

More hedging instruments may destabilize markets

William Brock^a, *Cars Hommes*^{b*} and *Florian Wagener*^b

This Draft: December 2008; first draft September 2006

Abstract

This paper formalizes the idea that more hedging instruments may destabilize markets when traders have heterogeneous expectations and adapt their behavior according to experience based reinforcement learning. In a simple asset pricing model with heterogeneous beliefs the introduction of additional Arrow securities may destabilize markets, and thus increase price volatility, and at the same time decrease average welfare. We also investigate whether a fully rational agent can employ additional hedging instruments to stabilize markets. It turns out that the answer depends on the composition of the population of non-rational traders and the information gathering costs for rationality.

keywords: financial innovation, asset pricing, hedging, reinforcement learning, bifurcations

Acknowledgments. Earlier versions of this paper have been presented at the SCE-conference on Computational Economics and Finance, Amsterdam, July 8-10, 2004, the workshop “Volatility of financial markets: theoretical models, forecasting and trading”, at the Lorentz Center Leiden, October 18-29, 2004, the workshop on “Complexity and Randomness in Economic Dynamical Systems”, Bielefeld, March 17-19, 2005, the ECB workshop on “Asset markets, expectations and learning”, November 6-7, 2006, the summer school on “Agent-based finance”, Trento, July 1-14, 2007, the “Complex markets” workshop, Warwick, April 4-5, 2008 and in various seminars. Stimulating discussions and helpful comments from Larry Blume, Wouter Den Haan, Doyne Farmer, Albert Marcet, Patrick Leoni, and other workshop and seminar participants are gratefully acknowledged. We are particularly grateful for the detailed comments by two anonymous referees and the Editor on an earlier draft, which have led to significant improvements. This research has been supported by the Netherlands Organization for Scientific Research (NWO), the NSF, the Vilas Trust and by a EU STREP-grant “Complex Markets”.

^a University of Wisconsin. Email: wbrock@ssc.wisc.edu.

^b University of Amsterdam, Email: C.H.Hommes@uva.nl (Hommes) (corresponding author); F.O.O.Wagener@uva.nl (Wagener).

“Our fundamental risks will thus be insured against, hedged, diversified, making for a safer world. By lightening the burden of risk, a new democratic finance will encourage all of us to be more venturesome, more inspired in our activities.”, Robert J. Shiller, *The New Financial Order: Risk in the 21st Century*, Princeton University Press, 2003.

1 Introduction

Robert J. Shiller advocates an expansion of the number of appropriately designed risk hedging instruments. Many writers, e.g., Rajan (2005), have recently raised concerns about the impact on market stability due to the explosive growth of innovations like “financial engineering” that has essentially amounted to a rapid growth in the number of risk hedging instruments. We agree with a commonly held view that increases in appropriately designed risk hedging instruments can increase welfare; but we are also concerned about the impact of addition of extra risk hedging instruments on the process of achieving (or not achieving) equilibrium.

This paper formalizes the idea that more hedging instruments or derivative securities may *destabilize* a market when traders are heterogeneous and learn from experience based on realized returns. Here is a sketch of the idea. Consider a heterogeneous agent intertemporal asset market where risk averse agents are learning the structure of asset prices in the economy by using, for example, different prediction strategies of future asset prices under some kind of reinforcement or evolutionary learning, for instance as in Brock and Hommes (1997). Let there be S states of the world and a finite number n of contingent claims or risk hedging instruments available. If $n < S - 1$ the market is incomplete. We model the risk hedging instruments as “Arrow” securities for state s , $1 \leq s \leq n < S - 1$, each paying 1 if state s occurs and 0 otherwise. Elementary Arrow securities are used here as a convenient analytical device, and a suitable combination of Arrow securities may serve as a proxy of more realistic financial instruments such as futures, derivatives or recently introduced collateralized debt obligations. Now

suppose that $n < S - 2$ and that a new risk hedging instrument, that is, a new Arrow security, is added for state $n + 1 < S - 1$. Then, since agents are risk averse, and since they can use the new Arrow security to hedge out “extra” risk, they will now tend to place bigger positions on the market. Thus if an agent’s forecasting tool turns out to be on the “right side” of the market, it will return a larger profit (because a larger position has been placed on the market), and therefore it will receive a stronger reinforcement and more individuals will switch to using that particular forecasting tool. This, in turn, implies that the learning system is now more likely to “overshoot”, i.e. to become unstable, and consequently market volatility increases. This intuitive idea will be formalized in a stylized model.

On the other hand it has been argued that an increasing multitude of derivative securities has made it possible for *rational* speculators to help *stabilize* markets since they can take bets on market imperfections and hedge their risk. A second contribution of our paper is to investigate the potential stabilizing role of rational traders in a market with co-existing non fully rational traders. Can a perfectly rational trader employ a growing number of hedging instruments to stabilize the market? It turns out that, when the information gathering costs for full rational expectations are large, rational traders can not prevent destabilization. However, we will also present conditions (depending on the composition of the co-existing population of non-rational traders) under which, as the number of hedging instruments increases, the benefits of “thinking hard” outweigh its costs, and as the market approaches completeness, rational agents may outperform non-rational traders, stabilize the market and limit welfare losses.

To formalize these ideas in the simplest setting we use the asset pricing model with heterogeneous beliefs of Brock and Hommes (1998), but the analysis can be generalized, for instance to a general equilibrium overlapping generations framework. We show that adding more Arrow securities may destabilize market dynamics and thus increase market volatility. In particular, we show that the primary bifurcation parameter, marking the onset of instability, occurs “earlier” when there are more Arrow securities.

Comparing our approach to the existing literature, it is probably fair to say that most research in finance leans towards the standard financial economics view that adding derivatives or futures markets increases welfare, reduces volatility and improves information revelation. The “General Equilibrium with Incomplete markets” (GEI) literature nevertheless contains a number of theoretical papers showing that introduction of new securities may *decrease* equilibrium welfare (e.g. Hart 1975, Elul 1995, Cass and Citanna 1998), or may *increase* price volatility (e.g. Citanna and Schmedders, 2005, Bhamra and Uppal, 2006; see the comprehensive survey by Mayhew (2000)¹). An important difference with our approach is that these papers investigate finite period, static GEI equilibrium models under rational expectations, while we attempt to model how learning dynamics and heterogeneous expectations affects the attainment of equilibrium; see Farmer and Geanakoplos (2008) for a recent discussion. An important contribution of our paper is that *reinforcement learning dynamics of non-fully rational agents* is a potentially amplifying force to price instability when the number of hedging instruments increases.

There is empirical evidence that experience based reinforcement learning, a key feature in our modeling framework, also plays an important role in investment decisions in real markets. For example, Ippolito (1992), Chevalier and Ellison (1997), Sirri and Tufano (1998), Rockinger (1996) and Karceski (2002) show for mutual funds data that money flows into past good performers, while flowing out of past poor performers, and that performance persists on a short term basis. Pension funds are less extreme in picking good performance but are tougher on bad performers (Del Guerico and Tkac, 2002). Recently, Benartzi and Thaler (2007) have shown that heuristics and biases

¹Mayhew (2000) also surveys empirical work on how futures and derivatives affect price volatility of the underlying. The empirical results on the introduction of futures are ambiguous, some authors finding a decrease, while others finding an increase in volatility. Gerlach et al. (2006) investigate the behavior of volatility of returns in bond and stock markets for a sample of eight countries using 150 years of data. Unsurprisingly, volatility has been high during episodes of economic and political turbulence. Interestingly, volatility has been high since the seventies, a time featuring an explosive growth of financial futures.

play a significant role in retirement savings decisions. For example, using data from Vanguard they show that the equity allocation of new participants rose from 58% in 1992 to 74% in 2000, following a strong rise in stock prices in the late 1990s, but dropped, back to 54% in 2002, following the strong fall in stock prices.

The asset pricing model with heterogeneous belief that we employ as our stylized framework is closely related to work in behavioral finance, evolutionary finance and adaptive learning. The reader is referred to a number of recent surveys: Barberis and Thaler (2003) give an extensive overview of *behavioral finance* (e.g. including the work on noise trader models as in DeLong et al., 1990). More general surveys on learning and bounded rationality and their role in enforcing convergence to rational expectations or creating excess volatility include Evans and Honkapohja (2001), Grandmont (1998) and Sargent (1993). Hens and Schenk-Hoppé (2008) provide stimulating and comprehensive survey chapters in evolutionary finance.

The paper is organized as follows. Section 2 extends the asset pricing model with heterogeneous beliefs to include Arrow securities. The main result here is that, when there are more Arrow securities, the primary bifurcation towards instability occurs earlier. Section 3 investigates the potential stabilizing role of rational traders. Conditions (depending on the composition of non-rational traders) under which rational agents can employ the Arrow securities to stabilize the market are given. Section 4 provides an example where financial innovation leads to an increase of market volatility as well as a decrease in welfare. The example also shows that rational traders can not always stabilize prices, when the market approaches completeness. Section 5 concludes. The paper closes with a summary, conclusions, a brief discussion of hedging strategies in real markets and suggestions for future research. An Appendix provides proofs of the results.

