1 - \pi_t^*$. One-day X^* and m^* then coincide (there is no uncertainty resolution), so, for example, the noise-adjusted mean of 0.0038 corresponds to a raw change of $\sqrt{0.0038} \approx 0.06$ (or $\pi_t^* = 0.47$, $\pi_{t+1}^* = 0.53$).

²⁸Recall that an observation for a given state pair is only included conditional on the realized return being in one of the two states. Longer contracts are likelier to generate greater absolute returns, explaining this positive covariance.

Table 1: Descriptive Statistics for Excess Movement

	\bar{X}^*/\bar{u}_0		\bar{X}^*/\bar{T}		\bar{u}_0	\bar{T}	N (Obs.)
	Raw	Noise-Adj.	Raw	Noise-Adj.			
Overall mean: <i>[Bootstrapped SE]</i>	1.89 [0.25]	1.23 [0.22]	0.0059 [0.0015]	0.0038 [0.0013]	0.18 [0.00]	56 [2]	1,809
<i>By return state:</i>							
1 (-20%)	5.83 [1.18]	4.83 [1.05]	0.0049 [0.0027]	0.0041 [0.0024]	0.17 [0.01]	200 [20]	26
2 (-15%)	11.61 [3.32]	5.70 [3.06]	0.0180 [0.0096]	0.0088 [0.0083]	0.22 [0.01]	141 [25]	19
3 (-10%)	5.76 [0.99]	2.37 [1.07]	0.0151 [0.0059]	0.0062 [0.0051]	0.21 [0.01]	81 [12]	49
4 (-5%)	2.67 [0.59]	1.39 [0.50]	0.0088 [0.0038]	0.0046 [0.0029]	0.14 [0.01]	42 [5]	272
5 (0%)	0.70 [0.16]	0.47 [0.14]	0.0045 [0.0019]	0.0030 [0.0017]	0.23 [0.00]	37 [2]	700
6 (+5%)	1.71 [0.35]	1.14 [0.34]	0.0039 [0.0015]	0.0026 [0.0014]	0.11 [0.01]	49 [3]	567
7 (+10%)	3.87 [1.00]	2.92 [1.03]	0.0053 [0.0023]	0.0040 [0.0023]	0.18 [0.01]	129 [9]	144
8 (+15%)	5.65 [1.48]	5.26 [1.48]	0.0060 [0.0027]	0.0056 [0.0027]	0.21 [0.01]	200 [11]	58
9 (+20%)	3.44 [0.89]	2.09 [1.27]	0.0032 [0.0015]	0.0020 [0.0020]	0.22 [0.01]	232 [9]	36
<i>By date:</i>							
1996–2000	10.89 [2.24]	9.67 [2.17]	0.0211 [0.0074]	0.0187 [0.0072]	0.21 [0.01]	107 [11]	109
2001–2005	1.75 [0.51]	0.55 [0.40]	0.0042 [0.0021]	0.0013 [0.0015]	0.22 [0.01]	90 [11]	112
2006–2010	1.25 [0.22]	0.68 [0.19]	0.0065 [0.0026]	0.0035 [0.0021]	0.17 [0.00]	32 [4]	502
2011–2015	1.75 [0.36]	1.09 [0.28]	0.0050 [0.0024]	0.0031 [0.0019]	0.19 [0.00]	67 [5]	530
2016–2018	0.36 [0.21]	-0.11 [0.14]	0.0011 [0.0017]	-0.0003 [0.0009]	0.16 [0.00]	50 [3]	556
<i>By π_0^*:</i>							
0–0.25	1.01 [0.60]	0.30 [0.55]	0.0055 [0.0048]	0.0017 [0.0047]	0.09 [0.01]	16 [2]	185
0.25–0.5	1.58 [0.21]	0.91 [0.17]	0.0067 [0.0017]	0.0039 [0.0014]	0.23 [0.00]	55 [3]	883
0.5–0.75	2.84 [0.61]	2.19 [0.58]	0.0053 [0.0020]	0.0041 [0.0019]	0.23 [0.00]	123 [7]	284
0.75–1	2.54 [0.95]	1.88 [0.90]	0.0048 [0.0031]	0.0036 [0.0029]	0.06 [0.00]	31 [2]	457

Notes: Empirical conditional means of risk-neutral excess movement $\bar{X}^* \equiv \hat{\mathbb{E}}[X_{i,j}^*]$ are calculated over all interior state pairs $j = 2, \dots, 8$, aside from averages by bin, which are calculated for each state pair separately. Standard errors are estimated using a block bootstrap for the normalized statistic \bar{X}^*/\bar{u}_0 or \bar{X}^*/\bar{T} , with a block size of one month (where contracts are classified by the month in which they expire) and 10,000 draws.

under different levels of ϕ for each value of π_0^* , as shown in gray. Panel (a) uses the tighter bound from [Proposition 8\(i\)](#), with Δ estimated using local averages for the conditional expectations $\mathbb{E}[X^* | \theta, \pi_0^*]$ in (9) over π_0^* .²⁹ Panel (b) uses the conservative bound from the second inequality in [Proposition 8\(ii\)](#), so the gray lines in this panel align with the solid bound lines in [Figure 3](#).

