

DISCUSSION PAPER SERIES

DP13857

SENTIMENT AND SPECULATION IN A MARKET WITH HETEROGENEOUS BELIEFS

Ian Martin and Dimitris Papadimitriou

**FINANCIAL ECONOMICS AND
MONETARY ECONOMICS AND
FLUCTUATIONS**



SENTIMENT AND SPECULATION IN A MARKET WITH HETEROGENEOUS BELIEFS

Ian Martin and Dimitris Papadimitriou

Discussion Paper DP13857

Published 13 July 2019

Submitted 01 July 2019

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **FINANCIAL ECONOMICS AND MONETARY ECONOMICS AND FLUCTUATIONS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Ian Martin and Dimitris Papadimitriou

SENTIMENT AND SPECULATION IN A MARKET WITH HETEROGENEOUS BELIEFS

Abstract

We present a dynamic model featuring risk-averse investors with heterogeneous beliefs. Individual investors have stable beliefs and risk aversion, but agents who were correct in hindsight become relatively wealthy; their beliefs are overrepresented in market sentiment, so “the market” is bullish following good news and bearish following bad news. Extreme states are far more important than in a homogeneous economy. Investors understand that sentiment drives volatility up, and demand high risk premia in compensation. Moderate investors supply liquidity: they trade against market sentiment in the hope of capturing a variance risk premium created by the presence of extremists.

JEL Classification: G11, G12, G13, G02, E44

Keywords: heterogeneous beliefs, Excess Volatility, sentiment, Speculation, target prices

Ian Martin - i.w.martin@lse.ac.uk
London School of Economics and CEPR

Dimitris Papadimitriou - D.Papadimitriou3@lse.ac.uk
London School of Economics

Acknowledgements

We thank Brandon Han, Nobu Kiyotaki, John Moore, Jeremy Stein, Andrei Shleifer, John Campbell, and seminar participants at LSE, Arizona State University, Harvard University, Shanghai Jiao Tong University (SAIF) and HKUST for comments. Ian Martin is grateful to the ERC for support under Starting Grant 639744.

Sentiment and speculation in a market with heterogeneous beliefs

Ian Martin Dimitris Papadimitriou*

June, 2019

Abstract

We present a dynamic model featuring risk-averse investors with heterogeneous beliefs. Individual investors have stable beliefs and risk aversion, but agents who were correct in hindsight become relatively wealthy; their beliefs are overrepresented in market sentiment, so “the market” is bullish following good news and bearish following bad news. Extreme states are far more important than in a homogeneous economy. Investors understand that sentiment drives volatility up, and demand high risk premia in compensation. Moderate investors supply liquidity: they trade against market sentiment in the hope of capturing a variance risk premium created by the presence of extremists.

*Both authors are at the London School of Economics. We thank Brandon Han, Nobu Kiyotaki, John Moore, Jeremy Stein, Andrei Shleifer, John Campbell, and seminar participants at LSE, Arizona State University, Harvard University, Shanghai Jiao Tong University (SAIF) and HKUST for comments. Ian Martin is grateful to the ERC for support under Starting Grant 639744.

In the short run, the market is a voting machine but in the long run it is a weighing machine.

—Attributed to Benjamin Graham by Warren Buffett.

In this paper, we study the effect of heterogeneity in beliefs on asset prices. We work with a frictionless dynamically complete market in which uncertainty evolves along a binomial tree. The model is populated by a continuum of risk-averse agents who differ in their beliefs about the probability of good news.

As a result, agents position themselves differently in the market. Optimistic investors make leveraged bets on the market; pessimists go short. If the market rallies, the wealth distribution shifts in favor of the optimists, whose beliefs become overrepresented in prices. If there is bad news, money flows to pessimists and prices more strongly reflect their pessimism going forward. At any point in time, one can define a representative agent who chooses to invest fully in the risky asset, with no borrowing or lending—our analog of Benjamin Graham’s “Mr. Market”—but the identity of the representative agent changes every period, with his or her beliefs becoming more optimistic following good news and more pessimistic following bad news. Thus market sentiment shifts constantly despite the stability of individual beliefs.

All agents understand the importance of sentiment and take it into account in pricing, so even moderate agents demand higher risk premia than they would in a homogeneous economy, as they correctly foresee that either good or bad news will be amplified by a shift in sentiment. The presence of sentiment induces speculation: agents take temporary positions, at prices they believe to be fundamentally incorrect, in anticipation of adjusting their positions in the future. In our model, speculation can act in either direction, driving prices up in some states and down in others. (In fact we show that for a broad class of assets, including the “lognormal” case in which asset payoffs are exponential in the number of up-moves, heterogeneity drives prices down and risk premia up.) This feature is emphasized by Keynes (1936, Chapter 12); in Harrison and Kreps (1978), by contrast, speculation only drives prices above fundamental value. In our setting it can also happen that an agent—even the representative agent—trades in one direction this period, in *certain* anticipation of reversing his or her position next period.

Extreme states are much more important than they are in a homogeneous-belief economy. Consider a stylized example. The riskless rate is 0%. A risky bond matures in 50 days, and will default (paying \$30 rather than the par value of \$100) only in the “bottom” state of the world, that is, only if there are 50 consecutive pieces of bad news. Investors’ beliefs about the probability, h , of an up-move are uniformly distributed between 0 and 1. Optimists therefore think default is almost impossible; a pessimistic agent with $h = 0.25$ thinks the default probability is less than 10^{-6} . Even an agent in the 95th percentile of pessimism, $h = 0.05$, thinks the default probability is less than 8%. Initially, the representative investor is the median agent, $h = 0.5$, who thinks the default probability is less than 10^{-15} . And yet we show that the bond trades at what might seem a remarkably low price: \$95.63. Moreover, almost half the agents—all agents with beliefs h below 0.478—initially go short at this price, though most will reverse their position within two periods.¹ The low price arises because all agents understand that if there is bad news next period, pessimists’ trades will have been profitable: their views will become overrepresented in the market, so the bond’s price will decline sharply *in the short run*. Only agents with $h < 0.006$ plan to stay short to the bitter end.

We start by solving the model in discrete time. Terminal payoffs are exogenously specified, and can be arbitrary, subject to being positive at every node so that expected utility is finite. We find the wealth distribution, prices, all agents’ investment decisions, and gross leverage at every node. We also characterize the cross-section of subjective perceptions of expected returns, volatilities, and Sharpe ratios. In general we do not take a stance on what the objectively correct beliefs are, nor even on whether there are objectively correct beliefs. But we can relate the equity premium perceived by the representative agent to an objectively measurable quantity, risk-neutral variance, that was proposed as a measure of the equity premium by [Martin \(2017\)](#).

After providing a formula for pricing in the general discrete-time case, we solve the model in a natural continuous-time limit in which the risky asset’s terminal payoffs are lognormally distributed. In this limit, the underlying asset price agrees with the corresponding price in the continuous-time model

¹Assuming there are two periods of bad news; if at any stage there is good news, the bond becomes riskless and disagreement vanishes.

of [Atmaz and Basak \(2018\)](#). As our framework is more tractable, we are able to study various issues that they do not (though, unlike us, they also price the underlying asset in the more general power utility case). We solve for agents’ subjective beliefs about expected returns and true (“ \mathbb{P} ”) volatility at all horizons; and for option prices at all maturities. Implied (“ \mathbb{Q} ”) volatility is higher at short horizons, due to the effect of sentiment; and lower at long horizons, due to the moderating influence of the terminal date at which pricing is dictated entirely by fundamentals. “In the short run, the market is a voting machine but in the long run it is a weighing machine.”

High implied volatility in the short run is also reflected in high physical measures of volatility (on which, in this continuous-time limit, all agents agree): there is no short-run variance risk premium. But physical measures of volatility decline more rapidly with horizon, so that there is a long-run variance risk premium.

As different investors have different beliefs but agree on asset prices, they have different stochastic discount factors (SDFs) whose properties help to reveal the interplay of beliefs, expected returns, and volatility. The volatility of any investor’s SDF equals the maximum Sharpe ratio that the investor perceives as achievable by trading dynamically in the market ([Hansen and Jagannathan, 1991](#)). By comparing this to the Sharpe ratio the investor perceives on the asset if it is statically held—or shorted—to maturity, we can measure the perceived benefit of dynamic trade (i.e., of speculation, as in our setting the only reason to trade dynamically is to exploit differences in beliefs: without belief heterogeneity, agents would hold a static position). We also solve for the entropies of investors’ SDFs ([Alvarez and Jermann, 2005](#)), which in our setting reveal the dollar value that different agents attach to being able to speculate.

All agents in our economy, particularly those with extreme beliefs, find speculation attractive. Extremists undertake conditional strategies that are increasingly aggressive as the market moves in their direction; in this sense, they are “long volatility.” We show that each investor can be thought of as having an investor-specific target price—the ideal outcome for the investor, given his or her beliefs and hence trading strategy—that can usefully be compared to what the investor expects to happen. The best possible outcome for

an extremist is that the market moves by even more than he or she expected.

Conversely, investors with more moderate beliefs are short volatility. Among moderates, there is a particularly interesting *gloomy investor* who perceives the lowest maximum attainable Sharpe ratio of all investors. The gloomy investor is slightly more pessimistic than the median investor and believes that the risky asset earns zero instantaneous risk premium. Nonetheless, he perceives that a sizeable Sharpe ratio can be attained via a short volatility position or, equivalently, via a contrarian market-timing strategy that exploits what he views as irrational exuberance on the up side and irrational pessimism on the down side. The gloomy investor can therefore be thought of as supplying liquidity to the extremists. He hopes to be proved right: in a sense that we make precise, the best outcome for him is the one that he expects.

As empirical researchers in finance often use high historical Sharpe ratios as a metric of success, we go on to study the properties of maximum-Sharpe-ratio strategies, and show that they feature short positions in out-of-the-money options. We view our exercise as a cautionary tale: while it is possible to earn very high Sharpe ratios via short option positions, these strategies are not remotely attractive to investors in our economy. Indeed, our investors would prefer to invest fully in cash than to rebalance, even slightly, toward a maximum-Sharpe-ratio strategy.

We make four key modelling choices. The first three are adopted from the model of [Geanakoplos \(2010\)](#) which inspired this paper. First, we assume that agents are dogmatic in their beliefs so that individuals do not experience changes in sentiment as time passes. The assumption is broadly consistent with one of the findings of [Giglio et al. \(2019\)](#), namely, that a substantial fraction of the variation in individual beliefs about expected returns, as reported in surveys of Vanguard clients, can be captured by individual fixed effects. If we allowed investors to learn over time, we believe that our mechanism would be amplified: that following good news, for example, optimistic agents would become relatively wealthier, as in our model, but all agents would also update their beliefs in an optimistic direction. (We formally prove that this intuition holds in a variant of the “risky bond” example described above. We assume that investors have heterogeneous priors about the true up-probability that they update fully rationally via Bayes’ rule, and show that the price of the

risky bond is *even lower* in the presence of learning.)

Second, we model uncertainty as evolving on a binomial tree so that the market is complete and agents can fully express their disagreement through trading. With an incomplete market, by contrast, agents may have strong differences in beliefs that are not revealed in prices. Market completeness also permits a clean interpretation of some of our results, as it generates a perfect correspondence between the cross-section and the time series. We exploit this fact to interpret our investors' trading behavior both in terms of conditional market-timing strategies and in terms of static positions in derivative securities.

Third, we allow for a continuum of beliefs, unlike papers including [Harrison and Kreps \(1978\)](#), [Scheinkman and Xiong \(2003\)](#), [Basak \(2005\)](#), [Banerjee and Kremer \(2010\)](#), and [Bhamra and Uppal \(2014\)](#). Aside from being realistic, this implies that the identities of the representative investor, and of the investor who chooses to sit out of the market entirely, are smoothly varying equilibrium objects that are determined endogenously in an intuitive and tractable way.

Fourth, and finally, our agents are risk-averse. In this respect we depart from several papers in the heterogeneous beliefs literature—including [Harrison and Kreps \(1978\)](#), [Scheinkman and Xiong \(2003\)](#) and [Geanakoplos \(2010\)](#)—that assume that agents are risk-neutral. Risk-neutrality simplifies matters in some respects, but complicates it in others. For example, short sales must be ruled out for equilibrium to exist. This is natural in some settings, but not if one thinks of the risky asset as representing, say, a broad stock market index. Moreover, the aggressive behavior of risk-neutral investors leads to extreme predictions: every time there is a down-move in the Geanakoplos model, all agents who are invested in the risky asset go bankrupt. From a technical point of view, short-sales constraints and risk-neutrality combine to give agents kinked indirect utility functions. Our agents have smooth indirect utility functions, and ultimately this is responsible for the tractability of our model and for our ability to study the dynamics described above.

1 Setup

We work in discrete time, with periods running from 0 to time T . Uncertainty evolves on a binomial tree, so that whatever the current state of the world, there are two possible successor states next period: “up” and “down.” There is a risky asset, whose payoffs at the terminal date T are specified exogenously. We normalize the net interest rate to 0%.

There is a unit mass of agents indexed by $h \in (0, 1)$. All agents have log utility over terminal wealth, zero time-preference rate, and are initially endowed with one unit of the risky asset, which we will think of as representing “the market.” Agent h believes that the probability of an up-move is h ; we often refer to h as the agent’s *belief*, for short. By working with the open interval $(0, 1)$, as opposed to the closed interval $[0, 1]$, we ensure that the investors’ beliefs are all absolutely continuous with respect to each other: that is, they all agree on *what events can possibly happen*. This means in particular that no investor will allow his or her wealth to go to zero in any state of the world.

The mass of agents with belief h follows a beta distribution governed by two parameters, α and β , such that the PDF is²

$$f(h) = \frac{h^{\alpha-1}(1-h)^{\beta-1}}{B(\alpha, \beta)}.$$

²The beta function $B(\cdot, \cdot)$ is defined by

$$B(x, y) = \int_{h=0}^1 h^{x-1}(1-h)^{y-1} dh.$$

If x and y are integers, then

$$B(x, y) = \frac{(x-1)!(y-1)!}{(x+y-1)!},$$

and more generally the beta function is related to the gamma function as follows:

$$B(x, y) = \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}.$$

We will repeatedly use basic facts about the beta function, such as that $B(x, y) = B(y, x)$, and that $B(x+1, y) = B(x, y) \cdot \frac{x}{x+y}$.




Figure 1: The distribution of beliefs for various choices of α and β .

The parameters α and β must be positive, but can otherwise be set arbitrarily. If $\alpha = \beta$ then the distribution of beliefs is symmetric with mean $1/2$. If $\alpha = \beta = 1$ then $f(h) = 1$, so that beliefs are uniformly distributed over $(0, 1)$; this is a useful case to keep in mind as one works through the algebra. The case $\alpha \neq \beta$ allows for asymmetric distributions with mean $\alpha/(\alpha + \beta)$ and variance $\alpha\beta/[(\alpha + \beta)^2(\alpha + \beta + 1)]$. Thus the distribution shifts toward 1 if $\alpha > \beta$ and toward 0 if $\alpha < \beta$, and beliefs are highly concentrated around the mean when α and β are large: if, say, $\alpha = 90$ and $\beta = 10$ then beliefs are concentrated around a mean of 0.9, with standard deviation 0.030. Figure 1 plots the distribution of beliefs, h , for a range of choices of α and β .

2 Equilibrium

The payoffs at terminal nodes of the binomial tree are specified exogenously. All agents have log utility over terminal wealth, so behave myopically; we can therefore consider each period in isolation. We start by taking next-period prices at the up- and down-nodes as given, and use these prices to determine the equilibrium price at the current node. This logic will ultimately allow us to solve the model by backward induction, and to express the price at time 0 in terms of the exogenous terminal payoffs. (See Result 2.)

Suppose, then, that the price of the risky asset will be either p_d or p_u next

period. Our problem, for now, is to determine the equilibrium price, p , at the current node; we assume that $p_d \neq p_u$ so that this pricing problem is nontrivial. Suppose also that agent h has wealth w_h at the current node. If he chooses to hold x_h units of the asset, then his wealth next period is $w_h - x_h p + x_h p_u$ in the up-state and $w_h - x_h p + x_h p_d$ in the down-state. So the portfolio problem is to solve

$$\max_{x_h} h \log [w_h - x_h p + x_h p_u] + (1 - h) \log [w_h - x_h p + x_h p_d].$$

The agent's first-order condition is therefore

$$x_h = w_h \left(\frac{h}{p - p_d} - \frac{1 - h}{p_u - p} \right). \quad (1)$$

The sign of x_h is that of $p - p_u$ for $h = 0$ and that of $p - p_d$ for $h = 1$. These must have opposite signs to avoid an arbitrage opportunity, so at every node there are some agents who are short and others who are long. The most optimistic agent³ levers up as much as possible without risking default, and correspondingly the most pessimistic agent takes on the largest short position possible that does not risk default if the good state occurs. For, the first-order condition (1) implies that as $h \rightarrow 1$, agent h holds $w_h/(p - p_d)$ units of stock. This is the largest possible position that does not risk default: to acquire it, the agent must borrow $w_h p/(p - p_d) - w_h = w_h p_d/(p - p_d)$. If the unthinkable (to this most optimistic agent!) occurs and the down state materialises, the agent's holdings are worth $w_h p_d/(p - p_d)$, which is precisely what the agent owes to his creditors.