2 An asset pricing model

In this section we extend the asset pricing model with heterogeneous beliefs of Brock and Hommes (1997, 1998) by adding contingent claims or Arrow securities and investigate how these hedging instruments affect market stability. For any time period t , there are S possible states s of the world in period $t+1$, occurring with probabilities α_s that are independent of time and common knowledge. Agents can buy risk free bonds and two types of risky assets, stocks and Arrow securities. Bonds are bought at a fixed price 1 and pay $R > 1$ in the next period. Stocks are bought at a market price p_t^0 in period t and next period in state of the world s they pay an amount

$$q_{t+1}^s = p_{t+1}^0 + y^s,$$

that is the new market price p_{t+1}^0 plus a dividend y^s depending on s . Finally Arrow securities for state i are bought at p_t^i and pay δ_i^s in state s , which is 1 if $s = i$ and 0 otherwise. Markets are incomplete: Arrow securities are only available for states $1, \dots, n$, where $n < S - 1$.

Let z_t^0 and z_t^i denote the demand of an agent for respectively the stock and the i 'th Arrow security. Introduce vector notation by setting $\tilde{z}_t = (z_t^1, \dots, z_t^n)$ and $z_t = (z_t^0, \tilde{z}_t)$; $\tilde{p}_t = (p_t^1, \dots, p_t^n)$ and $p_t = (p_t^0, \tilde{p}_t)$; $\delta = (\delta_1, \dots, \delta_n)$ and $\alpha = (\alpha_1, \dots, \alpha_n)$. Introduce moreover $\sigma^2 = \text{Var } q_{t+1}$; $\eta_i = \text{Cov}(q_{t+1}, \delta_i)$ and $\eta = (\eta_1, \dots, \eta_n)$, and $\Sigma = \text{Cov}(\delta)$. Finally, let $a > 0$ be the coefficient of risk aversion and let V_n denote the symmetric $(n+1) \times (n+1)$ variance-covariance matrix

$$V_n = a \text{Cov}((q_{t+1}, \delta)) = a \begin{pmatrix} \sigma^2 & \eta^T \\ \eta & \Sigma \end{pmatrix}. \quad (1)$$

Note that V_n is the variance-covariance matrix of the uncertain payments of the stock and the n Arrow securities multiplied by the coefficient of risk aversion a . The matrix is singular if and only if a riskless portfolio can be constructed out of stock and Arrow securities; this would for instance be possible if there were $n = S - 1$ Arrow securities

available. The elements of V_n can be computed. They read as

$$\sigma^2 = \sum (y^s - \bar{y})^2 \alpha_s, \quad \eta_i = \alpha_i (y^i - \bar{y}),$$

$$\Sigma_{ij} = \begin{cases} \alpha_i (1 - \alpha_i) & \text{if } i = j, \\ -\alpha_i \alpha_j & \text{if } i \neq j. \end{cases}$$

The inner product of two vectors v and w is denoted by $\langle v, w \rangle$. If W_t is the current wealth of an agent, his wealth next period in state s is

$$W_{t+1}^s = R (W_t - p_t^0 z_t^0 - \langle \tilde{p}_t, \tilde{z}_t \rangle) + q_{t+1}^s z_t^0 + z_t^s.$$

The excess profit in state s from trading the risky assets equals $\pi_{t+1}^s = W_{t+1}^s - RW_t$.

Utility is assumed to be of mean–variance type:

$$U_t = \mathbb{E}_t \pi_{t+1} - \frac{a}{2} \text{Var}_t \pi_{t+1} = \left\langle \begin{pmatrix} -Rp_t^0 + \mathbb{E}_t q_{t+1} \\ -R\tilde{p}_t + \mathbb{E}_t \delta \end{pmatrix}, z_t \right\rangle - \frac{1}{2} \langle z_t, V_n z_t \rangle. \quad (2)$$

The optimal demand vector is given by

$$z_t = V_n^{-1} \begin{pmatrix} -Rp_t^0 + \mathbb{E}_t q_{t+1} \\ -R\tilde{p}_t + \mathbb{E}_t \delta \end{pmatrix}. \quad (3)$$

2.1 Rational benchmark The case of all traders having rational expectations is the fundamental benchmark of the system. Arrow securities are endogenous to the system and therefore their total supply is zero. The total supply of the stock is ζ^0 .

Denote expected dividends by $\bar{y} = \sum y^s \alpha_s$. If all markets clear, then we obtain under rational expectations, using equations (1) and (3), the price dynamics

$$\begin{aligned} -Rp_t^0 + \mathbb{E}_t p_{t+1}^0 + \bar{y} &= a\sigma^2 \zeta^0, \\ -R\tilde{p}_t + \alpha &= a\eta \zeta^0. \end{aligned}$$

Imposing the transversality condition that prices remain bounded, these equations are solved by constant *fundamental* prices $p_t = p_* = (p_*^0, \tilde{p}_*)$, given as

$$p_*^0 = \frac{\bar{y} - a\sigma^2 \zeta^0}{R - 1}, \quad \tilde{p}_* = \frac{1}{R} (\alpha - a\eta \zeta^0). \quad (4)$$

The terms involving ζ^0 can be interpreted as the risk premium required by the investors to hold the risky assets.

2.2 Heterogeneous expectations. Consider now the case that agents are heterogeneous in their expectations or beliefs about next period's price of the stock, but homogeneous with respect to everything else.²

Demand of type h , $1 \leq h \leq H$, is given by equation (3), reading as

$$z_{ht} = V_n^{-1} \begin{pmatrix} -Rp_t^0 + \mathbb{E}_{ht}q_{t+1} \\ -R\tilde{p}_t + \mathbb{E}_t\delta \end{pmatrix} = V_n^{-1}B_{ht}. \quad (5)$$

Here B_{ht} is the *belief vector* of type h about the excess return of stock and Arrow securities; this belief vector determines the investment strategy through (5). We will refer to type h as a belief type or a prediction strategy. Since probabilities of states of the world are assumed to be common knowledge, the expectation $\mathbb{E}_t\delta$ is the same for all types. Note that agents differ in their assessment of $\mathbb{E}_{ht}q_{t+1}$, but agree on V_n . This simplifying assumption is made for analytical tractability of the heterogeneous agent case, but it is supported by the observation that there may be more agreement about the variance than about the mean³.

It will be convenient to work with price *deviations* $x_t = p_t - p_*$ from the fundamental

²Heterogeneous expectations play an increasingly important role in economics and finance. LeBaron (2006) and Hommes (2006) are up to date reviews, with each more than 100 references. Heterogeneity in forecasting future asset prices is supported by evidence from survey data, as discussed e.g. in Vissing-Jorgensen (2003) and Shiller (2000). Branch (2004) estimates a model with heterogeneous beliefs and time varying fractions, using survey data on inflation expectations, while Boswijk et al. (2007) estimate a simple two type asset pricing model with heterogeneous beliefs, fundamentalists versus trend extrapolators, on yearly S&P 500 data, 1871-2003.

³The observation that estimation of the variance or covariance from observed financial returns series will be much more accurate than estimation of the mean dates back to Merton (1980, especially Appendix A). The ARCH/GARCH literature has shown that, under regularity conditions, conditional variance is easier to estimate than conditional mean, see e.g. Bollerslev, Engle, and Nelson (1994, especially section 4).

benchmark prices p_* . We assume price expectations to be of the form

$$\mathbb{E}_{ht} p_{t+1}^0 = p_*^0 + f_{ht} = p_*^0 + f_h(x_{t-1}^0, \dots, x_{t-L}^0).$$

The “technical trading rule” f_h models how type h believes that the future price p_{t+1}^0 will deviate from the fundamental, given past prices.

2.2.1 Market clearing. Let the market share (or fraction) of type h agents in period t be denoted by n_{ht} . As before, Arrow securities are endogenous to the system, and their total supply is zero. Market clearing for stock and Arrow securities implies

$$\sum_h n_{ht} z_{ht}^0 = \zeta^0, \quad \sum_h n_{ht} \tilde{z}_{ht} = 0. \quad (6)$$

In deviations from the fundamental, the demand vector reads as

$$z_{ht} = \begin{pmatrix} \zeta^0 \\ 0 \end{pmatrix} + V_n^{-1} \begin{pmatrix} -R x_t^0 + f_{ht} \\ -R \tilde{x}_t \end{pmatrix}. \quad (7)$$

Adding these equations, weighted by fractions, and using equation (6) yields

$$R x_t^0 = \sum_h n_{ht} f_{ht}, \quad \tilde{x}_t = 0. \quad (8)$$

We make a couple of observations. First, according to (8), the price deviations of the Arrow securities are zero, implying that the Arrow securities are correctly priced. Secondly, if f_{ht} happens to be equal for all types h , beliefs are homogeneous, and there is no demand for Arrow securities. Only when beliefs are truly heterogeneous the demand for Arrow securities will be non-zero, as different types hedge their risk differently. Finally, under heterogeneous beliefs the market price of the stock will in general deviate from its fundamental benchmark. In fact, the expression $R x_t^0 = \sum_h n_{ht} f_{ht}$ in (8) is the same as in the asset pricing model without Arrow securities in Brock and Hommes (1998). However, as we will see below, the existence of Arrow securities will affect the magnitude of the fractions n_{ht} through reinforcement learning.