Across the two panels, the X^* values observed in the data exceed both sets of bounds, other than for high RN priors π_0^* and for high ϕ . For panel (a), we estimate $\hat{\Delta} = 0$ for $\pi_0^* \approx 0.5$, $\hat{\Delta} < 0$ for π_0^* below this cutoff, and $\hat{\Delta} > 0$ above it, as is evident from the bounds crossing zero at $\pi_0^* \approx 0.5$. As discussed in [Section 2.3](#), this indicates that the DGP is close to symmetric, with equally informative signals for the two states (θ_j, θ_{j+1}) on average (and thus roughly equal-sized upward and downward movements of π_t^*). As the bounds in [Proposition 8](#) apply for each possible π_0^* , the positive point estimates for $\mathbb{E}[X^* | \pi_0^*]$ clearly violate the bounds for $\pi_0^* < 0.5$, which are (at most) 0 for all ϕ . And while the empirical curves are closer to the more-conservative bounds in panel (b), the noise-adjusted estimates still exceed π_0^{*2} (the bound for $\phi = \infty$) for π_0^* less than about 2/3.

These figures do not, however, integrate over π_0^* , nor do they include any measures of statistical uncertainty necessary to make inferential statements or conduct hypothesis tests. To address these issues, we move on to our main estimation and results.

6.4 Main Results

We turn now to the empirical implementation of our theoretical bounds. Given our sample of noise-adjusted excess movement statistics and corresponding RN priors, each possible value of $\bar{\phi}$ leads to a *residual* excess movement value $e_i(\bar{\phi}) = X_i^* - \text{bound}(\pi_{0,i}^*, \bar{\phi})$ for contract i . We calculate two versions of this residual corresponding to the bound in part (i) and the unconditional bound in part (ii) of [Proposition 8](#) (equations (15) and (16), respectively):


$$\begin{aligned} e_i^\Delta(\bar{\phi}) &= X_i^* - \max\left\{0, \left(\pi_{0,i}^* - \frac{\pi_{0,i}^*}{\bar{\phi} + (1-\bar{\phi})\pi_{0,i}^*}\right) \hat{\Delta}_i\right\}, \\ e_i^{\text{main}}(\bar{\phi}) &= X_i^* - \left(\pi_{0,i}^* - \frac{\pi_{0,i}^*}{\bar{\phi} + (1-\bar{\phi})\pi_{0,i}^*}\right) \pi_{0,i}^*. \end{aligned} \tag{18}$$

The first version corresponds to the tighter bound from [Proposition 1](#), which requires a smoothed estimate of Δ_i as calculated for [Figure 5](#). This version accounts for the estimated DGP (through $\hat{\Delta}_i$) and thus conveys some useful preliminary information, but strictly speaking, part (i) of the proposition applies only conditional on a given $\pi_{0,i}^*$ and only under the unverifiable condition that $\text{Cov}(\pi_{0,i}, \Delta_i) = 0$. Thus the second version, which implements the more conservative bound from [Proposition 2](#) and which applies unconditionally, is the basis for our main set of results. Both residuals can be directly calculated using our noise-adjusted data for each possible value of $\bar{\phi}$.

We first present sample averages of these residual statistics for a range of values of $\bar{\phi}$, both for each individual state pair j and aggregated over all interior states. As these values have no

²⁹To emphasize that $\hat{\Delta} < 0$ for $\pi_0^* < 0.5$, the figure does not cut the bounds off at 0.

Figure 5: Excess Movement vs. RN Prior: Data and Theoretical Bounds



Notes: Empirical excess movement curves are kernel-weighted local averages (Epanechnikov kernel, bandwidth for π_0^* of 0.07) over all interior state pairs $j = 2, \dots, 8$. All statistics are estimates of conditional means $\tilde{\mathbb{E}}[\cdot]$ for RN beliefs $\tilde{\pi}_{i,j}^*$, and theoretical curves correspond to $\phi \equiv \tilde{\mathbb{E}}[\phi_{i,j}]$, with notation simplified for clarity. Bounds for (a) obtain $\hat{\Delta}$ using a kernel-weighted local average over π_0^* for each of the two terms in (9), with $\hat{\Delta}$ then plugged into the inequality in Proposition 8(i). Bounds for (b) use only the second inequality in Proposition 8(ii).

Table 2: Residual Excess Movement t -Statistics for Different $\bar{\phi}$

$\bar{\phi}$:	(a) Tighter Lower Bound from Proposition 8(i)						(b) Conservative Lower Bound from Proposition 8(ii)					
	1	2	5	10	50	∞	1	2	5	10	50	∞
Overall t -stat.:	5.19	5.15	4.98	4.64	3.03	0.78	5.19	4.08	2.68	1.75	0.07	-1.48
<i>By return state:</i>												
1 (-20%)	4.06	3.69	3.04	2.57	1.94	1.72	4.06	3.58	2.79	2.22	1.47	1.21
2 (-15%)	2.13	2.13	2.13	2.13	2.14	2.14	2.13	2.03	1.94	1.90	1.87	1.86
3 (-10%)	2.23	2.23	2.23	2.23	2.23	2.23	2.23	2.03	1.86	1.79	1.72	1.70
4 (-5%)	2.78	2.78	2.77	2.77	2.76	2.76	2.78	2.51	2.28	2.18	2.09	2.06
5 (0%)	3.09	3.09	3.08	3.08	3.06	3.06	3.09	1.86	0.74	0.26	-0.19	-0.31
6 (+5%)	3.26	2.66	1.27	0.04	-1.71	-2.42	3.26	1.96	-0.12	-1.77	-5.22	-8.39
7 (+10%)	2.82	2.73	2.51	2.29	1.70	1.23	2.82	2.38	1.73	1.28	0.50	0.06
8 (+15%)	3.53	3.49	3.42	3.37	3.29	3.25	3.53	3.20	2.77	2.51	2.21	2.10
9 (+20%)	1.63	1.58	1.52	1.48	1.44	1.43	1.63	1.36	1.05	0.89	0.72	0.68

