Dynamic Trading When You May Be Wrong

Remorov, Alexander

Massachusetts Institute of Technology

27 April 2015

Dynamic Trading When You May Be Wrong

Alexander Remorov*

April 27, 2015

Abstract

I analyze a model with heterogeneous investors who have incorrect beliefs about fundamentals. Investors think that they are right at first, but over time realize that they are wrong. The speed of the realization depends on investor confidence in own beliefs and arrival of new information. The model provides a tractable and clear link for how changing opinions translate into equilibrium dynamics for price, holdings, and expected profits. I am able to generate a wide range of realistic market behaviors, including momentum and reversals, as well as support and resistance levels in prices due to investors being reluctant to admit they are wrong.

Keywords: Asset Pricing, Learning, Being Wrong, Heterogeneous Beliefs, Behavioral Finance

JEL Classification: G11, G12, G14

*MIT Operations Research Center (e-mail: alexrem@mit.edu).

Table of Contents

"I'm only rich because I know when I'm wrong...I basically have survived by recognizing my mistakes."

— George Soros

1 Introduction

Investing is non-trivial. An investor has to carry out a thorough analysis of available information to make a prediction about the future stock price and to initiate the appropriate position in the stock. It will take some time to determine if the prediction is correct or not – as more information will be revealed, and the stock price will evolve over time. The investor has to appropriately update his forecast as new information arrives. It is possible that he had a very optimistic forecast initially, but then some bearish news about the company are revealed. Will the investor account for this and revise his expectations?

Personal experience and academic studies tell us that this may be a difficult task. People tend to suffer from "discomfirmation bias", whereby they are reluctant to agree with arguments opposing their prior opinions[1]. In various disciplines, admitting an error may be a difficult task due to personal behavioral traits, as well as due to institutional implications (e.g. making a serious mistake may lead to a loss of a job.) For example, physicians find it difficult to deal with human error (Leape (1994)), managers often continue to invest in losing projects (Shimizu and Hitt (2004)), while admitting mistakes may prove very costly for politicians (Sheldon and Sallot (2009)). Investors tend to hold on to losers for too long, possibly because they don't want to admit they made the wrong market call[2].

I study the implications of the reluctance to admit you are wrong on investor decision-making and the financial markets. I propose a model where investors are uncertain about the true state of the market. The investors observe signals which convey partial information about this state and appropriately update their beliefs. The updating is biased in the sense that investors tend to maintain their prior beliefs and initially do not revise their opinions much. Later the investors' beliefs change significantly enough, and they realize that they are wrong. In those situations they drastically change their opinions about the market.

[1] See Lord, Loss, and Lepper (1979) and Edwards and Smith (1996).

[2] This explanation has been offered by Shefrin and Statman (1985) and Barber et. al (2007), among others.

I derive closed-form expressions for the equilibrium in the last period. These allow me to carry out some initial analysis for how the relevant variables in the model depend on agent beliefs. I show that the agent who is more convinced in his beliefs pushes the price closer to his forecast for the asset payoff. As a result, he expects to earn lower profits in comparison to the agent whose conviction is not as strong. Equilibrium holdings and expected profits rise with greater difference of opinions between the agents. However, when an agent has very strong opinions, he pushes the price so much, that his expected profits decline if his opinions become even stronger.

I then look at a simplified structure of the model, where the signal can take on one of two values (low and high), and if the state is good, the signal always takes on the high value. If the state is bad, the signal can take on either the high or the low value. The agents disagree on the probability with which the signal takes on the low value in the bad state. I find that if uncertainty about the true state persists, price rises over time because agents become more convinced that the true state is good. The holdings of an agent tend to be increasing as time passes as long as the agent believes that if the state is bad, the low signal is less likely to occur in comparison to the other agent. The patterns for expected profits are non-trivial since they depend on how the difference of opinions and the conviction of the more bullish agent change over time.

Finally, I consider the full specification for the model. Here, the signal can still take a finite number of values, but its distribution is more general. Furthermore, I incorporate the key component, whereby if an agent's belief about the true state crosses a pre-specified threshold, then the agent not only drastically changes his opinions about the state, but also about the distribution of the signal. I then use simulations to analyze the dynamics for price, holdings, and expected profits in the case when the true state is bad, but the agents initially believe that it is good. I find that price initially decreases slowly, and then later decreases faster as agents realize they may be wrong. It is actually possible for the price to increase in some situations right before an agent realizes he is wrong which is consistent with the "support and resistance levels" phenomenon observed in the financial markets. By analyzing correlations in price movement I find that price can experience both momentum and reversals, depending on the prior beliefs of the agents and the probability thresholds for changing their opinions.

My work contributes to the Differences of Opinion (DO) literature studying the behavior of agents with heterogeneous beliefs. The idea with agents differently interpreting public signals has been explored over the past two decades, with recent models proposed by Banerjee and Kremer (2010) and Barbosa (2011). In those papers agents observe signals, which are noisy estimates of the terminal payoff. Investors interpret these signals differently, which results in different valuations of the risky asset, and generates dynamic trading.

In my model, investors observe signals about the state, which is a binary variable. Therefore, at any point in time, an agent's belief comes down to his perceived probability of the true state. This allows me to explicitly study how an agent's beliefs (perceived probabilities) change with response to new information, and how they are tied to his positions in the risky asset. To my knowledge, very few other papers have looked at the two-state set-up for studying Differences of Opinion. Siemroth (2014) looks at risk-neutral traders in a binary prediction market and their information acquisition problem. Palfrey and Wang (2012) and Ottaviani and Sørensen (2015) analyze traders with heterogeneous beliefs in a binary market. My paper uses a similar framework and incorporates a novel component whereby traders can realize they are wrong and drastically switch their beliefs.

In terms of big picture, I am looking at the interaction between bulls and bears – a concept which is popular among some technical analysts and professional traders. A portion of traders believe prices will rise, while another portion believe prices will fall. Prices move because new information arrives and/or the traders change their opinions about the fundamentals. While this seems like a plausible word description of what is happening in the financial markets, my paper provides a theoretical treatment of this concept, focusing on the changing of opinions piece of the story.

2 Literature Review

My framework combines two important aspects that have been modeled in the academic finance literature. The first is Differences of Opinion (DO) and is related to how agents with heterogeneous beliefs about certain aspects of the market will trade with one another. The second is Overconfidence, whereby some investors have excessive confidence in their own opinions, and may be reluctant to revise them.

The early DO models assumed agents have private information, which is a noisy signal of a value related to fundamentals. He and Wang (1995) analyze the implications for trading volume in a model with public and private information, and random supply of the stock. Kim and Verrecchia (1994) propose a dynamic model where investors have an opportunity to obtain private information about fundamentals at a cost; this may result in greater disagreement around earnings announcements.

A more recent paper, by Banerjee and Kremer (2010), assumes agents observe public signals, but interpret them differently. Their model is used to analyze how belief dispersion relates to volume and volatility. Investors are assumed to think only their beliefs are correct and ignore the beliefs of others. Barbosa (2011) addresses this potential shortcoming by allowing investors to adjust their beliefs after observing what others think.

A few models have considered agents with different priors about the signal. This way, even if all agents observe the same signal, their posterior about the fundamentals is different. Kandel and Pearson (1995) consider a simple model with traders using different likelihood functions to interpret public information. Hong and Stein (2003) introduce short-sales constraints into a DO framework and show that this leads to a delayed release of bearish information and potential market crashes. Cujean and Hasler (2014), and Andrei, Carlin, and Hasler (2014) develop continuous settings in which investors trade while using heterogenous models about the evolution of fundamentals.

The two-state set-up in my model has been employed in a few recent papers. Siemroth (2014) looks at risk-neutral traders in a prediction market with two possible terminal outcomes. Investors can pay for informative signals. Siemroth shows that more wealthy agents will be more likely to acquire information and as a result have better forecasts. Ottaviani and Sørensen (2015) analyze traders with heterogeneous beliefs in a binary market. In equilibrium there is price underreaction to information, which leads to short-run momentum and long-run reversals. Xiouros (2011) develops a model with two risk-averse agents in an endowment model, where the endowment growth depends on a binary state and is uncertain. Palfrey and Wang (2012) use a similar model to mine with two states and two possible values of the signal. They show that heterogeneous posteriors of the agents upon observing the same public information may lead to overpricing of the risky asset.

The new ingredient in my model is agents realizing they are wrong and changing their beliefs about the distribution of the signal. To my knowledge, my paper is the first to focus on this phenomenon in the context of how investor opinions change over time and translate into price and trade dynamics.

Overconfidence can be viewed as a case of Differences of Opinion. Some agents believe their forecasts are more accurate than they actually are, and furthermore, are reluctant to change their beliefs. Daniel, Hirshleifer, and Subrahmanyam (1998) propose a prominent model for handling this phenomenon. Investors exhibit overconfidence and biased self-attribution, which leads to momentum over short horizons and corrections over longer horizons. Gervais and Odean (2001) model traders who learn their ability while subject to biased self-attribution, and show that overconfidence, and increased trading volume, will rise at the end of bull markets and fall at the end of bear markets. Scheinkman and Xiong (2003) consider a continuous model with two types of agents who overestimate the precision of their forecasts, which may generate asset bubbles.

Finally, I summarize a few influential papers in the literature providing empirical evidence for the financial market phenomena generated by my model. The first is momentum. Jegadeesh and Titman (1993) were the first to document this anomaly, whereby stocks with superior recent returns tend to continue to outperform stock with poor recent returns. Asness, Moscowitz, and Pedersen (2013) show that momentum holds for a broad range of asset classes (not just stocks), while Moscowitz, Ooi, and Pedersen (2012) document momentum in the time series of returns for securities across various asset classes and different countries.

There is also evidence for reversal in stock returns, both at very long horizons (De Bondt, Thaler, 1985) and at very short horizons (Bremer, Sweeney, 1991 and Chan, 2003). In my model both short-term and long-term reversals are possible for particular cases of agent prior beliefs about fundamentals.

The concept of resistance levels in price has been well-established among practitioners using technical analysis, but has not received much support in academia. Nevertheless, a few papers have documented the success of using support and resistance levels. Brock, Lakonishok, and LeBaron (1992) find that the trading range break rule produces a profitable investment strategy, while Osler (2000) shows that support and resistance levels provided to customers by foreign exchange trading firms have short-term predictive power. The use of support and resistance has been advocated in popular trading books, including Murphy (1999) and Lefevre (1923).

3 The Model

There is one risky asset and one riskless asset that can be traded at dates $0, 1, \ldots, T$. The riskless asset is in perfectly elastic supply and pays a zero interest rate. The risky asset is in zero net supply and pays a liquidating dividend F at time $T+1$. The distribution of the dividend depends on the state of the world $\mathcal{S}$. The state $\mathcal{S}$ can take on two possible values: 1 (good state) or 2 (bad state). For now, the state will not change throughout the horizon.

The distribution of the liquidating dividend is as follows:

$$F = \begin{cases} \mu_1, & \text{if } \mathcal{S} = 1 \\ \mu_2, & \text{if } \mathcal{S} = 2 \end{cases}$$

Without loss of generality I assume $\mu_2 = 0$ and $\mu_1 = \mu$.

There are two agents trading the assets. The agents don't know what the true state is. At date 0, agent 1 believes the probability of a good state is $p_{1,0}$; agent 2 believes this probability is $p_{2,0}$. Throughout the investment horizon agents observe signals related to the state and update their beliefs accordingly. I will use the term *perceived probability* to define the probability that the true state is the good one, as perceived by an agent.

Before each date $t = 1, \ldots, T$ a public signal s_t is revealed. The signal s_t has distribution $F_k(z)$ if the state $\mathcal{S}$ is equal to k, for $k = 1, 2$. Traders disagree about these distributions; trader i believes that the distribution of the signal is $F_{i,k}(z)$ if the state $\mathcal{S}$ is equal to k.

The set-up so far assumes the agents do not change their perceived distributions of the signal in the states. Later I will incorporate the phenomenon whereby if the perceived probability of the good state changes significantly enough (in comparison to the initial perceived probability), then the agent drastically changes his opinions about the distribution of the signal.

Agents 1 and 2 have initial wealth $W_{1,0}$ and $W_{2,0}$, respectively. They maximize mean-variance utility over terminal wealth $W_{i,T+1}$:

$$U_i = \mathbb{E}_i(W_{i,T+1}) - \frac{\lambda}{2} var_i(W_{i,T+1})$$

where the mean and variance are computed using the beliefs of agent i.

I will restrict my equilibrium definition to assume that both agents know the beliefs of both themselves and the other agents, the way in which updating is done, and the preferences. Hence, at every date, an agent will submit his demand as if he knows what the beliefs of *both* agents are, and how the price will evolve in every possible scenario during the subsequent dates (because the agents behave as if there is perfect information). Since there are only two agents, an agent who knows his own demand and the price in the previous period can determine the demand of the other agent. Therefore, it is intuitive to assume that the agent knows the beliefs of the other agent as well. That being said, it may be possible that other equilibria exist; I have not explored this possibility in this paper and rather focus on a simple equilibrium which captures how investor opinions translate into prices, holdings, and profits.

I start with a simplified model, where the signal can take on one of two values (low and high), and if the state is good, the signal always takes on the high value. I derive closed-form recursive formulas for equilibrium price and expected profits of each trader. The model provides some basic intuition for how investor disagreement affects the market, and how it changes over time. I then consider a more general setting for the distributions of the signals.

For the rest of the paper, I will use the following conventions. Period t is the period of time between dates $t-1$ and t. $\mathcal{I}_t$ the information set at the end of period t (after the signal s_t is revealed). $\mathbb{E}_{i,t}(X)$ and $var_{i,t}(X)$ are the expected value and variance, respectively, of a random variable X as perceived by agent i at time t. For an event A, $\mathbb{P}_{i,t}(A)$ is the perceived probability of event A by agent i at time t.

4 Simple Model

4.1 Model Set-up

The distribution of the signal depends on the state as follows.

$$\text{if } \mathcal{S} = 1 \text{ then } s_t = 1 \text{ wp (with probability) } 1$$

$$\text{if } \mathcal{S} = 2 \text{ then } s_t = \begin{cases} 1, & \text{wp } 1 - r \\ -1, & \text{wp } r \end{cases}$$

Thus, if the state is good, the signal only takes on the *high value* 1, while if the state is bad, the signal may take on the *high value* 1 or the *low value* -1. Agents agree on the distribution of the signal in the good state, but disagree on the value of r in the bad state. If the state is bad, agent i believes that the signal is low with probability r_i, for $i = 1, 2$.

I assume the state does not change throughout the whole horizon. Furthermore, the agents maintain their perceptions r_1 and r_2 and do not adjust them due to new information.

4.2 Probability Updating

Agents update their beliefs about the true state based on the signals. Denote by $p_{i,t}$ the perceived probability for agent i during period t, after signal s_t has been revealed. More formally, $p_{i,t} = \mathbb{P}_{i,t}(\mathcal{S} = 1)$. After observing the value of the signal z, agent i updates this probability to $p_{i,t+1}$ according to the Bayes Rule:

$$p_{i,t+1} = \frac{\mathbb{P}_{i,t}(s_{t+1} = z|\mathcal{S} = 1)\mathbb{P}_{i,t}(\mathcal{S} = 1)}{\mathbb{P}_{i,t}(s_{t+1} = z|\mathcal{S} = 1)\mathbb{P}_{i,t}(\mathcal{S} = 1) + \mathbb{P}_{i,t}(s_{t+1} = z|\mathcal{S} = 2)\mathbb{P}_{i,t}(\mathcal{S} = 2)} \quad (1)$$

We now make the following observation. If at any time the agents observe a low signal then they immediately recognize the true state is bad with probability 1. In economic terms, these are situations when a big shock occurs. An agent who thought the true state is probably a bad one realizes he is correct, whereas an agent who thought the state is probably a good one realizes he is wrong.

Consider date t. If for all times $\tilde{t} \leq t$, the signal $s_{\tilde{t}}$ was equal to 1, the agents are still uncertain what the state is and only have some beliefs $p_{i,t}$ on the probability that the state is good. Now period $t+1$ begins and signal s_{t+1} is revealed. In view of (1) the updating of the probabilities $p_{1,t}, p_{2,t}$ is as follows. If a shock occurs:

$$p_{i,t+1} = 0 \text{ if } s_{t+1} = -1 \tag{2}$$

Otherwise, there is still uncertainty about the state, and we have:

$$p_{i,t+1} = \frac{p_{i,t}}{p_{i,t} + (1-p_{i,t})(1-r_i)} \text{ if } s_{t+1} = 1 \tag{3}$$

Thus the updating rule is different for the agents if uncertainty persists, but there is agreement on the state in case of a shock.

4.3 Equilibrium

I solve the model backwards to obtain closed-form formulas for the equilibrium.

First, I introduce some more notation. Let $W_{i,t}$ be the wealth of agent i at time t and $\Pi_{i,t}$ be the profits for the remaining horizon. Let P_t be the equilibrium price at time t. Because of each agent's belief at time t reduces to his perceived probability $p_{i,t}$ of the state being good, and the agents only update this probability over time, then at any time t the state of the system depends only on the time t and the perceived probabilities $p_{1,t}, p_{2,t}$. Therefore, the time t profits $\Pi_{i,t} = \Pi_{i,t}(p_{1,t}, p_{2,t})$ and price $P_t = P_t(p_{1,t}, p_{2,t})$ can be written as functions of $p_{1,t}$ and $p_{2,t}$ only.