2.2.2 Fitness. In order to close the model, the evolution of the market shares n_{ht} has to be specified. We assume that these shares are determined by the fitness u_{ht-1} of type h ; the subscript $t - 1$ indicates that fitness depends only on past observed prices. The fraction of agents using strategy type h will thus be driven by “experience” or “regret” through reinforcement learning. Given fitnesses u_{ht-1} , the fraction of agents using strategy type h is determined by a multinomial logit model:⁴

$$n_{ht} = \frac{e^{\beta u_{ht-1}}}{Z_t}, \quad Z_t = \sum_h e^{\beta u_{ht-1}}, \quad (9)$$

These fractions can be derived from a random utility model (Manski and McFadden (1981)). Note that as the fitness u_{ht-1} increases relative to the other type’s fitnesses, more agents will select trading strategy h . The *intensity of choice* parameter $\beta > 0$ in (9) measures how sensitive agents are to selecting the optimal prediction strategy. This parameter is inversely related to the variance of the noise in the observation of random utility. If $\beta = 0$, agents are insensitive to past performance and pick a strategy at random: all fractions will then be equal to $1/H$. In the other extreme case $\beta \rightarrow \infty$, all agents choose the forecast which performed best in the last period. An increase in the intensity of choice β can therefore be seen as to represent an increase in the degree of rationality with respect to evolutionary selection of strategies.

As fitness measure we use average risk-adjusted profit, that is, average profits corrected for the risk taken when buying risky assets:

$$u_{ht} = \left\langle \left(\begin{array}{c} -Rp_{t-1}^0 + p_t^0 + \bar{y} \\ -R\tilde{p}_{t-1} + \alpha \end{array} \right), z_{ht-1} \right\rangle - \frac{1}{2} \langle z_{ht-1}, V_n z_{ht-1} \rangle.$$

Notice that this measure is consistent with the utility (2) of the agents⁵. Using (5), the

⁴The results discussed below do not depend on the details of the specification of the updating rule (9). The exponential function may for instance be replaced by another increasing function. What is key in (9) is that, as the intensity of choice β moves from one extreme 0 to the other extreme $+\infty$, the distribution of types moves from uniform to a delta function with its peak at the best strategy.

⁵Another fitness measure that may be of interest is (non-risk adjusted) realized profits. The results for this alternative fitness measure are very similar to those presented below.

realized excess return vector

$$B_{t-1} = \begin{pmatrix} -Rp_{t-1}^0 + p_t^0 + \bar{y} \\ -R\tilde{p}_{t-1} + \alpha \end{pmatrix}, \quad (10)$$

and recalling that $z_{ht} = V_n^{-1}B_{ht}$, we rewrite risk-adjusted realized profits as

$$u_{ht} = \langle B_{t-1}, V_n^{-1}B_{h,t-1} \rangle - \frac{1}{2} \langle B_{h,t-1}, V_n^{-1}B_{h,t-1} \rangle.$$

In the special case where type h has rational expectations or perfect foresight, i.e. $B_{h,t-1} = B_{t-1}$, this expression simplifies to $u_t^R = \frac{1}{2} \langle B_{t-1}, V_n^{-1}B_{t-1} \rangle$. Now look at the difference between risk-adjusted profits of type h and fully rational agents, i.e.

$$\begin{aligned} u_{ht} - u_t^R &= -\frac{1}{2} \langle B_{t-1} - B_{h,t-1}, V_n^{-1}(B_{t-1} - B_{h,t-1}) \rangle \\ &= -\frac{1}{2} \langle (x_t^0 - f_{h,t-1})e_0, V_n^{-1}(x_t^0 - f_{h,t-1})e_0 \rangle \\ &= -\frac{1}{2} (V_n^{-1})_{00} (x_t^0 - f_{h,t-1})^2; \end{aligned} \quad (11)$$

here $e_0 = (1, 0, \dots, 0)$. Since u_t^R is independent of h and the fractions in the multinomial logit model are independent of the fitness level we conclude that risk-adjusted profits are equivalent, up to a constant factor, to (minus) squared prediction errors. In the case when there are no Arrow securities we have $(V_n^{-1})_{00} = 1/(a\sigma^2)$.

2.3 Adding Arrow securities. We can now address our main question: what happens to the price dynamics when Arrow securities are added?

2.3.1 General mechanism The previous subsection has shown that, when we add an extra Arrow security to the system, the dynamical behavior only changes through the term $(V_n^{-1})_{00}$ in the fitness measure (11). Moreover, using (7) and (8) we get

$$z_{ht} = \zeta^0 + (f_{ht} - Rx_t^0)V_n^{-1}e_0.$$

Let $r = \mathbb{C}ov(\delta_{n+1}, (q_t^0, \delta_1, \dots, \delta_n))$. We call the $n + 1$ -th Arrow security *irrelevant to the portfolio* z_{ht} if $\langle r, V_n^{-1}e_0 \rangle = 0$. It can be shown that the set of irrelevant configurations of y_s^0 and α_s is a closed set of measure zero. The following lemma is key.

LEMMA 1. *If the $(n + 1)$ -th Arrow security is relevant, the matrix V_n in (1) satisfies:*

$$(V_{n+1}^{-1})_{00} > (V_n^{-1})_{00}. \quad (12)$$

Appendix A.1 contains the proof of the lemma. Instead of working with $(V_n^{-1})_{00}$ it may be more intuitive to use

$$\sigma_n^2 = \frac{1}{a(V_n^{-1})_{00}}, \quad 0 \leq n < S - 1. \quad (13)$$

The quantity σ_n^2 may be viewed as a measure of risk when there are n Arrow securities. The inequalities (12) are then equivalent to

$$\sigma_0^2 > \sigma_1^2 > \dots > \sigma_{S-2}^2 > \sigma_{S-1}^2 = 0, \quad (14)$$

implying that the risk measure decreases when more Arrow securities are added to the market, as more risk can be hedged.

We are now ready to formulate our main result. A typical feature in reinforcement or evolutionary learning systems as in Brock and Hommes (1997,1998) is that the fundamental equilibrium may destabilize when the intensity of choice β to switch strategies increases. We claim that adding Arrow securities leads to earlier primary bifurcations:⁶

THEOREM 2.1. *Consider the asset price dynamics with reinforcement learning in (8–9) and fitness measure given by average risk-adjusted profits in (11). If in the market without Arrow securities the fundamental steady state is stable for $\beta < \beta_0^*$ but unstable for $\beta > \beta_0^*$, then for a generic full measure set of dividends y^s and probabilities α_s the primary destabilization values β_n^* for a situation with n Arrow securities and incomplete markets (i.e. $n < S - 1$) satisfy*

$$\beta_{n+1}^* < \beta_n^* < \beta_0^*, \quad 1 \leq n < S - 2. \quad (15)$$

This theorem implies that in the presence of *more* Arrow securities, the fundamental equilibrium destabilizes *earlier*. There is a simple economic intuition behind the theorem. When there are more Arrow securities, agents will use them to hedge out more

⁶The same type of results hold more generally for n assets that are linear combinations of Arrow securities.

risk and take bigger positions in the risky assets. The more Arrow securities there are, the higher will be the rewards for trading strategies that turn out to be on the right side of the market and, under reinforcement learning, successful strategies will attract even more followers. To see this, consider the demand of type h for the risky asset

$$z_{ht}^0 = \zeta^0 + (V_n^{-1})_{00}(f_{ht} - Rx_t^0) = \zeta^0 + \frac{(f_{ht} - Rx_t^0)}{a\sigma_n^2}. \quad (16)$$

A straightforward computation gives the demand of the j -th Arrow security as

$$\begin{aligned} z_{ht}^j &= V_{j0}^{-1}(f_{ht} - Rx_t^0) \\ &= -C \left((1 - \sum_{k=1}^n \alpha_k)(y_j - \bar{y}) + \sum_{k=1}^n \alpha_k(y_k - \bar{y}) \right) (f_{ht} - Rx_t^0), \end{aligned} \quad (17)$$

where $C = \alpha_1 \cdots \alpha_n / \det V_n > 0$ does not depend on j .

When the number of Arrow securities increases, the risk measure σ_n^2 decreases. Hence, it is clear from (16) that the introduction of additional Arrow securities forces optimistic (pessimistic) agents, with the same risk aversion coefficient a , to hold bigger long (short) positions in the stock. For example, optimistic traders who predict next period's asset price deviation f_{ht} from the fundamental price to grow faster than R times the current positive deviation, that is, for whom $f_{ht} - Rx_t^0 > 0$, will take larger positions when there are more Arrow securities. Moreover, from (17) (note the minus sign) we see that these optimistic traders take short positions in Arrows corresponding to above average dividends and long position in Arrows corresponding to below average dividends. Traders thus use the Arrow securities to hedge their risk and invest more in the stock if they expect its price to rise. This is a leverage effect. Moreover, strategies that forecasted the price movement better will attract more followers according to the risk-adjusted fitness measure (11) and inequality (12). Stated differently, using hedging portfolios strategies that turned out to be "right" will be rewarded more and attract more followers, while strategies that turned out to be "wrong" will lose more.

A major implication of theorem 2.1 is that, if all other parameters including the intensity of choice are fixed, adding Arrow securities may destabilize the market. For

when there are n Arrow securities, the fitness given by average risk-adjusted profits in (11), is proportional to $(V_n^{-1})_{00}$ or, equivalently, inversely proportional to the risk measure σ_n^2 in (13). In the presence of n Arrow securities, the *effective intensity of choice* of strategy selection in (9) is thus given by β_0/σ_n^2 , where β_0 is the intensity of choice without Arrow securities. Figure 1(c) plots the effective intensity of choice as the number of Arrow securities increases for a given dividend and probability structure. The effective intensity of choice increases past a critical value as the market approaches completeness, thus creating instability under reinforcement learning.