Notes: This table shows t -statistics for the average residuals $\overline{e_i(\bar{\phi})}$ in (18) for different values of $\bar{\phi}$. The columns for (a) use $e_i^{\Delta}(\bar{\phi})$, with Δ_i estimated as in panel (a) of Figure 5. The columns for (b) use $e_i^{\text{main}}(\bar{\phi})$. Standard errors are estimated using a block bootstrap with block size of one month and 10,000 draws. All statistics are calculated using conditional means of noise-adjusted X^* .

natural scaling, we present them as t -statistics, $t_{\bar{e}} = \overline{e_i(\bar{\phi})} / \widehat{SE}_{\bar{e}}$, where $\overline{e_i(\bar{\phi})}$ is the sample average of $e_i(\bar{\phi})$ and $\widehat{SE}_{\bar{e}}$ is its standard error. We calculate these standard errors using a block bootstrap with a block length of one month, with each block containing (i) raw excess movement statistics and priors for all contracts expiring in a given month, and (ii) noise variance estimates for any trading days in our intraday sample that fall in the same month. For each resampled data set, we use the set of $(X_i^*, \tilde{\pi}_{0,i}^*, \{\widehat{\text{Var}}(\epsilon_{t,i})\})$ values to recalculate noise-adjusted excess movement and residual values $e_i(\bar{\phi})$.³⁰ The bootstrap accordingly accounts for sampling uncertainty in all statistics used to calculate noise-adjusted X_i^* and $e_i(\bar{\phi})$. We conduct 10,000 such draws, from which we calculate standard errors as the standard deviation of $\overline{e_i(\bar{\phi})}$ across draws.

These residual t -statistics are presented in Table 2 for the two versions of the residual in (18), both overall and by state pair. Since $\mathbb{E}[e_i(\bar{\phi})] = 0$ under RE given a correctly specified $\bar{\phi}$, these t -statistics tell us how far the residuals are from being consistent with any hypothesized null for the SDF slope. Positive numbers correspond to the data exhibiting too much excess movement to be consistent with a given value of $\bar{\phi}$. In the first six columns, the only negative t -statistics for the tighter bound are for RN beliefs over $(\theta_j, \theta_{j+1}) = ([0\%, 5\%], [5\%, 10\%])$ (i.e., $j = 6$) for $\bar{\phi} > 10$; all other t -statistics, including the overall values, are positive (and mostly large in magnitude), indicating no value of $\bar{\phi}$ is consistent with the degree of observed excess RN movement. In the

³⁰For the residual $e_i^{\Delta}(\bar{\phi})$ corresponding to the tighter bound, we also re-estimate $\widehat{\Delta}_i$ in each bootstrap draw by calculating the same local average (with respect to $\pi_{0,i}^*$) as in Figure 5(a), and then evaluating it at the observed $\pi_{0,i}^*$.

remaining columns, the conservative bound t -statistics are somewhat smaller in magnitude but also generally positive, other than for large $\bar{\phi}$ and for returns in the middle of the distribution.

As admissible excess movement $\mathbb{E}[X_i^*]$ is monotonically increasing in the unobserved parameter $\bar{\phi}$, our main empirical exercise is to estimate the lower bound for this SDF slope such that the bound for $\mathbb{E}[X_i^*]$ is satisfied. This lower bound for $\bar{\phi}$ is estimated as the minimal value for which the average residual $\overline{e_i(\bar{\phi})}$ is zero, so that we are effectively finding the root of the function traced out in Table 2. Given that the tighter bound is generally not satisfied even for $\bar{\phi} = \infty$ (and since it holds only under restrictive assumptions), we now confine attention to the conservative bound from Proposition 8(ii). The estimated $\bar{\phi}$ is the minimal SDF slope for which the amount of observed excess movement in RN beliefs can be rationalized; it is thus an index of the restrictiveness of Assumptions 2–4 (or, given Proposition 7, Assumptions 2–3 and supermartingale ϕ_t). As above, we estimate this SDF slope both for each individual state pair ($\bar{\phi} = \bar{\phi}_j$) and overall ($\bar{\phi} = \bar{\phi}$). To make the estimates for $\bar{\phi}$ more interpretable, we also use Proposition 6 to translate them into local relative risk aversion values $\bar{\gamma} = \frac{(\bar{\phi}-1)}{0.05} = 20(\bar{\phi} - 1)$ for a representative agent who consumes the market.

Table 3 presents our main results. Whenever there is *no* value of $\bar{\phi}$ for which the bound for $\mathbb{E}[X_i^*]$ is satisfied — i.e., from Proposition 8(iii), when we estimate $\mathbb{E}[X_i^*] > \mathbb{E}[\pi_{0,i}^{*,2}]$ — we write $\bar{\phi} = \infty$. In brackets below each point estimate, we provide the lower bound of a one-sided 95% confidence interval (CI) for the parameter in question.³¹ Starting from the same bootstrap resampling procedure as described just above, we obtain these CIs by inverting a one-sided test. The CI lower bound $\hat{\phi}_{LB}$ is the minimal $\bar{\phi}$ such that $\overline{e_i^{\text{main}}(\bar{\phi})} = 0$ is not rejected at the 5% level using the bootstrapped data (see Online Appendix C.7 for further details).