At time t agent i is maximizing:

$$U_{i,t} = \mathbb{E}_{i,t}(W_{i,t} + \Pi_{i,t}) - \frac{\lambda}{2} var_{i,t}(W_{i,t} + \Pi_{i,t})$$

Since $W_{i,t}|\mathcal{I}_t$ is constant, the problem is equivalent to maximizing:

$$\mathbb{E}_{i,t}(\Pi_{i,t}) - \frac{\lambda}{2} var_{i,t}(\Pi_{i,t}) \tag{4}$$

If $s_t = -1$ at any time $t \leq T$, both agents know with perfect certainty that the true state is bad and the final payoff will be 0. Therefore the price at date t and all subsequent dates is zero; equilibrium holdings are zero, and PnL for the remaining horizon is also zero.

I now consider what happens at times when $s_t = 1$. Denote by $x_{i,t}$ the equilibrium holdings of agent i at time t.

4.4 Equilibrium – Last Period

Consider date T and suppose so far the signal has been equal to 1 all the time. Then the equilibrium price and holdings take the following simple form:

Proposition 4.1. *Suppose $s_t = 1$ for $t \leq T$. Then the equilibrium price is:*

$$P_T(p_{1,T}, p_{2,T}) = \frac{p_{1,T}p_{2,T}(2 - p_{1,T} - p_{2,T})}{p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T})}\mu \tag{5}$$

The expected value of profits for agent i, $i = 1, 2$, is:

$$\mathbb{E}_{i,T}(\Pi_{i,T}(p_{1,T}, p_{2,T})) = \frac{(p_{1,T} - p_{2,T})^2 p_{i,T}(1 - p_{i,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))^2} \tag{6}$$

Proof. See Appendix. □

Proposition 4.1 allows me to get some initial intuition about how the equilibrium depends on the difference in beliefs. I analyze how the risky asset price, agent holdings[3], and expected profits depend on the perceived probabilities $p_{1,T}$, $p_{2,T}$ of the good state.

I find that the agent with "the stronger" belief pushes the price in his direction. Because the price is closer to his expected payoff, the agent expects lower profits than the agent with beliefs that are not as strong. Consistent with general intuition, equilibrium holdings and expected profits increase with greater disagreement between the agents. However, I obtain an interesting result whereby if the beliefs of an agent are strong enough, the he expects to earn *lower* profits if his beliefs become even stronger.

The parameters involved are the payoff in good state μ, the risk aversion coefficient λ, and the two perceived probabilities $p_{1,T}, p_{2,T}$. The payoff in the good state enters as a multiplicative factor in the price and the holdings and so is not too important. Similarly,

[3] The formula for agent holdings is provided in the proof of Proposition 4.1 in the Appendix.

the risk aversion coefficient only enters as a multiplicative factor in the holdings and the expected price. Therefore, I fix $\mu = 1$ and $\lambda = 3$ and only vary the perceived probabilities. For notational convenience, I drop the time subscript and use $p_1 = p_{1,T}, p_2 = p_{2,T}$.

I consider p_1 varying in the interval $[0\%, 100\%]$ and three values for p_2: 20%, 50%, and 80%. Due to the symmetry in the model I need to consider the full range of values for only one of the probabilities, and I do this for probability p_1. Note that for values of p_1 equal to 0% and 100% I replace the values for the price, holdings, and expected profits with their limits; the formulas for these limits are provided in the Appendix.

Figure 4.1 plots the equilibrium price. We see that it is increasing in the perceived probabilities of the agents. The increase is non-linear, and more rapid for "strong" beliefs of agent 1, i.e. perceived probabilities that are close to 0% or to 100%. When agent 2 has weak beliefs about the true state with $p_2 = 50\%$, the price is increasing close to linearly in p_1.

To understand which agent has a larger impact on the price as a result of his beliefs, I look at which of the two expected payoffs of the agents the price is closest to. With $\mu = 1$ the expected payoff of agent i is just his perceived probability p_i. I can write the equilibrium price P as a weighted average $wp_1 + (1 - w)p_2$ of the probabilities, and analyze how w depends on these probabilities. A larger value of w signifies a larger weight on the beliefs of agent 1. From Proposition 4.1 it follows that the formula for w is very simple:

$$w = \frac{p_2(1 - p_2)}{p_1(1 - p_1) + p_2(1 - p_2)} \tag{7}$$

Formula (7) shows that the weight is symmetric for p_1 around the value of 0.5, and furthermore a stronger opinion by the agent results in a smaller value for $p_1(1 - p_1)$ in the denominator, and therefore a larger weight. I conclude that if an agent is more sure that the state is good (p_1 close to 100%) or that the state is bad (p_1 close to 0%) then he pushes the price more in the direction of his belief. Figure 4.2 plots the weight as a function of the perceived probabilities. We see that the weight on the beliefs of agent 1 is low and changes slowly for values of perceived probability around 0.5, but rises fast as this probability gets close to 0% and 100%. It also looks quite parabolic as a function of p_1 (although the actual

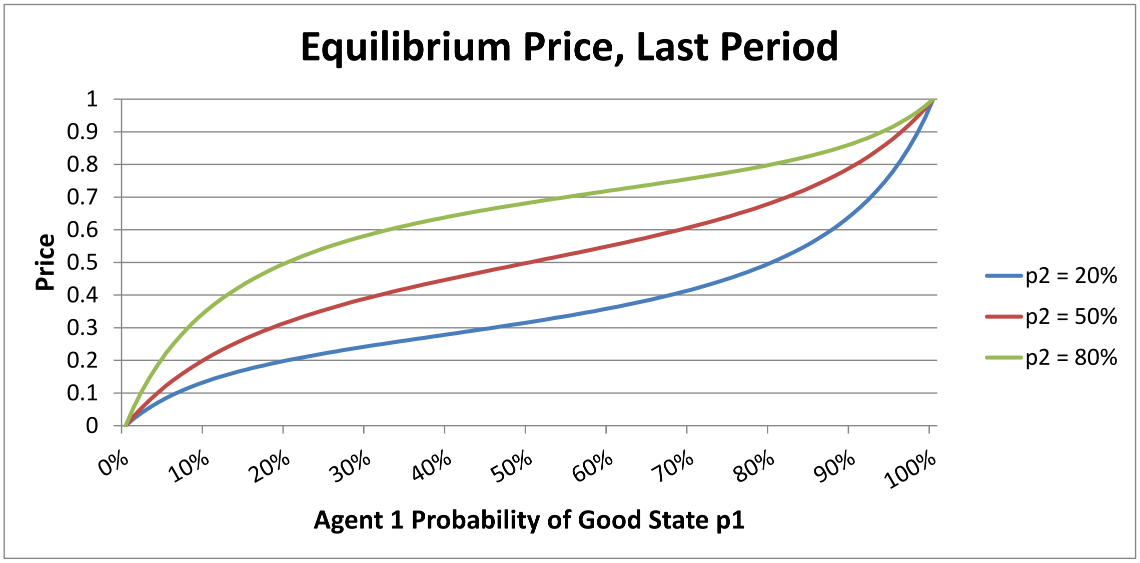

Figure 4.1: Equilibrium price in the last period, as a function of the agents' perceived probabilities p_1 and p_2 of the good state.

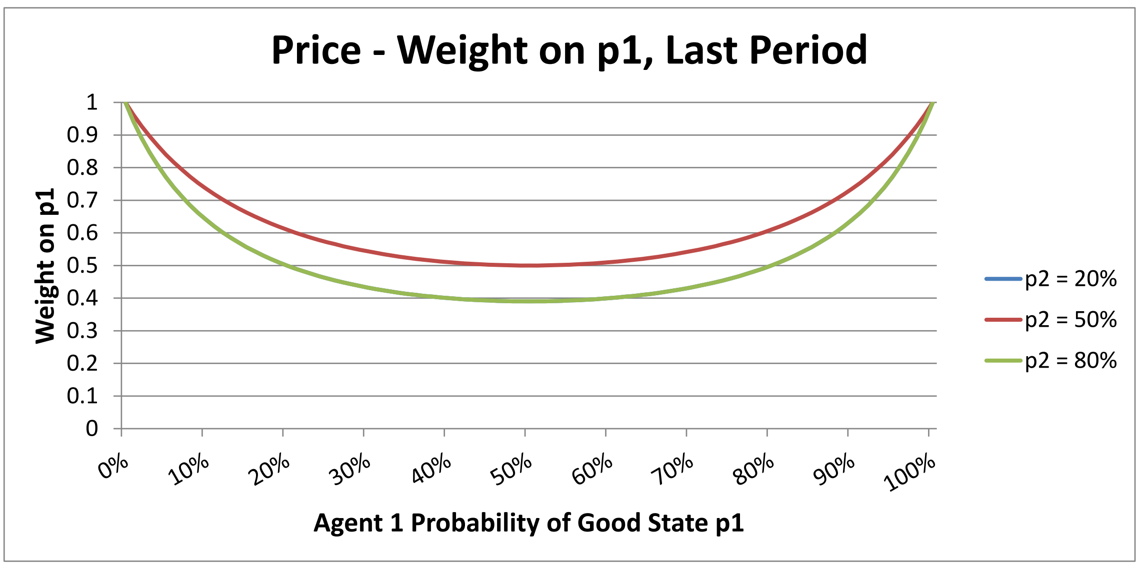

Figure 4.2: Weight w on the expected payoff p_1 of agent 1 so that the equilibrium price P is equal to the weighted average $wp_1 + (1 - w)p_2$ of the expected payoffs by the agents. The curves for $p_2 = 20\%$ and $p_2 = 80\%$ overlap exactly.

form is not a parabola).

I next look at the equilibrium holdings. Because net supply of the risky asset is zero, it suffices to only consider the holdings of agent 1, as the negative value of these holdings gives the holdings of agent 2. We see that the holdings are zero when agents have the same beliefs, and is positive if and only if agent 1 has a higher perceived probability of the good state in comparison to agent 2. The holdings are increasing in the beliefs of the agent, since as the agent becomes more bullish, he is inclined to take larger and more positive positions in the risky asset. Finally, as with the price, equilibrium holdings change more rapidly with p_1 as p_1 gets closer to 0% or 100%.

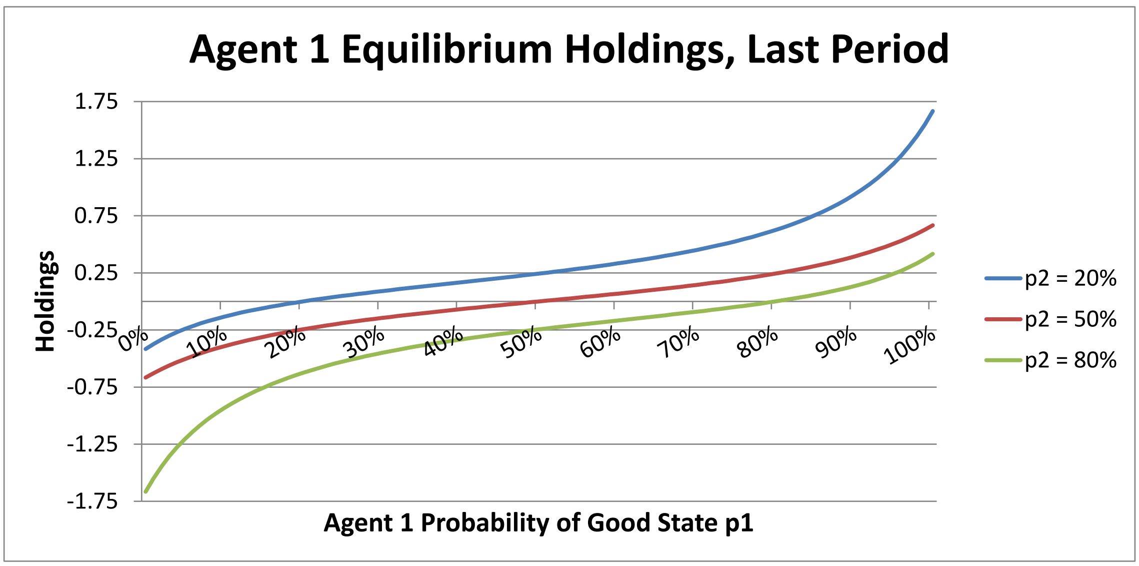

Figure 4.3: Equilibrium holdings of agent 1 in the last period. Note that the holdings of agent 2 are just the negative of the holdings of agent 1.

I also calculate equilibrium profits for the agents. Even though the agents have symmetric holdings, their expected PnL is not the same because they have different beliefs on the distribution of the payoff. Figure 4.4 shows the expected profits for agent 1. We see that the expected PnL is zero when the agents have the same beliefs, and rises as the difference in the opinions of the agents increases. However, I observe a very interesting effect. As the belief of agent 1 becomes strong enough (p_1 close enough to 0% or 100%), his expected profits actually start to fall. The fall is quite rapid when the difference of opinions is large (e.g. when $p_2 = 20\%$, and p_1 gets close to 100%, see the blue curve). The fall in profits is caused by the fact that agent 1 pushes the price in his direction so much, that it gets very close to his expected payoff, and so he expects to make much less *per unit* of holdings of the risky asset. Even though the magnitude of his holdings increases, it is not enough to offset the decrease in this expected PnL per unit of holdings, and so expected profits fall.

Figure 4.5 plots the profits for agent 2. As with agent 1, the expected profits are zero when the agents have the same beliefs and rise as these beliefs diverge more. This rise is slower than for agent 1 when p_1 is close to 50%. The intuition behind this is that in such situations the weight w of agent 1's opinion on the price is lower, so agent 2 is pushing the price more, and hence expecting lower profits. When p_1 gets further away from 50%, agent 1 is the one who starts pushing the price, so there is a large deviation between the price and the beliefs of agent 2, therefore agent 2 expects to make more money in such situations. For extreme values of p_1 close to 0% and 100%, the expected PnL of agent 2 rises drastically.

Note that the expected PnL of agent 2 never gets arbitrarily large. If $p_{1,T} = 0$, using (6) the expected profits for agent 2 are:

$$\mathbb{E}_{1,T}(\Pi_{1,T}(0, p_{2,T})) = \frac{p_{2,T}}{\lambda(1 - p_{2,T})} \tag{8}$$

The formula for the case when $p_{1,T} = 1$ is similar.

I conclude that during the last period price is increasing in the beliefs of the agents, and is changing more rapidly as these beliefs become stronger. Equilibrium holdings increase with a larger disagreement, and the agent with the more bullish forecast on the payoff holds a positive amount of the asset. Finally, each agent is expecting higher profits as his opinions diverge more in the comparison with the other agent. The exception to this phenomenon is when the agent has very strong opinions (perceived probability of the good state close to 0% or to 100%), in which case he pushes the price so much in the direction of his belief, that his expected profits go down.

4.5 Equilibrium - Full Horizon

I now show how to recursively calculate the equilibrium over the full time horizon. Consider date t (with $t \leq T - 1$) so that signal s_t has already been revealed. I again assume for

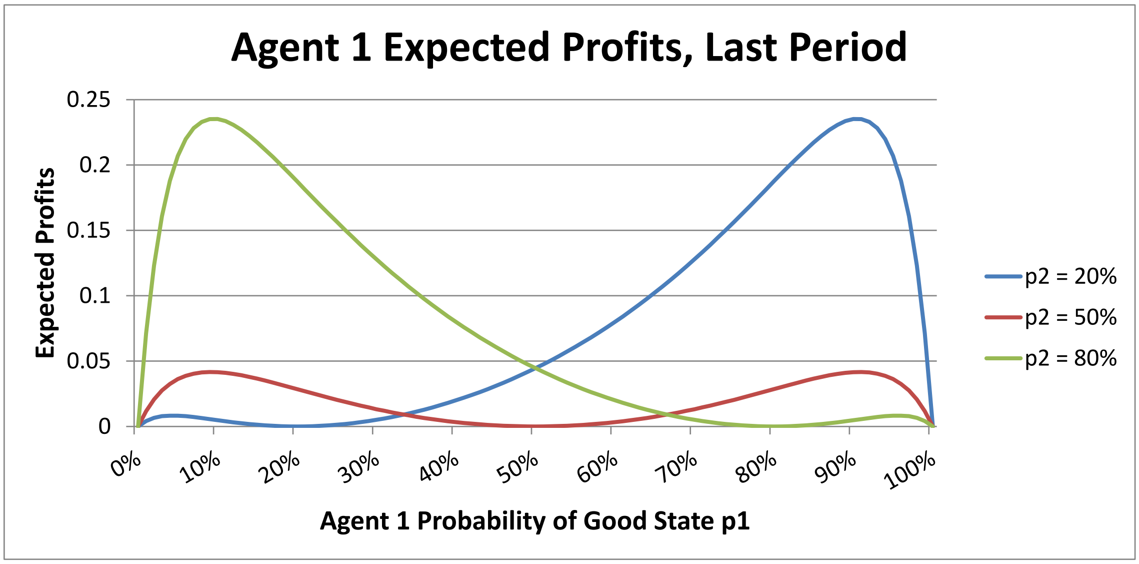

Figure 4.4: Expected profits of agent 1 in the last period, as a function of the agents' perceived probabilities p_1 and p_2 of the good state.