It will often be convenient to think in terms of the risk-neutral probability of an up-move, p^* , defined by the property that the price can be interpreted as a risk-neutral expected payoff, $p = p^* p_u + (1 - p^*) p_d$. (There is no discounting,

³This is an abuse of terminology: there is no 'most optimistic agent' since h lies in the open set $(0, 1)$. More formally, this discussion relates to the behavior of agents in the limit as $h \rightarrow 1$. An agent for whom $h = 1$ would want to take arbitrarily large levered positions in the risky asset, so there is a behavioral discontinuity at $h = 1$ (and similarly at $h = 0$).

as the riskless rate is zero.) Hence

$$p^* = \frac{p - p_d}{p_u - p_d}.$$

In these terms, the first-order condition (1) becomes

$$x_h = \frac{w_h}{p_u - p_d} \frac{h - p^*}{p^*(1 - p^*)},$$

for example. An agent whose h equals p^* will have zero position in the risky asset: by the defining property of the risk-neutral probability, such an agent perceives that the risky asset has zero expected excess return.

Agent h 's wealth next period is therefore

$$w_h + x_h(p_u - p) = w_h(p_u - p_d) \frac{h}{p - p_d} = w_h \frac{h}{p^*} \quad (2)$$

in the up-state, and

$$w_h - x_h(p - p_d) = w_h(p_u - p_d) \frac{1 - h}{p_u - p} = w_h \frac{1 - h}{1 - p^*} \quad (3)$$

in the down-state. In either case, all agents' returns on wealth are linear in their beliefs. Moreover, this relationship (which is critical for the tractability of our model) applies at every node. It follows that person h 's wealth at the current node must equal

$$\lambda_{\text{path}} h^m (1 - h)^n$$

where λ_{path} is a constant that is independent of h but which can depend on the path travelled to get to the current node, which we have assumed has m up and n down steps.

As aggregate wealth is equal to the value of the risky asset—which is in unit supply—we must have

$$\int_0^1 \lambda_{\text{path}} h^m (1 - h)^n f(h) dh = p.$$

This enables us to solve for the value of λ_{path} :

$$\lambda_{\text{path}} = \frac{B(\alpha, \beta)}{B(\alpha + m, \beta + n)} p.$$

(This expression can be written in terms of factorials if α and β are integers: for example, if $\alpha = \beta = 1$ then $\lambda_{\text{path}} = \frac{(m+n+1)!}{m!n!} p$. See footnote 2.)

Substituting back, agent h 's wealth equals

$$w_h = \frac{B(\alpha, \beta)}{B(\alpha + m, \beta + n)} h^m (1 - h)^n p. \quad (4)$$

This is maximized by $h \equiv m/(m + n)$: the agent whose beliefs turned out to be most accurate ex post ends up richest.

The wealth distribution—that is, the fraction of aggregate wealth held by type- h agents—is a probability distribution over h . Specifically, it is the beta distribution with parameters $\alpha + m$ and $\beta + n$,

$$\frac{w_h f(h)}{p} = \frac{h^{\alpha+m-1} (1-h)^{\beta+n-1}}{B(\alpha + m, \beta + n)}. \quad (5)$$

We can now revisit Figure 1 in light of this fact. For the sake of argument, suppose that $\alpha = \beta = 1$ so that wealth is initially distributed uniformly across investors of all types $h \in (0, 1)$. If, by time 4, there have been $m = 1$ up- and $n = 3$ down-moves, then equation (5) implies that the new wealth distribution follows the line denoted $\alpha = 2, \beta = 4$. (Investors with h close to 0 or to 1 have been almost wiped out by their aggressive trades; the best performers are moderate pessimists with $h = 1/4$, whose beliefs happen to have been vindicated ex post.) At time 8, following three more up-moves and one down-move, the new wealth distribution is marked by $\alpha = \beta = 5$. And if by time 12 there have been a further four up-moves then the wealth distribution is marked by $\alpha = 9, \beta = 5$. The changing wealth distribution in this example illustrates a key feature of our model: at any point in time, wealth is concentrated in the hands of investors whose beliefs appear correct in hindsight.

Now we solve for the equilibrium price using the first-order condition

$$x_h = \underbrace{\frac{B(\alpha, \beta)}{B(\alpha + m, \beta + n)}}_{w_h} h^m (1 - h)^n p \left(\frac{h}{p - p_d} - \frac{1 - h}{p_u - p} \right).$$

The price p adjusts to clear the market, so that in aggregate the agents hold one unit of the asset:

$$\begin{aligned} \int_0^1 x_h f(h) dh &= \frac{p [(m + \alpha)(p_u - p) - (n + \beta)(p - p_d)]}{(m + n + \alpha + \beta)(p_u - p)(p - p_d)} \\ &= 1. \end{aligned}$$

It follows that

$$p = \frac{(m + \alpha)p_d p_u + (n + \beta)p_u p_d}{(m + \alpha)p_d + (n + \beta)p_u}. \quad (6)$$

Equivalently, the risk-neutral probability of an up-move must satisfy

$$p^* = \frac{(m + \alpha)p_d}{(m + \alpha)p_d + (n + \beta)p_u}$$

in equilibrium.

These results can usefully be interpreted in terms of *wealth-weighted beliefs*. For example, at time t , after m up-moves and $n = t - m$ down-moves, the wealth-weighted cross-sectional average belief, $H_{m,t}$, can be computed with reference to the wealth distribution (5):

$$H_{m,t} = \int_0^1 h \frac{w_h f(h)}{p} dh = \frac{m + \alpha}{t + \alpha + \beta}. \quad (7)$$

In these terms we can write

$$p^* = \frac{H_{m,t} p_d}{H_{m,t} p_d + (1 - H_{m,t}) p_u}. \quad (8)$$

It follows that

$$\frac{p_u}{p} = \frac{H_{m,t}}{p^*} \quad \text{and} \quad \frac{p_d}{p} = \frac{1 - H_{m,t}}{1 - p^*}. \quad (9)$$

Hence p^* is smaller than $H_{m,t}$ if $p_u > p_d$ and larger than $H_{m,t}$ if $p_u < p_d$: in




Figure 2: The range of beliefs in the investor population.

either case, risk-neutral beliefs are more pessimistic than the wealth-weighted average belief.

The share of wealth an agent of type h invests in the risky asset is

$$\frac{x_h p}{w_h} = p \left(\frac{h}{p - p_d} - \frac{1 - h}{p_u - p} \right) = \frac{h}{1 - \frac{p_d}{p}} - \frac{1 - h}{\frac{p_u}{p} - 1}.$$

This can be rewritten in a more compact form using (9):

$$\begin{aligned} \frac{x_h p}{w_h} &\stackrel{(9)}{=} \frac{h}{1 - \frac{1 - H_{m,t}}{1 - p^*}} - \frac{1 - h}{\frac{H_{m,t}}{p^*} - 1} \\ &= \frac{h - p^*}{H_{m,t} - p^*}. \end{aligned} \tag{10}$$

So the agent with $h = H_{m,t}$ can be thought of as *the representative agent*: by equation (10), this is the agent who chooses to invest her wealth fully in the market, with no borrowing or lending.

The identity of the representative investor therefore moves around over time, as does the identity of the investor with $h = p^*$ who chooses to hold his or her wealth fully in the bond. Figure 2 illustrates in the case $p_u > p_d$, so that $p^* < H_{m,t}$. Pessimistic investors with $h < p^*$ choose to short the risky asset; moderate investors with $p^* < h < H_{m,t}$ hold a balanced portfolio with long positions in both the bond and the risky asset; and optimistic investors with $h > H_{m,t}$ take on leverage, shorting the bond to go long the risky asset.

In a homogeneous economy in which all agents agree on the up-probability,

$h = H$, it is easy to check that

$$p^* = \frac{Hp_d}{Hp_d + (1 - H)p_u}. \quad (11)$$

Comparing equations (8) and (11), we see that for *short-run* pricing purposes our heterogeneous economy looks the same as a homogeneous economy featuring a representative agent with belief $H_{m,t}$. But as the identity of the representative agent changes over time, the similarity will disappear when we study the pricing of multi-period claims.

For future reference, the risk-neutral variance of the asset is

$$p^* \left(\frac{p_u}{p} \right)^2 + (1 - p^*) \left(\frac{p_d}{p} \right)^2 - 1 = \frac{(H_{m,t} - p^*)^2}{p^*(1 - p^*)}. \quad (12)$$

(The risk-neutral expectation of the asset's return is uninteresting: it must, by definition, equal the gross riskless rate.) Below, we will compare this quantity with subjective expected returns, motivated by the results of [Martin \(2017\)](#).

We can also use equation (10) to calculate the *leverage ratio* of investor h , which we define as the ratio of funds borrowed, $x_h p - w_h$, to wealth, w_h :

$$\frac{x_h p - w_h}{w_h} = \frac{h - H_{m,t}}{H_{m,t} - p^*}. \quad (13)$$

If $p_u > p_d$ then $p^* < H_{m,t}$, by (9); in this case equation (13) shows that people who are optimistic relative to the representative investor borrow from pessimists. We can define *gross leverage* as the total dollar amount these

optimists borrow,⁴ scaled by aggregate wealth:

$$\begin{aligned}
\frac{\int_{H_{m,t}}^1 (x_h p - w_h) f(h) dh}{p} &= \int_{H_{m,t}}^1 \frac{w_h f(h)}{p} \frac{x_h p - w_h}{w_h} dh \\
&= \int_{H_{m,t}}^1 \frac{w_h f(h)}{p} \frac{h - H_{m,t}}{H_{m,t} - p^*} dh \\
&= \frac{H_{m,t}^{m+\alpha} (1 - H_{m,t})^{n+\beta}}{(m + \alpha + n + \beta) B(\alpha + m, \beta + n) (H_{m,t} - p^*)}.
\end{aligned}$$

Conversely, if $p_u < p_d$ then optimists are lenders and pessimists borrowers. In either case, we can define gross leverage as the absolute value of the above expression,

$$\frac{H_{m,t}^{m+\alpha} (1 - H_{m,t})^{n+\beta}}{(m + \alpha + n + \beta) B(\alpha + m, \beta + n) |H_{m,t} - p^*|}. \quad (14)$$

Alternatively, scaling by the wealth of the borrowers and assuming that $p_u > p_d$ for simplicity, we define *borrower fragility*

$$\frac{\int_{H_{m,t}}^1 (x_h p - w_h) f(h) dh}{\int_{H_{m,t}}^1 w_h f(h) dh} = \frac{\int_{H_{m,t}}^1 \frac{w_h f(h)}{p} \frac{x_h p - w_h}{w_h} dh}{\int_{H_{m,t}}^1 \frac{w_h f(h)}{p} dh}, \quad (15)$$

which equals gross leverage divided by the fraction of wealth held by borrowers.

Figure 3 gives a numerical example with uniformly distributed beliefs (i.e., $\alpha = \beta = 1$) and $T = 2$. Terminal payoffs are chosen so that (i) $p_u/p_d = 2$ at the penultimate nodes and (ii) the asset would initially trade at a price of 1 in a homogeneous economy with $H = 1/2$. Initially, sentiment in the heterogeneous belief economy is the same— $H_{0,0} = 1/2$ —but the price is lower, at 0.96, because of the anticipated effect of *future* sentiment. If bad news arrives, money flows to pessimists. The representative agent and risk-neutral beliefs become more pessimistic and the price declines, accompanied by increases in gross leverage and borrower fragility.

⁴The total dollar amount borrowed by *all* investors is zero, as the riskless asset is in zero net supply.




Figure 3: At each node, \bar{p} denotes the price in a homogeneous economy with $H = 1/2$; p is the price in a heterogeneous economy with $\alpha = \beta = 1$; and p^* and $H_{m,t}$ indicate the risk-neutral probability of an up-move and the identity of the representative agent in the heterogeneous economy. (In the homogeneous economy, the risk-neutral probability of an up-move is $1/3$ at every node.) GL and BF denote gross leverage and borrower fragility, respectively.

2.1 Subjective beliefs

Investors disagree on the properties of the asset. Consider first moments. Agent h 's subjectively perceived expected excess return on the market is

$$\frac{hp_u + (1-h)p_d}{p} - 1 = \frac{(h-p^*)(p_u - p_d)}{p} = \frac{(h-p^*)(H_{m,t} - p^*)}{p^*(1-p^*)}. \quad (16)$$

Hence the share of wealth invested by agent h in the market (10) equals the ratio of the subjectively perceived expected excess return on the market (16) to (objectively defined) risk-neutral variance (12). In particular, risk-neutral variance reveals the expected excess return perceived by the representative agent, which is given by equation (16) with $h = H_{m,t}$.

The cross-sectional average expected excess return is

$$\frac{\left(\frac{\alpha}{\alpha+\beta} - p^*\right) (H_{m,t} - p^*)}{p^*(1-p^*)}, \quad (17)$$

which may be positive or negative. But as the wealth-weighted average belief equals the representative investor's belief (by (7)), the wealth-weighted cross-sectional average expected excess return must be positive: it equals

$$\int_0^1 \frac{w_h (h - p^*)(H_{m,t} - p^*)}{p p^*(1 - p^*)} f(h) dh = \frac{(H_{m,t} - p^*)^2}{p^*(1 - p^*)}. \quad (18)$$

Note that this quantity can also be interpreted as the expected excess return perceived by the representative agent $h = H_{m,t}$. The cross-sectional standard deviation of return expectations is

$$\sqrt{\frac{\alpha\beta}{(\alpha + \beta)^2 (\alpha + \beta + 1)}} \frac{|H_{m,t} - p^*|}{p^*(1 - p^*)}, \quad (19)$$

using the formula for the standard deviation of the beta distributed random variable h in equation (16).

Next we consider second moments. Person h 's subjectively perceived variance of the asset's return is

$$h \left(\frac{p_u}{p} \right)^2 + (1 - h) \left(\frac{p_d}{p} \right)^2 - \left(\frac{hp_u + (1 - h)p_d}{p} \right)^2 = \frac{h(1 - h)(H_{m,t} - p^*)^2}{p^{*2}(1 - p^*)^2},$$

and person h 's perceived Sharpe ratio is therefore

$$\frac{h - p^*}{\sqrt{h(1 - h)}},$$

which is increasing in h for all p^* .

The left panel of Figure 4 shows the risk premium perceived by different investors h in each of the possible states in the two-period example of Figure 3. Within any period, optimists perceive higher risk premia than pessimists. Across periods, most agents think expected returns go up at the down-node (“bad times”) and down at the up-node (“good times”), though the picture is complicated by the fact that volatility declines sharply over time in our two-period example, which exerts a downward influence on risk premia. To correct for this fact, the right panel plots the corresponding Sharpe ratios. All investors believe that Sharpe ratios are high in bad times and low in good




Figure 4: Left: The risk premium perceived by different agents in the initial state (\cdot), down state (d), and up state (u). Right: The corresponding Sharpe ratios.

times. But the representative investor (whose identity, in each state, is indicated by dots in the right panel of Figure 4) is more optimistic—and perceives a *higher* Sharpe ratio—in good times than in bad times. Thus “Mr. Market” disagrees with *every* individual investor about the behavior of Sharpe ratios in good and bad states.

The figure also shows that extremists perceive extreme Sharpe ratios, reflecting the fact that they are extremely confident in their beliefs and perceive that true volatility is close to zero. This might seem surprising in view of the general heuristic that second moments of returns are relatively easy to measure empirically, which suggests that there should be less room for disagreement about volatility. Indeed this is, to a large extent, an artefact of the simple two-period setting of the present example. When we move to continuous time in Section 4, the conventional view will reemerge in a particularly stark form: in a diffusion example, we will see that there is no room at all for disagreement about second moments. All agents will perceive the same volatility in returns, but will disagree about expected returns.

The variance risk premium perceived by investor h (that is, subjective minus risk-neutral variance) is equal to

$$\frac{(H_{m,t} - p^*)^2}{p^*(1 - p^*)} \left[\frac{h(1 - h)}{p^*(1 - p^*)} - 1 \right].$$

This is maximized—and weakly positive—for investor $h = 1/2$, and negative for agents with beliefs h that are further from $1/2$ than p^* is.

The wealth return for agent h is h/p^* in the up state and $(1-h)/(1-p^*)$ in the down state, as shown in equations (2) and (3). So agent h 's subjective expected excess return on own wealth is

$$\frac{h^2}{p^*} + \frac{(1-h)^2}{1-p^*} - 1 = \frac{(h-p^*)^2}{p^*(1-p^*)}.$$

All agents expect to earn a nonnegative excess return on wealth, though they have very different positions. Only agent $h = p^*$ chooses to take no risk, fully invests in the bond, and so correctly anticipates zero excess return.

2.2 A risky bond

The dynamic that drives our model is particularly stark in the “risky bond” example outlined in the introduction. Suppose that the terminal payoff is 1 in all states apart from the very bottom one, in which the payoff is ε ; the price of the asset is therefore 1 as soon as an up-move occurs. Writing p_t for the price at time t following t consecutive down-moves we have, from equation (6),

$$p_t = \frac{\alpha p_{t+1} + (t + \beta)p_{t+1}}{\alpha p_{t+1} + t + \beta}.$$

Defining $y_t \equiv 1/p_t - 1$, this can be rearranged as

$$y_t = \frac{\beta + t}{\alpha + \beta + t} y_{t+1}. \quad (20)$$

We can interpret y_t as the inducement to invest in the risky asset at time t , following t consecutive down-moves: it is the realized excess return on the asset if there is an up-move from t to $t+1$. Equation (20) determines the rate at which this inducement must rise in equilibrium.