Starting with the overall estimates in the first row, the point estimate for the conservative lower bound for $\bar{\phi}$ is slightly greater than 50, in line with the t -statistic of 0.07 in Table 2(b) for $\bar{\phi} = 50$. This translates to an extraordinarily high estimated lower bound for $\bar{\gamma}$ of 1,075. Values of $\bar{\phi}$ below 9.8 (for $\bar{\gamma}$, below 175) are rejected at the 5% level. This is our main empirical result: under our maintained assumptions, extremely high risk aversion is needed to rationalize the large degree of excess movement in RN beliefs observed in the data.

For the individual return-state pairs, all but two of the point estimates are infinite, indicating that no amount of utility curvature (or SDF slope) can rationalize the observed excess movement such that the bounds are satisfied. For many of the states ($j = 2, 3, 4, 8$), their confidence intervals also have lower bounds of ∞ (or more precisely, are empty), indicating outright model rejection. Only for RN beliefs over state pairs in the middle of the distribution — i.e., for $j = 5$ and 6, with return midpoints of 0% and 5%, respectively — are there finite point estimates and confidence intervals that contain reasonable risk-aversion values of about 20. RN beliefs over these intermediate return states are thus comparatively well-behaved; for all other states, there is so much mean-reverting variation in RN beliefs that the bounds are only met for implausibly large values of $\bar{\phi}_j$, if at all.

As discussed in Section 5.2, the pricing kernel puzzle tends to emerge only in the range of positive returns. One can thus focus attention on the interior states with negative returns, $j = 2$ –5,

³¹The set identification implied by our theoretical bound motivates our use of these one-sided intervals.

Table 3: Main Estimation Results

	<i>Conservative Lower Bound for:</i>	
	SDF Slope $\bar{\phi}$	RRA $\bar{\gamma}$
Overall bound: <i>[95% CI Lower Bound]</i>	54.7 [9.8]	1,075 [175]
<i>By return state:</i>		
1 (-20%)	∞ [24.2]	∞ [464]
2 (-15%)	∞ [∞]	∞ [∞]
3 (-10%)	∞ [∞]	∞ [∞]
4 (-5%)	∞ [∞]	∞ [∞]
5 (0%)	19.4 [2.1]	368 [22]
6 (+5%)	4.8 [2.2]	75 [24]
7 (+10%)	∞ [4.6]	∞ [73]
8 (+15%)	∞ [∞]	∞ [∞]
9 (+20%)	∞ [1.0]	∞ [1]

Notes: The first column reports estimates for the minimal value of $\bar{\phi}$ satisfying the conservative bound for excess movement in [Proposition 8\(ii\)](#). These estimates are translated to relative risk aversion $\bar{\gamma}$ using [Proposition 6](#), as shown in the second column. Point estimates are obtained by finding the value $\bar{\phi}$ such that $e_i^{\text{main}}(\bar{\phi}) = 0$ in (18). Confidence interval lower bounds are obtained by inverting a test for $\bar{\phi}$ using bootstrapped data; see Online Appendix C.7 for details. All estimates use conditional means of noise-adjusted excess movement.

for estimates that are unlikely to be affected by the possibility that $\phi < 1$. For all four of these states, the null is fully rejected. It is thus unlikely that [Assumption 3](#) is driving our results. Similarly, by [Proposition 9](#), our findings cannot in general be produced solely by miscalibrated priors. This indicates that our rejection must be arising either from excessively volatile physical belief revisions, or a strong violation of CTI.

6.5 Suggestive Evidence on the Correlates of RN Excess Movement

What ingredients would a model need in order to get closer to explaining the observed RN excess movement? To provide some preliminary guidance on this question, we conclude our empirical analysis with reduced-form evidence on the correlates of excess movement. [Table 4](#) shows results from a set of time-series regressions to this end. The dependent variable in all cases is the monthly average of noise-adjusted RN excess movement $X_{t,t+1,i,j}^*$, and all variables are normalized to have unit standard deviation. See Online Appendix C.8 for details on variable construction.

Table 4: Regressions for Monthly Average of RN Excess Movement

	(1)	(2)	(3)	(4)	(5)	(6)
Option Bid-Ask Spread	0.24 [0.15]					-0.03 [0.11]
Option Volume	0.07 [0.09]					-0.05 [0.10]
RN Belief Stream Length		0.28 [0.14]			0.16 [0.05]	0.18 [0.07]
VIX ²			0.33 [0.16]		0.58 [0.32]	0.62 [0.36]
Variance Risk Premium			0.38 [0.24]			
Volatility of Risk-Aversion Proxy			0.06 [0.10]			
Repurchase-Adjusted $ pd_t - \overline{pd} $				0.37 [0.12]	0.17 [0.05]	0.18 [0.06]
12-Month S&P 500 Return				0.30 [0.16]	0.53 [0.22]	0.53 [0.21]
R^2	0.08	0.08	0.28	0.14	0.37	0.37
Obs.	264	264	264	264	264	264

Notes: Heteroskedasticity- and autocorrelation-robust standard errors are in brackets, calculated using the equal-weighted periodogram estimator with $0.4 \text{ Obs.}^{2/3} = 16$ degrees of freedom, as recommended by Lazarus, Lewis, Stock, and Watson (2018). Dependent variable in all regressions is the empirical average of noise-adjusted $\widehat{\mathbb{E}}[X_{t,t+1,i,j}^*]$, for all available expiration dates and interior state pairs, across all trading dates within each given month. All variables are normalized to have unit standard deviation, and all regressions include a constant. See Online Appendix C.8 for details on variable construction.