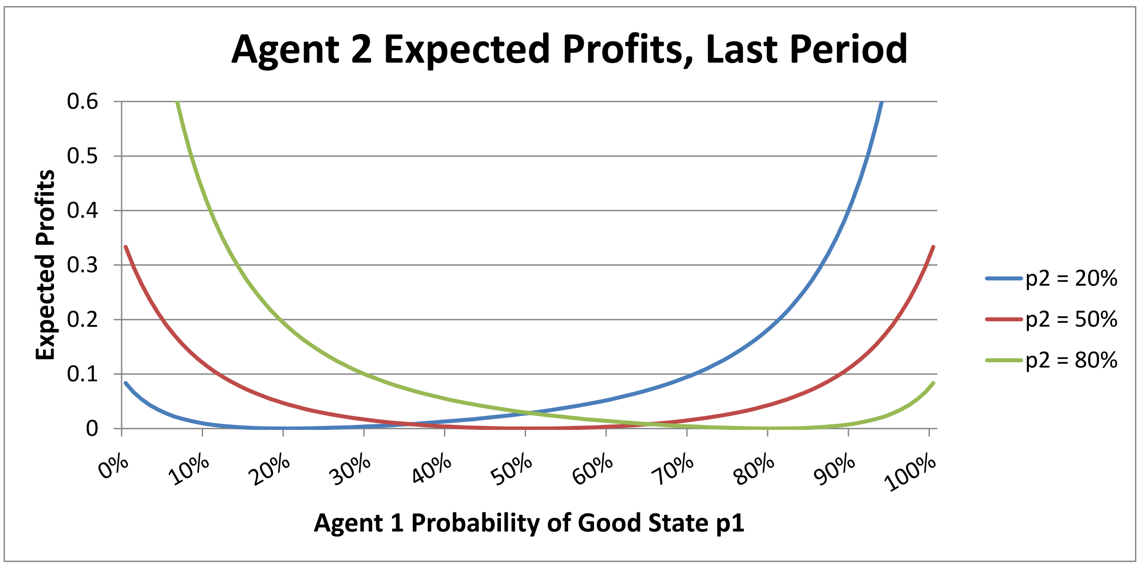

Figure 4.5: Expected profits of agent 2 in the last period, as a function of the agents' perceived probabilities p_1 and p_2 of the good state.

all times $\tilde{t} \leq t$, the signal $s_{\tilde{t}}$ was equal to 1, so that the agents are still uncertain what the state is.

Assume the price the agents face is P. Consider agent 1. He considers holding x units of the risky asset and cares about profits $\Pi_{1,t}$ for the remaining horizon. There are two cases of what can happen next period depending on what the value of the signal s_{t+1} is.

Case 1: $s_{t+1} = -1$. Then agents realize $\mathcal{S} = 2$ for sure, and price at time $t + 1$ is 0. The PnL of agent 1 is $x(0 - P)$.

Case 2: $s_{t+1} = 1$. Then the agents are not sure what the true state is, and update their perceived probabilities of the state. The PnL of the agent is:

$$\Pi_{1,t} = x(P_{t+1}(p_{1,t+1}, p_{2,t+1}) - P) + \Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1})$$

Since agent 1 believes $\mathcal{S} = 1$ with probability $p_{1,t}$, I can easily derive his perceived probabilities of occurrence of cases 1 and 2. From there, I can fully describe the distribution for this profits:

$$\Pi_{1,t} = \begin{cases} x(0 - P), \text{ wp } (1 - p_{1,t})r_1 \\ x(P_{t+1}(p_{1,t+1}, p_{2,t+1}) - P) + \Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1}), \text{ wp } 1 - (1 - p_{1,t})r_1 \end{cases} \tag{9}$$

The above relation shows how $\Pi_{1,t}$ recursively depends on $\Pi_{1,t+1}$. Using the mean-variance preferences of the agents I derive how his expected profits and the price depend on the beliefs of the agents, the next period expected profits, and the next period price. In particular, I obtain relatively simple formulas for how current period price and *sum* of expected profits depend on the next period price and sum of expected profits for the agents.

Define $D_{i,t}(p_{1,t}, p_{2,t})$ as the expected profits for agent i at time t:

$$D_{i,t}(p_{1,t}, p_{2,t}) = \mathbb{E}_{i,t}(\Pi_{i,t}(p_{1,t}, p_{2,t}))$$

Because the profit $\Pi_{i,t}$ depends only on time t and the perceived probabilities, then so does the expected profit $D_{i,t}$. Define $D_t(p_{1,t}, p_{2,t})$ as the sum of the expected profits of the two agents at time t:

$$D_t(p_{1,t}, p_{2,t}) = D_{1,t}(p_{1,t}, p_{2,t}) + D_{2,t}(p_{1,t}, p_{2,t})$$

Proposition 4.2. *Let $t \le T-1$ and suppose $s_{\tilde{t}} = 1$ for $\tilde{t} \le t$. Let $p_{1,t+1}$ and $p_{2,t+1}$ be the perceived probabilities of the good state if the signal s_{t+1} is equal to 1. Then the following relation holds for equilibrium price:*

$$P_t(p_{1,t}, p_{2,t}) = \frac{b_1 b_2(2-b_1-b_2) - \lambda b_1 b_2(1-b_1)(1-b_2)D_{t+1}(p_{1,t+1}, p_{2,t+1})}{b_1(1-b_1) + b_2(1-b_2)} P_{t+1}(p_{1,t+1}, p_{2,t+1}) \tag{10}$$

where:

$$b_1 = 1 - (1-p_{1,t})r_1 \; ; \; b_2 = 1 - (1-p_{2,t})r_2 \tag{11}$$

For the sum of expected profits, we have:

$$D_t(p_{1,t}, p_{2,t}) = \frac{(b_1-b_2)^2 + \lambda b_1 b_2(2-b_1-b_2)D_{t+1}(p_{1,t+1}, p_{2,t+1})}{\lambda(b_1(1-b_1) + b_2(1-b_2))} \tag{12}$$

Proof. See Appendix. □

The formulas (10) and (12) allow me to fully solve for the equilibrium prices in the model. In the appendix I also provide the recursive formulas for expected profits for each agent individually. They are slightly more complicated than the ones above, but are still closed-form.

I analyze how equilibrium prices, holdings, and expected profits change over time. I also look at how they depend on the disagreement of the agents, and their beliefs about the distribution of the signal.

I find that price tends to rise as agents become more convinced that the true state is the good one. The holdings of an agent i tend to be increasing or decreasing over time depending on if probability r_i is higher or lower, respectively, in comparison to the other agent. Expected profits are often non-monotone functions of time because they depend on the time-varying magnitude of the difference of beliefs, as well as by how much the more bullish agent is pushing the price.

As with the last period, I fix the values $\mu = 1$ and $\lambda = 3$ for the payoff in the good state and the risk aversion coefficient, respectively. There are four more parameters in the model: the initial perceived probabilities of the good state $p_{1,0}, p_{2,0}$ and the probabilities relating to the beliefs about the signal distribution r_1, r_2. For simplicity I use $p_1 = p_{1,0}$ and $p_2 = p_{2,0}$.

I fix the beliefs of the second agent at $p_2 = 20\%$, $r_2 = 20\%$, and only vary the beliefs of the first agent. I consider four cases: p_1 equal to 10% and 30%, and r_1 equal to 10% and 30%. I found that the behavior of the model is similar for other parameter values[4].

Note that I look at the relevant variables in the model *assuming* the value of the signal in each period is high. This way there is still uncertainty in the model as time progresses. Recall that if the signal ever takes on the low value, both agents realize the true state is bad, equilibrium price drops to 0, and there is no more trading.

Figure 4.6 shows price as a function of time for the four cases of parameters. Price is increasing with time, because for each subsequent period when a high signal is revealed, both agents become more confident that the true state is good, and hence expect a larger terminal payoff. The speed of this rise depends on how fast the perceived probabilities $p_{i,t}$ increase. For low values of r_1, equal to 10%, agent 1 faces more uncertainty about which state is good when he observes a high signal. Therefore his belief that the true state is good does not increase as much, and the corresponding equilibrium price rises slowly (see the blue and the green curves on the chart). For higher values of r_1, equal to 30%, this increase in beliefs occurs faster, leading to a more rapid rise in price (see the red and the purple curves).

I next analyze how holdings evolve over time. Figure 4.7 plots the equilibrium holdings of agent 1; the negative of these holdings gives the equilibrium holdings of agent 2 since net supply of the risky asset is zero. I also plot in figure 4.8 the perceived probabilities of the high signal occurring in the next period[5]. Comparing the two charts we see that the agent with the higher perceived probability for a high signal tends to hold a positive amount of the asset. For example, for the case $p_1 = 10\%, r_1 = 30\%$ (red curve), agent 1 starts out with a negative position in the early periods, and over time this position increases and turns positive around date 11. From figure 4.8 we see that agent 1 (red curve) initially has a lower

[4]Note that for these parameter values, the probabilities b_1 and b_2 never get too close to 0% or 100%, so we never need to deal with "degenerate" cases for the formulas in Proposition 4.2.

[5]These are the probabilities $b_{i,t}$ in Proposition 4.2.

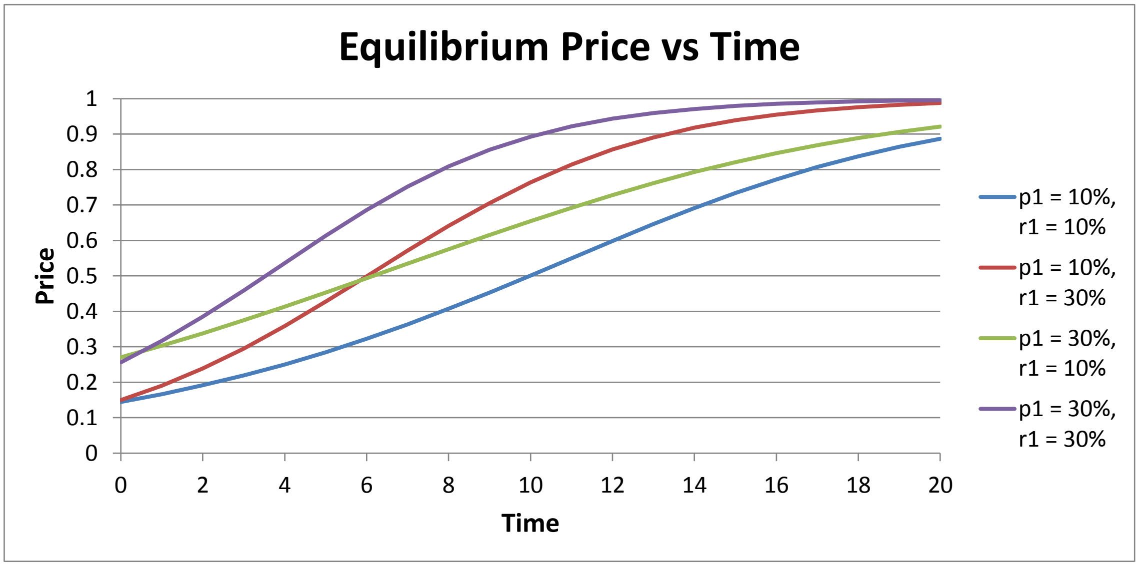

Figure 4.6: Equilibrium price as a function of time. The beliefs of agent 2 are fixed at $p_2 = 20\%$, $r_2 = 20\%$, while the beliefs p_1 and r_1 of agent 1 vary from 10% to 30%.

perceived probability for the high signal next period in comparison to agent 2 (light blue curve). The red curve rises faster and around date 11 becomes higher than the light blue curve. Thus the time when an agent becomes more bullish in terms of the signal coincides with the time when his position turns positive. The same phenomenon is observed for the other three cases.

Holdings for agent 1 are increasing with time for high values of r_1 (red and purple curves), and decreasing for low values of r_1 (blue and green curves). The intuition behind this phenomenon is as follows. When $r_1 = 10\%$, agent 1 realizes that his perceived probability of a low signal, conditional on the state being bad, is lower than for agent 2 (since $r_2 = 20\%$). Therefore, every time the agents observe a high signal, agent 1 will be less convinced the true state is good relative to his earlier belief, in comparison to agent 2. (This is assuming both agents have the same prior on the probability of the good state). As more high signals arrive, agent 2 will increase his perceived probability of the good state faster than agent 1, and towards the end of the horizon this probability will be higher for agent 2 than for agent 1, so agent 2 will be more bullish and expect to hold a positive amount of the stock. Furthermore, since the perceived probability $p_{2,t}$ will be close to 100% at that point, he will be demanding a large amount in the stock due to his very high confidence in beliefs. Both agents realize this, and therefore the holdings of agent 1 will decrease more and more with time, as observed in the chart. When $r_1 = 30\%$, the same logic applies, except now agent 1 becomes more bullish than agent 2 with time.

Finally, I look at the expected profits of the agents. Figure 4.9 shows the expected PnL of agent 1, while figure 4.10 plots the PnL of agent 2.

I first compare the expected profits between the two agents. As we get closer to the end of the time horizon, both agents are quite bullish, with at least one of them having a very strong opinion on the state being a good one ($p_{i,t}$ close to 100%). If agent 1 is the more bullish one (red and purple curves), he is pushing the price a lot and expecting lower profits than agent 2. If agent 2 is the more bullish one (blue and green curves), he is the one pushing the price and hence expecting lower profits than agent 1. Recall that whether an agent is more bullish or not towards the end of the time horizon depends on whether r_i for them is higher than for the other agent. We can therefore conclude that the agent with the higher value of r_i will be expecting lower profits than the other agent after a large amount of time has passed, and a lot of signals have been revealed. Since the agents "work backwards" to calculate expected profits in earlier periods, and their beliefs about the unconditional probability of the signal being equal to 2 are similar, it follows that the agent with the higher value of r_i will be expecting lower profits than the other agent *throughout the whole horizon*.

I also consider how expected PnL changes over time. I will look at each of the four cases individually. For the case $p_1 = 10\%$, $r_1 = 10\%$, agent 1 is initially less bullish than agent 2, and as time passes by remains less bullish (since $r_1 < r_2$). At the same time, agent 2 becomes even more convinced the true state is good, so that $p_{2,t}$ becomes close to 100%. He ends up pushing the price more, so agent 1 expects higher profits after more signals are revealed; we see this is in the blue curve in figure 4.9. As for agent 2, in the early periods, he becomes more bullish than agent 1, while the difference of opinions between the two agents increases. As a result, he expects higher profits as more signals are revealed. Towards the middle and end of the investment horizon the difference in opinions starts to decrease, while the belief $p_{2,t}$ of agent 2 becomes close to 100%, so that he pushes the price so much, that he now expects to make lower profits than in earlier periods. This behavior is evident from the blue curve in figure 4.10.

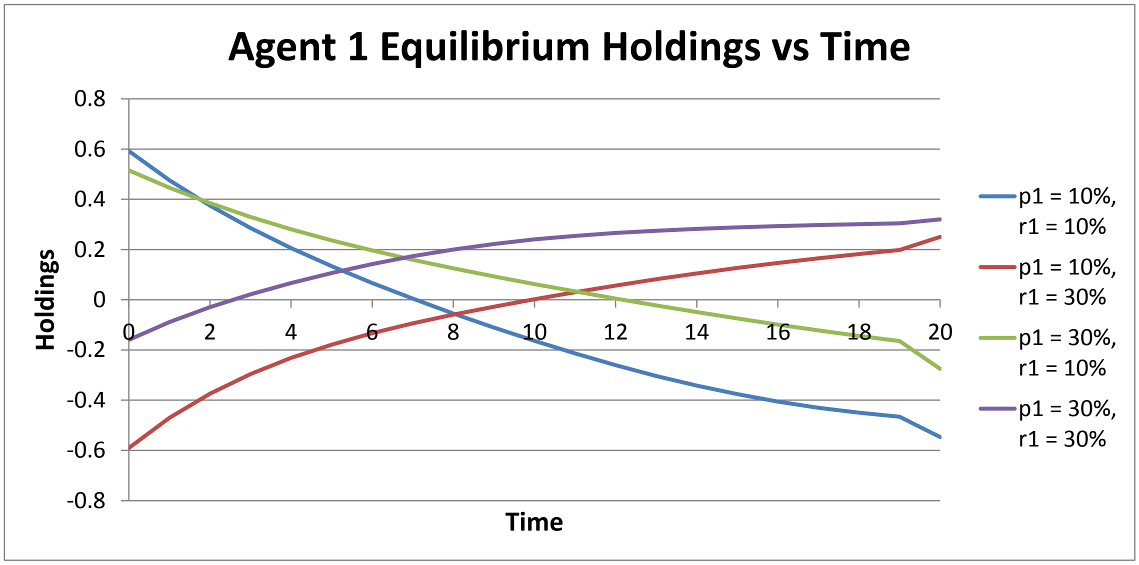

Figure 4.7: Equilibrium holdings of agent 1 as a function of time. Note that the holdings of agent 2 are just the negative of the holdings of agent 1.