Solving equation (20) forward,

$$y_t = \frac{(\beta + t)(\beta + t + 1) \cdots (\beta + T - 1)}{(\alpha + \beta + t)(\alpha + \beta + t + 1) \cdots (\alpha + \beta + T - 1)} y_T, \quad (21)$$

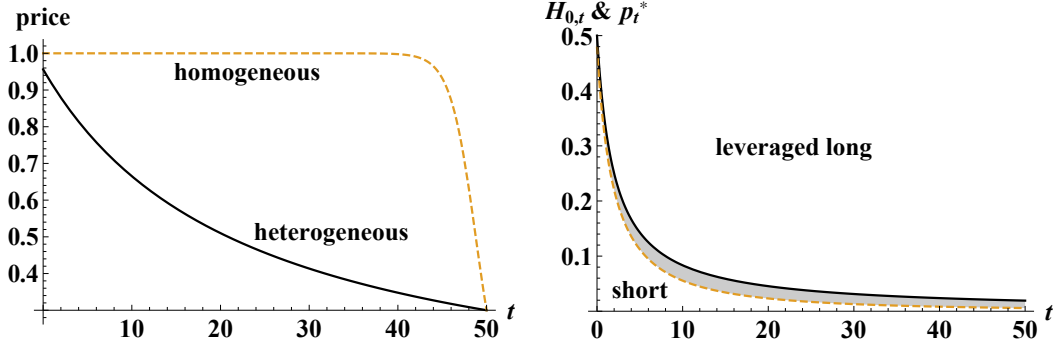


Figure 5: Left: The risky bond’s price over time in the heterogeneous and homogeneous economies following consistently bad news. Right: $H_{0,t}$ reveals the identity of the representative agent at time t following consistently bad news. Investors who are more optimistic, $h > H_{0,t}$, have leveraged long positions in the risky bond. The risk-neutral probability reveals the identity of the investor who is fully invested in the riskless bond at time t , with zero position in the risky bond. Investors who are more pessimistic, $h < p_t^*$, are short the risky bond. Investors with $p_t^* < h < H_{0,t}$ (shaded) are long both the risky and the riskless bond.

and the terminal condition dictates that $y_T = (1 - \varepsilon)/\varepsilon$. Thus, finally,

$$p_t = \frac{1}{1 + \frac{\Gamma(\beta+T)\Gamma(\alpha+\beta+t)}{\Gamma(\beta+t)\Gamma(\alpha+\beta+T)} \frac{1-\varepsilon}{\varepsilon}}.$$

If $\alpha = \beta = 1$, we can simplify further, to

$$p_t = \frac{1}{1 + \frac{1+t}{1+T} \frac{1-\varepsilon}{\varepsilon}}. \quad (22)$$

We can calculate the risk-neutral probability of an up-move at time t , which we (temporarily) denote by p_t^* , by applying (9) with $p = p_t$, $p_u = 1$, and $p_d = p_{t+1}$ to find that

$$p_t^* = H_{0,t} p_t = \frac{\alpha p_t}{\alpha + \beta + t}. \quad (23)$$

Figure 5 illustrates these calculations in the example described in the introduction, with $T = 50$ periods to go, and a recovery value of $\varepsilon = 0.30$. The panels show how the price and risk-neutral probability evolve if bad news arrives each period. Initially, the bond trades at what might seem a remarkably

low price of 0.9563.

By contrast, in a homogeneous economy with $H = 1/2$ the price, p_t , and risk-neutral probability, p_t^* , following t down-moves would be

$$p_t = \frac{1}{1 + \frac{1-\varepsilon}{\varepsilon}0.5^{T-t}} \quad \text{and} \quad p_t^* = \frac{p_t}{2},$$

respectively. Thus with homogeneous beliefs the bond price is approximately 1, and the risk-neutral probability of an up-move is approximately 1/2, until shortly before the bond's maturity.

From the perspective of time 0, the risk-neutral probability of default—call it δ^* —satisfies

$$p_0 = 1 - \delta^* + \delta^*\varepsilon, \quad \text{so} \quad \delta^* = \frac{1 - p_0}{1 - \varepsilon}.$$

In the homogeneous case, therefore,

$$\delta^* = \frac{1}{1 + \varepsilon(2^T - 1)} = O(2^{-T});$$

and in the heterogeneous case with $\alpha = 1$,

$$\delta^* = \frac{1}{1 + \varepsilon T} = O(1/T).$$

There is a qualitative difference between the homogeneous economy, in which default is exponentially unlikely, and the heterogeneous economy, in which default is only polynomially unlikely.⁵

To understand pricing in the heterogeneous economy, it is helpful to think through the portfolio choices of individual investors. We use equations (5), (7), and (10), together with the prices and risk-neutral probabilities given in (22) and (23) above, to find investors' holdings of the risky asset at each node.

The median investor, $h = 0.5$, thinks the probability that the bond will default—i.e., that the price will follow the path shown in Figure 5 all the way to the end—is $2^{-50} < 10^{-15}$. Even so, he believes the price is right at time

⁵This holds more generally for any $\alpha = \beta > 1$: it is easy to show that $\delta^* = O(T^{-\alpha})$ by Stirling's formula. It is also true if $\varepsilon > 1$, i.e. in the 'lottery ticket' case. Then, δ^* is interpreted as the probability of the lottery ticket paying off, which is exponentially small in the homogeneous economy but only polynomially small in the heterogeneous belief economy.




Figure 6: Left: The number of units of the risky bond held by different agents, $x_{h,t}$, plotted against time. Right: The evolution of leverage for the median investor under the optimal dynamic and static strategies. Both panels assume bad news arrives each period.

zero (in the sense that he *is* the representative agent) because of the short-run impact of sentiment. Meanwhile, a modestly pessimistic agent with $h = 0.25$ will choose to short the bond at the price of 0.9563—and will remain short at time $t = 1$ before reversing her position at $t = 2$ —despite believing that the bond’s default probability is less than 10^{-6} . (Recall from equation (10) that p_t^* is the belief of the agent who is neither long nor short the asset. More optimistic agents, $h > p^*$, are long, and more pessimistic agents, $h < p^*$, are short.) Following a few periods of bad news, almost all investors are long, but the most pessimistic investors have become extraordinarily wealthy.

The left panel of Figure 6 shows the holdings of the risky asset for a range of investors with different beliefs, along the trajectory in which bad news keeps on coming. The optimistic investor $h = 0.75$ starts out highly leveraged so rapidly loses almost all his money. The median investor, $h = 0.5$, initially invests fully in the risky bond without taking on leverage. If bad news arrives, this investor takes on leverage in order to be able to increase the size of her position despite her losses; after about 10 periods, the investor is almost completely wiped out. Moderately bearish investors start out short. For example, investor $h = 0.25$ starts out short about 10 units of the bond, despite believing that the probability it defaults is less than one in a million, but reverses her position after two down-moves. Investor $h = 0.01$, who thinks that there is more than a 60% chance of default, is initially extremely short but eventually reverses

position as still more bearish investors come to dominate the market.

The right panel of Figure 6 shows how the median investor's leverage changes over time if he follows the optimal dynamic and static strategies. If forced to trade statically, his leverage ratio is initially 0.457. This seemingly modest number is dictated by the requirement that the investor avoid bankruptcy at the bottom node (and in fact the leverage of all investors with $h \geq 0.2$ is visually indistinguishable at the scale of the figure). If the median investor can trade dynamically, by contrast, the optimal strategy is, initially, to invest fully in the risky bond without leverage. Subsequently, however, optimal leverage rises fast. Thus the dynamic investor keeps his powder dry by investing cautiously at first but then aggressively exploiting further selloffs.

All investors perceive themselves as better off if able to trade dynamically, of course. In Appendix A we analytically characterize the perceived advantage of dynamic versus static trade as a function of each investor's belief h .

The volume of trade (in terms of the number of units of the risky asset transacted) in the transition from time t to time $t + 1$ is⁶

$$\frac{1}{2} \int_0^1 \left| \frac{(1-h)^t}{\frac{1}{1+t}} \frac{h-p_t^*}{H_{0,t}-p_t^*} - \frac{(1-h)^{t+1}}{\frac{1}{2+t}} \frac{h-p_{t+1}^*}{H_{0,t+1}-p_{t+1}^*} \right| dh = \frac{4(1+t)^{1+t}}{(3+t)^{3+t}} \left(1+t + \frac{1+\varepsilon T}{1-\varepsilon} \right),$$

while gross leverage and borrower fragility, calculated from (14) and (15), equal

$$\left(\frac{1+t}{2+t} \right)^{2+t} \left(1 + \frac{1+T}{1+t} \frac{\varepsilon}{1-\varepsilon} \right) \quad \text{and} \quad \left(\frac{1+t}{2+t} \right) \left(1 + \frac{1+T}{1+t} \frac{\varepsilon}{1-\varepsilon} \right)$$

respectively.

The left panel of Figure 7 shows the time series of volume, gross leverage, and borrower fragility, assuming bad news arrives each period. (If good news arrives at any stage, volume drops permanently to zero.) In this stylized example there is a burst of trade at first: volume substantially exceeds the total supply of the asset initially, as agents with extreme views undertake highly leveraged trades, but declines rapidly over time as wealth becomes concentrated in the hands of investors with similar beliefs. The right panel shows the corresponding series if $\varepsilon = 0.9$. In this case disagreement generates

⁶We include the factor of 1/2 to avoid double-counting.




Figure 7: Volume (solid), gross leverage (dashed), and borrower fragility (dotted) over time, with $\varepsilon = 0.3$ (left) or $\varepsilon = 0.9$ (right). Heterogeneous case only: volume is zero in the homogeneous economy.

more aggressive trading, and more volume, because the relative safety of the asset permits agents to take on more leverage: extremists on both sides of the market are “picking up nickels in front of a steamroller.”

2.2.1 Bayesian learning in the risky bond example

We briefly generalize our previous analysis to allow investors to update their beliefs over time using Bayes’ rule. We continue to index investors by $h \in (0, 1)$, and we continue to assume that the distribution of h follows a beta distribution with parameters α and β . Now, however, investor h ’s prior belief is that the probability of an up-move is $\tilde{h} \sim \text{Beta}(\zeta h, \zeta(1-h))$. We have $\mathbb{E}\tilde{h} = h$, so that agent h ’s mean belief is h ; and $\text{var}\tilde{h} = h(1-h)/(1+\zeta)$, where $\zeta > 0$ is a constant that controls the uncertainty of investors about the true probability of an up-move. For large ζ , investor h ’s prior is sharply peaked around h , and the setting nests the dogmatic case we consider elsewhere, because as $\zeta \rightarrow \infty$ agent h becomes certain that the up-probability is h .

The following result formally confirms—for the risky bond example—our intuition that learning amplifies the effect of heterogeneity in beliefs.

Result 1. *For any α and β , and for any $\zeta > 0$, the initial price of the bond is lower with learning than without.*

Proof. See Appendix A. □




Figure 8: An example with late resolution of uncertainty. Heterogeneous-economy price (p), homogeneous-economy price (\bar{p}), and the cross-sectional average perceived excess return in the heterogeneous economy (ER).

2.3 An example with late resolution of uncertainty

Consider an example with an odd number of periods, T , and $\alpha = \beta = 1$; and let $0 < \varepsilon < 1$. If there have been an even number of up-moves at time T , the asset pays off $\frac{1}{1+\varepsilon}$; if there have been an odd number of up-moves, the asset pays $\frac{1}{1-\varepsilon}$.

In the homogeneous economy with $H = 1/2$, the asset trades at a price of 1 in every node, and at every period, until the terminal payoff: it is therefore riskless until the final period.

In the heterogeneous economy it follows immediately from Result 2, below, that the asset also trades at 1 initially. But the asset is now volatile: although the payoff of the asset is up in the air until the very last period, the effect of sentiment ripples back so that the asset is volatile throughout its lifetime, and its price therefore embeds a risk premium.⁷

Figure 8 shows an example with $T = 3$ and $\varepsilon = 1/2$. In a homogeneous

⁷There is also an equilibrium in which the asset's price is 1 until time $T - 1$, as in the homogeneous economy. Then the market is incomplete, and agents have no means of betting against one another. But this equilibrium is not robust to vanishingly small perturbations of the terminal payoffs, which would restore market completeness.

economy, the asset's price is completely stable until immediately before the terminal date. In the heterogeneous economy, the asset's price is volatile, and it embeds a time-varying risk premium.

3 The general case

Write $z_{m,t} = 1/p_{m,t}$, where m is the number of up moves that have taken place by time t . Equation (6) implies that the following recurrence relation holds at each node:

$$z_{m,t} = H_{m,t}z_{m+1,t+1} + (1 - H_{m,t})z_{m,t+1}. \quad (24)$$

That is, the price at each node is the weighted average harmonic mean of the next-period prices, with weights given by the beliefs of the currently representative agent. By backward induction, $z_{0,0}$ is a linear combination of the reciprocals of the terminal payoffs $z_{0,T}, z_{1,T}, \dots, z_{T,T}$:

$$z_{0,0} = \sum_{m=0}^T c_m z_{m,T}. \quad (25)$$

Pricing is not path-dependent in our economy. Indeed, as

$$\underbrace{\frac{m + \alpha}{t + \alpha + \beta}}_{H_{m,t}} \underbrace{\frac{t - m + \beta}{t + 1 + \alpha + \beta}}_{1 - H_{m+1,t+1}} = \underbrace{\frac{t - m + \beta}{t + \alpha + \beta}}_{1 - H_{m,t}} \underbrace{\frac{m + \alpha}{t + 1 + \alpha + \beta}}_{H_{m,t+1}},$$

the risk-neutral probability of going up and then down (from any starting node) equals the risk-neutral probability of going down and then up. That is, from (9),

$$p_{m,t}^*(1 - p_{m+1,t+1}^*) = (1 - p_{m,t}^*)p_{m,t+1}^*.$$

These observations allow us to find a general pricing formula that applies for arbitrary terminal payoffs $p_{m,T}$. (The payoffs must be positive so that the expected utility of any agent is well defined.) The proof of the result, and all subsequent results, is in the Appendix.

Result 2. *If the risky asset has terminal payoffs $p_{m,T}$ at time T (for $m =$*

$0, \dots, T$), then its initial price is

$$p_0 = \frac{1}{\sum_{m=0}^T \frac{c_m}{p_{m,T}}}, \quad (26)$$

where

$$c_m = \binom{T}{m} \frac{B(\alpha + m, \beta + T - m)}{B(\alpha, \beta)}. \quad (27)$$

The time 0 price of the Arrow–Debreu security that pays off if there have been m up-moves by time T is

$$q_m^* = c_m \frac{p_0}{p_{m,T}}.$$

The coefficients c_m have a so-called beta-binomial distribution, $BB(T, \alpha, \beta)$. This is a binomial distribution with a random probability of success in each trial given by a $Beta(\alpha, \beta)$ distribution.⁸ In the Appendix, we generalize equation (25) and Result 2 to price the risky asset at any node.

As a corollary of Result 2, we can find the effect of belief heterogeneity on prices for a broad class of assets.

Result 3. *If beliefs are symmetric, and the risky asset has terminal payoffs such that $\frac{1}{p_{m,T}}$ is convex if viewed as a function of m , then the asset’s time 0 price is decreasing in the degree of belief heterogeneity. In particular, it is sufficient (though not necessary) that $\log p_{m,T}$ be weakly concave for the asset’s price to be decreasing in the degree of belief heterogeneity.*

Result 3 applies if the terminal payoff is concave in m . But it also applies for some convex payoffs. If, for example, the asset’s payoffs increase or decrease geometrically in m , then the log payoffs are linear in m , so that the concavity condition (just) holds. We provide a more extensive analysis of this case in the next section.

⁸In fact, c_m can be interpreted as the *cross-sectional average* (among investors) perceived probability of reaching node (m, T) .

4 A diffusion limit

We consider a natural continuous time limit by allowing the number of periods to tend to infinity and specifying geometrically increasing terminal payoffs. This is the setting of [Cox, Ross and Rubinstein \(1979\)](#), in which the Black–Scholes formula emerges in the corresponding limit with homogeneous beliefs. We are able to solve for the asset price, risk-neutral probabilities, the volatility term structure, individuals’ trading strategies, and other quantities of interest.

Denote by $2N$ the total number of periods (corresponding to time T).⁹ We assume that

$$p_{m,T} = e^{2\sigma\sqrt{\frac{T}{2N}}(m-N)}. \quad (28)$$

As we will see, σ can be interpreted as the volatility of terminal payoffs (on which all agents will turn out to agree). If we set $\lambda = e^{\sigma\sqrt{\frac{T}{2N}}}$, then we see that $p_{m,2N} = \lambda^m (\frac{1}{\lambda})^{2N-m}$, where $\lambda = u = d^{-1}$ and u, d are the up and down percentage movements of the stock price in the Cox–Ross–Rubinstein model. If we now set $\psi = \frac{m-N}{\sqrt{N}}$ then $p_{m,T} = e^{\sigma\sqrt{2T}\psi}$. From the perspective of each agent, m has a binomial distribution; we show, in the Appendix, that in the limit as $N \rightarrow \infty$, ψ has an asymptotic normal distribution from the perspective of each investor.