The first column shows that proxies for option illiquidity and trading activity — namely, volume-weighted average monthly bid-ask spread in our options sample, and exponentially detrended option trading volume — are insignificant as predictors of X^* , which provides further evidence that option-market frictions are unlikely to be the main drivers of our results. Column (2) shows, however, that one economically meaningful factor specific to the option market *does* robustly predict excess movement: the average length of RN belief streams (i.e., \overline{T}_i for contracts i traded in the given month). As in Section 6.3, excess volatility seems to be concentrated at longer horizons, suggesting the possibility of overreaction to weak signals about events resolving relatively far in the future.³²

Column (3) considers volatility-related predictors. Excess movement has a significant positive relationship with the VIX; a weak positive relationship with the variance risk premium, calculated as VIX² minus realized variance following Bollerslev, Tauchen, and Zhou (2009); and essentially no relationship with the volatility of Bekaert, Engstrom, and Xu’s (2022) risk-aversion proxy.³³ This suggests that X^* comoves strongly with the quantity of market uncertainty; slightly less strongly

³²Interestingly, RN beliefs are *most volatile* (in terms of m^*) in the two weeks before maturity. But they also exhibit much higher uncertainty resolution and thus X^* close to 0, illustrating the importance of measuring *excess* movement.

³³Using daily data, they estimate time-varying relative risk aversion ra_t^{BEX} for a representative agent with habit-like preferences and preference shocks. We then take the sum of squared daily changes in ra_t^{BEX} in a given month.

with the price of this uncertainty; and not at all with the volatility of risk aversion, which can be thought of as a proxy for $\text{Var}(\phi_t)$. Column (4) considers proxies for (mis)valuation and return reversals, in the form of the absolute deviation of the log repurchase-adjusted price-dividend ratio (obtained from Nagel and Xu, 2022) and the trailing 12-month S&P return. Both are significantly positively related to X^* . As noted by Greenwood and Shleifer (2014), the trailing 12-month return predicts Gallup survey-based return expectations very well, suggesting a plausible role here for similar survey expectations to predict excess movement. Column (5) considers all four predictors from (1)–(4) that are significant separately at the 10% level and shows that they remain significant jointly, and explain 37% of the variation in X^* . Column (6) adds back the illiquidity and volume predictors; they remain insignificant, while the other predictors from (5) retain their significance.

Taken together, these results suggest again that RN excess movement is a real phenomenon: it is not attributable to option-specific frictions, and it comoves strongly with variables that are intuitively related to aggregate equity valuations and volatility. While we find limited evidence that variation in ϕ_t contributes substantially to X^* , this finding is specific to the particular risk-aversion proxy used here. But perhaps alternative models for time variation in ϕ_t are capable of generating significant excess movement, alongside alternative models of belief formation.

7. Conclusion

We derive new bounds on the admissible rational 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, and they allow for time variation in discount rates. Further, by using asset prices, we do not require direct measures of physical beliefs over future outcomes, and our bounds exploit intertemporal consistency requirements of rational beliefs without the need for the econometrician to know what agents' beliefs "should" be under RE.

When taken to the data using risk-neutral beliefs over the future return on the S&P 500 index, we find that these observed RN beliefs are so volatile that our bounds are routinely violated. This implies a violation of our joint assumptions. We provide evidence suggesting that RE violations are likely to be at least partly responsible for this violation, though we remain open to the possibility that alternative forms of CTI violations could play a part in explaining our results.

We believe that there are numerous feasible ways to make additional progress in identifying the specific causes of our bound violations. Conducting additional tests on the empirical correlates of excess movement, as well as generalizing our analysis to alternative asset classes, may provide useful additional information. Further, detailed data on changes in individual portfolios could allow for tests on the rationality of individual beliefs, which would help distinguish between micro and macro explanations for the observed excess movement in RN beliefs. We leave these possibilities to future work.