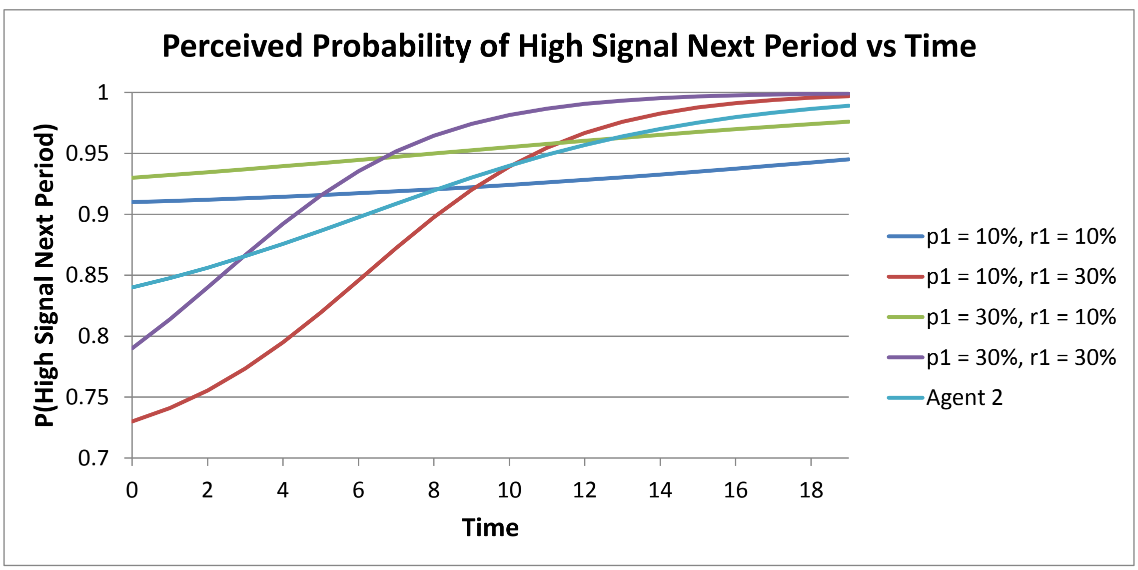

Figure 4.8: Probability of high signal occurring in the next period, as perceived by each agent. The dark blue, purple, red, and green curves correspond to probabilities as perceived by agent 1 for the different cases for p_1 and r_1. The light blue curve corresponds to the probability as perceived by agent 2 with $p_2 = 20\%$ and $r_2 = 20\%$.

When $p_1 = 30\%$, $r_1 = 10\%$, the behavior for agent 1 towards the end of the horizon is similar to when $p_1 = 10\%$ – namely agent 1 is less bullish than agent 2, while agent 2 pushes the price, so agent 1 expects larger profits. However, during the early periods agent 1 is more bullish than agent 2, and the difference between their opinions decreases in subsequent periods. Therefore in those periods expected profits for agent 1 fall. After that agent 2 becomes more bullish than agent 1, and expected profits start to rise again – as we see in figure 4.9 from the green curve. The expected PnL for agent 2, shown in figure 4.10 (green curve), follows a different pattern. Initially, agent 2 is less bullish, however the difference of opinions decreases with time, so expected profits also decrease. After that, agent 2 is more bullish and pushing the price a little, so his expected profits rise marginally or decrease.

If $p_1 = 10\%$, $r_1 = 30\%$, agent 1 is initially less bullish, but becomes more bullish with time. This is the same case as for agent 2 in the case $p_1 = 30\%$, $r_1 = 10\%$, so the same pattern is observed: a rapid decrease in expected profits (as difference of opinions decreases), then a slight increase or decrease in subsequent periods as agent 1 starts to push the price (red curve in figure 4.9). For agent 2 we actually observe the same pattern – because even though agent 2 is less bullish towards the end of the horizon, his beliefs are also very strong, so he is expecting lower profits.

The last case is $p_1 = 30\%$, $r_1 = 30\%$. Here, agent 1 is more bullish than agent 2 throughout the whole investing horizon; however the difference in opinions first increases, reaches a maximum at date 6, and then decreases. Both agents become very bullish with time and so expect lower profits towards the end of the horizon. Initially, the expected profits for both of them rise (as difference of opinions increases), and then start to fall. We see this in the purple curves in figures 4.9 and 4.10.

The simple model gives us a clear picture of how difference in beliefs and strength of opinions influence equilibrium price, holdings, and expected profits. Price rises over time, and rises more rapidly when the agents become more bullish faster. The holdings of an agent tend to be increasing (decreasing) if the probability of a low signal in the bad state for the agent is higher (lower) than for the other agent. This is because towards the end of the horizon the agent with the higher value of this probability will be the more bullish one.

The patterns for the expected profits vary significantly depending on the parameter

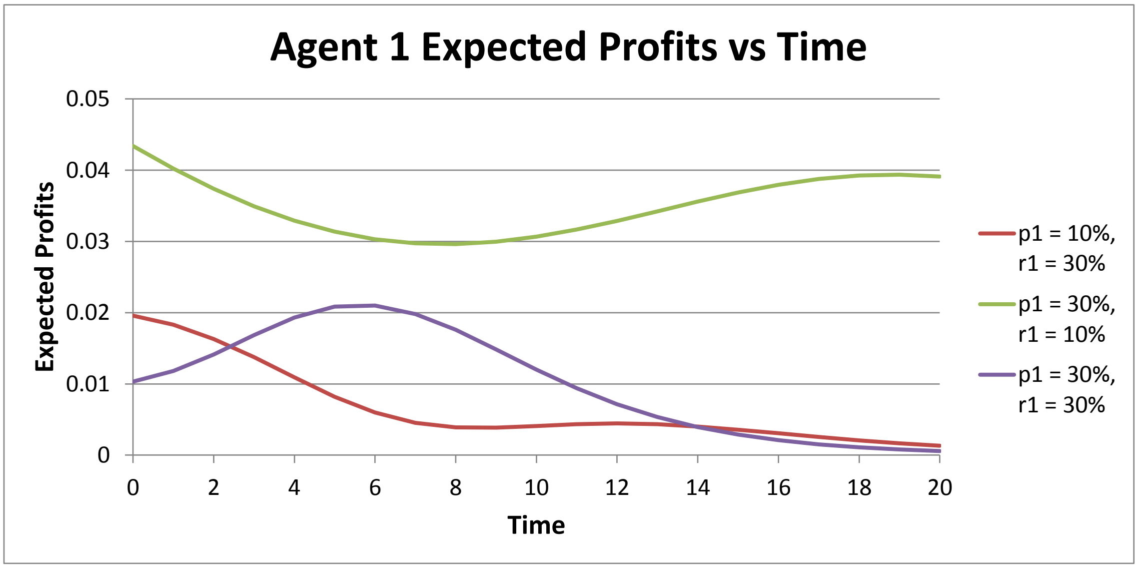

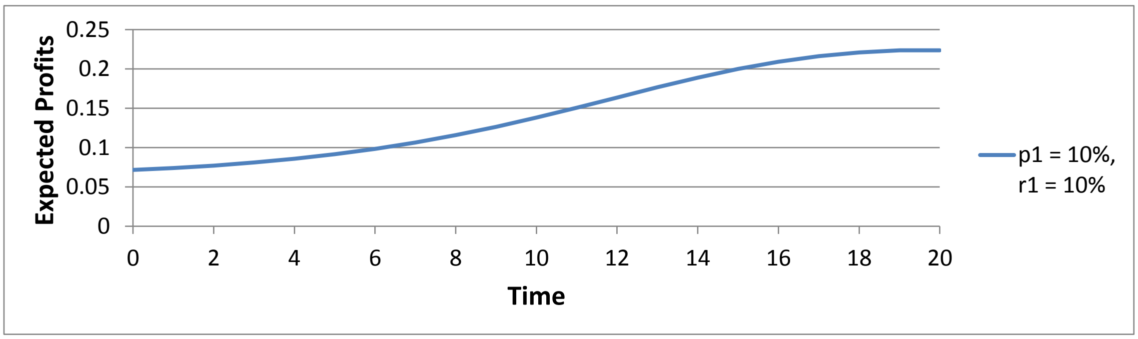

Figure 4.9: Expected profits of agent 1 as a function of time. The top chart plots the cases for $p_1 = 10\%, r_1 = 30\%$, as well as $p_1 = 30\%$ and $r_1 = 10\%, 30\%$. The case $p_1 = 10\%, r_1 = 10\%$ is shown on a separate chart since the expected profits in that case are significantly larger.

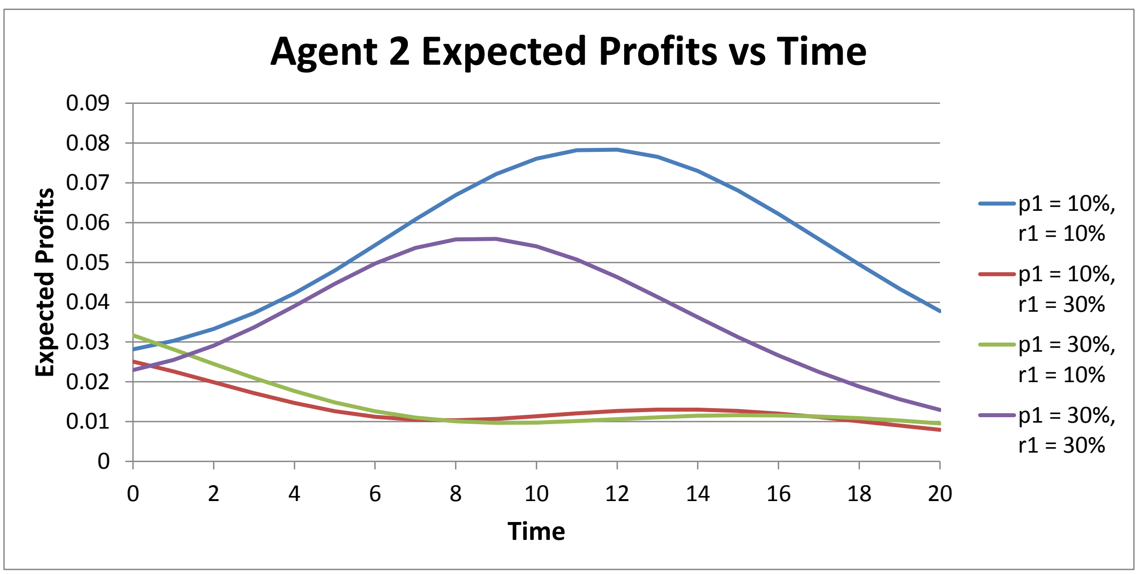

Figure 4.10: Expected profits of agent 2 as a function of time.

values. However, they are largely consistent with our results for the last period profits: as difference of opinions increases, agents expect higher profits; when an agent starts to become very convinced the true state is good, he expects lower profits. Over time, the difference in opinions may increase or decrease depending on the agent priors p_1, p_2 and beliefs about signal distributions r_1, r_2. Thus it is possible that expected profits initially rise (as difference of opinions increases) and then fall (as an agent becomes very bullish or the difference decreases). On the other hand, it is possible that the difference of opinions decreases from the initial date, and/or an agent becomes very bullish early in the time horizon, so expected profits decrease (and sometimes marginally increase) throughout the whole period.

While I already get some interesting results with the current set-up, the model is still quite limited, because agents can only become more bullish with time (unless a low signal arrives, so that price drops to zero). I next extend the model to allow for situations when agents may both increase and decrease their perceived probabilities of the good state depending on the realization of the signal.

5 Discrete Distribution of the Signal

I extend the simple model to incorporate a distribution of the signal that is still discrete, but now includes more than two values. The signal can take the values $z_1, z_2, \ldots, z_n$, with different probabilities depending on the state. The two agents disagree on the distribution of the signal in each state. This is the piece that will drive the model. Agent i believes that signal s_t takes on the value j with probability $f_{i,k}(z_j)$ if state $\mathcal{S}$ is equal to k. The beliefs $f_{i,k}$ about the signal are not updated throughout the investing horizon; later I allow the agents to update the beliefs about the signal distribution if they realize they are wrong.

The rest of the model is the same as before.

Agents still update their beliefs $p_{1,t}, p_{2,t}$ about the probability of the good state according to the Bayes Rule formula (1). With the current set-up, the formula for updating probability $p_{i,t}$ for agent i after observing signal s_{t+1}, is:

$$p_{i,t+1} = \frac{p_{i,t} f_{i,1}(s_{t+1})}{p_{i,t} f_{i,1}(s_{t+1}) + (1 - p_{i,t}) f_{i,2}(s_{t+1})} \tag{13}$$

The update depends on the current belief $p_{i,t}$ and the relative likelihood ratio:

$$\frac{f_{i,1}(s_{t+1})}{f_{i,2}(s_{t+1})} \tag{14}$$

We can model an agent who is reluctant to update his beliefs as follows. Suppose agent 1 is quite convinced the true state is the good one. Then for all the possible values of the signal s_{t+1}, the likelihood ratio is greater than 1, or smaller than 1, but still quite close to 1. This way, upon observing a new signal, the agent either becomes even more convinced that the true state is good, or he only marginally decreases his belief $p_{i,t}$ for the probability of the good state. As time passes, it is possible that the agent's perceived probability of the good state becomes low enough for the agent to realize he is wrong.

I solve for equilibrium using the same approach as in the simple model. The structure of the final payoff is the same, therefore Proposition (4.1) still holds. However, the recursive formulas are more complicated, because there are more cases for the signal value next period, and there is no "degenerate" case when after a certain realization of the signal there is no more uncertainty.

Consider date $t \leq T-1$. Both agents currently know each other's beliefs $p_{1,t}, p_{2,t}$. They also both know how these beliefs will change next period when signal s_{t+1} arrives. Let $p_{1,t+1,j}, p_{2,t+1,j}$ be these updated beliefs if the value of this signal is z_j. Define the following, for agent $i = 1, 2$ and signal s_{t+1} value z_j, $j = 1, 2, \ldots, n$:

$$a_{i,j} = p_{i,t} f_{i,1}(z_j) + (1 - p_{i,t}) f_{i,2}(z_j)$$

$$R_j = P_{t+1}(p_{1,t+1,j}, p_{2,t+1,j})$$

$$D_{i,j} = \mathbb{E}_{i,t+1}(\Pi_{i,t+1}(p_{1,t+1,j}, p_{2,t+1,j}))$$

These parameters are enough to pin down the relation for the price and expected profits between the current period and the next period.

Proposition 5.1. *Let $t \leq T-1$. Then the equilibrium price is:*

$$P = \frac{w_{2,t}\mu_{1,t} + w_{1,t}\mu_{2,t}}{w_{1,t} + w_{2,t}} \tag{15}$$

where:

$$\mu_{i,t} = \sum_{j=1}^{n} a_{i,j} R_j - \lambda \sum_{j=1}^{n} a_{i,j} R_j D_{i,j} + \lambda (\sum_{j=1}^{n} a_{i,j} D_{i,j})(\sum_{j=1}^{n} a_{i,j} R_j) \tag{16}$$

$$w_{i,t} = \sum_{j=1}^{n} a_{i,j} R_j^2 - (\sum_{j=1}^{n} a_{i,j} R_j)^2 \tag{17}$$

The expected PnL of agent i is:

$$\mathbb{E}_{i,t}(\Pi_{i,t}) = \frac{\mu_{i,t} - P}{\lambda w_{i,t}} \Big((\sum_{j=1}^{n} a_{i,j} R_j) - P \Big) + \sum_{j=1}^{n} a_{i,j} D_{i,j} \tag{18}$$

Proof. See Appendix. □

6 Realizing You Are Wrong

I now introduce the realization of being wrong into the model. Agents are usually quite convinced of their own beliefs and marginally update their opinions in response to the signals. However, if over time their opinions change significantly enough, then they realize their beliefs are incorrect and drastically change them. The agents don't just update their perceived probability of the good state, but also update their perceived distribution of the signal.

I assume that at any point in time each agent belongs to a particular type $m = 1, 2, \ldots, M$. This type uniquely determines the beliefs of the agent about the probability distribution of the signal. For most types, the agents are quite stubborn, so that they don't significantly update their perceived probability of the state being a true one. However, if this perceived probability changes drastically enough over time, the agent changes his type as well as his perceived distribution of the signal.

The evolution of the agent types is modeled as follows. Let $m_{i,t}$ be the type of agent i at date t. I assume that there is a one-to-one relation between the agent type $m_{i,t}$ and his perceived probability of the good state $p_{i,t}$.

Each agent i has probability thresholds $\gamma_{i,0} = 0, \gamma_{i,1}, \ldots, \gamma_{i,M} = 1$ which do not change throughout the time horizon. The agent type is determined as follows:

$$m_{i,t} = m \text{ iff } \gamma_{i,l-1} \leq p_{i,t} < \gamma_{i,l} \text{ for } m = 1, 2, \ldots, M \tag{19}$$

Finally, I describe how the perceived distribution of the signal is updated over time. Each agent of type m has a *baseline* distribution $g_{i,k}^{m}(z)$ of the signal in state k At date $t+1$, the agent uses his perceived signal distribution $g_{i,k}^{m_{i,t}}(z)$ and probability $p_{i,t}$ of good state from time t to determine the perceived probability $p_{i,t+1}$ of the good state at time $t+1$. Using (19) the agent updates his type $m_{i,t+1}$; this type could be the same as in the previous period, but could also change if the perceived probability changes drastically enough from time t to $t+1$. The perceived signal distribution at time $t+1$ is then $g_{i,k}^{m_{i,t+1}}(z)$.