3 The role of rational agents

In this section, we investigate the consequences of adding a fully rational agent, that is, a trader type that forecasts future prices perfectly, to the ecology of traders. The rational trader type will earn the highest average risk-adjusted profit, but will have to bear information gathering costs to obtain perfect foresight. We will investigate whether or not rational agents are able to drive out boundedly rational types, when more Arrow securities become available.

3.1 Introducing rational agents. Introduce a rational agent type who is a price taker and a mean-variance maximizer, as are all other agents, but who has a perfect forecast. The perfect foresight predictor is not freely available, but can only be obtained at per period information gathering costs $C \geq 0$, as in Brock and Hommes (1997). A rational agent, indexed by $h = 0$, is thus characterized by the fact that his belief about the future price of the asset is correct:

$$f_{0t} = x_{t+1}^0. \quad (18)$$

In line with (7), the demand function of the rational type is given by

$$z_{0t} = \begin{pmatrix} \zeta^0 \\ 0 \end{pmatrix} + V_n^{-1} \begin{pmatrix} -Rx_t^0 + x_{t+1}^0 \\ -R\tilde{x}_t \end{pmatrix}. \quad (19)$$

Adding a rational trader type to the market, the pricing equation (8) becomes

$$Rx_0^t = n_{0t}x_{t+1}^0 + \sum_h n_{ht}f_{ht}. \quad (20)$$

Recall from Subsection 2.2.2 that the fitness of the rational forecasting strategy is equal to

$$u_{0t}^R = u_t^R - C = \frac{1}{2} \langle B_{t-1}, V_n^{-1} B_{t-1} \rangle - C = \frac{1}{2a\sigma_n^2} (x_t^0 - Rx_{t-1}^0)^2 - C, \quad (21)$$

where u_{0t}^R is the risk-adjusted profit of a rational trader type, and where C represents the cost of obtaining the rational forecast. Recall also from (11) that

$$u_{ht} - u_t^R = -\frac{1}{2a\sigma_n^2} (x_t^0 - f_{h,t-1})^2, \quad (22)$$

implying that, ignoring the information gathering costs C , rational agents achieve the highest fitness, that is, rational agents earn the highest average risk-adjusted profits. Notice also that the difference in fitness becomes higher, the larger are the forecasting errors of type h . Since the risk measure σ_n^2 decreases with the number n of Arrow securities and tends to 0 as the market approaches completeness, it follows that the differences in fitness between type h and a rational forecasting rule increases with the number of Arrow securities and blows up as the market approaches completeness.

In order to get some more intuition, it is useful to rewrite the model. Introduce the ‘‘average expectation’’ $\langle x_{t+1}^0 \rangle^e$ of the boundedly rational traders

$$\langle x_{t+1}^0 \rangle^e = \frac{\sum_{h=1}^H n_{ht}f_{ht}}{\sum_{h=1}^H n_{ht}} = \frac{\sum_{h=1}^H n_{ht}f_{ht}}{1 - n_{0t}}.$$

The pricing equation (20) then takes the form

$$Rx_t^0 = n_{0t} (x_{t+1}^0 - \langle x_{t+1}^0 \rangle^e) + \langle x_{t+1}^0 \rangle^e. \quad (23)$$

Note that if $n_{0t} = 0$, we retrieve equation (8). Using equation (23), for $t - 1$, we rewrite the expression for u_t^R in (21) to obtain

$$u_t^R = \frac{(1 - n_{0t})^2}{2a\sigma_n^2} (x_t^0 - \langle x_t^0 \rangle^e)^2. \quad (24)$$

From this expression it becomes clear that the rational trader exploits deviations of the “average expectation of boundedly rational traders” from the realized market price. The larger the average forecasting errors of boundedly rational agents, the higher the benefits for rational traders. Moreover, the benefits for rational traders are higher, the more boundedly rational agents are present in the market (i.e. the smaller n_{0t}).

To analyze the evolution of a market with a fraction of rational traders, we finally rewrite the pricing equation as an evolution equation:

$$n_{0t}x_{t+1}^0 = Rx_t^0 - \sum_{h=1}^H n_{ht}f_{ht}. \quad (25)$$

Note that in the limit $C \rightarrow \infty$, the fraction of rational agents vanishes, $n_{0t} \rightarrow 0$, and we recover the case without rational agents.

3.2 Analyzing the dynamics. It is now convenient to introduce the parameter $\varepsilon = e^{-\beta C}$; then $\varepsilon \rightarrow 0$ as $C \rightarrow \infty$. The evolution equation (25) is a singularly perturbed nonlinear difference equations of the form

$$\varepsilon \Gamma x_{t+1}^0 = \Phi \quad (26)$$

where Φ and Γ do not depend on x_{t+1}^0 . We have

$$\Phi(x_t^0, \dots, x_{t-(L+2)}^0) = Rx_t^0 - \sum_{h=1}^H n_{ht}f_{ht}$$

and

$$\Gamma(x_t^0, \dots, x_{t-(L+2)}^0) = \frac{1}{\varepsilon + \sum_{h=1}^H e^{\beta(u_{ht} - u_t^R)}};$$

Let ξ_0 denote a steady state of the unperturbed equation, that is, let $\Phi(\xi_0, \dots, \xi_0) = 0$. The associated characteristic values of this steady state are the roots of the L 'th order polynomial

$$Q_0(\lambda) = \frac{\partial \Phi}{\partial x_t^0} \lambda^L + \frac{\partial \Phi}{\partial x_{t-1}^0} \lambda^{L-1} + \dots + \frac{\partial \Phi}{\partial x_{t-L}^0},$$

where the partial derivatives are evaluated at (ξ_0, \dots, ξ_0) .

If $\lambda = 1$ is not a root of Q_0 , then the steady state can be continued as a function ξ_ε for small but positive values of ε . The characteristic values associated to ξ_ε are roots of a polynomial Q_ε , which is of the form

$$Q_\varepsilon(\lambda) = \varepsilon \Gamma(\xi_\varepsilon, \dots, \xi_\varepsilon) \lambda^{L+1} + Q_0(\lambda) + \varepsilon Q_1(\lambda; \varepsilon),$$

where Q_1 is also a polynomial of order L .

Denote the characteristic values of ξ_ε by $\lambda_j(\varepsilon)$, $j = 0, \dots, L$. It follows from standard arguments that one characteristic value, say $\lambda_0 = \lambda_0(\varepsilon)$, tends to infinity as $1/\varepsilon$ as $\varepsilon \rightarrow 0$, while the other characteristic values tend to the characteristic values of the system without perfectly foresighted agents. The solutions that diverge from ξ_ε at the rate λ_0^t are local rational bubbles. The following theorem shows that it is possible to exclude these by a judicious choice of the initial prediction x_0^0 of the perfectly foresighted type.

THEOREM 3.1. *Assume that the steady state ξ_0 of system (8) without perfect foresight agents has no characteristic value equal to 1, and let λ_1 be the characteristic value of ξ_0 with the largest absolute value $m_1 = |\lambda_1|$. There is a small neighborhood U of the steady state ξ_ε of system (25), and a cost level $C_0 > 0$, such that for all initial conditions $x_{-\ell}$, $\ell = 1, \dots, L$ in U , for any $\delta > 0$, and for any $C > C_0$, the following holds.*

There is an initial prediction $x_1^0 = \psi^e(x_{-1}^0, \dots, x_{-L}^0)$ of the perfect foresight type, and an implied initial market clearing price $x_0^0 = \psi(x_{-1}^0, \dots, x_{-L}^0)$ such that for all t the dynamics

$$x_t^0 = \psi(x_{t-1}^0, \dots, x_{t-L}^0) \tag{27}$$

is well defined, C^k , and restricted to an invariant hypersurface of (25). For all $x_t^0 \in U$ the dynamics on the invariant manifold satisfies

$$|x_t^0| \leq |x_0^0| (m_1 + \delta)^t.$$

Moreover, as $C \rightarrow \infty$, the linearization of the evolution (27) at $x = 0$ tends to the linearization of the system (8) without perfect foresight agents.

Theorem 3.1 implies that perfect foresight agents can choose their initial forecast x_1^0 in such a way that the dynamics in (25) is restricted to an invariant manifold, so that rapidly exploding rational bubbles are avoided. The theorem thus implies that in a heterogeneous agents setting, a “transversality condition” avoiding bubble solutions can be imposed. If the costs for perfect foresight are high, so that the fraction of perfectly foresighted agents is small, the dynamics on the invariant manifold is similar to the dynamics without perfect foresight in (8). In particular, the local stability of the fundamental steady state is described by a characteristic equation Q_C which tends to the characteristic equation Q_0 in the case without perfect foresight agents, when $C \rightarrow \infty$. This implies that, when the information gathering costs for perfect foresight are high, the first bifurcation to local instability in the presence of a *small* fraction of perfectly foresighted agents is close to the first bifurcation in the case without these agents. But more can be said. The invariant manifold persists after the first bifurcation so that additional steady states and/or cycles created immediately after a first bifurcation also persist. This implies, for example, that if the system without perfect foresight agents exhibits a generic saddle-node, a period doubling or a Hopf bifurcation, the system with a sufficiently small fraction of these traders exhibits the same bifurcation at almost the same critical value.