References

- AÏT-SAHALIA, Y. AND A. W. LO (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.
- ALVAREZ, F. AND U. J. JERMANN (2005): "Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth," *Econometrica*, 73, 1977–2016.
- AUGENBLICK, N. AND M. RABIN (2021): "Belief Movement, Uncertainty Reduction, and Rational Updating," *Quarterly Journal of Economics*, 136, 933–985.
- BANSAL, R. AND A. YARON (2000): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *NBER Working Paper 8059*.
- (2004): "Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles," *Journal of Finance*, 59, 1481–1509.
- 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.
- BASAK, S. (2000): "A Model of Dynamic Equilibrium Asset Pricing With Heterogeneous Beliefs and Extraneous Risk," *Journal of Economic Dynamics & Control*, 24, 63–95.
- BEKAERT, G., E. C. ENGSTROM, AND N. R. XU (2022): "The Time Variation in Risk Appetite and Uncertainty," *Management Science*, 68, 3975–4004.
- VAN BINSBERGEN, J. H., W. F. DIAMOND, AND M. GROTTERRIA (2022): "Risk-Free Interest Rates," *Journal of Financial Economics*, 143, 1–29.
- BLUME, L., T. COURY, AND D. EASLEY (2006): "Information, Trade and Incomplete Markets," *Economic Theory*, 29, 379–394.
- BOLLERSLEV, T., G. TAUCHEN, AND H. ZHOU (2009): "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.
- 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.
- 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.
- CHRISTOFFERSEN, P., S. HESTON, AND K. JACOBS (2013): "Capturing Option Anomalies with a Variance-Dependent Pricing Kernel," *Review of Financial Studies*, 26, 1962–2006.
- COCHRANE, J. H. (2011): "Presidential Address: Discount Rates," *Journal of Finance*, 66, 1047–1108.
- CONSTANTINIDES, G. M., J. C. JACKWERTH, AND A. SAVOV (2013): "The Puzzle of Index Option Returns," *Review of Asset Pricing Studies*, 3, 229–257.
- DE BONDT, W. F. M. AND R. THALER (1985): "Does the Stock Market Overreact?" *Journal of Finance*, 40, 793–805.
- DE LA O, R. AND S. MYERS (2021): "Subjective Cash Flow and Discount Rate Expectations," *Journal of Finance*, 76, 1339–1387.

- DEW-BECKER, I. AND S. GIGLIO (2016): "Asset Pricing in the Frequency Domain: Theory and Empirics," *Review of Financial Studies*, 29, 2029–2068.
- DRIESSEN, J., J. KOËTER, AND O. WILMS (2022): "Horizon Effects in the Pricing Kernel: How Investors Price Short-Term Versus Long-Term Risks," *Working Paper*.
- 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.
- (1991): "Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis," *Journal of Political Economy*, 99, 263–286.
- FAMA, E. F. (1991): "Efficient Capital Markets: II," *Journal of Finance*, 46, 1575–1617.
- GABAIX, X. (2012): "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," *Quarterly Journal of Economics*, 127, 645–700.
- GANDHI, M., N. J. GORMSEN, AND E. LAZARUS (2022): "Does the Market Understand Time Variation in the Equity Premium?" *Working Paper*.
- 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.
- GIGLIO, S. AND B. KELLY (2018): "Excess Volatility: Beyond Discount Rates," *Quarterly Journal of Economics*, 133, 71–127.
- GREENWOOD, R. AND A. SHLEIFER (2014): "Expectations of Returns and Expected Returns," *Review of Financial Studies*, 27, 714–746.
- HANSEN, L. P. AND R. JAGANNATHAN (1991): "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99, 225–262.
- JACKWERTH, J. C. (2000): "Recovering Risk Aversion from Option Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.
- KLEIDON, A. W. (1986): "Variance Bounds Tests and Stock Price Valuation Models," *Journal of Political Economy*, 94, 953–1001.
- KÜBLER, F. AND H. POLEMARCHAKIS (2017): "The Identification of Beliefs From Asset Demand," *Econometrica*, 85, 1219–1238.
- LAZARUS, E. (2019): "Horizon-Dependent Risk Pricing: Evidence from Short-Dated Options," *Working Paper*.
- LAZARUS, E., D. J. LEWIS, J. H. STOCK, AND M. W. WATSON (2018): "HAR Inference: Recommendations for Practice," *Journal of Business & Economic Statistics*, 36, 541–559.
- LEROY, S. F. AND R. D. PORTER (1981): "The Present-Value Relation: Tests Based on Implied Variance Bounds," *Econometrica*, 49, 555–574.
- LI, Z. M. AND O. B. LINTON (2022): "A ReMeDI for Microstructure Noise," *Econometrica*, 90, 367–389.
- MALZ, A. M. (2014): "A Simple and Reliable Way to Compute Option-Based Risk-Neutral Distributions," *Federal Reserve Bank of New York Staff Report No. 677*.
- 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.

- MILGROM, P. AND N. STOKEY (1982): “Information, Trade, and Common Knowledge,” *Journal of Economic Theory*, 26, 17–27.
- NAGEL, S. AND Z. XU (2022): “Asset Pricing with Fading Memory,” *Review of Financial Studies*, 35, 2190–2245.
- POLKOVNICHENKO, V. AND F. ZHAO (2013): “Probability Weighting Functions Implied in Options Prices,” *Journal of Financial Economics*, 107, 580–609.
- RADNER, R. (1979): “Rational Expectations Equilibrium: Generic Existence and the Information Revealed by Prices,” *Econometrica*, 47, 655.
- ROSS, S. (2015): “The Recovery Theorem,” *Journal of Finance*, 70, 615–648.
- SCHREINDORFER, D. AND T. SICHERT (2022): “Volatility and the Pricing Kernel,” *Working Paper*.
- SHILLER, R. J. (1981): “Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?” *American Economic Review*, 71, 421–436.
- STEIN, J. (1989): “Overreactions in the Options Market,” *Journal of Finance*, 44, 1011–1023.
- ZHANG, L., P. A. MYKLAND, AND Y. AÏT-SAHALIA (2005): “A Tale of Two Time Scales,” *Journal of the American Statistical Association*, 100, 1394–1411.