7 Numerical Results for Discrete Model

The discrete model offers a clear framework that generates many interesting patterns for price and agent holdings as they change over time. I present two sets of results obtained using my framework. The first are exact values for the variables in the model for a particular market setting. Agents initially believe the true state is good, but observe a low signal in each period; thus over time they realize they are wrong. I am able to obtain clear intuition on how agents' changing beliefs influence their holdings and equilibrium price. The second set of results are generated by simulating the evolution of signal values through time. I look at the behavior of prices most prevalent across the simulated price paths. I find that the model can replicate a lot of the important phenomena in real-world markets, including short run momentum and reversals over longer horizons.

7.1 Persistent Bad News

I want to analyze how the agents in the model react to persistent bad news. Recall that the agents don't like to admit they are wrong. Therefore, if they are initially optimistic about the true state, it will take them some time to realize their beliefs are incorrect. Once they do realize they may be wrong, they will adjust their opinions more rapidly as well as change

Probability of signal taking High/Low value conditional on state			State	
			Good	Bad
Type = 1	Signal Value	High	80%	75%
		Low	20%	25%
Type = 2		High	60%	40%
		Low	40%	60%
Type = 3		High	25%	20%
		Low	75%	80%

Table 7.1: Perceived probability distributions for the signal by each type of agent. Each cell lists the probability of the signal taking on the high or low value, conditional on the state being high or low. The signal value is specified across the rows, while the state is specified across the columns.

their beliefs about the distribution of the signal. I find that price experiences different rates of decrease over time depending on the changing investor opinions. Agents tend to exhibit "support and resistance" behavior when they get close to their belief thresholds whereby they sometimes hold a positive (negative) position in the risky asset even though they are more bearish (more bullish) on the fundamentals. In extreme situations this can lead to an increase in price in spite of arrival of bad news.

To keep the model more tractable and the results easier to interpret, there will be only two possible values the signal can take: 1 (high value) and -1 (low value).

I will consider a time horizon with $T = 15$ periods, followed by another period after which the liquidating dividend is paid out. At each date the low signal is released. so that $s_t = -1$ for $t = 1, 2, \ldots, T$.

Each agent can belong to one of three different types. The first type is *bullish*. These agents are convinced the good state is the true one and think the signal value will most likely be a high one. Upon seeing a low signal they revise their perceived probability of the good state downwards, but quite marginally. The second type of the agent is uncertain. An agent of this type drastically changes the perceived probability upwards if he sees a high signal, and downwards if he sees a low signal. The third, and last, type of agent is bearish. These guys are convinced the true state is bad, and being stubborn, do not significantly increase their perceived probability upon observing a high signal.

Table 7.1 lists the probability distributions of the signal as perceived by each type of agent.

I fix the final payoff $\mu = 1$ and risk aversion coefficient $\lambda = 3$. The initial perceived probabilities are $p_{1,0} = 90\%$ and $p_{2,0} = 70\%$. These probabilities, combined with the thresholds for the agent types, pin down the initial types of the agents.

The only thing I vary in the model are the probability thresholds γ. I assume the thresholds are the same for each agent so that $\gamma_{i,1} = \gamma_1$, which varies between 80% and 60%, and $\gamma_{i,2} = \gamma_2$ varying between 40% and 20%. I will refer to γ_1 as the *high threshold* and to γ_2 as the *low threshold*.

Figure 7.1 plots price as a function of time for the four cases, while figure 7.2 plots the agent types. We see that in each case price initially declines slowly, then drops more rapidly, and finally decreases slowly again towards the end of the horizon. This is caused by the fact that initially both agents are of type 1, or agent 1 is of type 1 and agent 2 is of type 2. The agents of type 1 marginally decrease their perceived probabilities upon observing a low signal. After a few periods the perceived probabilities drop enough to cross the high threshold, so agents of type 1 become type 2, and drastically decrease their perceived probabilities in subsequent periods. A little later the probabilities drop further and cross the low threshold. At that point the agents become type 3, and again don't significantly revise their perceived probabilities.

The patterns for the rate of price decrease depend on the agent types and their perceived probabilities. If the agent just switched from type 2 to type 1, but is very bullish, then he believes there is a large probability a high signal will occur in the next period. If that happens, the agent will significantly revise his perceived probability upwards and have an even higher probability than in the previous period, and become type 1 again. Thus observing a low signal will not result in a large drop in price (see for example periods 0 and 1 for the blue and green curves, where agent 2 has type 2 but is bullish). On the other hand, if the agent is bearish, he anticipates a high likelihood of a low signal in the next period, and further evidence of a bad state, resulting in a lower price. Observing a low signal produces a substantial price fall in this case, as can be observed from periods 3 and 4 in the blue and green curves.

The beliefs of both agents matter here. Consider the red curve in periods 7 and 8, where we see a substantial price drop while agent 1 has type 1, and agent 2 has type 3. Agent 2 is

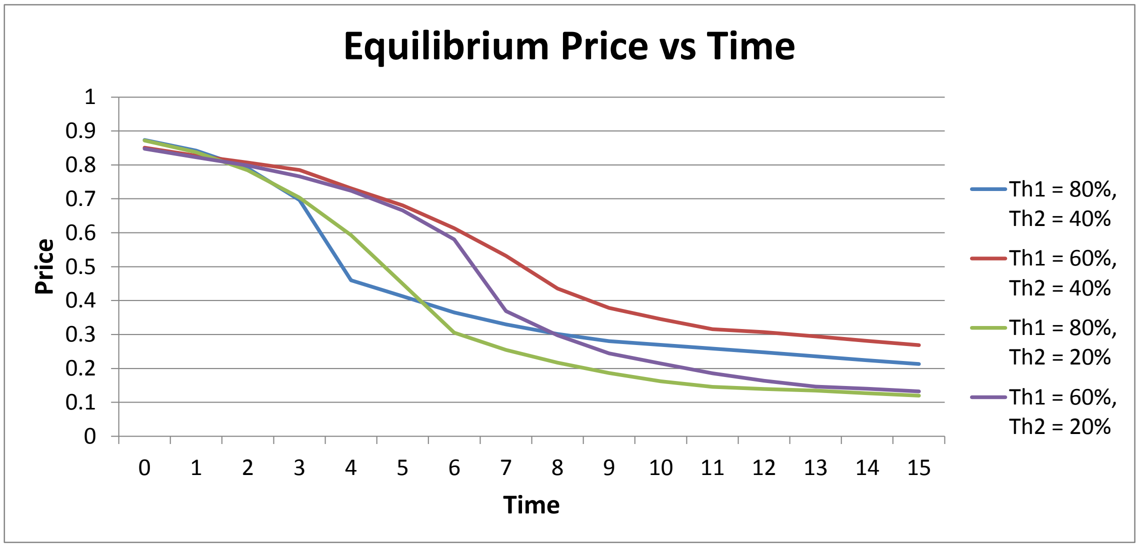

Figure 7.1: Equilibrium price as a function of time. The initial perceived probabilities are $p_{1,0} = 90\%$ for agent 1 and $p_{2,0} = 70\%$ for agent 2. The signal value is low in each period. The thresholds are the same for the agents and vary, with $\gamma_1 = 80\%, 60\%$ and $\gamma_2 = 40\%, 20\%$.

	Agent 1 Type vs Time																
Agent 1 Type	1	1	1	1	2	2	2	2	2	3	3	3	3	3	3	3	**Th1 = 80%, Th2 = 40%**
	1	1	1	1	1	1	1	1	1	2	2	3	3	3	3	3	**Th1 = 60%, Th2 = 40%**
	1	1	1	1	2	2	2	2	2	2	2	3	3	3	3	3	**Th1 = 80%, Th2 = 20%**
	1	1	1	1	1	1	1	1	1	2	2	2	2	3	3	3	**Th1 = 60%, Th2 = 20%**
Time	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	

	Agent 2 Type vs Time																
Agent 2 Type	2	2	2	2	3	3	3	3	3	3	3	3	3	3	3	3	**Th1 = 80%, Th2 = 40%**
	1	1	2	2	3	3	3	3	3	3	3	3	3	3	3	3	**Th1 = 60%, Th2 = 40%**
	2	2	2	2	2	2	3	3	3	3	3	3	3	3	3	3	**Th1 = 80%, Th2 = 20%**
	1	1	2	2	2	2	2	3	3	3	3	3	3	3	3	3	**Th1 = 60%, Th2 = 20%**
Time	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**	**11**	**12**	**13**	**14**	**15**	

Figure 7.2: Agent types as they change over time. Each row corresponds to a different case for the threshold values, listed on the right. Green color represents type 1, yellow represents type 2, and red represents type 3.

very convinced the true state is bad, and is anticipating a further decrease in price. Agent 1 is not yet of type 2, but is close to the high threshold. Because his perceived probability is now quite low (around 60%), then a subsequent low signal would turn him into a type 2 agent who is bearish – and expecting a further price decrease. Thus, the price experiences a large drop, before agent 1 even realizes he is wrong in his bullish beliefs.

I observe a surprising phenomenon whereby price is initially *higher* for higher threshold values. Even though with a larger high threshold an agent is more likely to switch from type 1 to type 2 and thus become more bearish than before, there is still a large probability a high signal will follow, pushing the agent to become type 1 again (and actually more bullish than before). Thus observing a low signal still maintains the bullishness of the agents – and in periods 0 and 1 the blue and green curves with larger values of the high threshold (80%) are above the red and purple curves with a smaller high threshold (60%).

At the same time, a larger low threshold means the agent realizes he is wrong faster, and stops revising his perceived probability downwards so drastically with each subsequent low signal realization. At that point price does not drop as substantially, and towards the end of the horizon the blue and red curves with a large low threshold (40%) are above the green and purple curves with a small low threshold (20%).

I next analyze how holdings evolve over time. Figure 7.3 plots the equilibrium holdings for agent 1; the negative of these holdings gives the equilibrium holdings of agent 2. The patterns for the four cases are quite different, and I discuss them one at a time.

For the case $\gamma_1 = 80\%, \gamma_2 = 40\%$ (blue curve) holdings start out quite large and positive (around 2 shares of the risky asset), and then drop close to zero in date 9 and stay there until the end of the horizon. Intuitively this makes sense: agent 1 is more bullish than agent 2 throughout the whole horizon and thus wants to go long the asset. Furthermore, in dates 0 to 3 agent 1 has type 1, while agent 2 has type 2, while in dates 4 to 8 agent 1 has type 2, while agent 2 has type 3. Thus agent 1 has a different type from agent 2 and one that anticipates a higher probability of a high signal. In date 9 and later, both agents are of type 3 and thus have the same perceived distribution of the signal. Even though they still disagree on the perceived probability of the good state, this difference is not too large. As a result, holdings are close to zero in those periods.

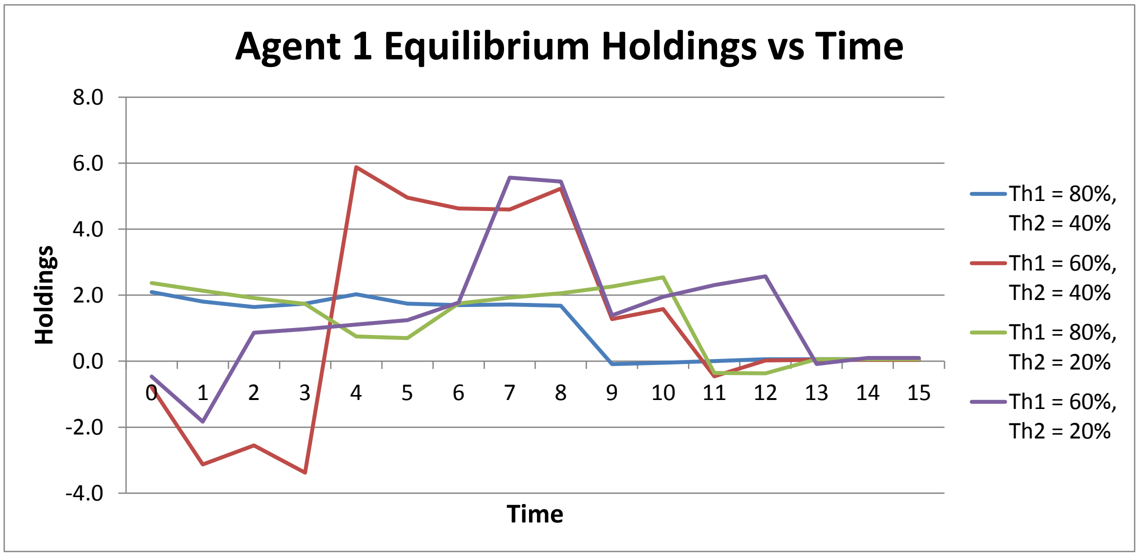

Figure 7.3: Equilibrium holding of agent 1 in the last period. Note that the holding of agent 2 is just the negative of the holding of agent 1.

Note that in date 9 the position is slightly negative even though agent 1 is more bullish. This is because, while the two agents have similar expectations for next period's price, the first agent is expecting a larger differential in profits between if the signal high and if this signal is low in the next period. The agents have similar price expectations since by this time their perceived probabilities are quite close, *and* since they are of the same type. In terms of expected profits, both agents, being type 3, recognize that a high signal would make them revise their perceived probability upwards and lead to a greater difference of opinions, thus providing greater potential profits. Since agent 1 is more bullish than agent 2, a high signal would make his perceived probability further away from the "extreme" opinion of 0% and hence provide higher PnL relative to the case of a low signal, than for agent 2.

In economic terms, the above phenomenon makes sense. Agent 1 has just crossed his psychological "breaking point" and realized he is wrong. He is thus very bearish just like agent 2, but furthermore expects high profits in case of a high signal, with more bullish expectations than agent 2; thus profits covary more with price, so utility has a lower coefficient for holdings in terms of the expected profits component, as well as a lower coefficient overall since price expectations are similar. Agent 1 will therefore hold a negative position in equilibrium.

The green curve plots agent 1's holdings for the case $\gamma_1 = 80\%, \gamma_2 = 20\%$. It looks similar to the blue curve – agent 1 holds a relatively large positive position, until he switches to type 3 at date 11. At that point he holds a negative position for a few periods (for the same reasons as for the blue curve at date 9), and then again a small positive one for the remaining horizon.

There are a few minor differences between the blue and the green curve. In dates 4 and 5 agent 1's holding drops quite significantly. This is because in those dates both agents have the "uncertain" type and expect large price volatility the next period; being risk-averse, they choose to hold smaller magnitudes of the risky asset.

The other difference we observe is in date 10, when agent 1 increases his holdings right before he realizes he is wrong in date 11. This is a nice desirable phenomenon that the model generates – people increasing their stakes when they are close to the limit when they realize they are wrong. Here, this happens because both agents forecast a significantly lower variance in the next period, and hence take on larger positions.

The two cases for $\gamma_1 = 60\%$ are also interesting. Even though agent 1 is more bullish than agent 2, he holds a large negative position in the risky asset in the early periods. Just like for date 9 for the blue curve, here both agents have similar perceived probabilities but disagree on the expected profits they can potentially earn. Agent 1 realizes that a low signal would lead to a greater difference of opinions and his opinion being not as extreme, and thus expects to make greater profits with a low signal next period in comparison to a high signal. Since agent 1 has perceived probability that is closer to 100% than for agent 2, he therefore expects a smaller profit differential between how much he makes if a low signal occurs in comparison to if a high signal occurs in the next period. Thus profits covary more with price, so agent 1 holds a negative position in equilibrium.

In later periods for $\gamma_1 = 60\%$ the pattern for holdings is similar to the cases for $\gamma_1 = 80\%$ we already discussed. Namely, agent 1 holds a positive position as he still has type 1, while agent 2 has type 2 or 3. Later he still holds a positive (but smaller) position as he has type 2, while agent 2 has type 3. In the last few periods, both agents become type 3 and holdings are close to zero since there is little difference in opinions at that point.

I thus make the following conclusions about agent holdings:

- If the agents have a large deviation in their perceived probabilities, the more bullish agent holds a positive position.
- If both agents are of type 1 or 3 and not close to their thresholds, the more bullish agent holds a small positive position.
- If both agents are of type 2 and not close to their thresholds, the more bullish agent holds a large positive position.
- If the agents have similar perceived probabilities, and the more bullish agent is close to the low threshold but is quite bearish, he holds a negative position[6].
- If the agents have similar perceived probabilities, and the more bearish agent is close to the high threshold but is quite bullish, then he holds a positive position[7].

In the Appendix I also consider the expected profits of the agents. Figure 9.1 shows the expected PnL for agent 1, while figure 9.2 plots the PnL for agent 2. The general pattern for each agent and for each of the four cases is the same: initial low expected profits, then higher profits sometime in the middle of the horizon, and then again low profits (in fact, close to zero) close to the liquidation date. This is due to the fact that towards the middle of the investment horizon the agents have a large difference of opinions (and usually different beliefs about the signal distribution), whereas in the early and late periods this difference is much smaller.

I now consider a situation where agents have "extreme overconfidence" in their beliefs about the distribution of the signal when they belong to the bullish and the bearish types. This way they very marginally update their perceived probabilities upon observing new information. The perceived distributions of the signal for each type are listed in table 7.2. The rest of the framework and model parameters are the same as before. Agents still initially believe the true state is good, and a low signal arrives in each period.