3.3 Can rational agents employ more hedging instruments to stabilize markets?

In Section 2 we have seen that in a heterogeneous world with only boundedly rational agents, destabilization comes earlier if more Arrow securities are added to the system. We are now ready to explore what happens in the presence of perfectly foresighted agents.

3.3.1 A small fraction of rational traders cannot prevent destabilization. Let $\bar{x}_n(C)$ be a steady state of the system (25) with a (small) fraction of rational agents,

with n Arrow securities and costs C for perfect foresight, such that the limit

$$\bar{x}_n^\infty = \lim_{C \rightarrow \infty} \bar{x}_n(C)$$

exists. Let $\beta = \beta_n^\infty$ be the bifurcation value at which the steady state \bar{x}_n^∞ of the system (8) without rational agents and with n Arrow securities first loses stability. According to our earlier Theorem 2.1, in a world with only boundedly rational traders $\beta_{n+1}^\infty < \beta_n^\infty$, that is, the primary bifurcation to instability occurs earlier when there are more Arrow securities. The following result extends this result to the case with a small fraction of rational agents:

THEOREM 3.2. *Assume that the stability losing bifurcation β_n^∞ in the case without rational agents and n Arrow securities is of co-dimension one (e.g. a generic Hopf, period doubling or saddle-node bifurcation). There is a constant C_0 such that for $C > C_0$, the system (25) in the presence of a rational type and with n Arrow securities first loses stability at $\beta = \beta_n(C)$, and*

$$\beta_{n+1}(C) < \beta_n(C),$$

for $n = 1, \dots, S - 2$.

The heuristic content of this result is that *when perfect foresight is costly, rational traders cannot prevent destabilization*. When the costs for perfect rationality are high, so that the fraction of rational agents remains small, the first bifurcation still comes earlier when there are more Arrow securities.

3.3.2 Rational traders may eventually stabilize markets. Theorem 3.2 implies that, as long as the fraction of rational agents remains small, the market may destabilize when the number of Arrow securities increases. But what happens if the number of Arrow securities keeps increasing and the market approaches completeness? Will perfectly rational agents be able to use a sufficiently large set of Arrow securities to hedge out their risk, outperform the other strategies, grow in number and eventually

stabilize the market? Stated differently, will the benefits of an almost complete market for perfectly rational traders outweigh the costs of “thinking hard”?

The answer to this question depends on the *composition* of the population of heterogeneous, boundedly rational agents. There are two alternatives. The first occurs if all boundedly rational types are *biased* at the fundamental steady state. As the risk measure σ_n^2 decreases towards zero, rational agents will drive out all biased types, stabilizing the market and forcing prices to converge to fundamental value. This is the content of theorem 3.3 below.

We obtain the second alternative, if there are boundedly rational types who at the fundamental steady state predict that the price remains at steady state. Rational traders are unable to drive out these unbiased boundedly rational types and the market may remain unstable, even when approaching market completeness. An example of this phenomenon is given in section 4.

THEOREM 3.3. *Assume that all boundedly rational types are biased in the fundamental steady state $x_t = 0$, that is, assume that*

$$f_{ht}^0(0, \dots, 0) \neq 0, \quad h = 1, \dots, H.$$

Then for every $C > 0$ and every $0 \leq n < S - 1$, there is a $\bar{\sigma}^2 > 0$ such that if the risk measure satisfies $0 < \sigma_n^2 < \bar{\sigma}^2$, then the system (25) with a fraction of rational agents has a locally attracting stable steady state.

The proof of this result is given in A.5. The idea of the proof is simple: the prediction errors of the boundedly rational traders lead to huge positions, and consequently huge losses, if the risk measure σ_n^2 is very small. If the losses are much bigger than the costs C of acquiring a perfect forecast, the rational type dominates the market and (locally) stabilizes the fundamental price.

3.4 Welfare. So far we have focussed on its potential (de)stabilizing effect, but now consider how the introduction of additional hedging instruments affects welfare.

Welfare at time t averaged over the population of trader types is given by

$$W_t = \sum_{h=0}^H n_{h,t-1} u_{ht}. \quad (28)$$

Notice the timing here: welfare is averaged over the fractions $n_{h,t-1}$ of type h whose realized utility is u_{ht} . Using (11) and $u_t^R = \frac{1}{2} \langle B_{t-1}, V^{-1} B_{t-1} \rangle$, with the realized return vector B_{t-1} as in (10), we can rewrite welfare as

$$\begin{aligned} W_t &= \sum_{h=0}^H n_{h,t-1} u_{ht} = \sum_{h=0}^H n_{h,t-1} (u_{ht} - u_t^R) + u_t^R \\ &= -\frac{1}{2a\sigma_n^2} \sum_{h=1}^H n_{h,t-1} (x_t^0 - f_{h,t-1})^2 - n_{0,t-1} C \\ &\quad + \frac{1}{2a\sigma_n^2} (x_t^0 - Rx_{t-1}^0)^2 + (x_t^0 - Rx_{t-1}^0) a\zeta^0 + \frac{1}{2} (a\zeta^0)^2 \sigma^2. \end{aligned} \quad (29)$$

The first term represents (minus) the squared forecasting errors averaged over the population of non-rational agents. Substituting $Rx_{t-1}^0 = \sum_{h=0}^H n_{h,t-1} f_{h,t-1}$ into the third term, and merging it with the first term into a “variance”, average welfare simplifies to

$$W_t = \frac{1}{2} (a\zeta^0)^2 \sigma^2 - n_{0,t-1} C + (x_t^0 - Rx_{t-1}^0) a\zeta^0 - \frac{1}{2a\sigma_n^2} \sigma_{\epsilon_{ht}}^2, \quad (30)$$

where $\sigma_{\epsilon_{ht}}^2$ is the “variance” of the forecasting errors $\epsilon_{ht} = x_t^0 - f_{h,t-1}$ of non-rational types, with “mean” $\mu_{\epsilon_{ht}} = x_t^0 - \sum_{h=0}^H n_{h,t-1} f_{h,t-1}$. The first term reflects the “risk premium” for the population of traders to hold the risky assets, and the second term the costs of rational agents. If rational agents drive out all non-rational types and force prices to their fundamental benchmark (i.e. $x_t \equiv 0$), welfare becomes $\frac{1}{2} (a\zeta^0)^2 \sigma^2 - C$, the risk premium net of the costs of rationality. The third term in (30) reflects a (temporary) “irrationality bias” in population averaged welfare, which e.g. is positive (negative) if a positive deviation from fundamental price grows faster (slower) than a rational bubble (i.e. if $x_t^0 - Rx_{t-1}^0 > 0$ (< 0)). If prices fluctuate around its fundamental value, the time average of this “irrationality bias” will be close to 0. The last term, $-\sigma_{\epsilon_{ht}}^2 / (2a\sigma_n^2)$, captures the effect of the spread of the forecasting errors of non-rational agents. Since the risk measure σ_n^2 decreases with the number of Arrow securities, non-zero forecasting errors of non-rational types will blow up when the number of hedging instruments increases. In fact, when non-rational agents and their

forecasting errors persist, average welfare may blow up to minus infinity as the market approaches completeness (see Figure 1d). On the other hand, if rational agents can use the hedging instruments to drive out non-rational agents and stabilize the market, welfare losses will be limited to the costs of “thinking hard”.

4 Example

This subsection presents a simple example illustrating that adding Arrow securities destabilizes the system and may lead to cycles and even chaos, that average welfare decreases, and that rational agents can not drive out unbiased traders when the market approaches completeness. There are three types of traders with forecasting rules (in deviations from the fundamental benchmark):

$$f_{0t} = x_{t+1}^0, \quad f_{1t} = 1, \quad f_{2t} = x_{t-1} + g(x_{t-1} - x_{t-2}). \quad (31)$$

Rational traders (type 0) have perfect foresight. Type 1 agents use information about economic fundamentals and predict that the price of the risky asset will be equal to its fundamental value, but they make an (small) error (normalized to 1). Type 2 are trend followers who do not use fundamental information, but extrapolate the latest observed price change by an extrapolation factor g . Note that this is an unbiased forecasting rule at the fundamental steady state. Taking $a\sigma^2 = 1$, the fitnesses of the strategies read as

$$u_{0t} = -C, \quad u_{1t} = -(x_{t-1} - 1)^2, \quad u_{2t} = -(x_{t-1} - (1 + g)x_{t-3} + x_{t-4})^2. \quad (32)$$

4.1 Dynamics without rational traders. First consider the case that $C = \infty$, so that the fully rational traders are absent from the market ecology.

The following bifurcation scenario occurs⁷. For $\beta = 0$, the steady state $x_* = 1/(2R - 1) = 1/(1+2r) \approx 1$ (recall that $r = R-1$). This steady state is close to 1, the predicted

⁷See for mathematical treatments of bifurcation theory e.g. Kuznetsov (1995).