We use [Result 2](#) to price the asset at each node of the tree, then take the limit as N tends to infinity. As the number of up/down steps increases with N , the extent of disagreement over any individual step must decline to generate sensible limiting results—that is, we allow the parameters α, β , which control the belief dispersion in the market, to tend to infinity with N . In particular we will write $\alpha = \theta N + \eta\sqrt{N}$ and $\beta = \theta N - \eta\sqrt{N}$. Small values of θ correspond to a high belief heterogeneity, while the limit $\theta \rightarrow \infty$ corresponds to the homogeneous case; we will refer to $\frac{1}{\theta}$ as capturing the degree of heterogeneity in the market. The level of optimism in the market is captured by η .

To be more precise, we will introduce a cross-sectional expectation operator $\widetilde{\mathbb{E}}[\cdot]$. So, for example, the cross-sectional mean of h satisfies $\widetilde{\mathbb{E}}[h] = \frac{\alpha}{\alpha+\beta} = \frac{1}{2} + \frac{\eta}{2\theta\sqrt{N}}$ and $\widetilde{\text{var}}[h] = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)} = \frac{1}{8\theta N+1} + O(\frac{1}{N^2})$. As $\widetilde{\mathbb{E}}\left[\mathbb{E}^{(h)}[\psi]\right] = \frac{\eta}{\theta}$,

⁹The choice of an even number of periods is unimportant, but it simplifies the notation in some of our proofs.

we can interpret η as controlling the cross-sectional mean expected terminal payoff.

In the work of Cox, Ross, and Rubinstein, the central limit theorem is used to approximate a binomial distribution with a normal random variable. A similar, though slightly more convoluted, situation arises in our setting. The argument starts by rewriting equation (25) as¹⁰

$$p_0^{-1} = \mathbb{E}_m \left[e^{-\sigma\sqrt{2T}\frac{m-N}{\sqrt{N}}} \right] = M_\psi \left(-\sigma\sqrt{2T}\psi \right).$$

where we write \mathbb{E}_m to indicate that the expectation is taken over m which, by Result 2, can be viewed as a random variable following the beta-binomial distribution with parameters $2N$, α , and β ; and $M_\psi(\cdot)$ denotes the moment generating function (MGF) of $\psi = \frac{m-N}{\sqrt{N}}$. As ψ is asymptotically normal by a result of Paul and Plackett (1978), $M_\psi(\cdot)$ converges to the MGF of a Normal distribution—a known, and simple, function. We provide full details in the Appendix.

Result 4. *The price of the asset at time 0 is given by:*

$$p_0 = \exp \left(\frac{\eta}{\theta} \sigma \sqrt{2T} - \frac{\theta + 1}{2\theta} \sigma^2 T \right). \quad (29)$$

If $\eta = 0$, so that the cross-sectional distribution of beliefs is symmetric around $h = 1/2$, then the price at time 0 is decreasing in the degree of heterogeneity, θ^{-1} , consistent with Result 3. But if the cross-sectional average belief is sufficiently optimistic—that is, if η is sufficiently positive—then the price may be increasing in the heterogeneity of beliefs.

We now study what this price implies for different agents' expectations about returns. We parametrize an agent by the number of standard deviations, z , by which his or her belief deviates from the mean: $h = \tilde{\mathbb{E}}[h] + z\sqrt{\widetilde{\text{var}}[h]} \approx \frac{1}{2} + \frac{\eta}{2\theta\sqrt{N}} + \frac{z}{\sqrt{8\theta N}}$.¹¹ Thus an agent with $z = 2$ is two standard deviations more optimistic than the mean agent. When we use this parametrization, we write

¹⁰From now on we suppress the explicit dependence of price on state in our notation and write, for example, p_0 rather than $p_{0,0}$.

¹¹Note that $\tilde{\mathbb{E}}[h] = \frac{\alpha}{\alpha+\beta} = \frac{1}{2} + \frac{\eta}{2\theta\sqrt{N}}$ and $\widetilde{\text{var}}[h] = \frac{\alpha\beta}{(\alpha+\beta)^2(\alpha+\beta+1)} = \frac{1}{8\theta N+1} + O(\frac{1}{N^2})$. The lower order terms, $O(1/N^2)$, will not play any role as N approaches infinity.

superscripts z rather than h (for example, $\mathbb{E}^{(z)}$ rather than $\mathbb{E}^{(h)}$).

Result 5. *The return of the asset from time 0 to time t , from the perspective of agent $h = \tilde{\mathbb{E}}[h] + z\sqrt{\widetilde{\text{var}}[h]}$ has a lognormal distribution with*

$$\begin{aligned}\mathbb{E}^{(z)} \log R_{0 \rightarrow t} &= \frac{\theta + 1}{\theta + \frac{t}{T}} \left(\frac{z\sigma}{\sqrt{\theta T}} + \frac{\theta + 1}{2\theta} \sigma^2 \right) t \\ \text{var}^{(z)} \log R_{0 \rightarrow t} &= \left(\frac{\theta + 1}{\theta + \frac{t}{T}} \right)^2 \sigma^2 t.\end{aligned}$$

The expected return on the asset follows immediately.

Result 6. *The (annualized) expected return of the asset from 0 to t is*

$$\frac{1}{t} \log \mathbb{E}^{(z)} R_{0 \rightarrow t} = \frac{\theta + 1}{\theta + \frac{t}{T}} \left[\frac{z\sigma}{\sqrt{\theta T}} + \frac{\theta + 1}{2\theta} \frac{2\theta + \frac{t}{T}}{\theta + \frac{t}{T}} \sigma^2 \right].$$

In particular, the instantaneous expected return is

$$\lim_{t \rightarrow 0} \frac{1}{t} \log \mathbb{E}^{(z)} R_{0 \rightarrow t} = \frac{\theta + 1}{\theta} \frac{z\sigma}{\sqrt{\theta T}} + \left(\frac{\theta + 1}{\theta} \right)^2 \sigma^2$$

and the expected return to maturity is

$$\frac{1}{T} \log \mathbb{E}^{(z)} R_{0 \rightarrow T} = \frac{z\sigma}{\sqrt{\theta T}} + \frac{2\theta + 1}{2\theta} \sigma^2. \quad (30)$$

Thus, although different agents perceive different expected returns, all agents agree on the volatility of log returns (though not on the volatility of simple returns).

We note in passing that if dynamic trade were shut down entirely, so that all agents had to trade once at time 0 and then hold their positions statically to time T , then equilibrium would not exist in the limit.¹²

¹²To see this, write ψ_z for the share of wealth invested by agent z in the risky asset. Given any positive time 0 price, $R_{0 \rightarrow T}$ is lognormal from every agent's perspective by Result 5 (which applies even in the static case at horizon T , because the terminal payoff is specified exogenously). Hence we must have $\psi_z \in [0, 1]$ for all z to avoid the possibility of terminal wealth becoming negative if $R_{0 \rightarrow T}$ is sufficiently close to zero or is sufficiently large. Market clearing requires that $\psi_z = 1$ on average across agents, so we must in fact have $\psi_z = 1$ for

Result 7. Recall that $\tilde{\mathbb{E}}$ is the cross-sectional expectation operator. The cross-sectional mean (or median) expected return is

$$\tilde{\mathbb{E}} \left[\frac{1}{t} \log \mathbb{E}^{(z)} R_{0 \rightarrow t} \right] = \frac{(\theta + 1)^2 \left(\theta + \frac{t}{2T} \right)}{\theta \left(\theta + \frac{t}{T} \right)^2} \sigma^2.$$

Disagreement is the standard deviation of expected returns $\frac{1}{t} \log \mathbb{E}^{(z)} R_{0 \rightarrow t}$:

$$\sqrt{\widetilde{\text{var}} \left[\frac{1}{t} \log \mathbb{E}^{(z)} R_{0 \rightarrow t} \right]} = \frac{\theta + 1}{\theta + \frac{t}{T}} \frac{\sigma}{\sqrt{\theta T}}.$$

Our next result characterizes option prices at all maturities $t \leq T$ and all strikes K . As always, options can be quoted in terms of the Black–Scholes formula. What is unusual is that, in our setting, implied volatilities $\tilde{\sigma}_t$ can be expressed in a simple but non-trivial closed form at all maturities t .

Result 8. The time 0 price of an option with maturity t and strike price K is

$$C(t, K) = p_0 \Phi(d_1) - K \Phi(d_1 - \tilde{\sigma}_t \sqrt{t}), \quad (31)$$

where

$$d_1 = \frac{\log(p_0/K) + \frac{1}{2} \tilde{\sigma}_t^2 t}{\tilde{\sigma}_t \sqrt{t}} \quad \text{and} \quad \tilde{\sigma}_t = \frac{\theta + 1}{\sqrt{\theta \left(\theta + \frac{t}{T} \right)}} \sigma.$$

In particular, short-dated options (with $t \rightarrow 0$) have $\tilde{\sigma}_0 = \frac{\theta+1}{\theta} \sigma$, and long-dated options (with $t = T$) have $\tilde{\sigma}_T = \sqrt{\frac{\theta+1}{\theta}} \sigma$.

Implied volatility is increasing in the degree of heterogeneity, θ^{-1} , and the term structure of implied volatility is downward-sloping. As θ^{-1} tends to 0, we recover the conventional Black–Scholes formula with constant implied volatility σ . For comparison, recall from Result 5 that all agents agree on physical volatility, which is

$$\frac{1}{\sqrt{t}} \sigma^{(z)} (\log R_{0 \rightarrow t}) = \frac{\theta + 1}{\theta + \frac{t}{T}} \sigma = \sqrt{\frac{\theta}{\theta + \frac{t}{T}}} \tilde{\sigma}_t.$$

all z . But there is no positive price for which all agents will choose to invest their wealth fully in the risky asset.




Figure 9: The term structures of implied volatility and of annualized physical volatility.

In a homogeneous belief economy, both implied and physical volatilities would be constant, at σ , at all maturities. The sentiment and speculation induced by heterogeneous beliefs boosts implied and physical volatility in the short run, and also generates a variance risk premium, as shown in Figure 9. For example, the annualized variance risk premium to horizon T takes the simple form

$$\frac{1}{T} (\text{var}^* \log R_T - \text{var} \log R_T) = \tilde{\sigma}_T^2 - \sigma^2 = \frac{\sigma^2}{\theta}.$$

Two illustrative calibrations.—In the figures below, we set the horizon over which disagreement plays out to $T = 10$ years, and we set σ , which equals the volatility of log fundamentals (i.e. payoffs), to 12%. The belief heterogeneity parameter θ dictates the amount of disagreement, the level of short-run volatility, and the size of the long-run variance risk premium. In our baseline calibration, we set $\theta = 1.8$, which implies that one-month, one-year, and two-year implied volatilities are 18.6%, 18.2%, and 17.7%, respectively, as shown in Table 1. These numbers are close to their empirically observed counterparts, which are indicated with solid dots in Figure 10a: in the data of [Martin and Wagner \(2019\)](#), implied volatility averages 18.6%, 18.1%, and 17.9% at the one-month, one-year, and two-year horizons.

With this value of θ , the model-implied cross-sectional mean expected returns are 3.3% and 1.9% at the one- and 10-year horizons. For comparison,

	Model	Data
1mo implied vol	18.6%	18.6%
1yr implied vol	18.2%	18.1%
2yr implied vol	17.7%	17.9%
1yr disagreement	4.4%	4.8%
10yr disagreement	2.9%	2.9%
1yr mean risk premium	3.3%	3.8%
10yr mean risk premium	1.9%	3.6%

Table 1: Observables in the baseline calibration, and the corresponding time-averaged moments in the data.

in the survey data of [Ben-David, Graham and Harvey \(2013\)](#), the time-series average levels of cross-sectional average expected returns are 3.8% and 3.6% at these two horizons, as indicated with green dots in [Figure 10a](#). The cross-sectional standard deviations of expected returns (“disagreement”) at the one- and 10-year horizons are 4.4% and 2.9% in the model and 4.8% and 2.8%, on average, in the data of [Ben-David, Graham and Harvey \(2013\)](#), as indicated with red dots in [Figure 10a](#).

We also consider a calibration in which $\theta = 0.2$ to explore the behavior of asset prices under conditions with substantial disagreement, and to discuss some interesting qualitative features of equilibrium that arise once θ is less than one. The resulting term structures of physical implied volatility, and of average perceived risk premia and disagreement, are shown in [Figure 10b](#). Heightened belief heterogeneity generates steeply downward-sloping term structures of physical and implied volatility and of risk premia.

4.1 The perceived value of speculation

An agent’s stochastic discount factor (SDF) links his or her perceived true probabilities of events to the associated risk-neutral probabilities. As individuals disagree on true probabilities but agree on risk-neutral probabilities—equivalently, on asset prices, which are directly observable—they have differ-

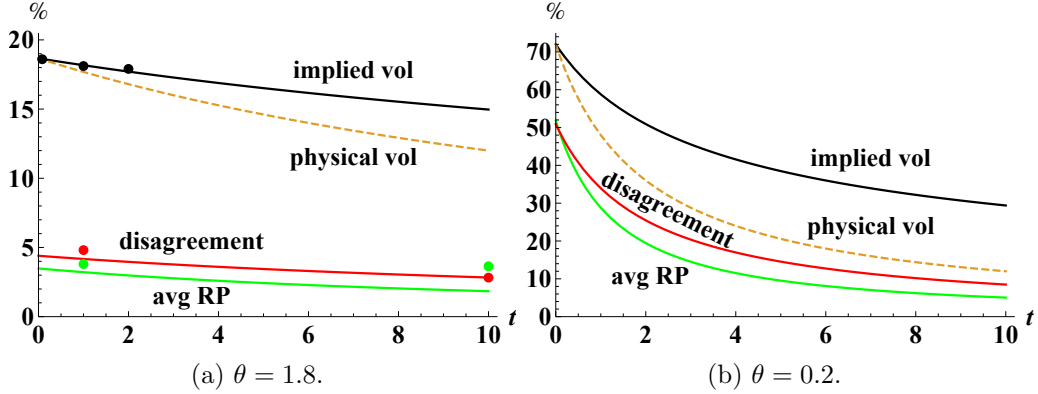


Figure 10: Term structures of implied and physical volatility, mean expected returns and disagreement in the baseline (left) and crisis (right) calibrations.

ent stochastic discount factors. We now analyze the properties of individuals' SDFs, and hence explore agents' attitudes to speculation.

Result 9. *The variance of the SDF of investor z is finite for $\theta > 1$ and is equal to*

$$\text{var}^{(z)} M_{0 \rightarrow t}^{(z)} = \frac{\theta}{\sqrt{\theta^2 - (\frac{t}{T})^2}} \exp \left\{ \frac{\left[z \sqrt{\frac{\theta t}{T}} + (\theta + 1) \sigma \sqrt{t} \right]^2}{\theta \left(\theta - \frac{t}{T} \right)} \right\} - 1. \quad (32)$$

By the [Hansen and Jagannathan \(1991\)](#) bound, this result supplies the maximum Sharpe ratio as perceived by agent z , $\text{MSR}_{0 \rightarrow t}^{(z)}$:

$$\text{MSR}_{0 \rightarrow t}^{(z)} \equiv \max_{\tilde{R}_{0 \rightarrow t}} \frac{\mathbb{E}^{(z)} \tilde{R}_{0 \rightarrow t} - R_{f,0 \rightarrow t}}{\sigma^{(z)}(\tilde{R}_{0 \rightarrow t})} = \sigma^{(z)}(M_{0 \rightarrow t}^{(z)}),$$

where $\tilde{R}_{0 \rightarrow t}$ is an arbitrary gross return and we write $\sigma^{(z)}(\cdot) = \sqrt{\text{var}^{(z)}(\cdot)}$ for the standard deviation perceived by investor z , and we have used the fact that the net interest rate is zero. As the market is complete, there is a strategy that attains the maximal Sharpe ratio (MSR) implied by the Hansen–Jagannathan bound for any agent—and of course different agents will perceive different maximal Sharpe ratios, and different associated trading strategies.

It follows from (32) that the annualized maximum Sharpe ratio perceived by agent z over very short horizons is

$$\lim_{t \rightarrow 0} \frac{1}{\sqrt{t}} \text{MSR}_{0 \rightarrow t}^{(z)} = \left| \frac{\theta + 1}{\theta} \sigma + \frac{z}{\sqrt{\theta T}} \right|. \quad (33)$$

(We annualize, here and in the figures below, by scaling the Sharpe ratio by $\frac{1}{\sqrt{t}}$.) This equals the annualized instantaneous Sharpe ratio of the risky asset; setting $z = 0$ it implies that the median investor perceives a maximal Sharpe ratio equal to short-dated implied volatility $\tilde{\sigma}_0$. Over longer horizons, however, all agents believe that there are dynamic strategies with Sharpe ratios strictly exceeding that of the risky asset.

Minimizing (32) with respect to z , we find that the investor who perceives the smallest MSR (at all horizons t) has $z = z_g$, where

$$z_g = -\frac{\theta + 1}{\sqrt{\theta}} \sigma \sqrt{T}. \quad (34)$$

Definition 1. We refer to investor $z = z_g$ as the gloomy investor. The gloomy investor perceives that the instantaneous Sharpe ratio on the risky asset is exactly zero, by Result 6 or equation (33).

The maximum Sharpe ratio to maturity perceived by the gloomy investor satisfies

$$\text{MSR}_{0 \rightarrow T}^{(z_g)} = \sqrt{\frac{\theta}{\sqrt{\theta^2 - 1}} - 1}.$$

There are, of course, more *pessimistic* investors ($z < z_g$), but we think of them as being less gloomy in the sense they perceive attractive trading opportunities associated with shorting the risky asset.