Appendix A. Selected Proofs of Main Results

Some preliminaries are useful before proceeding to the proofs. Start by defining the *RN measure* as

$$\mathbb{P}^*(H_T) \equiv \begin{cases} \mathbb{P}(H_T) \frac{\pi_0^*}{\pi_0} & \text{if } \pi_T(H_T) = 1 \\ \mathbb{P}(H_T) \frac{1-\pi_0^*}{1-\pi_0} & \text{if } \pi_T(H_T) = 0, \end{cases} \quad (\text{A.1})$$

where $\mathbb{P}(H_T)$ is the probability of observing history H_T under *DGP*. As shown in Lemma B.1 in the Online Appendix, (A.1) follows from the usual definition of the RN measure in a general asset-pricing setting. It suffices to think of \mathbb{P}^* as representing the change of measure that adjusts the frequency of each path of signal realizations such that a person with RN beliefs has RE. Define $\mathbb{E}^*[\cdot]$ to be the expectation under \mathbb{P}^* . Then the following results hold under [Assumption 1](#).

LEMMA A.1. For any H_T and θ , $\mathbb{P}^*(H_T|\theta) = \mathbb{P}(H_T|\theta)$.

LEMMA A.2. For any H_t , $\pi_t^*(H_t) = \mathbb{E}^*[\pi_{t+1}^*(H_{t+1})|H_t]$.

LEMMA A.3. For any *DGP*, $\mathbb{E}^*[X^*] = 0$.

Proof of Lemmas A.1–A.3. For the physical measure,

$$\begin{aligned} \mathbb{P}(H_T) &= \mathbb{P}(\theta = 1) \cdot \mathbb{P}(H_T|\theta = 1) + \mathbb{P}(\theta = 0) \cdot \mathbb{P}(H_T|\theta = 0) \\ &= \pi_0 \cdot \mathbb{P}(H_T|\theta = 1) + (1 - \pi_0) \cdot \mathbb{P}(H_T|\theta = 0), \end{aligned} \quad (\text{A.2})$$

where the second line uses that $\pi_0 = \mathbb{E}[\pi_T] = \mathbb{E}[\theta] = \mathbb{P}(\theta = 1)$ by our assumption of RE. Meanwhile, for the RN measure, we have from (A.1) and (A.2) that

$$\begin{aligned}\mathbb{P}^*(H_T) &= \frac{\pi_0^*}{\pi_0} \cdot \pi_0 \cdot \mathbb{P}(H_T|\theta = 1) + \frac{1 - \pi_0^*}{1 - \pi_0} \cdot (1 - \pi_0) \cdot \mathbb{P}(H_T|\theta = 0) \\ &= \pi_0^* \cdot \mathbb{P}(H_T|\theta = 1) + (1 - \pi_0^*) \cdot \mathbb{P}(H_T|\theta = 0).\end{aligned}\tag{A.3}$$

For any H_T such that $\pi_T = 1$, the definition of the RN measure (A.1) gives that $\mathbb{P}^*(H_T) = \frac{\pi_0^*}{\pi_0} \mathbb{P}(H_T)$, which implies $\mathbb{P}^*(\theta = 1) = \frac{\pi_0^*}{\pi_0} \mathbb{P}(\theta = 1)$. So using the definition of conditional probability, $\mathbb{P}^*(H_T|\theta = 1) = \mathbb{P}(H_T|\theta = 1)$. Similarly, $\mathbb{P}^*(H_T|\theta = 0) = \mathbb{P}(H_T|\theta = 0)$, proving Lemma A.1.

Given this, (A.3) becomes

$$\mathbb{P}^*(H_T) = \pi_0^* \cdot \mathbb{P}^*(H_T|\theta = 1) + (1 - \pi_0^*) \cdot \mathbb{P}^*(H_T|\theta = 0).$$

Summing over all possible H_T for which $\theta = 1$ gives $\pi_0^* = \mathbb{P}^*(\theta = 1)$, so \mathbb{P}^* is a valid probability distribution for which the law of iterated expectations (LIE) holds. Then noting $\mathbb{P}^*(\theta = 1) = \mathbb{E}^*[\theta] = \mathbb{E}^*[\pi_T] = \mathbb{E}^*[\pi_T^*]$, Lemma A.2 follows.

Finally, given Lemma A.2, the same proof for Lemma 1 (see Online Appendix B.1) implies that $\mathbb{E}^*[X^*] = 0$, as the expectation of RN movement under the RN measure must equal RN initial uncertainty. This proves Lemma A.3. \square

Lemma A.1 says that compared to the physical measure, the RN measure places higher likelihood of all signal histories resolving in state 1, but does so proportionally, so that likelihoods of signal histories *conditional on state 1* do not change. This implies that conditional expectations under the two respective measures are equal. Therefore, $\mathbb{E}^*[X^*|\theta] = \mathbb{E}[X^*|\theta]$ (see Online Appendix Lemma B.2 for the analogue of this result in the general asset-pricing setting). This implies

$$\begin{aligned}\mathbb{E}^*[X^*] &= \pi_0^* \cdot \mathbb{E}^*[X^*|\theta = 1] + (1 - \pi_0^*) \cdot \mathbb{E}^*[X^*|\theta = 0] \\ &= \pi_0^* \cdot \mathbb{E}[X^*|\theta = 1] + (1 - \pi_0^*) \cdot \mathbb{E}[X^*|\theta = 0] = 0,\end{aligned}\tag{A.4}$$

where the last equality applies Lemma A.3. For $\mathbb{E}[X^*]$, it is useful to similarly write

$$\mathbb{E}[X^*] = \pi_0 \cdot \mathbb{E}[X^*|\theta = 1] + (1 - \pi_0) \cdot \mathbb{E}[X^*|\theta = 0].\tag{A.5}$$

Now, recall that $\Delta \equiv \mathbb{E}[X^*|\theta = 0] - \mathbb{E}[X^*|\theta = 1]$. It will be useful to bound Δ :

LEMMA A.4. For any DGP, $\Delta \leq \pi_0^*$.