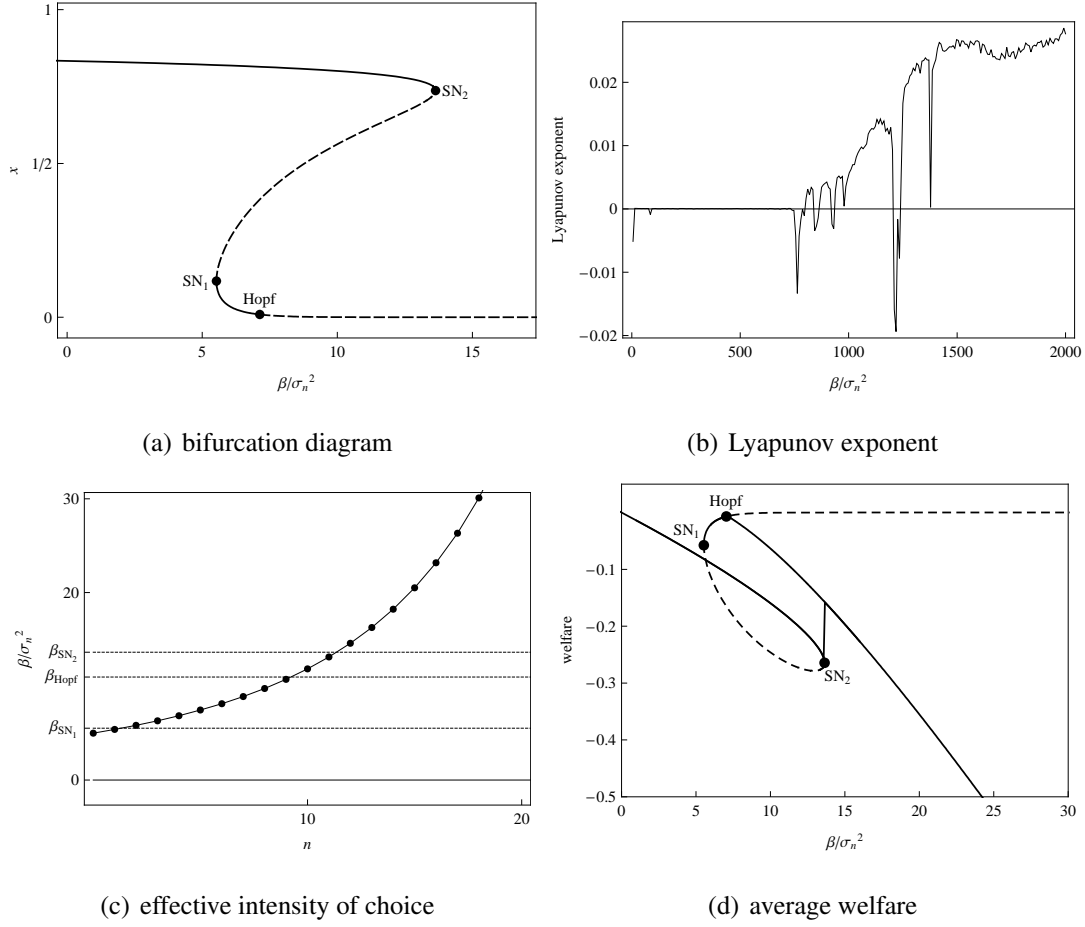


Figure 1: Bifurcation diagrams and average welfare in the 2-type example with biased traders versus trend extrapolators ($R = 1.1$, $g = 1.101$). Panel 1(a): bifurcations of steady states. The curve shown is the locus of the steady state equilibria x_* . Two saddle node (SN) bifurcations and one Hopf bifurcation occur, and $x_* \rightarrow 0$, the true fundamental, as $\beta \rightarrow \infty$. Panel 1(b): largest Lyapunov exponent. As the intensity of choice β increases the system loses stability in a Hopf bifurcation, after which cycles and chaos (with positive largest Lyapunov exponent) arises. Panel 1(c): effective intensity of choice $\beta_n = \beta / (\alpha \sigma_n^2)$ as a function of the number n of Arrow securities. Panel 1(d): welfare averaged over the population and over time, as a function of the effective intensity of choice. Dotted curves correspond to unstable steady states. The lower branch of the solid curve corresponds to the “biased” steady state (the upper branch in panel 1(a)), while the upper branch (between SN_1 and Hopf) corresponds to the stable near fundamental steady state (the lower branch in Panel 1(a)) and, after the Hopf bifurcation, to the quasi-periodic or chaotic attractor.

steady state of type 1. As β increases, the steady state $x_*(\beta)$ moves along the upper part of the curve in Figure 1a, and this steady state is stable. For $\beta = \beta_{SN_1} \approx 5.5$ two additional steady states are created in a saddle-node bifurcation, one stable (the lower one) and one unstable (the middle one)⁸. These two steady states are closer to the fundamental value $x \equiv 0$. As β increases, the steady state closest to the fundamental value loses stability through a Hopf bifurcation at $\beta_{Hopf} \approx 7.0$. At $\beta_{SN_2} \approx 13.6$ a second saddle-node bifurcation occurs, and the two upper steady states disappear. For $\beta_{Hopf} < \beta < \beta_{SN_2}$ a stable steady state co-exists with an attractor around the fundamental steady state. Figure 1c shows a Lyapunov exponent plot, illustrating the dynamical behavior after the Hopf bifurcation. After the Hopf bifurcation quasi-periodic behavior occurs with a Lyapunov exponent close to 0. For large values of β the dynamics becomes chaotic, with positive Lyapunov exponent. Introduction of additional Arrow securities has the same effect as increasing the parameter β . For example, with $S = 40$ states of the world with probabilities $\alpha_s = 1/S$ and dividends $y_s = s - 1$, fixing $\beta = 1$ yields the following dynamics depending upon the number n of Arrow securities (see Figure 1c): (i) unique stable steady state for $n = 0$ and $n = 1$; (ii) co-existence of two stable steady states for $2 \leq n \leq 9$; (iii) co-existence of stable steady state and (quasi-)periodic attractor for $n = 10$ and $n = 11$; (iv) (quasi-)periodic attractor, for $12 \leq n \leq 32$, and (v) chaotic behavior, for $33 \leq n < S = 39$.

Figure 1(d) plots average welfare, averaged over the population and over time, as a function of the effective intensity of choice β/σ_n^2 , in the case *without rational traders*. Welfare *decreases* when the number of Arrow securities *increases*. Only between the first saddle-node bifurcation SN_1 and the Hopf bifurcation, when a stable near fundamental steady state exists, welfare increases, but it decreases again after the Hopf bifurcation (averaged over the quasi-periodic or chaotic attractor). In particular, welfare explodes (to minus infinity) as the market approaches completeness. In this ecology of traders, rational agents can not stabilize welfare, because they can not drive out trend

⁸Hens and Pilgrim (2003) show that new financial securities may change the number of equilibria in a perfect foresight model.

following strategies.

4.2 Rational traders can not always eventually stabilize. We now add rational traders to the system with biased traders and trend extrapolators, that is, we take $0 \leq C < \infty$. Our objective is to show that, in contrast to Theorem 3.3, in the presence of (unbiased) trend extrapolators rational traders cannot stabilize the fundamental equilibrium, even locally. System (26) has a steady state equilibrium $x_* \in \mathbb{R}$ if $x = x_*$ satisfies the equation

$$F(x) = \varepsilon x - \left(\varepsilon + e^{-\eta(x-1)^2} + 1 \right) Rx + e^{-\eta(x-1)^2} + x = 0.$$

Note that this steady state equation is independent of the trend extrapolation factor g . From this equation, the value η can be solved as a function of x , yielding

$$\eta = \eta_*(x) = \frac{1}{(x-1)^2} \log \frac{1-Rx}{(1+\varepsilon)(R-1)x}.$$

Note that $\eta \rightarrow \infty$ as $x \rightarrow 0$; that is, for large values of η , there is a single equilibrium that tends to the fundamental equilibrium (see Figure 1(a)). We are interested in the stability of this near-fundamental steady state when the market approaches completeness, that is, when $\eta = \beta/\sigma_n^2 \rightarrow \infty$. The next result, which is proved by linearization around the near-fundamental steady state, gives conditions when rational agents can stabilize the market when the number of Arrow securities is large:

THEOREM 4.1. *Let x_* be the near-steady state equilibrium of (26) with specifications (31) and (32). Let $\eta = \beta/\sigma_n^2$. If η is sufficiently large and*

1. *if $0 < g < R$, then x_* is locally stable;*
2. *if $g > 2R - 1$, then x_* is locally unstable;*
3. *if $R < g < 2R - 1$, and C is sufficiently small (close to 0), then x_* is locally stable;*

The stability in the limiting case of a complete market depends on the magnitude of the trend extrapolation factor g . When the trend parameter is small ($0 < g < R$) the market will always be stable. For intermediate parameter values ($R < g < 2R - 1$), if the costs for rationality are small and the number of Arrow securities is large, rational agents can still stabilize the market. However, when the trend extrapolation factor is large ($g > 2R - 1 = 1 + 2r \approx 1$), the price dynamics will remain unstable even if the market approaches completeness⁹.

5 Concluding Remarks

In the last decade we have seen an explosive growth of risk hedging instruments in financial markets. There is also empirical evidence that investment decisions are (partly) driven by relative performance. It has been argued recently, e.g., by Rajan (2005), that under such conditions markets may be exposed to more financial-sector turmoil than in the past. We have formalized this idea in a stylized asset pricing model with heterogeneous beliefs. Hedging instruments are represented by Arrow securities, which may be viewed as proxies for more complicated financial instruments. When agents adapt their behavior based upon reinforcement learning, a general mechanism for potential instability applies. Adding Arrow securities to the market may destabilize price dynamics, and thus increase volatility, and at the same time decrease average welfare.