The dashed lines in the panels of Figure 11 plot the subjective Sharpe ratio of a static position in the risky asset (calculated from Results 5 and 6) against investor type, z . The solid lines plot the maximum attainable Sharpe ratio against investor type, z . The top panels use the baseline calibration, $\theta = 1.8$, and the bottom panels use the high-disagreement calibration, $\theta = 0.2$. The left panels show perceived Sharpe ratios over the next year; the right panels show annualized Sharpe ratios over the entire 10-year horizon.

The solid lines lie strictly above the dashed lines, indicating that all investors must trade dynamically to achieve their perceived MSR. In the baseline calibration, the annualized MSR perceived by the gloomy investor z_g , is 0.04 at the one-year horizon and 0.14 at the 10-year horizon. All investors perceive attainable Sharpe ratios at least as large as this. Recall that the gloomy investor believes that the risky asset is priced to earn precisely zero risk premium. Loosely speaking, the gloomy investor’s maximal-Sharpe-ratio strategy is to go long if the market sells off, and short if the market rallies, thereby exploiting what he views as irrational exuberance on the upside and irrational pessimism on the downside. This is a contrarian, “short vol” strategy. We will expand on this interpretation shortly.

If there is substantial disagreement, as in our calibration with $\theta = 0.2$, then agents perceive substantially higher attainable Sharpe ratios. At the one-year horizon depicted in Figure 11c, even the gloomy investor perceives an MSR of 0.39, while the median investor perceives an MSR of 1.50. Sharpe ratios increase very rapidly for investors with extreme beliefs, and especially so for optimists with extreme beliefs: an investor who is only moderately optimistic, with beliefs one standard deviation above the mean ($z = 1$), perceives an MSR of 8.2.

Perhaps more surprisingly, at the 10-year horizon shown in Figure 11d, all investors perceive that *arbitrarily* high Sharpe ratios are attainable. At first sight, this might seem obviously unreasonable. Surely very high Sharpe ratios should not be possible in equilibrium? But our investors are not mean-variance optimizers, so Sharpe ratios do not adequately summarize investment opportunities. (And indeed, Sharpe ratios are not considered sufficient measures of the attractiveness of a trading strategy in practice: investors appear to monitor performance measures such as max drawdowns, value at risk, and Sortino ratios, among other things.)

In order to measure the attractiveness of dynamic trading strategies in a theoretically motivated way, we calculate the maximum fraction of wealth, $\xi^{(z)}$, that an individual investor z would be prepared to sacrifice in order to avoid being shut out of the market. (We assume that other investors continue to trade, so that prices are unaffected by the absence of investor z .) When the investor is shut out, he is forced to hold his original position in the risky

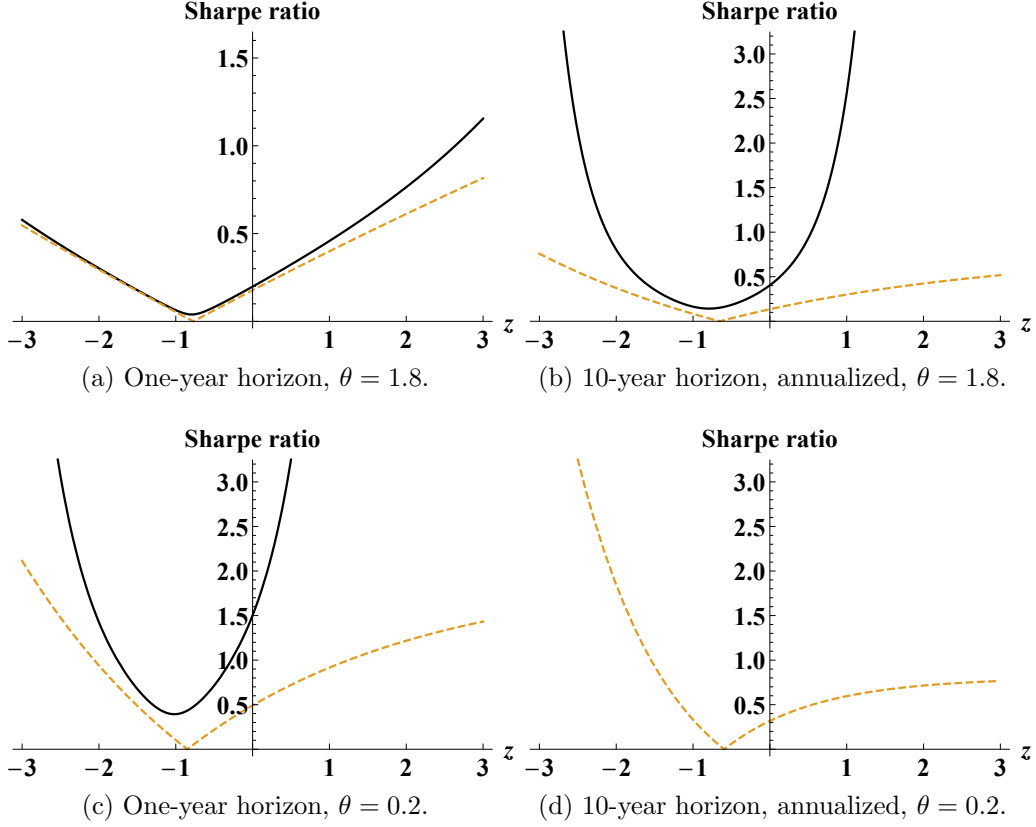


Figure 11: Maximal Sharpe ratios attainable through dynamic (solid) or static (dashed) trading, as perceived by investor z . All investors perceive that arbitrarily high Sharpe ratios are attainable dynamically in panel d.

asset, earning the return $R_{0 \rightarrow t}$ up to time t . Thus $\xi^{(z)}$ satisfies

$$\max_{R_{0 \rightarrow t}^{(z)}} \mathbb{E}^{(z)} \log \left[(1 - \xi^{(z)}) W_0^{(z)} R_{0 \rightarrow t}^{(z)} \right] = \mathbb{E}^{(z)} \log \left[W_0^{(z)} R_{0 \rightarrow t} \right]. \quad (35)$$

The [Alvarez and Jermann \(2005\)](#) bound states that

$$\max_{R_{0 \rightarrow t}^{(z)}} \mathbb{E}^{(z)} \log R_{0 \rightarrow t}^{(z)} = L^{(z)} \left[M_{0 \rightarrow t}^{(z)} \right], \quad (36)$$

where the entropy of the SDF, as perceived by investor z , is $L^{(z)} \left[M_{0 \rightarrow t}^{(z)} \right] = \log \mathbb{E}^{(z)} M_{0 \rightarrow t}^{(z)} - \mathbb{E}^{(z)} \log M_{0 \rightarrow t}^{(z)}$. The bound is attained because the market is

complete; we are using the fact that $\log R_{f,0 \rightarrow t} = 0$ in equation (36). Combined with equation (35), this implies that

$$\log(1 - \xi^{(z)}) = \mathbb{E}^{(z)} \log R_{0 \rightarrow t} - \mathbb{L}^{(z)} \left[M_{0 \rightarrow t}^{(z)} \right]. \quad (37)$$

Result 10. *The subjective entropy of the SDF is*

$$\mathbb{L}^{(z)} \left[M_{0 \rightarrow t}^{(z)} \right] = \frac{\left[z \sqrt{\frac{\theta t}{T}} + (\theta + 1) \sigma \sqrt{t} \right]^2}{2\theta \left(\theta + \frac{t}{T} \right)} + \frac{1}{2} \left(\log \frac{\theta + \frac{t}{T}}{\theta} - \frac{\frac{t}{T}}{\theta + \frac{t}{T}} \right),$$

so that the gloomy investor perceives the minimal SDF entropy.

We can also write

$$\mathbb{L}^{(z)} \left[M_{0 \rightarrow t}^{(z)} \right] = \underbrace{\frac{\theta + 1}{\theta + \frac{t}{T}} \left(\frac{z\sigma}{\sqrt{\theta T}} + \frac{\theta + 1}{2\theta} \sigma^2 \right) t}_{\mathbb{E}^{(h)} \log R_{0 \rightarrow t}} + \frac{z^2 \frac{t}{T}}{2 \left(\theta + \frac{t}{T} \right)} + \underbrace{\frac{1}{2} \left(\log \frac{\theta + \frac{t}{T}}{\theta} - \frac{\frac{t}{T}}{\theta + \frac{t}{T}} \right)}_{>0}.$$

It follows that

$$\xi^{(z)} = 1 - \exp \left\{ -\frac{z^2 \frac{t}{T}}{2 \left(\theta + \frac{t}{T} \right)} - \frac{1}{2} \left(\log \frac{\theta + \frac{t}{T}}{\theta} - \frac{\frac{t}{T}}{\theta + \frac{t}{T}} \right) \right\}.$$

The median investor perceives the minimal $\xi^{(z)}$.

Figure 12 plots $\xi^{(z)}$ against z with parameters $\sigma = 0.12$, $T = 10$, and $t = 1$ or $t = 10$. The left panel shows the baseline calibration with $\theta = 1.8$; the right panel shows the high disagreement calibration with $\theta = 0.2$.

4.2 Investor behavior and the wealth distribution

We now study how the distribution of terminal wealth varies across agents as a function of the terminal payoff of the risky asset. To do so, it is convenient to introduce the notion of an investor-specific *target price* $K^{(z)}$ defined via¹³

$$\log K^{(z)} = \mathbb{E}^{(z)} \log p_T + (z - z_g) \sigma \sqrt{\theta T}. \quad (38)$$

¹³If desired, the expected log price, $\mathbb{E}^{(z)} \log p_T = \log p_0 + \mathbb{E}^{(z)} \log R_{0 \rightarrow T}$, can be written in terms of the fundamental parameters of the model using Results 4 and 5.




Figure 12: The proportion of wealth investor z would sacrifice to avoid being prevented from trading dynamically for one or 10 years.

For example, the median and gloomy investors' target prices can be written in terms of the fundamental parameters as

$$\log K^{(0)} = \frac{\eta}{\theta} \sigma \sqrt{2T} + (\theta + 1) \sigma^2 T \quad \text{and} \quad \log K^{(z_g)} = \frac{\eta}{\theta} \sigma \sqrt{2T} - \frac{\theta + 1}{\theta} \sigma^2 T.$$

For comparison, $\log p_0 = \frac{\eta}{\theta} \sigma \sqrt{2T} - \frac{1}{2} \frac{\theta+1}{\theta} \sigma^2 T$, so the median and gloomy investors' target prices are, respectively, above and below the spot price.

As our next result shows, the target price represents the ideal outcome for investor z : the value of p_T that maximizes wealth, and hence utility, ex post.

Result 11. *The time T wealth of agent z can be expressed as a function of p_T as $W^{(z)}(p_T) = p_0 \times R_{0 \rightarrow T}^{(z)}$, where the wealth return $R_{0 \rightarrow T}^{(z)}$ is given by*

$$R_{0 \rightarrow T}^{(z)} = \sqrt{\frac{\theta + 1}{\theta}} \exp \left\{ \frac{1}{2} (z - z_g)^2 - \frac{1}{2(1 + \theta) \sigma^2 T} [\log (p_T / K^{(z)})]^2 \right\}. \quad (39)$$

Thus $W^{(z)}(p_T)$ is maximized when $p_T = K^{(z)}$.

This can also be written as a quadratic relationship between an investor's log wealth return, $r_{0 \rightarrow T}^{(z)} = \log R_{0 \rightarrow T}^{(z)}$, and the log return on the risky asset,

$r_{0 \rightarrow T} = \log R_{0 \rightarrow T}$:

$$r_{0 \rightarrow T}^{(z)} = \frac{1}{2} \log \frac{\theta + 1}{\theta} + \frac{1}{2} (z - z_g)^2 - \frac{1}{2(1 + \theta)} \left[\frac{r_{0 \rightarrow T} - \mathbb{E}^{(z)} r_{0 \rightarrow T}}{\sigma \sqrt{T}} - \sqrt{\theta} (z - z_g) \right]^2. \quad (40)$$

It follows that the expected elasticity of an investor's wealth return with respect to the risky asset return, $\mathbb{E}^{(z)}(\partial r_{0 \rightarrow T}^{(z)} / \partial r_{0 \rightarrow T})$, satisfies

$$\mathbb{E}^{(z)} \frac{\partial r_{0 \rightarrow T}^{(z)}}{\partial r_{0 \rightarrow T}} = 1 + \frac{z}{|z_g|}.$$

In particular, the median investor has an expected elasticity of one and the gloomy investor has an expected elasticity of zero.

If the risky asset goes nowhere, in the sense of having zero realized excess return (i.e., if $R_{0 \rightarrow T} = 1$), then the median and gloomy investors have the same return on wealth: $R_{0 \rightarrow T}^{(0)} = R_{0 \rightarrow T}^{(z_g)}$.

In our model, there is a useful distinction between what investors expect to happen and what they would like to happen. (The distinction also exists, but is uninteresting, in models in which a representative agent statically holds the market, as the target price is infinity in such models.) The gloomy investor would like to be proved right in logs: his log target price equals his expected log price. But targets and expectations differ for all other investors. More optimistic investors have a (log) target price that exceeds their expectations—i.e., they are best off if the risky asset modestly outperforms their expectations—while more pessimistic investors are best off if the risky asset modestly underperforms their expectations. (But any investor does very poorly if the asset performs *far* better or worse than anticipated.)

Differentiating the expression (39) twice with respect to p_T , we find

$$R_{0 \rightarrow T}^{(z)''} = \frac{R_{0 \rightarrow T}^{(z)}}{p_T^2} \left\{ \left(\frac{\log(p_T / K^{(z)})}{(1 + \theta)\sigma^2 T} + \frac{1}{2} \right)^2 - \frac{1}{4} - \frac{1}{(1 + \theta)\sigma^2 T} \right\}. \quad (41)$$

It immediately follows that any investor's wealth is concave in p_T near their target price, and convex far from their target price. A moderate investor's




Figure 13: Return on wealth against return on the market for different agents. Dots indicate the expected return on the risky asset perceived by each investor.

wealth is also concave near their *expected* log price. But an extreme investor's wealth is convex near their expected log price. These facts follow because if $\log p_T = \mathbb{E}^{(z)} \log p_T$ then we have, after some algebra,

$$\text{sign} \left[W^{(z)''}(p_T) \right] = \text{sign} \left[z^2 - z_g z - \frac{\theta + 1}{\theta} \right],$$

which is negative for moderate investors (including all investors with z between z_g and zero), and positive if $|z|$ is sufficiently large.

Figure 13 shows how different investors' outcomes depend on the risky asset's realized return. Dots indicate the expected gross return on the risky asset perceived by each of the investors. The median investor's wealth is a concave function of the risky asset return in the neighbourhood of the investor's expected outcome, while more extreme investors have wealth that is convex in the risky asset return in the neighbourhood of their expected outcome.

Equation (39) can be rewritten as

$$R_{0 \to T}^{(z)} = \sqrt{\frac{\theta + 1}{\theta}} \exp \left\{ -\frac{1}{2} \left[\frac{\log p_T - \mathbb{E}^{(z)} \log p_T}{\sqrt{\text{var}^{(z)} \log p_T}} \right]^2 + \frac{1}{2(1 + \theta)} \left[\sqrt{\theta} \frac{\log p_T - \mathbb{E}^{(z)} \log p_T}{\sqrt{\text{var}^{(z)} \log p_T}} + z - z_g \right]^2 \right\}.$$

This characterization shows that you get richer if you are an extremist (large $|z|$) whose expectations are realized than you do if you have conventional beliefs (z close to zero) that are realized: it is cheap to purchase claims to states of the world that extremists consider likely, because few people are extremists. As a result, there is substantially more wealth inequality in states in which the asset has an extreme positive or negative realized return.

Informally, extreme investors are “long volatility” near the outcome that they expect, while moderate investors are “short volatility” in their corresponding region. To formalize this intuition, we introduce a general result that holds in any frictionless arbitrage-free model in which options are traded. It is in the spirit of the famous result of Breeden and Litzenberger (1978), but the logic operates at the level of payoffs rather than of prices.

Result 12. *Let $W(\cdot)$ be such that $W(0) = 0$. Then choosing terminal wealth $W(p_T)$ is equivalent to holding the following portfolio:*

- Long $W'(K_0)$ units of the underlying asset (whose price is p_T at time T)
- Long bonds with face value $W(K_0) - K_0W'(K_0)$
- Long $W''(K) dK$ put options with strike K , for every $K < K_0$
- Long $W''(K) dK$ call options with strike K , for every $K > K_0$

The constant $K_0 > 0$ can be chosen arbitrarily.

Proof. Start from $W(p_T) = \int_0^{p_T} W'(K) dK = \int_0^\infty W'(K) \mathbf{1}_{\{p_T > K\}} dK$ and integrate by parts. \square

We now specialize Result 12 to our setting to identify a static portfolio whose payoffs replicate the investment strategy followed by an arbitrary investor h .

Result 13. *Investor z 's investment strategy is equivalent to the following:*

- a long position in bonds with face value $W^{(z)}(K^{(z)}) = p_0 \sqrt{\frac{\theta+1}{\theta}} e^{\frac{1}{2}(z-z_g)^2}$;
- short positions in options with strikes at and near her target level $K^{(z)}$;

- long positions in options with strikes far from $K^{(z)}$.