Figure 7.4 shows how price moves over time, and figure 7.5 shows the types of the agents. The general pattern for the price is the same as with the earlier situation I considered.

[6]In particular this happens when both agents are of type 3, and the more bullish agent is close to the low threshold; see the red and green curves at date 11.

[7]This occurs, for example, when both agents are of type 1, and the more bearish agent is close to the high threshold; see the red and purple curves at dates 0 and 1.

Probability of signal taking High/Low value conditional on state			State	
			Good	Bad
Type = 1	Signal Value	High	80%	78%
		Low	20%	22%
Type = 2		High	60%	40%
		Low	40%	60%
Type = 3		High	22%	20%
		Low	78%	80%

Table 7.2: Perceived probability distributions for the signal by each type of agent, when agents of type 1 and 3 have extreme overconfidence in their beliefs. The highlighted cells are the only cases when the perceived distribution is different from the distributions considered earlier in Table 7.1.

In particular, price starts to drop slowly, then this drop accelerates as agents realize they are wrong and become type 2, and then price declines slowly again towards the end of the horizon. The rate of the decline during the early and late periods is smaller than for the earlier set of parameters because now if agents are of type 1 or 3, they revise their perceived probabilities at a slower rate upon observing a low signal.

The intriguing phenomenon I want to draw attention to is the fact that sometimes price *increases* with the arrival of a low signal. If we look at the blue curve ($\gamma_1 = 80\%, \gamma_2 = 40\%$), price rises in date 3 and date 10. This also happens for the green curve ($\gamma_1 = 80\%, \gamma_2 = 20\%$) in dates 5 and 10. I discuss why these price increases occur.

For the blue curve in date 3, agent 2 is of type 2 and is still quite bullish on the state. He is close to the low threshold and will become type 3 if another low signal occurs. At the same time, agent 1 will remain type 1 regardless of the signal value in the next period. Therefore, a low signal will lead to a greater difference in opinions and in beliefs about the signal distribution for the agents, and result in much higher expected profits, than if a high signal occurs. Price will still be higher if a high signal occurs next period, therefore the next period price and expected profits have a large negative covariance. This leads to a high equilibrium price in the current period, which is in fact higher than the price in the period before. The same situation occurs for the green curve in date 5. We thus see that as agent 2 is about to realize he is wrong (by switching from type 2 to type 3), then he exhibits very strong resistance and causes the price to rise.

Price increases in dates 10 for both the blue and the green curve. I argue that this happens because of an "extra" drop in price in date 9. For both curves agent 2 has already been of type 3 for some time, and so will not switch to type 2 upon arrival of either a high or a low signal. On the other hand, agent 1 has just switched to type 2 after being type 1. A high signal in the next period will make him type 1 again, and in fact having a higher perceived probability than in the period before. A high signal will then produce a larger difference in opinions as well as in the agent types (agent 1 will be type 1 and agent 2 will be type 3). Hence a high signal next period will lead to larger expected profits than a low signal. Price will still be higher in the case of a high signal; therefore the covariance between next period price and expected profits is positive and large. As a result, price drops significantly in date 9. In date 10, after there is no longer the possibility of agent 1 switching to type 2 with a high signal, the covariance is much smaller in magnitude, while price expectations are similar. This causes the date 10 price to be higher than the price in date 9. I conclude that in this situation agent 1, after realizing he is wrong, depresses the price substantially.

7.2 Price Behavior

I next analyze price behavior when both good and bad news can arrive. I simulate the realizations of the signal in each period and look at predictability in price movement. Different price patterns are possible depending on the initial beliefs of the agents, as well as the thresholds for switching these beliefs. When one agent is very bearish on the state, while the other is very bullish, prices tend to experience reversals, which are particularly strong at longer horizons. On the other hand, when agents have similar beliefs and have high conviction about them, significant trending is observed. Finally, if one agent is convinced in his opinions and the other is very uncertain, then price movement is highly dependent on the agent belief thresholds.

I use the same parameter values as when looking at the persistent bad news in the previous section. The terminal payoff is $\mu = 1$, risk aversion coefficient $\lambda = 3$, and the agents' beliefs about the distribution of the signal are listed in Table 7.1. As before, I assume both agents have the same thresholds with the high threshold $\gamma_{i,1} = \gamma_1$ varying from 60% to 80% and the low threshold $\gamma_{i,2} = \gamma_2$ varying from 20% and 40%.

There are three cases for the initial beliefs of the agents. The first is when agents have very

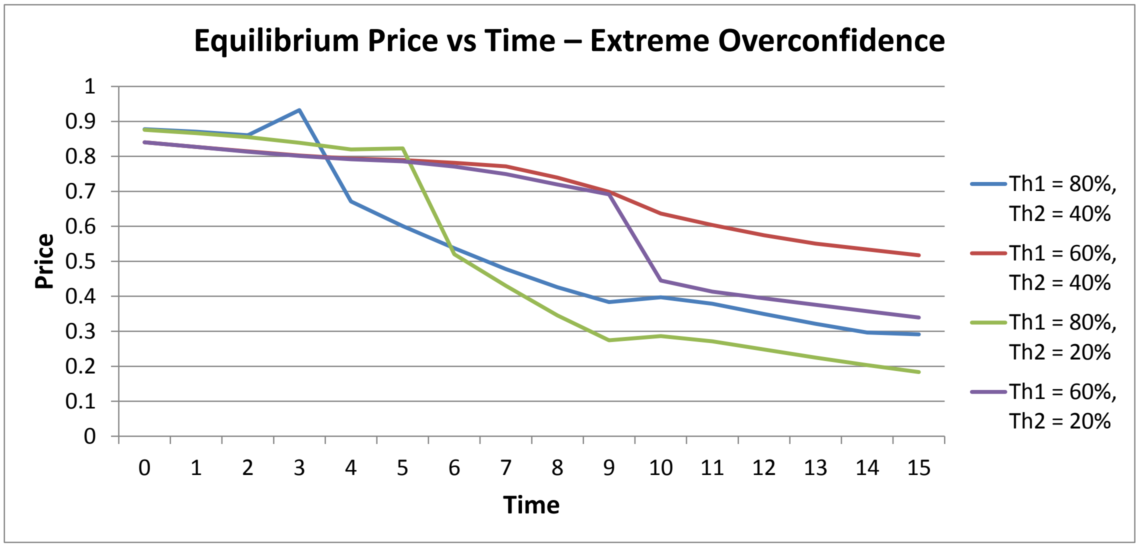

Figure 7.4: Equilibrium price as a function of time for the case of extreme overconfidence among the agents. As before, the initial perceived probabilities are $p_{1,0} = 90\%$ for agent 1 and $p_{2,0} = 70\%$ for agent 2, and in each period a low signal is released.

	Agent 1 Type vs Time – Extreme Overconfidence																
Agent 1 Type	1	1	1	1	1	1	1	1	1	2	2	2	2	2	3	3	Th1 = 80%, Th2 = 40%
	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	Th1 = 60%, Th2 = 40%
	1	1	1	1	1	1	1	1	1	2	2	2	2	2	2	2	Th1 = 80%, Th2 = 20%
	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	1	Th1 = 60%, Th2 = 20%
Time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	

	Agent 2 Type vs Time – Extreme Overconfidence																
Agent 2 Type	2	2	2	2	3	3	3	3	3	3	3	3	3	3	3	3	Th1 = 80%, Th2 = 40%
	1	1	1	1	1	2	2	3	3	3	3	3	3	3	3	3	Th1 = 60%, Th2 = 40%
	2	2	2	2	2	2	3	3	3	3	3	3	3	3	3	3	Th1 = 80%, Th2 = 20%
	1	1	1	1	1	2	2	2	2	2	3	3	3	3	3	3	Th1 = 60%, Th2 = 20%
Time	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	

Figure 7.5: Agent types at each date when agents are subject to extreme overconfidence. Green color represents type 1, yellow represents type 2, and red represents type 3.

different opinions: agent 1 is bullish with $p_1 = 90\%$ and agent 2 is bearish with $p_2 = 10\%$.[8] The second case is when agent 1 has high conviction in his beliefs ($p_1 = 90\%$) while agent 2 is very uncertain about the true state ($p_2 = 50\%$). The final case has the two agents with very similar beliefs and quite convinced they are right, so that $p_1 = 90\%, p_2 = 85\%$.

The evolution of the signal is now stochastic. I consider the most "basic" distribution whereby the signal in each period is high with probability 50% and low with probability 50%. This way, if the agents knew the true distribution of the signal, they would realize that the signal contains no information about the true state, and hence not change their beliefs throughout the horizon. Therefore equilibrium price would stay constant. However, since agents have beliefs about the signal distribution which are very different from the true distribution, they revise their opinions over time, and the price moves. I want to understand this price movement for different cases of agent beliefs and belief thresholds.

The particular statistics of interest are the serial correlation in price changes from the start of the investment horizon. For various values of the number of lags k, I calculate the following correlation:

$$\rho_k = \rho(P_k - P_0, P_{2k} - P_k) \tag{20}$$

Here, ρ is the correlation operator. In the above expression it measures the correlation between price move over the first k periods and the price move over the k periods after that. I look at two specific values for k: $k = 1$ representing short-run correlation, and $k = 5$ corresponding to long-run correlation.

The correlation values are listed in table 7.3. We see that correlation exhibits different patterns depending on the number of lags used and depending on the case for the initial beliefs of the agents.

When agents have very different beliefs (top two sub-tables), we observe price reversals at both the short and long horizons. The short-run correlation in price moves is negative and small. The main explanation for this is as follows. The agent who is further away from the threshold closest to him in terms of perceived probabilities is the one pushing the price[9].

[8] For notational convenience I denote by $p_1 = p_{1,0}$ and $p_2 = p_{2,0}$ for the initial perceived probabilities of the agents.

[9] It is possible that the agents are equally close to their closest thresholds, e.g. if $p_1 = 90\%, \gamma_1 = 80\%$ and $p_2 = 10\%, \gamma_2 = 10\%$.

p1 = 90%, p2 = 10%; Lag 1 Serial Correlation							
		High Threshold					
		60%	64%	68%	72%	76%	80%
Low Threshold	20%	-2%	-3%	-5%	-7%	-9%	-10%
	24%	-5%	-6%	-7%	-9%	-9%	-9%
	28%	-6%	-7%	-8%	-9%	-9%	-7%
	32%	-6%	-6%	-7%	-8%	-7%	-5%
	36%	-5%	-6%	-6%	-7%	-6%	-3%
	40%	-4%	-5%	-5%	-6%	-5%	-3%

p1 = 90%, p2 = 10%; Lag 5 Serial Correlation							
		High Threshold					
		3%	64%	68%	72%	76%	80%
Low Threshold	20%	-42%	-45%	-47%	-54%	-60%	-63%
	24%	-41%	-44%	-46%	-51%	-57%	-60%
	28%	-35%	-38%	-39%	-45%	-51%	-53%
	32%	-27%	-32%	-33%	-39%	-46%	-47%
	36%	-25%	-30%	-32%	-38%	-44%	-45%
	40%	-21%	-25%	-27%	-35%	-41%	-42%

p1 = 90%, p2 = 50%; Lag 1 Serial Correlation							
		High Threshold					
		60%	64%	68%	72%	76%	80%
Low Threshold	20%	7%	7%	9%	12%	15%	15%
	24%	9%	6%	9%	15%	18%	18%
	28%	-7%	-13%	-7%	6%	17%	20%
	32%	33%	35%	39%	43%	43%	39%
	36%	56%	57%	58%	57%	55%	50%
	40%	-3%	-3%	0%	3%	3%	1%

p1 = 90%, p2 = 50%; Lag 5 Serial Correlation							
		High Threshold					
		60%	64%	68%	72%	76%	80%
Low Threshold	20%	23%	16%	10%	-2%	-17%	-21%
	24%	25%	18%	10%	-5%	-22%	-31%
	28%	12%	4%	-4%	-19%	-34%	-39%
	32%	17%	9%	1%	-13%	-27%	-35%
	36%	13%	4%	-4%	-18%	-31%	-36%
	40%	15%	7%	-1%	-14%	-30%	-39%

p1 = 90%, p2 = 85%; Lag 1 Serial Correlation							
		High Threshold					
		60%	64%	68%	72%	76%	80%
Low Threshold	20%	5%	5%	4%	4%	4%	6%
	24%	5%	5%	4%	4%	4%	6%
	28%	6%	5%	4%	4%	3%	6%
	32%	5%	5%	4%	4%	3%	6%
	36%	5%	5%	4%	4%	4%	5%
	40%	5%	5%	4%	4%	3%	6%

p1 = 90%, p2 = 85 %; Lag 5 Serial Correlation							
		High Threshold					
		60%	64%	68%	72%	76%	80%
Low Threshold	20%	16%	15%	14%	11%	2%	-5%
	24%	16%	15%	14%	11%	3%	-4%
	28%	16%	15%	14%	11%	3%	-4%
	32%	16%	15%	14%	11%	4%	-3%
	36%	16%	15%	14%	11%	4%	-4%
	40%	16%	15%	14%	12%	4%	-4%

Table 7.3: Correlations in price change at various lags and for various initial beliefs of the agents. I list the correlation between the price move in the k next periods (starting from date 0), and the price move in the k periods after that, where k is the specified number of lags. The values for the number of lags considered are 1 (short-run) and 5 (long-run). The initial perceived probabilities of the agents are listed at the top of each sub-table. The cells with the correlation numbers contain bars, that are green if the value is negative, and red if the value is positive; the size of the bar corresponds to its magnitude in the $[0\%, 65\%]$ range.
For each case, I perform 1,000 runs of 10,000 simulations of signal paths, measuring the correlation using each sample of the simulated paths. The resulting 1,000 values are then averaged out to get the estimate of correlation. For each case, the standard errors do not exceed 0.25%.

Without loss of generality assume the bullish agent is the one pushing the price. If a bad signal arrives in the first period, price decreases; however the potential moves in the price will not be as different depending on if a bad or a good signal arrives in the next period, because at that point the bullish agent is not pushing the price as much. On the other hand, if a good signal arrives in the first period, price increases and the bullish agent is pushing the price even more. Then, a bad signal in the next period will result in a greater change in perceived probability of the bullish agent, than if another good signal comes – and therefore price will move substantially more in that case, in the direction opposite to its move in the first period. This results in a negative correlation. The magnitude of this correlation is small because over a single period the relative differences between potential moves in the price are not large.

Over longer horizons the above effect amplifies significantly. After the first five periods, one of the agents is pushing the price; again, assume it is the bullish agent. Over the next five periods, his perceived probability would move a lot more on the downside than on the upside because of his perceived distributions in the signal *and* the potential change in beliefs because of crossing a belief threshold. Therefore, on average price will experience a much greater move in direction opposite to the one over the first five periods, which leads to a large negative correlation.

I next look at the case when agent 1 has high conviction in his beliefs ($p_1 = 90\%$) while agent 2 does not ($p_2 = 50\%$). The corresponding correlation values are shown in the middle two sub-tables of table 7.3. At short horizons price usually experiences momentum; it is particularly strong for low threshold values of $\gamma_2 = 32\%, 36\%$. The reason for this is the fact that price tends to experience a particularly large move in period 2 if a bad signal arrives in that period, following a bad signal in the first period. This way price drops for two subsequent periods, with these moves being larger than for the other cases of signal values. Price drops so much in period 2 because agent 2 switches from type 2 to type 3 (this happens for $\gamma_2 = 32\%, 36\%$), and/or because price just generally moves more on the downside due to agent 1 pushing the price.

At long horizons we observe price reversals for high values of the high threshold ($72\% \leq \gamma_1 \leq 80\%$). Here, agent 1 is pushing the price and will often cross the high threshold in terms of his beliefs over the first five periods, or at least get close to it, causing a large price move on the downside. This would usually be followed by a large move on the upside, since either agent 1 will become type 1 again (after being type 2), and/or agent 2 will become type 2 again (after being type 3). The combination of these large moves in different directions over five-period intervals leads to a negative correlation in price movement; sometimes this correlation is large in magnitude when the agents start out closer to the thresholds.

For lower values of the high threshold price tends to experience reversals. This is due to the fact that agent 1 is now further away from the high threshold. As a result, following a decrease in perceived probability in the first five periods, he would experience a larger further decrease in this probability over the next periods (in comparison to a potential increase). As for agent 2, this move in the next five periods is also usually negative, especially if the low threshold is small. This combination of subsequent negative moves gives a positive correlation value.

The final case for the initial agent beliefs is when they are both very bullish on the true state (bottom two sub-tables in table 7.3). Price experiences weak momentum over short horizons and strong momentum over long horizons (except for the case of a large high threshold). The momentum in price can again be explained by the fact, that a bullish agent, after revising his beliefs downwards, will revise them further downwards more upon arrival of a low signal, than upwards upon arrival of a high signal in the next period. This phenomenon is stronger over long horizons, where the bullish agent often ends up switching his type from 1 to 2 after crossing the high threshold. The exception is for the value of high threshold of $\gamma_1 = 80\%$, where agents often already switch from type 1 to 2 in the first five periods, and then can potentially switch back to type 1, resulting in price reversals.

I conclude that the model generates both momentum and reversals in price. Price experiences momentum when both agents have high conviction in their beliefs, and these beliefs are similar. On the other hand, when these beliefs are drastically different, reversals are observed. Finally, if only one agent has high conviction, then there is usually trending over short horizons as well as over longer horizons when the agent who is unsure about his beliefs is close to his belief threshold and is likely to switch his type. When instead the high conviction-agent is close to his belief threshold, reversals tend to occur.