We have also investigated whether the benefits of “thinking hard” can outweigh its costs: can fully rational traders use the extra hedging instruments to drive out non-rational agents, stabilize the market and limit welfare losses? As long as their fraction is small, e.g. due to high information gathering costs, rational traders can *not* prevent destabilization when more Arrow securities are introduced. However, under some con-

⁹Hommes et al. (2005) estimated trend extrapolation factors in learning to forecast experiments with human subjects in the same asset pricing setting (without Arrow securities). Many individuals used trend following forecasting rules, with estimated trend parameters ranging from $0.4 \leq g \leq 1.3$, covering all three cases of Theorem 4.1.

ditions rational agents can stabilize prices. For example, when *all* non-rational trader types are *biased* (i.e. make a small error) at the fundamental steady state, then rational agents eventually stabilize prices and limit welfare losses to the costs of “thinking hard” as the market approaches completeness. As more and more hedging instruments are introduced, rational agents take bigger positions in their hedging portfolio and their benefits outweigh the information gathering costs for rationality. On the other hand, if one of the non-rational strategies uses an unbiased strategy such as a simple trend extrapolating forecast rule, and trend extrapolation is sufficiently strong, rational agents are unable to stabilize prices even when the market approaches completeness and there are no costs for rationality.

Our model is very stylized and much too simple to capture all aspects of financial complexity in real markets. But it is tempting to compare our main results to some stylized features of speculative trading, e.g. due to large hedge funds, in real markets. For example, Prabhu (2001) describes the LTCM “convergence trade” investment strategy to take a leveraged position to profit from an expected narrowing of the spread between the yields of “on the run” and “off the run” bonds as follows:

“For example, in August 1993, before Long-Term entered the market, 30-year bonds yielded 7.24%, while 29 1/2 year bonds yielded 7.36%. This 12 basis point spread would not allow it to earn the type of returns that its investors expected, so the traders at LTCM needed to leverage their trade in order to magnify this return. On this particular trade, such magnification was very easy. LTCM received cash when it shorted the on-the-run bond, and it could then use that cash to buy the off-the-run. This meant that it needed to put up very little cash in order to finance this pair of transactions, and could easily leverage the tiny arbitrage profit into large gains. This type of trade was reportedly often leveraged thirty to forty times in order to generate high returns on equity.”

The hedging portfolio computed in (16) and (17) has similar features. It contains both long and short positions, similar to many actual hedge fund strategies like the

130/30 type positions discussed in Lo and Patel (2007) on the new "Long-Only". When traders think that the stock is underpriced, they approximate it with the Arrow securities available as well as they can, go short in the approximate portfolio, and use the money to buy the stock. As the stock and its approximation have almost the same dividend structure, the traders have almost no dividend risk: the dividends they have to pay on the Arrow security portfolio, they pay out of the dividend revenues from the stock. However, they bet on making a gain out of the price movement. In toto, they "put very little cash in order to finance this pair of transactions".

How general are these results? Clearly, they will not always hold. We have derived them in a simple asset pricing model with heterogeneous beliefs, but the results can be generalized, for example, to a general equilibrium overlapping generations setting¹⁰. One could also think of a more general model, for example by including stabilizing forces such as an increase of the time horizons of agents (e.g Levine and Zame, 2002) or a decrease of the rates at which agents discount the future (e.g. Blume and Easley, 2006). What would happen in a more general model taking these stabilizing forces into account as well as the potentially destabilizing effect of learning? Which force will "win"? The search for an answer in more elaborate models will be an exciting area for future research.

References

- [1] Barberis, N. and Thaler, R. (2003), A survey of behavioral finance, *Handbook of the Economics of Finance*, In: Constantinidis, G.M. et al. (Eds.), Elsevier.
- [2] Benartzi, S. and Thaler, R.H. (2007), Heuristics and biases in retirement savings behavior, *Journal of Economic Perspectives* 21, 81-104.

¹⁰Our earlier working paper version CeNDEF WP 06-12, available at <http://www.ase.uva.nl/cendef/publications>, shows that the asset pricing model with heterogeneous beliefs is in fact a linearization of a general equilibrium two period overlapping generations model. It also shows that similar results hold in a third model setting, a noisy rational expectations model with heterogeneous costly information signals.

- [3] Bhamra, H.S. and Uppal, R. (2006), The effects of introducing a non-redundant derivative on the volatility of stock-market returns, *CEPR Working Paper* 5726.
- [4] Blume, L. and Easley, D. (2006), If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets, *Econometrica* 74, 929-1012.
- [5] Bollerslev, T., Engle, R., Nelson, D. (1994), Arch models, In: Engle, R., McFadden, D., (Eds.), *Handbook of Econometrics*, Volume IV, North Holland.
- [6] Boswijk, H.P., Hommes, C.H. and Manzan, S. (2007), Behavioral heterogeneity in stock prices, *Journal of Economic Dynamics and Control* 31, 1938-1970.
- [7] Branch, W.A. (2004), The theory of rationally heterogeneous expectations: evidence from survey data on inflation expectations, *Economic J.* 114, 592-621.
- [8] Brock, W.A. and Hommes, C.H. (1997), A rational route to randomness, *Econometrica* 65, 1059-1095.
- [9] Brock, W.A., and Hommes, C.H. (1998), Heterogeneous beliefs and routes to chaos in a simple asset pricing model, *J. Econ. Dyn. & Contr.* 22, 1235-1274.
- [10] Cass, D. and Citanna, A. (1998), Pareto improving financial innovation in incomplete markets, *Economic Theory* 11, 467-494.
- [11] Chevalier, J. and Ellison, G., (1997), Risk taking by mutual funds as a response to incentives, *Journal of Political Economy* 105, 1167-1200.
- [12] Citanna, A. and Schmedders, K. (2005), Excess price volatility and financial innovation, *Economic Theory* 26, 559-587.
- [13] Del Guerico, D. and Tkac, P.A., (2002), The determinants of the flow of funds of managed portfolios: mutual funds versus pension funds, *Journal of Financial and Quantitative Analysis* 37, 523-557.
- [14] DeLong, J.B., Shleifer, A., Summers, L.H. and Waldmann, R.J. (1990), Noise trader risk in financial markets, *Journal of Political Economy* 98, 703-738.
- [15] Elul, R. (1995), Welfare effects of financial innovation in incomplete markets economies with several consumption goods, *J. of Economic Theory* 65, 43-78.

- [16] Evans, G.W. and Honkapohja, S. (2001), Learning and expectations in macroeconomics, Princeton University Press, Princeton.
- [17] Farmer, J.D. and Geanakoplos, J. (2008), The virtues and vices of equilibrium and the future of financial economics, Working Paper Santa Fe Institute.
- [18] Gerlach, S., Ramaswamy, S. and Scatigna, M. (2006), 150 years of financial market volatility, *BIS Quarterly Review*, September 2006.
- [19] Grandmont, J.-M. (1998), Expectation formation and stability in large socio-economic systems, *Econometrica* 66, 741-781.
- [20] Hart, O.D. (1975), On the optimality of equilibrium when the market structure is incomplete, *Journal of Economic Theory* 11, 418-443.
- [21] Hens, T. and Schenk-Hoppé, K.R. (Eds.) (2008), *Handbook of Financial Markets: Dynamics and Evolution*, Elsevier, forthcoming.
- [22] Hens, T. and Pilgrim, B. (2003), General equilibrium foundations of finance, Kluwer, Dordrecht.
- [23] Hirsch, M.W., Pugh, C.C. and Shub, M. (1977), Invariant manifolds, *Lecture Notes in Mathematics* 583, Springer, Heidelberg.
- [24] Hommes, C.H. (2006), Heterogeneous agent models in economics and finance, In: Tesfatsion, L. and Judd, K.L., (eds), *Handbook of Computational Economics, Volume 2: Agent-Based Comp. Econ.*, chapter 23, North-Holland, 1109-1186.
- [25] Ippolito, R. A. (1989), Efficiency with costly information: a study of mutual fund performance, 1965-84, *Quarterly Journal of Economics* 105, 1-23.
- [26] Karceski, J. (2002), Returns-chasing behavior, mutual funds, and betas death, *Journal of Financial and Quantitative Analysis* 37, 559-594.
- [27] Kuznetsov, Y. (1995), Elements of applied bifurcation theory, Springer.
- [28] LeBaron, B. (2006), Agent-based Computational Finance, *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*, chapter 24, L. Tesfatsion and K.L. Judd (Eds.), North-Holland, Amsterdam, 1187-1233.