More precisely, the investor holds $W^{(z)''}(K) dK$ put options with strike K for all $K < K^{(z)}$, and $W^{(z)''}(K) dK$ call options with strike K , for all $K \geq K^{(z)}$, where $W^{(z)''}(K)$ is as defined in (41). (Note that $W^{(z)''}(K^{(z)}) < 0$, and that $W^{(z)''}(K) > 0$ if K is sufficiently far from $K^{(z)}$.)

The best possible payoff is $W^{(z)}(K^{(z)})$. This occurs if the asset hits its target price, $p_T = K^{(z)}$, in which case all the options expire worthless. Conversely, the investor's wealth approaches zero as $p_T \rightarrow 0$ or $p_T \rightarrow \infty$.

Proof. It follows from the definition (38) of $K^{(z)}$, and a direct calculation, that $W^{(z)'}(K^{(z)}) = 0$. The claims in the first paragraph then follow on setting $K_0 = K^{(z)}$ in Result 12. The fact that the best possible payoff is $W^{(z)}(K^{(z)})$ follows from equation (39). The payoff on the option portfolio must therefore be nonpositive. \square

4.3 Maximum-Sharpe-ratio strategies

Result 9 characterized the maximum Sharpe ratios perceived by different investors. We now turn to the question of which strategies achieve these maximal Sharpe ratios. We do so in two steps.

First, we know from the work of Hansen and Richard (1987) that the return with minimal second (subjective, uncentered) moment, R_* , is proportional to $M_{0 \rightarrow T}^{(z)}$, and, moreover, that this return achieves the minimal possible Sharpe ratio. It follows that a portfolio that is long the riskless asset and short R_* has the maximal possible Sharpe ratio. A Sharpe-ratio-maximizing strategy for investor z therefore must take the form $a - bM_{0 \rightarrow T}^{(z)}$ for some constants $a > 0$ and $b > 0$. For the strategy to be a gross return—that is, a payoff with price 1—we must have $a = 1 + b \mathbb{E}^{(z)} \left[M_{0 \rightarrow T}^{(z)2} \right]$. Second, the return on wealth chosen by investor z , which we derived in Result 11, reveals the SDF perceived by that investor: $M_{0 \rightarrow T}^{(z)} = 1/R_{0 \rightarrow T}^{(z)}$.

We can therefore write a Sharpe-ratio-maximizing return as

$$R_{\text{MSR},0 \rightarrow T}^{(z)} = 1 + b \left(\text{var}^{(z)} M_{0 \rightarrow T}^{(z)} + 1 \right) - \frac{b}{R_{0 \rightarrow T}^{(z)}}, \quad (42)$$




Figure 14: Return on a max-Sharpe-ratio strategy (solid) and return on wealth (dashed) against return on the market for agents $z = 0$ and 1 . Log scale on x -axis.

where b can be any positive constant (reflecting the fact that any strategy can be combined with a position in the riskless asset without altering its Sharpe ratio).

Result 11 showed that the optimally chosen wealth return $R_{0 \rightarrow T}^{(z)}$ is maximized when p_T equals investor z 's target return $K^{(z)}$. It follows from equation (42) that the maximal-Sharpe-ratio return $R_{\text{MSR},0 \rightarrow T}^{(z)}$ is also maximized when p_T equals $K^{(z)}$. But the maximal-Sharpe ratio return becomes negative if p_T is far from $K^{(z)}$, whereas the wealth return remains positive. Figure 14 illustrates using parameters from the baseline calibration.

We can also study the convexity of the MSR return, considered as a function of p_T . It follows from equations (39), (41), and (42) that

$$R_{\text{MSR},0 \rightarrow T}^{(z)''} = \frac{-b}{p_T^2 R_{0 \rightarrow T}^{(z)}} \left\{ \left(\frac{\log(p_T/K^{(z)})}{(1+\theta)\sigma^2 T} - \frac{1}{2} \right)^2 - \frac{1}{4} + \frac{1}{(1+\theta)\sigma^2 T} \right\}. \quad (43)$$

This is negative, and arbitrarily large in magnitude, for p_T far from $K^{(z)}$. Using Result 12 in the same way as it was used to derive Result 13, we see that the maximal-Sharpe-ratio strategy features extremely short positions in deep-out-of-the-money call and put options.¹⁴ If beliefs are sufficiently heterogeneous

¹⁴Specifically, let $g(K)$ denote the right-hand side of (43) with all occurrences of p_T (including the implicit dependence on p_T via $R_{0 \rightarrow T}^{(z)}$, which we do not write out explicitly




Figure 15: Dashed lines indicate the critical (annualized) log returns on the risky asset at which investor z 's MSR strategy has zero realized excess return in the baseline calibration. The unshaded area represents the range of returns on the risky asset over which the investor's MSR strategy outperforms the riskless asset; the shaded areas indicate where the MSR strategy underperforms.

that $(1 + \theta)\sigma^2 T < 4$, then the right-hand side of equation (43) is negative everywhere, so that the strategy calls for short positions in options of *all* strikes.

A given investor's MSR strategy realizes a positive excess return over a particular range of risky asset returns $R_{0 \rightarrow T}$. We can characterize this range by asking for which $R_{0 \rightarrow T}$ investor z 's MSR strategy has return exactly equal to the riskless rate, that is, $R_{\text{MSR},0 \rightarrow T}^{(z)} = 1$. (This question has an answer that is independent of b , because combining with the riskless asset does not affect the MSR return when it happens to earn the same as the riskless asset.) Straightforward algebra shows, using (32), (40), and (42), that the critical log returns for investor z satisfy

$$\frac{r_{0 \rightarrow T} - \mathbb{E}^{(z)} r_{0 \rightarrow T}}{\sigma \sqrt{T}} = \sqrt{\theta} (z - z_g) \pm \sqrt{\frac{(\theta + 1)^2}{\theta - 1} (z - z_g)^2 + (\theta + 1) \log \frac{\theta}{\theta - 1}}. \quad (44)$$

for lack of space) replaced by K . Applying Result 12 with $K_0 = K^{(z)}$, we find that the maximum-Sharpe-ratio strategy from the perspective of agent z can be implemented by holding $g(K) dK$ calls at each strike $K > K^{(z)}$ and $g(K) dK$ puts at each strike $K < K^{(z)}$, plus a bond position. As $g(K) \ll 0$ for K far from $K^{(z)}$, this strategy is (extremely) short out-of-the-money options.

When the risky asset’s log return equals $r_{0 \rightarrow T}$, as defined by the above equation, investor z ’s MSR strategy has zero excess return.

Figure 15 illustrates in the baseline calibration. The dashed lines show the log returns on the risky asset associated with zero excess return on investor z ’s MSR strategy. If the risky asset’s log return lies in the shaded region, the MSR strategy has a negative realized excess return; if the risky asset’s log return lies between the lines, the MSR strategy has a positive realized excess return. The range is narrowest (as can be seen directly from equation (44)) for the gloomy investor, who perceives that the Sharpe-ratio-maximizing strategy can be implemented by shorting relatively near-the-money call and put options.

We note, however, that although it is possible to earn high Sharpe ratios via short option positions, these strategies are not remotely attractive to investors in our economy. Indeed, as maximum-Sharpe-ratio strategies feature the possibility of unboundedly negative gross returns, our investors would prefer to invest fully in cash than to rebalance, even slightly, toward a maximum-Sharpe-ratio strategy.

5 Conclusions

We have presented a frictionless model in which individuals have stable beliefs and risk aversion. All investors are risk-averse; short sales are allowed; all agents avoid bankruptcy; and all agents are on their first-order conditions at all times.

Even so, the model generates a rich set of predictions. Heterogeneity in beliefs gives rise to sentiment, which induces speculation and drives up realized and implied volatility, particularly in the short run. All agents understand these facts, so expected returns are higher than in an otherwise identical homogeneous economy, and securities with payoffs in extreme states of the world are far more highly valued than in otherwise similar economies with homogeneous beliefs. Moderate investors are suppliers of liquidity: they trade in a contrarian manner—they are “short vol”—and capture a variance risk premium created by the presence of extremists.

References

- Alvarez, Fernando, and Urban J. Jermann.** 2005. “Using Asset Prices to Measure the Persistence of the Marginal Utility of Wealth.” *Econometrica*, 73(6): 1977–2016.
- Atmaz, Adem, and Suleyman Basak.** 2018. “Belief dispersion in the stock market.” *The Journal of Finance*, 73(3): 1225–1279.
- Banerjee, Snehal, and Ilan Kremer.** 2010. “Disagreement and learning: Dynamic patterns of trade.” *The Journal of Finance*, 65(4): 1269–1302.
- Basak, Suleyman.** 2005. “Asset pricing with heterogeneous beliefs.” *Journal of Banking & Finance*, 29(11): 2849–2881.
- Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey.** 2013. “Managerial Miscalibration.” *Quarterly Journal of Economics*, 128(4): 1547–1584.
- Bhamra, Harjoat S., and Raman Uppal.** 2014. “Asset Prices with Heterogeneity in Preferences and Beliefs.” *Review of Financial Studies*, 27(2): 519–580.
- Cox, J., S. Ross, and M. Rubinstein.** 1979. “Option Pricing: A Simplified Approach.” *Journal of Financial Economics*, 7: 229–263.
- Geanakoplos, John.** 2010. “The Leverage Cycle.” In *NBER Macroeconomic Annual 2009*. Vol. 24, , ed. Daron Acemoglu, Kenneth Rogoff and Michael Woodford, 1–65. Chicago:University of Chicago Press.
- Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus.** 2019. “Five facts about beliefs and portfolios.” Working paper.
- Hansen, Lars P., and Scott F. Richard.** 1987. “The Role of Conditioning Information in Deducing Testable Restrictions Implied by Dynamic Asset Pricing Models.” *Econometrica*, 55(3): 587–613.
- Hansen, Lars Peter, and Ravi Jagannathan.** 1991. “Implications of Security Market Data for Models of Dynamic Economies.” *Journal of Political Economy*, 99(2): 225–262.
- Harrison, J. Michael, and David M. Kreps.** 1978. “Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations.” *Quarterly Journal of Economics*, 92(2): 323–336.

- Keynes, John Maynard.** 1936. *The General Theory of Employment, Interest and Money*. Palgrave Macmillan.
- Martin, Ian.** 2017. “What is the Expected Return on the Market?” *Quarterly Journal of Economics*, 132(1): 367–433.
- Martin, Ian W. R., and Christian Wagner.** 2019. “What is the Expected Return on a Stock?” *Journal of Finance*. Forthcoming.
- Scheinkman, José A., and Wei Xiong.** 2003. “Overconfidence and Speculative Bubbles.” *Journal of Political Economy*, 111(6): 1183–1219.

A The risky bond example

This section contains some further calculations in the risky bond example of Section 2.2. Specifically, we ask what happens if agents are not allowed to trade dynamically. Agent h perceives a probability $1 - (1 - h)^T$ that the bond pays 1, and $(1 - h)^T$ that the bond pays ε , so solves

$$\max_{x_h} (1 - (1 - h)^T) \log(w_h - x_h p + x_h) + (1 - h)^T \log(w_h - x_h p + x_h \varepsilon).$$

The first-order condition (after setting $w_h = p$ to account for the fact that all agents are initially endowed with a unit of the risky asset) is

$$x_h = p \left(\frac{1 - (1 - h)^T}{p - \varepsilon} - \frac{(1 - h)^T}{1 - p} \right).$$

If T is reasonably large, most agents will have $(1 - h)^T \approx 0$, and so will choose $x_h \approx \frac{p}{p - \varepsilon}$; their wealth in the bad state of the world is then approximately zero. Thus, if forced to trade statically most agents will lever up (almost) as much as possible without risking bankruptcy.

For the market to clear, we require $\int_0^1 x_h dh = 1$, which implies that $p = \frac{(1+T)\varepsilon}{1+T\varepsilon}$. This is the same as the time-0 price in the case with dynamic trade. It follows that agent h 's demand for the asset is

$$x_h = 1 + (1 - (1 + T)(1 - h)^T) \frac{1 + T\varepsilon}{T(1 - \varepsilon)}.$$

If an individual investor is forced to trade statically (while everyone else is trading dynamically, so that the price at time t is observed) then the investor's leverage at time t , defined as debt-to-wealth ratio, is

$$\text{leverage}_t = \frac{p_0(x_h - 1)}{x_h p_t + p_0 - p_0 x_{h,0}} = \frac{1 - (1 + T)(1 - h)^T}{T - t(1 - (1 + T)(1 - h)^T)} \frac{1 + t - t\varepsilon + T\varepsilon}{1 - \varepsilon}.$$

For comparison, in the dynamic case investor h 's time- t demand will be

$$x_{h,t} = (1 - h)^t + \frac{(1 - h)^t}{1 - \varepsilon} [h(2 + t)(1 + t(1 - \varepsilon) + T\varepsilon) - 1 - T\varepsilon]$$

and the investor's leverage at time t , defined as in equation (13), is

$$\text{leverage}_t = \frac{x_{h,t} p_t - w_{h,t}}{w_{h,t}} = \frac{(h(2+t) - 1)(1 + t(1 - \varepsilon) + T\varepsilon)}{(1+t)(1-\varepsilon)}.$$

This strategy delivers the dynamic investor higher expected utility. An investor who follows the static strategy has wealth

$$\frac{p_0 (1 - (1 - h)^T)}{1 - (1 - p_0^*) \cdots (1 - p_{T-1}^*)}$$

if the bond does not default—which, in the investor's opinion, occurs with probability $1 - (1 - h)^T$. If the bond does default, the investor ends up with

$$\frac{p_0(1-h)^T}{(1-p_0^*) \cdots (1-p_{T-1}^*)} = \frac{p_0(1-h)^T(1-\varepsilon)}{1-p_0}.$$

This occurs with probability $(1 - h)^T$. The static investor therefore has expected utility

$$EU_{static} = \left[1 - (1 - h)^T\right] \log \left(\frac{p_0 (1 - (1 - h)^T)}{1 - (1 - p_0^*) \cdots (1 - p_{T-1}^*)} \right) + (1-h)^T \log \left(\frac{p_0(1-h)^T(1-\varepsilon)}{1-p_0} \right).$$

Conversely, a dynamic investor ends up with wealth

$$\frac{p_0(1-h)^t h}{(1-p_0^*) \cdots (1-p_{t-1}^*) p_t^*}$$

if the first up move occurs after t successive down-moves, where $t \in \{0, \dots, T-1\}$.

This outcome has probability $(1 - h)^t h$. If the bond defaults, his terminal wealth is

$$\frac{p_0(1-h)^T}{(1-p_0^*) \cdots (1-p_{T-1}^*)} = \frac{p_0(1-h)^T(1-\varepsilon)}{1-p_0}.$$

Thus his expected utility is

$$EU_{dynamic} = \sum_{t=0}^{T-1} (1-h)^t h \log \left(\frac{p_0(1-h)^t h}{(1-p_0^*) \cdots (1-p_{t-1}^*) p_t^*} \right) + (1-h)^T \log \left(\frac{p_0(1-h)^T(1-\varepsilon)}{1-p_0} \right).$$




Figure 16: The attractiveness of dynamic strategies relative to static strategies, for investors of differing levels of optimism h .

It follows that

$$\begin{aligned}
 EU_{dynamic} - EU_{static} &= \sum_{t=0}^{T-1} (1-h)^t h \log \left(\frac{h(1-h)^t [1 - (1-p_0^*) \cdots (1-p_{T-1}^*)]}{[1 - (1-h)^T] (1-p_0^*) \cdots (1-p_{t-1}^*) p_t^*} \right) \\
 &= \sum_{t=0}^{T-1} (1-h)^t h \log \left(\frac{(1-h)^t h (1+t)(2+t)T}{(1 - (1-h)^T) (1+T)} \right),
 \end{aligned}$$

which is independent of ε .

To convert this logic into dollar terms, suppose an investor is indifferent between wealth of $\omega_h w_h$ and being constrained to invest statically, and wealth of w_h and being allowed to invest dynamically. Then ω_h must satisfy $\mathbb{E}_{static} \log(\omega_h w_h R) = \mathbb{E}_{dynamic} \log(w_h R)$ which implies that $\omega_h - 1 = \exp\{EU_{dynamic} - EU_{static}\} - 1$. Figure 16 plots this quantity for $\varepsilon = 0.3$ and $T = 50$, as in the example in the main text.