Proof of Lemma A.4. First write:

$$\begin{aligned}\Delta &\equiv \mathbb{E}^*[X^*|\theta = 0] - \mathbb{E}^*[X^*|\theta = 1] \\ &= \mathbb{E}^*[m^*|\theta = 0] - u_0^* - (\mathbb{E}^*[m^*|\theta = 0] - u_0^*)\end{aligned}$$

$$= \mathbb{E}^*[m^*|\theta = 0] - \mathbb{E}^*[m^*|\theta = 1]. \quad (\text{A.6})$$

Further, using [equation \(A.4\)](#),

$$\begin{aligned} 0 &= \pi_0^* \cdot \mathbb{E}[X^*|\theta = 1] + (1 - \pi_0^*) \cdot \mathbb{E}[X^*|\theta = 0] \\ &= \pi_0^* \cdot (\mathbb{E}[m^*|\theta = 1] - u_0^*) + (1 - \pi_0^*) \cdot (\mathbb{E}[m^*|\theta = 0] - u_0^*), \end{aligned}$$

so using the definition of u_0^* ,

$$\pi_0^* \cdot \mathbb{E}[m^*|\theta = 1] + (1 - \pi_0^*) \cdot \mathbb{E}[m^*|\theta = 0] = \pi_0^*(1 - \pi_0^*). \quad (\text{A.7})$$

Solving for $\mathbb{E}[m^*|\theta = 0]$ gives $\mathbb{E}[m^*|\theta = 0] = \pi_0^* - \frac{\pi_0^*}{1 - \pi_0^*} \cdot \mathbb{E}[m^*|\theta = 1]$. Using this in [\(A.6\)](#),

$$\Delta = \pi_0^* - \frac{\pi_0^*}{1 - \pi_0^*} \cdot \mathbb{E}[m^*|\theta = 1] - \mathbb{E}^*[m^*|\theta = 1] = \pi_0^* - \frac{1}{1 - \pi_0^*} \cdot \mathbb{E}[m^*|\theta = 1]. \quad (\text{A.8})$$

Given that $\frac{1}{1 - \pi_0^*} \geq 0$ and $\mathbb{E}[m^*|\theta = 1] \geq 0$, Δ is bounded above by π_0^* . \square

Proof of Proposition 1. Start from [equation \(A.5\)](#) and apply [equation \(A.4\)](#):

$$\begin{aligned} \mathbb{E}[X^*] &= \pi_0 \cdot \mathbb{E}[X^*|\theta = 1] + (1 - \pi_0) \cdot \mathbb{E}[X^*|\theta = 0] - 0 \\ &= \pi_0 \cdot \mathbb{E}[X^*|\theta = 1] + (1 - \pi_0) \cdot \mathbb{E}[X^*|\theta = 0] - (\pi_0^* \cdot \mathbb{E}[X^*|\theta = 1] + (1 - \pi_0^*) \cdot \mathbb{E}[X^*|\theta = 0]) \\ &= (\pi_0^* - \pi_0)(\mathbb{E}[X^*|\theta = 0] - \mathbb{E}[X^*|\theta = 1]) = (\pi_0^* - \pi_0)\Delta, \end{aligned} \quad (\text{A.9})$$

as stated. Then the second equality holds using [equation \(8\)](#) and the definition of Δ . \square

Proof of Proposition 2. From [Lemma A.4](#) above, we have $\Delta \leq \pi_0^*$. Further, [equation \(8\)](#) implies

$$\begin{aligned} \pi_0^* - \pi_0 &= \pi_0^* - \frac{\pi_0^*}{\pi_0^* + \phi(1 - \pi_0^*)} \\ &= \pi_0^* \left(1 - \frac{1}{\pi_0^* + \phi(1 - \pi_0^*)} \right) \geq 0, \end{aligned} \quad (\text{A.10})$$

where the last inequality uses $\pi_0^* + \phi(1 - \pi_0^*) \geq 0$ since $\phi \geq 1$. Using these two inequalities in the expression for $\mathbb{E}[X^*]$ in [\(A.9\)](#),

$$\mathbb{E}[X^*] = (\pi_0^* - \pi_0)\Delta \leq (\pi_0^* - \pi_0)\pi_0^*. \quad (\text{A.11})$$

Plugging in the expression for $\pi_0^* - \pi_0$ in [\(A.10\)](#) then gives [equation \(10\)](#). \square

Proof of Corollary 1. This is an immediate implication of [\(A.11\)](#) and $\pi_0 \geq 0$. \square

Proof of Corollary 2. As in [\(A.10\)](#), we have $\pi_0^* - \pi_0 \geq 0$. Using this in the equality in [\(A.11\)](#) alongside the assumption that $\Delta = \mathbb{E}^*[m^*|\theta = 0] - \mathbb{E}^*[m^*|\theta = 1] \leq 0$ gives $\mathbb{E}[X^*] \leq 0$. \square