- [29] Lo, A.W. and Patel, P.N. (2007), “130/30: The new long only”, (December 11, 2007). Available at SSRN: <http://ssrn.com/abstract=1074622>.
- [30] Levine, D. and Zame, W. (2002), Does market incompleteness matter?, *Econometrica* 70, 1805-1840.
- [31] Manski, C. and McFadden, D., (eds.) (1981), Structural analysis of discrete data with econometric applications, MIT Press, Cambridge MA.
- [32] Mayhew, S. (2000), The impact of derivatives on cash markets: what have we learned?, Dep. of Banking and Finance, University of Georgia, February 2000.
- [33] Merton, R.C. (1980), On estimating the expected return on the market. An exploratory investigation, *Journal of Financial Economics* 8, 323-361.
- [34] Prabhu, S. (2001), Long-Term Capital Management: the dangers of leverage, *Duke Journal of Economics*, May 2001.
- [35] Rajan, G.R. (2005), Has financial development made the world riskier? *NBER Working Paper Series*, No. 11728, November 2005.
- [36] Rockinger, M. (1996), Determinants of capital flows to mutual funds, Working Paper, HEC School of Management.
- [37] Sargent, T.J. (1993): Bounded rationality in macroeconomics, Clarendon Press.
- [38] Shiller, R.J., 2000, Measuring bubble expectations and investor confidence, *Journal of Psychology and Financial Markets* 1, 49–60.
- [39] Shiller, R. J. (2003), The new financial order: risk in the 21st century, Princeton University Press.
- [40] Sirri, E.R. and Tufano, P., (1998), Costly search and mutual fund flows, *Journal of Finance* 53, 1589-1621.
- [41] Skogestad, S. and Postlethwaite, I. (1996), Multivariable feedback control, John Wiley and Sons, New York.

- [42] Vissing-Jorgensen, A. (2003), Perspective on behavioral finance: does ‘irrationality’ disappear with wealth? Evidence from expectations and actions, In: Gertler, M., Rogoff, K. (Eds.) *NBER Macroeconomics Annual*, MIT Press.

A Proofs

A.1 Proof of Lemma 1. We first state and prove a more general matrix lemma:

LEMMA 2. Let Q_n be a symmetric (n, n) -matrix and Q_{n+1} a symmetric $(n+1, n+1)$ -matrix of the form $\begin{pmatrix} Q_n & r \\ r^T & s \end{pmatrix}$, where r is an n -vector and s a scalar, and let $\tilde{w} = (w, w_0)$, with w an n -vector and w_0 a scalar. Then

$$\langle \tilde{w}, Q_{n+1}^{-1} \tilde{w} \rangle = \langle w, Q_n^{-1} w \rangle + \frac{(w_0 - \langle r, Q_n^{-1} w \rangle)^2}{s - \langle r, Q_n^{-1} r \rangle}.$$

The proof of the first part of this lemma can be established by a variation on the use of the formula for the inverse of a partitioned matrix which uses the notion of Schur complement of a submatrix of a matrix (Skogestad and Postlethwaite (1996, p. 499)). The second part can be established using Schur’s formula for the determinant of a partitioned matrix (Skogestad and Postlethwaite (1996, p. 500)).

To prove Lemma 1, in Lemma 2 we set $Q_n = V_n$,

$$w = e_0 \in \mathbb{R}^n, \quad w_0 = 0, \quad \tilde{w} = \tilde{e}_0 = (e_0, 0) \in \mathbb{R}^{n+1},$$

$$s = 0 \quad \text{and} \quad r = \text{Cov}(\delta_{n+1}, (q_t^0, \delta_1, \dots, \delta_n)).$$

Since the $(n + 1)$ -th Arrow security is assumed to be relevant,

$$(w_0 - \langle r, Q_n^{-1} w \rangle)^2 = \langle r, V_n^{-1} e_0 \rangle^2 > 0.$$

Moreover, since V_n^{-1} is positive definite,

$$s - \langle r, Q_n^{-1} r \rangle = -\langle r, V_n^{-1} r \rangle < 0$$

and

$$\langle \tilde{e}_0, V_{n+1}^{-1} \tilde{e}_0 \rangle = \langle e_0, V_n^{-1} e_0 \rangle - \frac{\langle r, V_n^{-1} e_0 \rangle^2}{\langle r, V_n^{-1} r \rangle} > \langle e_0, V_n^{-1} e_0 \rangle.$$

This proves Lemma 1.

A.2 Proof of Theorem 2.1. The proof follows immediately from inequality (12) or equivalently, the inequalities (14). The fitness given by average risk-adjusted profits (11) is proportional to $(V_n^{-1})_{00}$ or, equivalently, inversely proportional to σ_n^2 . Let β_0^* be the first bifurcation value when there are no Arrow securities. Then the system with n Arrow securities will undergo its first bifurcation if

$$\frac{\beta}{\sigma_n^2} = \frac{\beta_0^*}{\sigma_0^2}, \quad \text{that is, if } \beta = \beta_0^* \frac{\sigma_n^2}{\sigma_0^2} \stackrel{\text{def}}{=} \beta_n^*.$$

From (14) we infer that $\beta_0^* > \beta_1^* > \dots > \beta_{S-2}^*$. Consequently, with more Arrow securities the primary bifurcation comes earlier. ■

A.3 Proof of Theorem 3.1 Choose C_0 so large that for $C > C_0$ and $\varepsilon = e^{-\beta C} < e^{-\beta C_0}$ the “rational” eigenvalue λ_0 satisfies $\lambda_0(\varepsilon) > |\lambda_1(\varepsilon)|^k + \delta$. Then it follows from a straightforward application of the theorem on pseudo-hyperbolic maps of Hirsch, Pugh and Shub (1977) that all orbits of (26) diverging from the origin at a speed at most $(m_1 + \delta)^t$ form a k -times continuously differentiable hypersurface in phase space, tangent to the corresponding eigenspace. The map ψ parameterizes this surface.

A.4 Proof of Theorem 3.2 We are interested in bifurcations of system (25), as the parameter β is varied, for large values of C . The most direct approach is to introduce a new parameter $\varepsilon = e^{-\beta C}$, and to study the resulting equation for small values of ε . However, as ε not only depends on C , but also on the bifurcation parameter β , this leads to certain technical problems.

We therefore take an idea from singularity theory and study an *unfolding* of equation (25), where ε is now a second free bifurcation parameter:

$$\varepsilon x_{t+1}^0 = Rx_t^0 \left(\varepsilon + \sum_{h=1}^H e^{\beta u_{ht-1}} \right) - \sum_{h=1}^H e^{\beta u_{ht-1}} f_{ht}.$$

The original system (25) is the *subfamily* of the new system that is obtained by restricting to the curve $\gamma_C(\beta) = (\beta, e^{-\beta C})$ in parameter space.

The assumption that β_n^∞ is the destabilizing bifurcation value of β , and that the bifurcation is a codimension one bifurcation of (8), implies that there is a curve of bifurcation points $h_n(\varepsilon) = (\tilde{\beta}_n(\varepsilon), \varepsilon)$ that intersects the curve $\varepsilon = 0$ transversally at $(\beta_n^\infty, 0)$. Since $\beta_{n+1}^\infty < \beta_n^\infty$, there are constants $\varepsilon_0 > 0$ and $b_n > 0$ such that

$$0 < b_0 < \dots < b_{n-1} < \beta_{n+1}(\varepsilon) < b_n < \beta_n(\varepsilon) < \dots .$$

for all $0 \leq \varepsilon \leq \varepsilon_0$. Let $C_0 > 0$ such that $0 < e^{-\beta C} < \varepsilon_0$ if $C > C_0$ and $\beta > b_0$. As the first destabilizing bifurcation value $\beta = \beta_{n+1}^C$ of (25) in the presence of n Arrow securities is first coordinate of the intersection point of γ_C with h_n , it follows that $\beta_{n+1}^C < b_n < \beta_n^C$, as claimed. ■

A.5 Proof of Theorem 3.3 Recall that

$$u_{ht-1} = -\frac{1}{\sigma_n^2} (x_{t-1} - f_{ht-2})^2$$

Introducing $\eta = \beta/\sigma_n^2$ and $\varepsilon = e^{-\beta C}$, the evolution equation reads, after division by ε , as

$$x_{t+1}^0 = Rx_t^0 \left(1 + \frac{1}{\varepsilon} \sum_{h=1}^H e^{-\eta(x_{t-1} - f_{ht-2})^2} \right) - \frac{1}{\varepsilon} \sum_{h=1}^H e^{-\eta(x_{t-1} - f_{ht-2})^2} f_{ht}.$$

Expanding this equation around the steady state $x_t = 0$ yields an expression of the form

$$x_{t+1}^0 = Rx_t^0 + \delta \left(Rx_t^0 + a_0 + \sum_{\ell} a_{\ell} x_{t-\ell} \right) + \mathcal{O}(2), \quad (33)$$

where $\delta = \varepsilon^{-1} \sum_{h=1}^H e^{-\eta(f_h^0)^2}$ and where $O(2)$ collects the terms of higher than first order in the x_t .

It follows from the implicit function theorem that there is a unique steady state $x_*(\delta)$, depending smoothly on $\delta \in [0, \delta_0]$, such that $x_*(0) = 0$. The characteristic polynomial of (33) has for $\delta = 0$ one root $\lambda = R$ and $L+2$ roots 0. If necessary by decreasing $\delta_0 > 0$, we have that for all $0 \leq \delta \leq \delta_0$, one characteristic root is outside the unit circle, while the others are within.

As all f_h^0 at $x = 0$ are bounded away from zero, the condition $\delta < \delta_0$ can be satisfied if η is sufficiently large, or, equivalently, σ_n^2 sufficiently small.

As before, by choosing the initial price appropriately, the dynamics are restricted to the center-stable manifold of the steady state, which here is a purely stable manifold. But on the stable manifold, the steady state is locally asymptotically stable. ■