Proof of Result 1:

Proof. The probability distribution that each investor assigns to the event of m up moves at time t is a BetaBinomial($\zeta h, \zeta(1-h), t$) distribution. Therefore

$$\mathbb{P}(m \text{ up moves at time } t) = \binom{t}{m} \frac{B(\zeta h + m, \zeta(1-h) + t - m)}{B(\zeta h, \zeta(1-h))}$$

From this, we can see how each agent updates their beliefs after a down move. Then, by using Bayesian updating, we find that after t down moves (at time t), the updated belief of the investor satisfies

$$\tilde{h} \mid t \text{ down moves} \sim \text{Beta}(\zeta h, \zeta(1-h) + t).$$

We will denote this posterior belief by \tilde{h}_t , which has the property that $\mathbb{E} \tilde{h}_t = \frac{h}{1+t/\zeta}$. We can now proceed by replicating the steps of the benchmark model. The agent's first-order condition at time t becomes

$$x_{h,t} = w_{h,t} \left(\frac{\frac{h}{1+t/\zeta}}{p-p_d} - \frac{1 - \frac{h}{1+t/\zeta}}{p_u - p} \right). \quad (45)$$

It follows that the wealth of an investor after t down moves is

$$w_{h,t} = \underbrace{\lambda_{path} (1-h) \left(1 - \frac{h}{1+1/\zeta}\right) \cdots \left(1 - \frac{h}{1+(t-1)/\zeta}\right)}_{I_t(h)},$$

where we can set $I_0(h) = 1$ since the initial wealth does not depend on h . We can then find λ_{path} by equating aggregate wealth to the value of the risky asset, which is in unit supply:

$$\lambda_{path} \int_0^1 I_t(h) f(h) dh = p_t.$$

Finally, in order to clear the market we must have

$$1 = \lambda_{path} \left[\int_0^1 I_t(h) \left(\frac{\frac{h}{1+t/\zeta}}{p-p_d} - \frac{1 - \frac{h}{1+t/\zeta}}{p_u - p} \right) f(h) dh \right]. \quad (46)$$

We can now let

$$G_t = \frac{\int_0^1 I_t(h) h f(h) dh}{\left(1 + \frac{t}{\zeta}\right) \int_0^1 I_t(h) f(h) dh},$$

where $G_0 = \frac{1}{2}$. Then equation (46) can be rewritten as $\frac{1}{p} = \frac{G_t}{p-p_d} - \frac{1-G_t}{p_u-p}$. As G_t can easily be calculated via numerical integration, this allows us to compute the price path of the risky bond for arbitrary learning parameters ζ .

In order to prove analytically that learning drives the price of the risky bond even further below the case previously analyzed, note that

$$y_{m,t} = G_t y_{m+1,t+1} + (1 - G_t) y_{m,t+1}$$

where, as in the benchmark model, we have defined $y_{m,t} = 1/p_{m,t} - 1$.

In the risky bond case $y_{m+1,t+1} = 0$, so we can drop subscripts m and write y_t for $y_{0,t}$. Then $y_t = (1 - G_t)y_{t+1}$, from which it follows that $y_0 = (1 - G_0)(1 - G_1) \cdots (1 - G_{T-1})y_T$. Now,

$$1 - G_t = \frac{\int_0^1 I_{t+1}(h)f(h)dh}{\int_0^1 I_t(h)f(h)dh},$$

and so we have a telescoping product

$$\begin{aligned} (1 - G_0)(1 - G_1) \cdots (1 - G_{T-1}) &= \int_0^1 I_T(h)f(h)dh \\ &\geq \int_0^1 (1 - h)^T f(h)dh = \frac{B(\alpha, \beta + T)}{B(\alpha, \beta)}. \end{aligned}$$

Comparing this result with equation (21) we see that the present model gives a higher value for y_0 and hence a lower value of the risky bond, p_0 , confirming our original intuition. \square

B Proofs of results

Proof of Result 2:

Proof. Observe from the recurrence relation (24) that a pricing formula in the form (25) holds. Each constant c_m is a sum of products of terms of the form $H_{j,s}$ and $1 - H_{j,s}$ over appropriate j and s . We noted in the text that $H_{m,t}(1 - H_{m+1,t+1}) = (1 - H_{m,t})H_{m,t+1}$: that is, pricing is path-independent.

Fix m between 0 and T . By path independence, all the possible ways of getting from the initial node to node m at time T make an equal contribution to c_m . By considering the path that travels down for $T - m$ periods and then up for m periods, and then multiplying by the number of paths, $\binom{T}{m}$, we find

that

$$\begin{aligned}
c_m &= \binom{T}{m} (1 - H_{0,0}) \cdots (1 - H_{0,T-m-1}) H_{0,T-m} H_{1,T-m+1} \cdots H_{m-1,T-1} \\
&= \binom{T}{m} \frac{\beta}{\alpha + \beta} \cdot \frac{\beta + 1}{\alpha + \beta + 1} \cdots \frac{\beta + T - m - 1}{\alpha + \beta + T - m - 1} \cdot \frac{\alpha}{\alpha + \beta + T - m} \cdots \frac{\alpha + m - 1}{\alpha + \beta + T - 1} \\
&= \binom{T}{m} \frac{B(\alpha + m, \beta + T - m)}{B(\alpha, \beta)}.
\end{aligned}$$

The risk-neutral probability q_m^* can be determined using the facts that $p_{m,t}^* = H_{m,t} p_{m,t} / p_{m+1,t+1}$ and $1 - p_{m,t}^* = (1 - H_{m,t}) p_{m,t} / p_{m,t+1}$. (We are restating (9) with subscripts to keep track of the current node.) Thus—using again path-independence in the first line—

$$\begin{aligned}
q_m^* &= \binom{T}{m} (1 - p_{0,0}^*) \cdots (1 - p_{0,T-m-1}^*) \cdot p_{0,T-m}^* p_{1,T-m+1}^* \cdots p_{m-1,T-1}^* \\
&= \binom{T}{m} (1 - H_{0,0}) \frac{p_{0,0}}{p_{0,1}} \cdots (1 - H_{0,T-m-1}) \frac{p_{0,T-m-1}}{p_{0,T-m}} \cdot H_{0,T-m} \frac{p_{0,T-m}}{p_{1,T-m+1}} \cdots H_{m-1,T-1} \frac{p_{m-1,T-1}}{p_{m,T}} \\
&= c_m \frac{p_{0,0}}{p_{m,T}}. \quad \square
\end{aligned}$$

We also have the following generalization of Result 2. We omit the proof, which is essentially identical to the above.

Lemma 1. *For any node m, t :*

$$z_{m,t} = \sum_{j=0}^{T-t} c_{m,t,j} z_{m+j,T}$$

where j represents the number of further up-moves after time t , and

$$c_{m,t,j} = \binom{T-t}{j} \frac{B(m + \alpha + j, T - m + \beta - j)}{B(m + \alpha, t - m + \beta)}.$$

Moreover, the risk neutral probability of ending up at j, T starting from node m, t is given by

$$q_{m,t,j}^* = c_{m,t,j} \frac{p_{m,t}}{p_{m+j,T}}.$$

Proof of Result 3:

Proof. We will start by proving the following Lemma.

Lemma 2. *If $Y_1 \sim BB(\bar{\alpha}, \bar{\alpha}, T)$ and $Y_2 \sim BB(\underline{\alpha}, \underline{\alpha}, T)$, for $\bar{\alpha} > \underline{\alpha}$ then Y_1 second order stochastically dominates Y_2 .*

Proof. A sufficient condition for second order stochastic dominance, for variables with the same expectation, is the single crossing dominance. That is, it is sufficient to prove that:

$$F_{\underline{\alpha}}(s) \geq F_{\bar{\alpha}}(s) \iff s \leq c^* \quad (47)$$

for some c^* , where $F_{\bar{\alpha}}(s), F_{\underline{\alpha}}(s)$ are the cdfs of Y_1, Y_2 respectively.¹⁵ Because of symmetry c^* will be just $T/2$. To prove the above it is sufficient to prove that $f_{\underline{\alpha}}(k) - f_{\bar{\alpha}}(k)$ is decreasing in $[0, T/2]$, where $f(\cdot)$ denotes the probability mass function. Then $F_{\underline{\alpha}}(s) - F_{\bar{\alpha}}(s)$ would be decreasing (as a sum of decreasing functions) and the proof of the lemma would be completed, since this would imply equation 47. Hence, we need to show that:

$$\binom{T}{k} \left[\frac{B(k + \underline{\alpha}, T - k + \underline{\alpha})}{B(\underline{\alpha}, \underline{\alpha})} - \frac{B(k + \bar{\alpha}, T - k + \bar{\alpha})}{B(\bar{\alpha}, \bar{\alpha})} \right]$$

is decreasing in k (in the interval $[0, T/2]$). Equivalently:

$$\Gamma(k + \underline{\alpha})\Gamma(T - k + \underline{\alpha}) \left[\frac{1}{\Gamma(T + 2\underline{\alpha})B(\underline{\alpha}, \underline{\alpha})} - \frac{\Gamma(k + \bar{\alpha})\Gamma(T - k + \bar{\alpha})}{\Gamma(k + \underline{\alpha})\Gamma(T - k + \underline{\alpha})} \frac{1}{\Gamma(T + 2\bar{\alpha})B(\bar{\alpha}, \bar{\alpha})} \right]$$

is decreasing.

But the above holds because of the following 2 facts:

First, $h(k) = \Gamma(k + \underline{\alpha})\Gamma(T - k + \underline{\alpha})$ is decreasing because

$$[\log h(k)]' = \psi(k + \underline{\alpha}) - \psi(T - k + \underline{\alpha}) < 0$$

where $\psi(\cdot)$ is the digamma function, which is an increasing function since $\Gamma(\cdot)$ is log-convex (and $k < T - k$).

Second, $\frac{\Gamma(k + \bar{\alpha})\Gamma(T - k + \bar{\alpha})}{\Gamma(k + \underline{\alpha})\Gamma(T - k + \underline{\alpha})}$ is increasing. Indeed, assume $k_1 > k_2$. Then, we

¹⁵See, for instance, Osband & Roy (2018) "Gaussian-Dirichlet Posterior Dominance in Sequential Learning".

want:

$$\frac{\Gamma(k_1 + \bar{\alpha})\Gamma(T - k_1 + \bar{\alpha})}{\Gamma(k_1 + \underline{\alpha})\Gamma(T - k_1 + \underline{\alpha})} > \frac{\Gamma(k_2 + \bar{\alpha})\Gamma(T - k_2 + \bar{\alpha})}{\Gamma(k_2 + \underline{\alpha})\Gamma(T - k_2 + \underline{\alpha})}$$

Equivalently:

$$\frac{\Gamma(k_1 + \bar{\alpha})\Gamma(T - k_1 + \bar{\alpha})}{\Gamma(k_2 + \bar{\alpha})\Gamma(T - k_2 + \bar{\alpha})} > \frac{\Gamma(k_1 + \underline{\alpha})\Gamma(T - k_1 + \underline{\alpha})}{\Gamma(k_2 + \underline{\alpha})\Gamma(T - k_2 + \underline{\alpha})}$$

Now using the property that $\Gamma(z + 1) = z\Gamma(z)$ for any z and that $k_1, k_2 \in \mathbb{Z}$, we get:

$$\frac{(\bar{\alpha} + k_2)(\bar{\alpha} + k_2 + 1) \dots (\bar{\alpha} + k_1 - 1)}{(\bar{\alpha} + T - k_1)(\bar{\alpha} + T - k_1 + 1) \dots (\bar{\alpha} + T - k_2 - 1)} > \frac{(\underline{\alpha} + k_2)(\underline{\alpha} + k_2 + 1) \dots (\underline{\alpha} + k_1 - 1)}{(\underline{\alpha} + T - k_1)(\underline{\alpha} + T - k_1 + 1) \dots (\underline{\alpha} + T - k_2 - 1)}$$

which is true since for example $\frac{(\bar{\alpha} + k_2)}{(\bar{\alpha} + T - k_1 + 1)} > \frac{(\underline{\alpha} + k_2)}{(\underline{\alpha} + T - k_1 + 1)}$.

Therefore this proves that Y_1 single-crossing dominates Y_2 and hence it also second order stochastically dominates Y_2 and the lemma has been proved. \square

Having established the above Lemma, we can now go back to proving Result 3. It is well known that if Y_1 second order stochastically dominates Y_2 then for any concave function $u(\cdot)$:

$$\mathbb{E}_{Y_1}[u(m)] \geq \mathbb{E}_{Y_2}[u(m)].$$

Pick $u(m) = -\frac{1}{p_{m,T}}$. Then we get: $\mathbb{E}_{Y_1}[\frac{1}{p_{m,T}}] \leq \mathbb{E}_{Y_2}[\frac{1}{p_{m,T}}]$ and therefore:

$$\frac{1}{\mathbb{E}_{Y_1}[\frac{1}{p_{m,T}}]} \geq \frac{1}{\mathbb{E}_{Y_2}[\frac{1}{p_{m,T}}]}$$

That is if p_1, p_2 are the corresponding prices (where p_1 corresponds to the case with less heterogeneity, as $\bar{\alpha} > \underline{\alpha}$), we have that $p_1 > p_2$.

To show that log-concavity of p implies that $1/p$ is convex, note that log-concavity is equivalent to $(p')^2 \geq pp''$. \square

Proof of Result 4:

Proof. As shown in equation (25),

$$p_{0,0}^{-1} = \sum_{j=1}^{2N} c_j z_{j,T}$$

From Result 2, c_j equals the probability that a $BB(2N, \alpha, \beta)$ random variable takes the value j . Therefore we can equivalently write

$$p_{0,0}^{-1} = \mathbb{E}_j [z_{j,T}] = \mathbb{E}_j \left[e^{-\sigma\sqrt{2T} \frac{j-N}{\sqrt{N}}} \right] \quad (48)$$

where the random variable j has a beta-binomial distribution, $BB(2N, \alpha, \beta) \equiv BB(2N, \theta N + \eta\sqrt{N}, \theta N - \eta\sqrt{N})$.

The Paul and Plackett theorem (see online Appendix for more details) states that j , appropriately shifted and scaled, converges in distribution and in moment generating function to a Normal distribution. More specifically,

$$\Psi_N \equiv \frac{j - N - \frac{\eta}{\theta}\sqrt{N}}{\sqrt{\frac{1+\theta}{2\theta}N}} \rightarrow N(0, 1)$$

where $\mathbb{E}[j] = N + \frac{\eta}{\theta}\sqrt{N}$ and $\text{var}[j] = \frac{1+\theta}{2\theta}N$. As

$$\frac{j - N}{\sqrt{N}} = \Psi_N \sqrt{\frac{1+\theta}{2\theta}} + \frac{\eta}{\theta},$$

we have

$$\begin{aligned} p_{0,0}^{-1} &= \mathbb{E} \left[e^{-\sigma\sqrt{2T} \left(\Psi_N \sqrt{\frac{1+\theta}{2\theta}} + \frac{\eta}{\theta} \right)} \right] \\ &\rightarrow \mathbb{E} \left[e^{-\sigma\sqrt{2T} \left(Z \sqrt{\frac{1+\theta}{2\theta}} + \frac{\eta}{\theta} \right)} \right] \\ &= \exp \left(-\frac{\eta}{\theta} \sigma \sqrt{2T} + \frac{\theta+1}{2\theta} \sigma^2 T \right). \end{aligned}$$

From the first to the second line, convergence of expectations follows from the fact that the beta-binomial converges to Normal in moment generating functions (for more details, see the Online Appendix). \square

Proof of Result 5:

Proof. We want to find the perceived expectation and variance of returns from 0 to t . In order to achieve that, we need to first compute $p_{m,t}$, following the lines of the proof of Result 4, and then find the limiting distribution that it has from the perspective of any investor h . We outline the main steps here, and present further details in the Online Appendix.

Define $\phi = \frac{t}{T}$ and set $m = \phi N + \psi_t \sqrt{\phi N}$, so that ψ_t is a convenient parametrization of m . Given that $z_{m+j,2N} = \lambda^{-2(m+j-N)}$, we have, similarly to equation (48)

$$p_{m,t}^{-1} = \mathbb{E}_j \left[e^{-\sigma \sqrt{2T} \frac{m+j-N}{\sqrt{N}}} \right] \quad (49)$$

where we view j as a random variable with beta-binomial distribution

$$BB \left(2(1-\phi)N, (\phi+\theta)N + (\psi_t \sqrt{\phi} + \eta) \sqrt{N}, (\phi+\theta)N - (\psi_t \sqrt{\phi} + \eta) \sqrt{N} \right).$$

By the Paul and Plackett theorem, the standardized version of j converges in distribution and in moment generating function to a standard Normal random variable. Therefore we can find the (limiting) expectation on the right hand side of (49), by just considering the expectation under a Normal distribution, with the corresponding mean and variance (for detailed calculations see the proof in the Online Appendix). As N tends to infinity, we will write $p_{\psi_t} := p_{m,t}$ (where, $\psi_t = \frac{m-\phi N}{\sqrt{\phi N}}$), to emphasize that we are considering the continuous time limit, in which ψ_t becomes the relevant state variable. We get:

$$p_{\psi_t} = b_t \cdot e^{\frac{\theta+1}{\phi+\theta} \sigma \sqrt{2\phi T} \psi_t} \quad (50)$$

where $b_t = e^{-\frac{1-\phi}{2} \frac{\theta+1}{\phi+\theta} \sigma^2 T + \frac{1-\phi}{\phi+\theta} \eta \sigma \sqrt{2T}}$.