8 Conclusion

I propose a new model for analyzing how changing investor opinions influence dynamic trading and prices. I focus on a particular behavioral aspect, whereby investors are reluctant to admit they are wrong about their perceptions of market fundamentals. The difference of opinions between agents in the model and the eventual realization of being wrong lead to interesting and realistic dynamics for prices, investor holdings, and expected profits.

I start with a simple framework for how information about fundamentals arrives and how agents revise their beliefs using this information. I derive closed-form expressions for variables of interest in the last period. These allow for some initial intuition with regards to how difference of opinions and agent conviction influence prices. The main result is that the agent who is more confident in his beliefs end up "pushing" the price in the direction of his beliefs, sometimes so much, that he actually expects to earn lower profits.

The model allows for an explicit formula of computing equilibrium prices over the full investment horizon. The risky asset price tends to increase over time as agents become more convinced that the terminal payoff on the asset is high. Equilibrium agent holdings increase over time whenever that agent revises upwards his perceived probability of the true state by a greater amount than the other agent upon observing a good signal.

I next consider an advanced model with a more complicated structure of agent opinions about the signal distribution. If an agent revises his perceived probability of the true state substantially enough, he realizes he is wrong and completely changes his perceived distribution of the signal. I look at price behavior when agents are initially bullish on the fundamentals, but persistent bad news arrive, causing them to gradually reduce their forecasts for the terminal payoff. Price initially declines slowly due to overconfidence in agent beliefs, and then more rapidly as agents admit they are wrong and revise their opinions faster. When agents are close to realizing they are wrong, they often exhibit support and resistance behavior; for extreme cases of their beliefs this generates a price increase even though bad news keep arriving.

Finally, the model can produce both momentum and reversals in prices. When agents have very different opinions about the true state, there is significant negative correlation in price movement, especially over longer horizons. When investors have similar beliefs and high conviction in them, price tends to experience trending. If only one agent has high conviction in his beliefs, price behavior is very sensitive to belief thresholds, with momentum being the more prevalent pattern.

It is common for a person to be convinced they are right and to be unwilling to accept they are wrong in their beliefs. This paper provides a tractable framework for understanding this phenomenon and demonstrates that it can have a substantial impact on investor trading and price behavior.

9 Appendix

9.1 Proof of Proposition 4.1

Consider the optimization problem for agent 1. So far, the value of the signal has been high in each period, hence there is still uncertainty about what the final payoff will be. Agent 1 believes that the distribution of this payoff is:

$$F = \begin{cases} \mu, \text{ wp } p_{1,T} \\ 0, \text{ wp } 1 - p_{1,T} \end{cases}$$

Suppose the current price (at date T) is P. If the agent considers holding x units of the risky asset, his expected value and variance of profits are:

$$\mathbb{E}_{1,T}(\Pi_{1,T}) = \mathbb{E}_{1,T}(x(F - P)) = p_{1,T}\mu x - xP$$

and:

$$var_{1,T}(\Pi_{1,T}) = var_{i,t}(x(F - P)) = var_{1,T}(xF)$$
$$= (p_{1,T}\mu^2 - (p_{1,T}\mu)^2)x^2 = p_{1,T}(1 - p_{1,T})\mu^2 x^2$$

The optimization problem in (4) becomes:

$$\max_x \; p_{1,T}\mu x - xP - \frac{\lambda}{2}(p_{1,T}(1 - p_{1,T})\mu^2 x^2)$$

The optimal demand of agent 1 is:

$$x_{1,T} = \frac{p_{1,T}\mu - P}{\lambda p_{1,T}(1 - p_{1,T})\mu^2}$$

Similarly, the optimal demand of agent 2 is:

$$x_{2,T} = \frac{p_{2,T}\mu - P}{\lambda p_{2,T}(1 - p_{2,T})\mu^2}$$

Then, the equilibrium price at time T is:

$$P_T(p_{1,T}, p_{2,T}) = \frac{1}{\frac{1}{\lambda p_{1,T}(1-p_{1,T})\mu^2} + \frac{1}{\lambda p_{2,T}(1-p_{2,T})\mu^2}} \Big(\frac{p_{1,T}\mu}{\lambda p_{1,T}(1-p_{1,T})\mu^2} + \frac{p_{2,T}\mu}{\lambda p_{2,T}(1-p_{2,T})\mu^2} \Big)$$

This simplifies to:

$$P_T(p_{1,T}, p_{2,T}) = \frac{p_{1,T}p_{2,T}(2-p_{1,T}-p_{2,T})}{p_{1,T}(1-p_{1,T}) + p_{2,T}(1-p_{2,T})}\mu \tag{21}$$

The expected value of profits for agent 1 is:

$$\mathbb{E}_{1,T}(\Pi_{1,T}(p_T, q_T)) = \frac{(p_{1,T}\mu - P)^2}{\lambda p_{1,T}(1-p_{1,T})\mu^2} \tag{22}$$

We have:

$$p_{1,T}\mu - P = p_{1,T}\mu - \frac{p_{1,T}p_{2,T}(2-p_{1,T}-p_{2,T})}{p_{1,T}(1-p_{1,T}) + p_{2,T}(1-p_{2,T})}\mu = \frac{(p_{1,T}-p_{2,T})p_{1,T}(1-p_{1,T})}{p_{1,T}(1-p_{1,T}) + p_{2,T}(1-p_{2,T})}\mu$$

Substituting into (22), the expected value of profits for agent i, $i = 1, 2$, is:

$$\mathbb{E}_{i,T}(\Pi_{i,T}(p_{1,T}, p_{2,T}))) = \frac{(p_{1,T}-p_{2,T})^2 p_{i,T}(1-p_{i,T})}{\lambda(p_{1,T}(1-p_{1,T}) + p_{2,T}(1-p_{2,T}))^2} \tag{23}$$

9.2 Limits for Variables in the Last Period

When one of the perceived probabilities $p_{1,T}, p_{2,T}$ in the last period is equal to 0 or 1, the values for some of the variabels in the model become degenerate. In order to incorporate such situations in the model, I replace these values with their limits, which are calculated below. Recall that the variables of interest are the price, the holdings, and the expected profits.

I consider the cases when $p_{1,T} \in \{0, 1\}$; the cases when $p_{2,T} \in \{0, 1\}$ give analogous results due to symmetry in the model.

<u>Case 1.1:</u> $p_{1,T} = 0, p_{2,T} \in (0, 1)$.

Equation (21) gives:

$$P = \lim_{p_{1,T}\to 0} \frac{p_{1,T}p_{2,T}(2 - p_{1,T} - p_{2,T})}{p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T})}\mu = 0 \tag{24}$$

Equilibrium holdings of the agents are:

$$x_{1,T} = \lim_{p_{1,T}\to 0} \frac{(p_{1,T} - p_{2,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))\mu} = -\frac{1}{\lambda(1 - p_{2,T})\mu} \tag{25}$$

$$x_{2,T} = -x_{1,T} \tag{26}$$

Expected profits are:

$$\mathbb{E}_{1,T}(\Pi_{1,T}) = \lim_{p_{1,T}\to 0} \frac{(p_{1,T} - p_{2,T})^2 p_{1,T}(1 - p_{1,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))^2} = 0 \tag{27}$$

$$\mathbb{E}_{2,T}(\Pi_{2,T}) = \lim_{p_{2,T}\to 0} \frac{(p_{1,T} - p_{2,T})^2 p_{2,T}(1 - p_{2,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))^2} = \frac{p_{2,T}}{\lambda(1 - p_{2,T})} \tag{28}$$

Case 1.2: $p_{1,T} = 0, p_{2,T} = 0$.

I assume $p_{1,T} = p_{2,T} = q$ with q approaching 0. Equation (21) gives:

$$P = \lim_{q\to 0} \frac{2q^2(1 - q)}{2q(1 - q)}\mu = 0 \tag{29}$$

Equilibrium holdings of the agents are:

$$x_{1,T} = \lim_{q\to 0} \frac{(q - q)}{\lambda(2q(1 - q))\mu} = 0 \tag{30}$$

$$x_{2,T} = -x_{1,T} = 0 \tag{31}$$

Expected profits are:

$$\mathbb{E}_{1,T}(\Pi_{1,T}) = \lim_{q\to 0} \frac{(q - q)^2 q(1 - q)}{\lambda(2q(1 - q))^2} = 0 \tag{32}$$

$$\mathbb{E}_{2,T}(\Pi_{2,T}) = 0 \tag{33}$$

Case 2.1: $p_{1,T} = 1, p_{2,T} \in (0,1)$.

Equation (21) gives:

$$P = \lim_{p_{1,T}\to 1} \frac{p_{1,T}p_{2,T}(2 - p_{1,T} - p_{2,T})}{p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T})}\mu = \mu \tag{34}$$

Equilibrium holdings of the agents are:

$$x_{1,T} = \lim_{p_{1,T}\to 1} \frac{(p_{1,T} - p_{2,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))\mu} = \frac{1}{\lambda p_{2,T}\mu} \tag{35}$$

$$x_{2,T} = -x_{1,T} \tag{36}$$

Expected profits are:

$$\mathbb{E}_{1,T}(\Pi_{1,T}) = \lim_{p_{1,T}\to 1} \frac{(p_{1,T} - p_{2,T})^2 p_{1,T}(1 - p_{1,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))^2} = 0 \tag{37}$$

$$\mathbb{E}_{2,T}(\Pi_{2,T}) = \lim_{p_{2,T}\to 1} \frac{(p_{1,T} - p_{2,T})^2 p_{2,T}(1 - p_{2,T})}{\lambda(p_{1,T}(1 - p_{1,T}) + p_{2,T}(1 - p_{2,T}))^2} = \frac{1 - p_{2,T}}{\lambda p_{2,T}} \tag{38}$$

Case 2.2: $p_{1,T} = 1, p_{2,T} = 1$.

I assume $p_{1,T} = p_{2,T} = q$ with q approaching 1. Equation (21) gives:

$$P = \lim_{q\to 1} \frac{2q^2(1-q)}{2q(1-q)}\mu = \mu \tag{39}$$

Equilibrium holdings of the agents are:

$$x_{1,T} = \lim_{q\to 1} \frac{(q - q)}{\lambda(2q(1-q))\mu} = 0 \tag{40}$$

$$x_{2,T} = -x_{1,T} = 0 \tag{41}$$

Expected profits are:

$$\mathbb{E}_{1,T}(\Pi_{1,T}) = \lim_{q\to 0} \frac{(q-q)^2 q(1-q)}{\lambda(2q(1-q))^2} = 0 \tag{42}$$

$$\mathbb{E}_{2,T}(\Pi_{2,T}) = 0 \tag{43}$$

Note that I do not consider the cases $p_{1,T} = 0, p_{2,T} = 1$ and $p_{1,T} = 1, p_{2,T} = 0$. In those cases holdings and expected profits become arbitrarily large.

9.3 Proof of Proposition 4.2

I solve the optimization problem of agent 1; the result for agent 2 is similar. Define:

$$a = (1 - p_{1,t})r_1 \ ; \ b = 1 - (1 - p_{1,t})r_1$$
$$R = P_{t+1}(p_{1,t+1}, p_{2,t+1})$$
$$D = \mathbb{E}_{1,t+1}(\Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1}))$$
$$G = (1 - (1 - p_{1,t})r_1)\mathbb{E}_{1,t+1}(\Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1}))$$

Since the agent has mean-variance preferences, I need to calculate his expected value and variance of profits. Using (9) the expected value of PnL is:

$$\mathbb{E}_{1,t}(\Pi_{1,t}) = ax(0 - P) + bx(R - P) + bD = x(bR - P) + G \tag{44}$$

The variance is more complicated. We use the conditional variance formula:

$$\text{var}(\Pi_{1,t}|\mathcal{I}_t) = \mathbb{E}(var(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) + var(\mathbb{E}(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) \tag{45}$$

where all the variances and expectations are calculated using the beliefs of agent 1. We have:

$$\begin{aligned} var(\Pi_{1,t}|\mathcal{I}_{t+1}) &= var(x(P_{1,t+1}(p_{1,t+1}, p_{2,t+1}) - P) + \Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1})|\mathcal{I}_{t+1}) \\ &= var(\Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1})|\mathcal{I}_{t+1}) \end{aligned}$$

which is independent of x. The other term is more important:

$$\begin{aligned} var(\mathbb{E}(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) &= var(\mathbb{E}(x(P_{t+1}(p_{1,t+1},p_{2,t+1})-P)+\Pi_{1,t+1}(p_{1,t+1},p_{2,t+1})|\mathcal{I}_{t+1})|\mathcal{I}_t)) \\ &= var_{1,t}(xP_{t+1}(p_{1,t+1},p_{2,t+1})+\mathbb{E}_{1,t+1}(\Pi_{1,t+1}(p_{1,t+1},p_{2,t+1}))) \end{aligned}$$

We know the term in brackets is equal to 0 with probability a, and is equal to $xR+D$ with probability b. Therefore:

$$var(\mathbb{E}(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) = b(xR+D)^2 - b^2(xR+D)^2 = b(1-b)(xR+D)^2 \tag{46}$$

Using (44), (45), (46), the optimization problem for agent 1 reduces to:

$$\max_x \; U_{1,t}(x) = x(bR-P) - \frac{\lambda}{2}b(1-b)(xR+D)^2 + K$$

where K is some constant independent of x. Differentiating with respect to x we get:

$$\frac{\partial}{\partial x}U_{i,t} = bR - P - \lambda b(1-b)(xR+D)R$$

Then, the optimal demand of agent 1 is:

$$x_{1,t} = \frac{b_1R - \lambda b_1(1-b_1)RD_1 - P}{\lambda b_1(1-b_1)R^2}$$

where $b_i = 1-(1-p_{i,t})r_i$, $D_i = \mathbb{E}_{i,t+1}(\Pi_{i,t+1}(p_{1,t+1},p_{2,t+1}))$. Similarly, the demand of agent 2 is:

$$x_{2,t} = \frac{b_2R - \lambda b_2(1-b_2)RD_2 - P}{\lambda b_2(1-b_2)R^2}$$

Then equilibrium price is:

$$P = \frac{1}{\dfrac{1}{b_1(1-b_1)R^2} + \dfrac{1}{b_2(1-b_2)R^2}} \times \Big(\frac{b_1R - \lambda b_1(1-b_1)RD_1}{b_1(1-b_1)R^2} + \frac{b_2R - \lambda b_2(1-b_2)RD_2}{b_2(1-b_2)R^2}\Big)$$

Simplifying, we get:

$$P = \frac{b_1 b_2(2 - b_1 - b_2) - \lambda b_1 b_2(1 - b_1)(1 - b_2)(D_1 + D_2)}{b_1(1 - b_1) + b_2(1 - b_2)} R \tag{47}$$

The expected value for PnL for agent 1 is:

$$\mathbb{E}_{1,t}(\Pi_{1,t}) = x_{1,t}(b_1 R - P) + b_1 D_1 = \frac{b_1 R - \lambda b_1(1 - b_1) R D_1 - P}{\lambda b_1(1 - b_1) R^2}(b_1 R - P) + b_1 D_1$$

We have:

$$b_1 R - P = \frac{(b_1 - b_2) + \lambda b_2(1 - b_2)(D_1 + D_2)}{b_1(1 - b_1) + b_2(1 - b_2)} b_1(1 - b_1) R$$

And also:

$$b_1 R - \lambda b_1(1 - b_1) R D_1 - P = \frac{(b_1 - b_2) + \lambda b_2(1 - b_2) D_2 - \lambda b_1(1 - b_1) D_1}{b_1(1 - b_1) + b_2(1 - b_2)} b_1(1 - b_1) R$$

Therefore we get a recursive formula for expected PnL of agent 1:

$$\begin{aligned} \mathbb{E}_{1,t}(\Pi_{1,t}) = {} & \frac{(b_1 - b_2) + \lambda b_2(1 - b_2) D_2 - \lambda b_1(1 - b_1) D_1}{\lambda(b_1(1 - b_1) + b_2(1 - b_2))} \\ & \times \frac{(b_1 - b_2) + \lambda b_2(1 - b_2)(D_1 + D_2)}{b_1(1 - b_1) + b_2(1 - b_2)} \times b_1(1 - b_1) + b_1 D_1 \end{aligned}$$

The formula for agent 2 is similar:

$$\begin{aligned} \mathbb{E}_{2,t}(\Pi_{2,t}) = {} & \frac{(b_2 - b_1) + \lambda b_1(1 - b_1) D_1 - \lambda b_2(1 - b_2) D_2}{\lambda(b_1(1 - b_1) + b_2(1 - b_2))} \\ & \times \frac{(b_2 - b_1) + \lambda b_1(1 - b_1)(D_1 + D_2)}{b_1(1 - b_1) + b_2(1 - b_2)} \times b_2(1 - b_2) + b_2 D_2 \end{aligned}$$

The sum of the profits is:

$$\begin{aligned} \mathbb{E}_{1,t}(\Pi_{1,t}) + \mathbb{E}_{2,t}(\Pi_{2,t}) = {} & \frac{(b_1 - b_2) + \lambda b_2(1 - b_2) D_2 - \lambda b_1(1 - b_1) D_1}{\lambda(b_1(1 - b_1) + b_2(1 - b_2))} \\ & \times \frac{(b_1 - b_2) b_1(1 - b_1) + (b_1 - b_2) b_2(1 - b_2)}{b_1(1 - b_1) + b_2(1 - b_2)} + b_1 D_1 + b_2 D_2 \end{aligned}$$

We can further simplify this to:

$$\mathbb{E}_{1,t}(\Pi_{1,t}) + \mathbb{E}_{2,t}(\Pi_{2,t}) = \frac{(b_1 - b_2)^2 + \lambda b_2(1-b_2)(b_1-b_2)D_2 - \lambda b_1(1-b_1)(b_1-b_2)D_1}{\lambda(b_1(1-b_1) + b_2(1-b_2))} + b_1 D_1 + b_2 D_2$$

Simplifying further, we get:

$$\mathbb{E}_{1,t}(\Pi_{1,t}) + \mathbb{E}_{2,t}(\Pi_{2,t}) = \frac{(b_1 - b_2)^2 + \lambda b_1 b_2(2 - b_1 - b_2)(D_1 + D_2)}{\lambda(b_1(1-b_1) + b_2(1-b_2))} \tag{48}$$

as desired.