We then view p_{ψ_t} as a function of ψ_t , for which we care about each limiting distribution. We know that $m (= \phi N + \psi_t \sqrt{\phi N})$ has a binomial distribution with mean $2\phi N h$ and variance $2\phi N h(1-h)$ from the perspective of agent h . Indeed by the Central Limit Theorem (or by De Moivre's theorem), a standardized version of m converges to a standard Normal distribution:

$$\frac{m - 2\phi N h}{\sqrt{2\phi N h(1-h)}} \rightarrow N(0, 1). \quad (51)$$

Equivalently, we have:

$$\frac{\psi_t - (2h - 1)\sqrt{\phi N}}{\sqrt{2h(1 - h)}} \rightarrow N(0, 1), \quad (52)$$

where $(2h - 1)\sqrt{N} = \frac{\eta}{\theta} + \frac{z}{\sqrt{2\theta}}$ and $h(1 - h) = \frac{1}{4} + O(\frac{1}{N})$. Therefore, the expectation and variance of $\log(p_t)$ are

$$\begin{aligned} \mathbb{E}^{(z)} \log p_t &= \frac{t(\theta + 1) \frac{z}{\sqrt{\theta}} \sigma \sqrt{T} - \frac{1}{2}(T - t)(\theta + 1)\sigma^2 T}{\theta T + t} + \frac{\eta}{\theta} \sigma \sqrt{2T} \\ \text{var}^{(z)} \log p_t &= \sigma^2 t \left(\frac{\theta + 1}{\theta + \frac{t}{T}} \right)^2. \quad \square \end{aligned}$$

Proof of Result 6:

Proof. We are interested in finding

$$\mathbb{E}^{(z)} [R_{0 \rightarrow t}] = \mathbb{E}^{(z)} \left[\begin{array}{c} p_{\psi_t} \\ p_{0,0} \end{array} \right],$$

where as in the proof of Result 5 we use the notation $p_{\psi_t} \equiv p_{m,t}$, which we have already computed in equation (50)

$$p_{0,0}^{-1} \cdot b_t \cdot \mathbb{E}^{(z)} \left[e^{\frac{\theta+1}{\phi+\theta} \sigma \sqrt{T} \sqrt{2\phi} \psi_t} \right];$$

and we have established, in equation (52), that ψ_t converges in distribution and in moment generating function to a Normal (as m does too). Hence asymptotically, the above is the expectation of a log-normal variable. In particular, after some algebra,

$$\mathbb{E}^{(z)} [R_{0 \rightarrow t}] = e^{\frac{\phi(\theta+1)}{\theta+\phi} \left[\frac{z}{\sqrt{\theta}} \sigma \sqrt{T} + \frac{\theta+1}{2} \left(\frac{1}{\theta} + \frac{1}{\theta+\phi} \right) \sigma^2 T \right]}. \quad (53)$$

Setting $\phi = \frac{t}{T}$, the proof is complete. Finally, note that by substituting $\phi = 1$ and $h = \frac{1}{2} + \frac{\eta}{2\theta\sqrt{N}} + \frac{z}{\sqrt{8\theta N}}$ we obtain equation (30). \square

Proof of Result 8:

Proof. Note that $2\phi N$ is the number of periods corresponding to $t = \phi T$.

Writing $q_{m,t}$ for the risk neutral probability of going from node $(0, 0)$ to node (m, t) , we have (as in Lemma 1) $q_{m,t} = \frac{p_{0,0}}{p_{m,t}} c_{m,t}$, where

$$c_{m,t} = \binom{2\phi N}{m} \frac{B(\alpha + m, \beta + 2\phi N - m)}{B(\alpha, \beta)}.$$

As the risk-free rate is 0, it follows that the time zero price of a call option with strike K , maturing at time t , is

$$\begin{aligned} C(0, t; K) &= \sum_{m=0}^{2\phi N} q_{m,t} (p_{m,t} - K)^+ \\ &= p_{0,0} \sum_{m=0}^{2\phi N} c_{m,t} \left(1 - \frac{K}{p_{m,t}}\right)^+ \\ &= p_{0,0} \mathbb{E} \left[\left(1 - \frac{K}{b_t} e^{-\frac{\theta+1}{\phi+\theta} \sigma \sqrt{2\phi T} \psi_t}\right)^+ \right] \end{aligned}$$

where the expectation is taken with respect to the random variable m which follows a $BB(2N\phi, \alpha, \beta)$ distribution and in the last line we have substituted $p_{m,t}$ with its (continuous time limit) value computed at equation (50) (remember, $\psi_t = \frac{m - \phi N}{\sqrt{\phi N}}$). By the result of Paul and Plackett, the asymptotic distribution of m satisfies

$$\frac{m - \phi N - \frac{\eta}{\theta} \phi \sqrt{N}}{\sqrt{\frac{\phi + \theta}{2\theta} \phi N}} \rightarrow \Psi \sim N(0, 1)$$

as $N \rightarrow \infty$. Equivalently:

$$\frac{1}{\sqrt{\frac{\phi + \theta}{2\theta}}} \left(\psi_t - \frac{\eta}{\theta} \sqrt{\phi} \right) \rightarrow \Psi \sim N(0, 1)$$

Thus

$$C(0, t; K) = p_{0,0} \cdot \mathbb{E} \left[\left(1 - \frac{K}{b_t} e^{-\frac{\theta+1}{\phi+\theta} \sigma \sqrt{2\phi T} \left(\Psi \sqrt{\frac{\phi+\theta}{2\theta}} + \frac{\eta\sqrt{\phi}}{\theta}\right)}\right)^+ \right].$$

(Note that convergence in distribution implies convergence of the expectation by the Helly-Bray theorem, since the function of Ψ inside the expectation is bounded and continuous.) This expectation is now standard, and we have

$$C(0, t; K) = p_{0,0} \left[\Phi \left(-\frac{\log(X)}{\tilde{\sigma}\sqrt{t}} \right) - e^{\frac{\tilde{\sigma}^2 t}{2}} \frac{K}{b_t} e^{-\frac{\theta+1}{\theta+\phi} \sigma \sqrt{2T} \frac{\eta\phi}{\theta}} \Phi \left(-\frac{\log(X) + \tilde{\sigma}^2 t}{\tilde{\sigma}\sqrt{t}} \right) \right]$$

where $X = \frac{K}{b_t} e^{-\frac{\theta+1}{\theta+\phi} \sigma \sqrt{2T} \frac{\eta\phi}{\theta}}$ and

$$\tilde{\sigma}^2 t = \frac{(\theta + 1)^2}{\theta(\theta + \phi)} \sigma^2 t = \text{var} \left[\log \left(\frac{K}{b_t} e^{-\frac{\theta+1}{\theta+\phi} \sigma \sqrt{2T} (\Psi \sqrt{\frac{\phi+\theta}{2\theta} \phi + \frac{\eta\phi}{\theta}})} \right) \right]$$

Finally, noting that $p_{0,0} = e^{\frac{\tilde{\sigma}^2 t}{2}} \frac{K}{b_t} e^{-\frac{\theta+1}{\theta+\phi} \sigma \sqrt{2T} \frac{\eta\phi}{\theta}}$, we arrive at the Black-Scholes formula

$$C(0, t; K) = p_{0,0} \Phi(d_1) - K \Phi(d_1 - \tilde{\sigma}\sqrt{t})$$

where

$$d_1 = \frac{\log \left(\frac{p_{0,0}}{K} \right) + \frac{1}{2} \tilde{\sigma}^2 t}{\tilde{\sigma}\sqrt{t}}$$

and volatility is determined endogenously via

$$\tilde{\sigma} = \frac{\theta + 1}{\sqrt{\theta(\theta + \frac{t}{T})}} \sigma. \quad \square$$

Proof of Result 9:

Proof. An agent's SDF links his or her perceived true probabilities to the objectively observed risk-neutral probabilities. Thus

$$M_t^{(h)}(m) = \frac{p_{0,0}}{p_{m,t}} \frac{c_{m,t}}{\pi_t^{(h)}(m)}$$

where $\pi_t^{(h)}(m)$ is the probability that we will end up at node (m, t) , as perceived by agent h . As $c_{m,t}$ has a beta-binomial distribution and $\pi_t^{(h)}(m)$ has a binomial distribution, they are each asymptotically Normal¹⁶ and we have the following

¹⁶Note that the price at 0, is given by Result 4. Moreover the asymptotic distributions of c_m and $\pi^{(h)}(m)$ are given in the proof of Result 5.

characterization for the SDF M_T :

$$M_t^{(h)}(m) \sim \sqrt{\frac{4h(1-h)\theta}{\phi+\theta}} p_{0,0} b_t^{-1} e^{-\frac{\theta+1}{\theta+\phi} \sigma \sqrt{2\phi T} \psi_t - \frac{\theta(m-\phi N - \frac{\eta}{\theta} \phi \sqrt{N})^2}{(\phi+\theta)\phi N} + \frac{(m-2\phi N h)^2}{4h(1-h)\phi N}} \quad (54)$$

where $\psi_t = \frac{m-\phi N}{\sqrt{\phi N}}$ is asymptotically Normal from the perspective of any agent h by the De Moivre–Laplace theorem.¹⁷ Parametrizing further h with z such that $h = \frac{1}{2} + \frac{\eta}{2\theta\sqrt{N}} + \frac{z}{\sqrt{8\theta N}}$, the right hand side can be rewritten as

$$M_t^{(z)}(\psi_t) \sim \sqrt{\frac{\theta}{\phi+\theta}} p_{0,0} b_t^{-1} e^{-\frac{\theta+1}{\theta+\phi} \sigma \sqrt{2\phi T} \psi_t - \frac{\theta}{(\phi+\theta)} (\psi_t - \frac{\eta}{\theta} \sqrt{\phi})^2 + (\psi_t - \sqrt{\phi}(\frac{\eta}{\theta} + \frac{z}{\sqrt{2\theta}}))^2}. \quad (55)$$

Thus $M_t^{(z)}(\psi_t)$ is asymptotically equivalent to a function of the random variable ψ_t , and hence of the variable $\Psi^{(z)} = \sqrt{2}(\psi_t - \sqrt{\phi}(\frac{\eta}{\theta} + \frac{z}{\sqrt{2\theta}}))$ which converges in distribution to a standard normal (as $\Psi^{(z)} = \frac{m-2\phi N h}{\sqrt{2\phi N h(1-h)}}$). By the continuous mapping theorem, since this function is continuous, it converges in distribution to $f(Z)$ (where $f(\cdot)$ is the corresponding function).

In order to be able to take expectations of M_t^2 (for the rest of the proof, we suppress the dependence on z in our notation) we need one additional condition. In particular we will prove that the above sequence of random variables is uniformly integrable.

For that, rewrite equation (55) as $(M_t^2)^{(N)} := D e^{A(\psi_t^{(N)})^2 + B\psi_t^{(N)} + C}$ to denote a sequence of random variables whose limiting expectation we want to find (we write $\psi_t^{(N)}$, $(M_t^2)^{(N)}$ instead of ψ_t , M_t^2 , to emphasize the dependence on N). We want to prove that there exists an $\varepsilon > 0$ such that

$$\sup_N \mathbb{E}[(e^{A(\psi_t^{(N)})^2 + B\psi_t^{(N)} + C})^{1+\varepsilon}] < \infty.$$

As L_p convergence for $p > 1$ implies uniform integrability, this will give us the result we want.

¹⁷The notation $A \sim B$ is used to denote A being asymptotically equivalent to B , or in other words: $\lim_{N \rightarrow \infty} \frac{A}{B} = 1$.

By Hoeffding's inequality,¹⁸

$$\mathbb{P}\left(\left|\frac{m - \phi N}{\sqrt{\phi N}}\right| \geq k\right) \leq 2e^{-k^2} \quad (56)$$

for any $k > 0$. As the coefficient, A , on ψ_t^2 in M_t^2 satisfies $A = \frac{2\phi}{\phi + \theta} < 1$, we can set $\varepsilon > 0$ such that $A = 1 - \varepsilon$. Then inequality (56) implies that

$$\mathbb{P}\left(e^{\frac{1}{1+\varepsilon^2} \frac{(m_t - \phi N)^2}{\phi N}} \geq x\right) \leq 2 \frac{1}{x^{1+\varepsilon^2}} \quad (57)$$

for $x > 0, \gamma > 0$.

Using this inequality together with the fact that $e^{\frac{1}{1+\varepsilon^2} \frac{(m - \phi N)^2}{\phi N}} \geq 1$ we have

$$\begin{aligned} \mathbb{E}[e^{\frac{1}{1+\varepsilon^2} \psi_t^2}] &= \mathbb{E}[e^{\frac{1}{1+\varepsilon^2} \frac{(m_t - \phi N)^2}{\phi N}}] \leq \int_0^\infty \mathbb{P}\left(e^{\frac{1}{1+\varepsilon^2} \frac{(m_t - \phi N)^2}{\phi N}} \geq x\right) dx \\ &\leq 1 + \int_1^\infty \mathbb{P}\left(e^{\frac{1}{1+\varepsilon^2} \frac{(m_t - \phi N)^2}{\phi N}} \geq x\right) dx \\ &\leq 1 + 2 \int_1^\infty \frac{1}{x^{1+\varepsilon^2}} dx \\ &= 1 + \frac{2}{\varepsilon^2} < \infty. \end{aligned}$$

Finally note that $(1 + \varepsilon)A = 1 - \varepsilon^2 < \frac{1}{1+\varepsilon^2}$. Hence there exists a constant, K , such that $(1 + \varepsilon)(A\psi_t^2 + B\psi_t + C) < \frac{1}{1+\varepsilon^2}\psi_t^2 + K$, and therefore $\mathbb{E}[e^{A\psi_t^{(N)} + B\psi_t^{(N)} + C}] < \mathbb{E}[e^{\frac{1}{1+\varepsilon^2}\psi_t^2 + K}] < \infty$. Thus our sequence is uniformly integrable, and hence there is convergence of expectations.¹⁹

We can now work towards finding the variance of M_t from the perspective of agent h . The results above imply that this problem reduces, in the limit as $N \rightarrow \infty$, to finding the moment generating function (the expectation of an exponential) of a chi-squared random variable. By computing this expectation

¹⁸Hoeffding's inequality states that if Z_1, Z_2, \dots, Z_n are i.i.d. random variables, with $Z_i \in [a, b]$, and $\bar{X} = \frac{1}{n} \sum_{i=1}^n Z_i$, then $\mathbb{E}[|\bar{X} - \mathbb{E}[\bar{X}]| \geq k] \leq 2e^{-\frac{2nk^2}{(b-a)^2}}$. In our case, m_t is the sum of $2\phi N$ i.i.d. Bernoulli variables, so the theorem can be applied.

¹⁹From equation (57) one could deduce that our sequence of random variables is dominated by the tail of a Pareto distribution, which has a finite expectation, and then use the dominated convergence theorem to reach the conclusion that there is convergence of expectations.

we find that

$$\mathbb{E}[M_t^2] = \frac{\theta}{\sqrt{\theta^2 - \phi^2}} \exp \left\{ \frac{[z\sqrt{\theta\phi} + (\theta + 1)\sigma\sqrt{\phi T}]^2}{\theta(\theta - \phi)} \right\}. \quad \square$$

Proof of Result 10:

Proof. We follow the logic of the proof of Result 9. Note, from equation (55), that $\log M_t$ is a quadratic function of ψ_t . Let us assume this quadratic has the form $F\psi_t^2 + G\psi_t + H$ for some constants F, G, H . Then this sequence of random variables converges in distribution to the corresponding quadratic of a Normal variable. By the Hoeffding inequality (56), $\mathbb{P}(2F\psi_t^2 \geq x) = \mathbb{P}(|\psi_t| \geq \sqrt{x/2F}) \leq 2e^{-x/2F}$. Thus $\mathbb{E}[2F\psi_t^2] \leq 2 \int_0^\infty e^{-x/2F} dx = 4F < \infty$, and hence $\mathbb{E}[F\psi_t^2 + G\psi_t + H] < \mathbb{E}[2F\psi_t^2 + c] < \infty$ for some constant c , which implies that the sequence is uniformly integrable. We can thus take the expectation under the corresponding normal distribution. In particular, $\frac{m-2\phi Nh}{\sqrt{2\phi Nh(1-h)}}$ converges to a standard Normal. We can then write ψ_t in terms of this random variable (as in the proof of the previous result) to find

$$\mathbb{E} \log(M_t) = \frac{[z\sqrt{\theta\phi} + (\theta + 1)\sigma\sqrt{t}]^2}{2\theta(\theta + \phi)} + \frac{1}{2} \left(\log \frac{\theta + \phi}{\theta} - \frac{\phi}{\theta + \phi} \right). \quad \square$$

Proof of Result 11:

Proof. Note that $W_T^{(z)} = W_0 \cdot R_{0 \rightarrow T}^{(z)}$, where $R_{0 \rightarrow T}^{(z)}$ is the growth optimal return from 0 to T as perceived by investor z , and W_0 is the initial endowment which equals $p_{0,0}$. As $N \rightarrow \infty$,

$$W_T^{(h)} = (M_T^{(h)})^{-1} p_{0,0} \sim p_T \sqrt{\frac{\theta + 1}{\theta}} e^{\frac{\theta}{(1+\theta)}(\psi - \frac{\eta}{\theta})^2 + (\psi - (\frac{\eta}{\theta} + \frac{z}{\sqrt{2\theta}}))^2}.$$

Substituting $\psi = \frac{m-N}{\sqrt{N}}$ and parametrizing $\sqrt{N}(2h - 1) = \frac{\eta}{\theta} + \frac{z}{\sqrt{2\theta}}$, we have

$$W_T^{(z)} = \sqrt{\frac{\theta + 1}{\theta}} \exp \left(-\frac{\psi^2}{\theta + 1} + \psi \left(\frac{2\eta}{\theta(\theta + 1)} + \frac{2z}{\sqrt{2\theta}} + \sigma\sqrt{2T} \right) - \frac{z^2}{2\theta} - \frac{2z\eta}{\sqrt{2\theta\theta}} - \frac{\eta^2}{\theta^2(\theta + 1)} \right).$$

Finally, substituting $\log(p_T) = \sigma\sqrt{2T}\psi$, we obtain Result 11. \square