9.4 Proof of Proposition 5.1

The proof is similar to the one for the recursive relations in the simple model. I solve the optimization problem for agent 1; for agent 2 the result is similar.

Suppose the current price (at date t) is P, and agent 1 considers holding x units of the risky asset. The distribution of the profit $\Pi_{1,t}$ for the agent, conditional on signal s_{t+1}, is as follows:

$$\Pi_{1,t}(p_{1,t}, p_{2,t}) = x(R_j - P) + \Pi_{1,t+1}(p_{1,t,j}, p_{2,t,j}), \text{ if } s_{t+1} = z_j \tag{49}$$

At time t agent 1 believes the probability of the good state is $p_{1,t}$, and therefore, he believes that $s_{t+1} = z_j$ with probability $a_{1,j}$:

$$\begin{aligned} \mathbb{P}_{1,t}(s_{t+1} = j) &= \mathbb{P}_{1,t}(\mathcal{S} = 1)\mathbb{P}_{1,t}(s_{t+1} = z_j | \mathcal{S} = 1) + \mathbb{P}_{1,t}(\mathcal{S} = 2)\mathbb{P}_{1,t}(s_{t+1} = z_j | \mathcal{S} = 2) \\ &= p_{1,t} f_{1,1}(z_j) + (1 - p_{1,t}) f_{1,2}(z_j) = a_{1,j} \end{aligned} \tag{50}$$

Since the agent has mean-variance preferences, I need to calculate his expected value and variance of profits. Using (49) and (50) the expected value of PnL is:

$$\begin{aligned} \mathbb{E}_{1,t}(\Pi_{1,t}) &= \sum_{j=1}^{n} a_{1,j} \Big(x(R_j - P) + D_{1,j} \Big) \\ &= x\Big((\sum_{j=1}^{n} a_{1,j} R_{1,j}) - P \Big) + \sum_{j=1}^{n} a_{1,j} D_{1,j} \end{aligned} \tag{51}$$

For the variance I use the conditional variance formula as in the simple model:

$$\mathrm{var}(\Pi_{1,t}|\mathcal{I}_t) = \mathbb{E}(var(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) + var(\mathbb{E}(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) \tag{52}$$

where all the variances and expectations are calculated using the beliefs of agent 1. Following the same derivation as with the simple model, only the second term depends on x, and can be written as:

$$var(\mathbb{E}(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) = var_{1,t}(xP_{t+1}(p_{1,t+1}, p_{2,t+1}) + \mathbb{E}_{1,t+1}(\Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1})))$$

The term in brackets is constant conditional on the value of s_{t+1}:

$$xP_{t+1}(p_{1,t+1}, p_{2,t+1}) + \mathbb{E}_{1,t+1}(\Pi_{1,t+1}(p_{1,t+1}, p_{2,t+1})) = xR_j + D_{1,j} \quad \text{if } s_{t+1} = z_j$$

From (50) we know the distribution of s_{t+1}. This is enough to calculate the variance:

$$var(\mathbb{E}(\Pi_{1,t}|\mathcal{I}_{t+1})|\mathcal{I}_t) = \sum_{j=1}^{n} a_{1,j}(xR_j + D_{1,j})^2 - (\sum_{j=1}^{n} a_{1,j}(xR_j + D_{1,j}))^2 \tag{53}$$

Using (51), (52), (53), the optimization problem for agent 1 reduces to:

$$\begin{aligned}\max_x \; U_{1,t}(x) = x\Big((\sum_{j=1}^{n} a_{1,j}R_{1,j}) - P\Big) - \frac{\lambda}{2}\Big(\sum_{j=1}^{n} a_{1,j}(xR_j + D_{1,j})^2 \\ - (\sum_{j=1}^{n} a_{1,j}(xR_j + D_{1,j}))^2\Big) + K\end{aligned}$$

where K is some constant independent of x. Differentiating with respect to x we get:

$$\begin{aligned}\frac{\partial}{\partial x}U_{1,t} &= (\sum_{j=1}^{n} a_{1,j}R_j) - P - \lambda\sum_{j=1}^{n} a_{1,j}R_j(xR_j + D_{1,j}) + \lambda\Big(\sum_{j=1}^{n} a_{1,j}(xR_j + D_{1,j})\Big)(\sum_{j=1}^{n} a_{1,j}R_j) \\ &= \mu_{1,t} - P - \lambda w_{1,t}x\end{aligned}$$

where for $i = 1, 2$:

$$\mu_{i,t} = \sum_{j=1}^{n} a_{i,j} R_j - \lambda \sum_{j=1}^{n} a_{i,j} R_j D_{i,j} + \lambda (\sum_{j=1}^{n} a_{i,j} D_{i,j})(\sum_{j=1}^{n} a_{i,j} R_j) \tag{54}$$

$$w_{i,t} = \sum_{j=1}^{n} a_{i,j} R_j^2 - (\sum_{j=1}^{n} a_{i,j} R_j)^2 \tag{55}$$

The optimal demand of agent 1 is:

$$x_{1,t} = \frac{\mu_{1,t} - P}{\lambda w_{1,t}}$$

Similary, the demand of agent 2 is:

$$x_{2,t} = \frac{\mu_{2,t} - P}{\lambda w_{2,t}}$$

Then equilibrium price is:

$$P = \frac{w_{2,t}\mu_{1,t} + w_{1,t}\mu_{2,t}}{w_{1,t} + w_{2,t}} \tag{56}$$

Using (51) the expected PnL of agent i is:

$$\begin{aligned} \mathbb{E}_{i,t}(\Pi_{i,t}) &= x_{i,t}\Big((\sum_{j=1}^{n} a_{i,j} R_j) - P) \Big) + \sum_{j=1}^{n} a_{i,j} D_{i,j} \\ &= \frac{\mu_{i,t} - P}{\lambda w_{i,t}} \Big((\sum_{j=1}^{n} a_{i,j} R_j) - P) \Big) + \sum_{j=1}^{n} a_{i,j} D_{i,j} \end{aligned} \tag{57}$$

9.5 Expected Profits for Persistent Bad News

I present and analyze the patterns for expected profits of the agents for the discrete model in the case when a bad signal arrives in each period. Figures 9.1 and 9.2 plot the expected profits of agents 1 and 2, respectively, as they change over time.

Each of the curves on the expected profits charts has exactly one peak. This peak occurs at different times depending on the agent and the case. For agent 1, the cases for large high threshold (blue and green curves) offer high expected profits in early periods. Expected PnL then continues to rise, more so for the blue curve – since in dates 4 and 5 agent 1 is bullish and of type 2, while agent 2 already has type 3. So there is a large difference of opinions, while agent 2 is less likely to adjust his opinion to potential high signals, leading to greater expected profits for agent 1, than if agent 2 was of type 2. After that expected profits start to drop as the agents start to converge in their beliefs. The drop is more rapid since agent 1 becomes type 3 faster due to the higher low threshold at which point expected profits are very low.

Continuing with agent 1, the cases for small high threshold (red and purple curves) offer lower profits in early periods. The profits then continue to rise; more rapidly for the red curve – again because agent 2 becomes type 3 faster. For the purple curve, profits first rise slowly, and then grow very rapidly, as agent 1 still has type 1, but agent 3 has type 3, so there is a large difference in opinions. As with the blue and green curves, the expected profits then drop as the beliefs converge, and agent 1 becomes type 3 along with agent 2 who by that time has switched to type 3 several periods earlier.

The pattern for agent 2 is similar. It is interesting that the curves with similar high (low) threshold are a lot more similar for agent 2 than for agent 1. This is likely because the peak occurs for agent 2 when he has type 2 (or has just become type 3) – since after that he pushes the price too much and is too convinced in his beliefs to expect a large difference in opinions later on. Thus the location of the high threshold matters a lot more for agent 2.

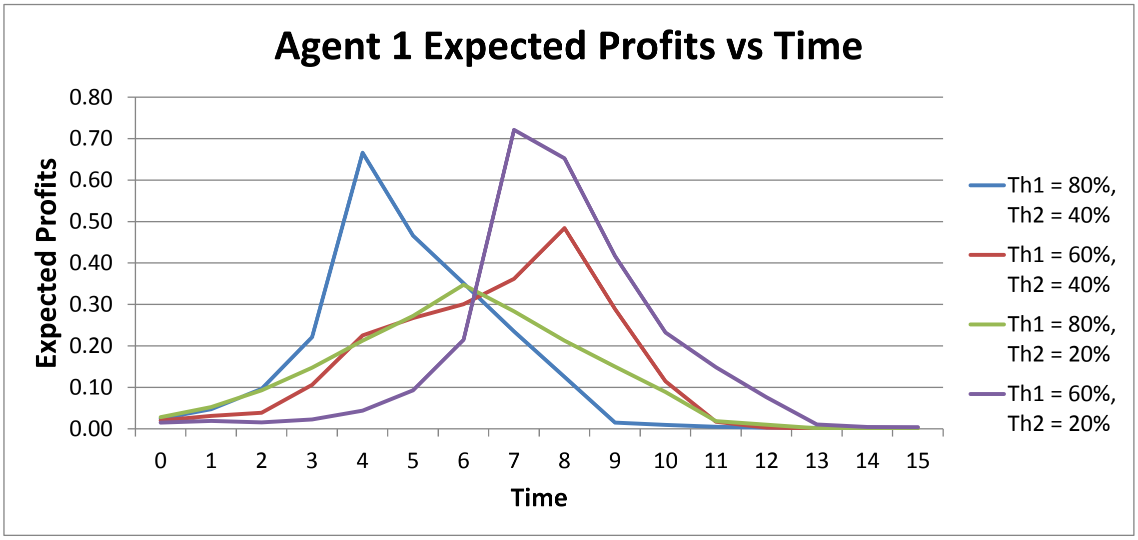

Figure 9.1: Expected profits of agent 1 as a function of time.

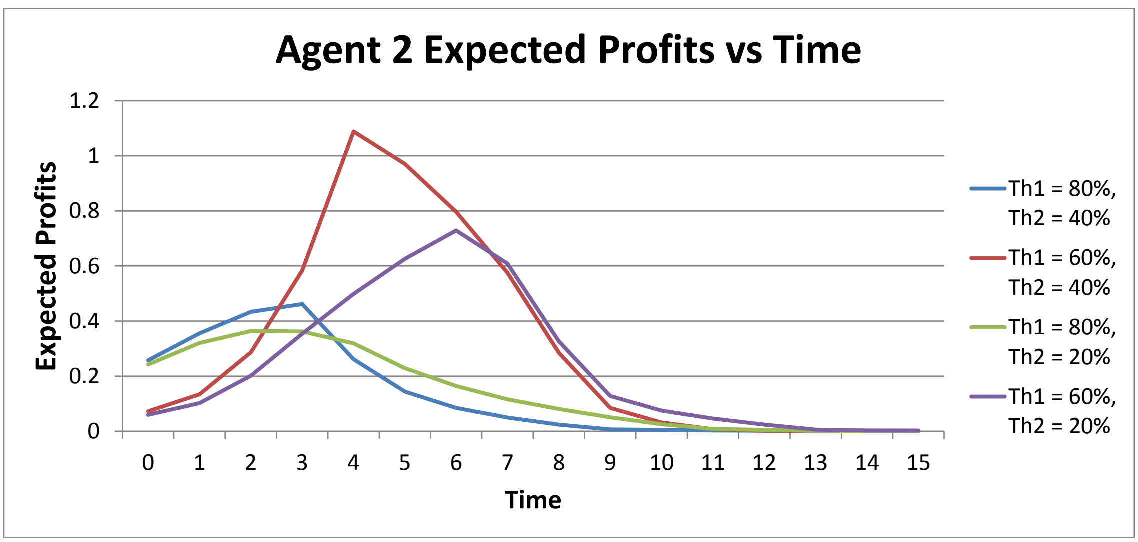

Figure 9.2: Expected profits of agent 2 as a function of time.

References

[1] Daniel Andrei, Bruce Carlin, and Michael Hasler. Model disagreement and economic outlook. Technical report, National Bureau of Economic Research, 2014.

[2] Clifford S Asness, Tobias J Moskowitz, and Lasse Heje Pedersen. Value and momentum everywhere. *The Journal of Finance*, 68(3):929–985, 2013.

[3] Snehal Banerjee and Ilan Kremer. Disagreement and learning: Dynamic patterns of trade. *The Journal of Finance*, 65(4):1269–1302, 2010.

[4] Brad M Barber, Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. Is the aggregate investor reluctant to realise losses? evidence from taiwan. *European Financial Management*, 13(3):423–447, 2007.

[5] António MRG Barbosa. Differential interpretation of information and the post-announcement drift: A story of consensus learning. *Available at SSRN 1932654*, 2011.

[6] Marc Bremer and Richard J Sweeney. The reversal of large stock-price decreases. *The Journal of Finance*, 46(2):747–754, 1991.

[7] William Brock, Josef Lakonishok, and Blake LeBaron. Simple technical trading rules and the stochastic properties of stock returns. *The Journal of Finance*, 47(5):1731–1764, 1992.

[8] Wesley S Chan. Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of Financial Economics*, 70(2):223–260, 2003.

[9] Julien Cujean and Michael Hasler. Disagreement and predictability over the business cycle. Technical report, Working paper, University of Maryland, 2014.

[10] Werner FM De Bondt and Richard Thaler. Does the stock market overreact? *The Journal of Finance*, 40(3):793–805, 1985.

[11] Kari Edwards and Edward E Smith. A disconfirmation bias in the evaluation of arguments. *Journal of Personality and Social Psychology*, 71(1):5, 1996.

[12] Simon Gervais and Terrance Odean. Learning to be overconfident. *Review of Financial Studies*, 14(1):1–27, 2001.

[13] Hua He and Jiang Wang. Differential informational and dynamic behavior of stock trading volume. *Review of Financial Studies*, 8(4):919–972, 1995.

[14] Harrison Hong and Jeremy C Stein. Differences of opinion, short-sales constraints, and market crashes. *Review of Financial Studies*, 16(2):487–525, 2003.

[15] Narasimhan Jegadeesh and Sheridan Titman. Returns to buying winners and selling losers: Implications for stock market efficiency. *The Journal of Finance*, 48(1):65–91, 1993.

[16] Eugene Kandel and Neil D Pearson. Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy*, pages 831–872, 1995.

[17] Oliver Kim and Robert E Verrecchia. Market liquidity and volume around earnings announcements. *Journal of Accounting and Economics*, 17(1):41–67, 1994.

[18] Lucian L Leape. Error in medicine. *Jama*, 272(23):1851–1857, 1994.

[19] Edwin Lefevre. *Reminiscences of a Stock Market Operator*. 1923.

[20] Charles G Lord, Lee Ross, and Mark R Lepper. Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence. *Journal of personality and social psychology*, 37(11):2098, 1979.

[21] Tobias J Moskowitz, Yao Hua Ooi, and Lasse Heje Pedersen. Time series momentum. *Journal of Financial Economics*, 104(2):228–250, 2012.

[22] John J Murphy. *Technical analysis of the financial markets: A comprehensive guide to trading methods and applications*. Penguin, 1999.

[23] Carol L Osler. Support for resistance: technical analysis and intraday exchange rates. *Economic Policy Review*, 6(2), 2000.

[24] Marco Ottaviani, Peter Norman Sørensen, et al. Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal. *American Economic Review*, 105(1):1–34, 2015.

[25] Thomas R Palfrey and Stephanie W Wang. Speculative overpricing in asset markets with information flows. *Econometrica*, 80(5):1937–1976, 2012.

[26] Hersh Shefrin and Meir Statman. The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of finance*, 40(3):777–790, 1985.

[27] Catherine A Sheldon and Lynne M Sallot. Image repair in politics: Testing effects of communication strategy and performance history in a faux pas. *Journal of Public Relations Research*, 21(1):25–50, 2009.

[28] Katsuhiko Shimizu and Michael A Hitt. Strategic flexibility: Organizational preparedness to reverse ineffective strategic decisions. *The Academy of Management Executive*, 18(4):44–59, 2004.

[29] Costas Xiouros. Disagreement, habit and the dynamic relation between volume and prices. *Available at SSRN 1744082*, 2011.