Diverse Beliefs

A.A. Brown *

Statistical Laboratory,
University of Cambridge

L.C.G. Rogers †

Statistical Laboratory,
University of Cambridge

May 14, 2009

First version: July 2008

Abstract

This paper presents a general framework for studying diverse beliefs in dynamic economies. Within this general framework, the characterization of a central-planner general equilbrium turns out to be very easy to derive, and leads to a range of interesting applications. We show how for an economy with log investors holding diverse beliefs, rational overconfidence is to be expected; volume-of-trade effects are effectively modelled; a range of sample moments from macroeconomic growth data can be closely approximated; and the Keynesian 'beauty contest' can be modelled and analysed. We remark that models where agents receive private information can formally be considered as models of diverse beliefs.

1 Introduction.

Dynamic general equilibrium models provide us with perhaps our best hope of understanding how markets and prices evolve, but are often frustratingly difficult to solve. Representative agent models are an exception, but the limitations of the representative agent assumption are only too plain. Stepping up to models with many heterogeneous agents drastically reduces the available range of tractable examples, but is a necessary approach to realism. The simplest form of heterogeneity one could consider is one where agents have different preferences, and perhaps different endowments, but such models

*Wilberforce Road, Cambridge CB3 0WB, UK (phone = +44 1223 337969 , email = A.A.Brown@statslab.cam.ac.uk)

†Wilberforce Road, Cambridge CB3 0WB, UK (phone = +44 1223 766806, email = L.C.G.Rogers@statslab.cam.ac.uk)

are not immediately suited to explaining effects arising from different information, or from different beliefs, since the causes are not being modelled.

In a recent survey, Kurz [27] discusses the literature on models with different information or beliefs, presents a compelling critique of models with private information, and expounds his own theory of how to handle diverse beliefs. Models where agents receive private signals about random quantities of interest have been extensively studied, but are in general hard to work with; see, for example, Lucas [31], Townsend [38], Grossman & Stiglitz [14], Diamond & Verrecchia [11], Singleton [37], Brown & Jennings [6], Grundy & McNichol [15], Wang [41], He & Wang [19], Judd & Bernardo [24], Morris & Shin [32], [33], Hellwig [20], [21], Angeletos & Pavan [2]. Problems such as the Grossman-Stiglitz paradox, and the Milgrom-Stokey no-trade theorem necessitate the introduction of exogenous noise into the models, but nonetheless the treatment of private information is only tractable under very restricted modelling assumptions. There are also problems at a conceptual level, as Kurz points out. Firstly, what is private information? In reality, the majority of agents' information is common, such as macroeconomic indicators or the past performance of the stock, so we have to accept that a very small amount of private information might have a significant impact. Secondly, if private information does exist, what could we say about it? The private nature of the information would make it very difficult for us to verify any model that relied upon it.

For these reasons, we prefer to examine the class of models where all agents have the same information, but interpret that information differently. Although Kurz distinguishes such models from private information models, we can make the simple but important observation[1]: *a private information model can be considered as a model where all agents have common information, but have different beliefs about that information.* Indeed, given a model where different agents receive private signals, we could regard this as a model where all agents receive the same information but interpret it differently: every agent gets to see *all* the private signals, but believes that the signals received by the others are independent of everything else in the economy!

In our treatment, the agents' different beliefs are modelled as different *probability measures* $\mathbb{P}^j$ defined over the same stochastic base $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0})$. Contrast this with the situation of private information, where all agents share the same probability $\mathbb{P}$, but work over different stochastic bases $(\Omega, \mathcal{F}^j, (\mathcal{F}^j_t)_{t\geq 0})$. The diverse beliefs setting is far easier to work with, and, as we shall show, leads to simple but effective analyses. The literature on diverse beliefs is surveyed by Kurz, and includes the papers of Harrison & Kreps [18], Leland [29], Varian [39], [40], Harris & Raviv [17], Detemple & Murthy [10], Frankel & Rose [13], Kandel & Pearson [25], Cabrales & Hoshi [8], Wu & Guo [43], [44], Buraschi & Jiltsov [7], Fan [12], Jouini & Napp [23]. Among these, we would pick out the elegant early contribution of Leland [29], which treats a static problem[2], and the more recent contributions of Buraschi & Jiltsov [7], and of Jouini & Napp [23]. Buraschi

[1] We do not claim that understanding diverse beliefs models as presented here renders irrelevant all studies on private information. The inclusion remarked on is too general to permit useful conclusions.

[2] Agents aim to maximise their expected utility of terminal wealth in a complete-markets model. Agents act as price takers; Leland does not derive equilibrium prices.

& Jiltsov propose a continuous-time model with a single asset whose log-dividend is a Brownian motion with Ornstein-Uhlenbeck drift, together with a signal process whose drift depends linearly on the drift in the asset. Two agents start with different priors for the drift, and filter from the observations. Lengthy calculations with the specific form of the model lead to equilibrium prices for stock and option. In view of this, the derivation of the multi-agent equilibrium which we present in the next section is disquieting; how can it be so easy? The reason is that we are handling the equilibrium problem at a higher level of generality, expressing the solution in terms of the likelihood-ratio martingales of the individual agents. The likelihood-ratio martingales which Buraschi & Jiltsov need are given by quite lengthy expressions, but we do not need to write them down explicitly in order to see the form of the equilibrium solution. Indeed, the form of the state-price density process given by (2.8) is one of the main contributions of this paper. The paper of Jouini & Napp [23] uses a similar[3] formulation and the same first-order conditions for equilibrium that we use, but with a somewhat different objective. Their aim is to interpret the equilibrium arising in terms of a representative agent whose beliefs are to be discovered, and who assigns consumption to the different agents in some possibly different way. Our goal is rather to study the equilibrium as it stands, without trying to assign a particular interpretation to it.

Our model supposes mutiple agents take positions in a single[4] asset which pays a continuous dividend stream, and is in unit net supply. There is a riskless asset, in zero net supply. The agents have different beliefs, represented as different probability measures, which we assume with no loss of generality[5] are absolutely continuous with respect to some reference measure. Though we have diverse beliefs, we stress that we do not take a continuum of stochastically identical agents; agents' diversities do not just get replaced by an average. The form of the agents' beliefs is otherwise unrestricted: the agents could be stubborn bigots who assume they know the true distribution of the processes they observe and never change their views, they could be Bayesians updating their beliefs as time evolves, there could be linkages between the beliefs of the different agents - all such structure is irrelevant at the first pass.

Having derived the central-planner equilibrium in Section 2, we immediately show how this framework gives with no effort the result that all agents are 'rationally over-confident' - they all think that the particular consumption stream that they have chosen is better than those chosen by the others.

Obtaining explicitly-soluble examples with diverse beliefs is no easier than in the situation where all beliefs are the same, and from Section 3 onwards we make the simplifying assumption that all the agents have log utilities. This allows us to identify the state-price density process quite explicitly, and to obtain expressions for the equilibrium price of the asset, and for the riskless rate of return. We are also able to identify explicitly the portfolio of the risky asset over time for each of the agents. In contrast

[3]We endow each agent with an initial wealth, Jouini & Napp gives each an endowment stream.

[4]The restriction to a single asset is for notational convenience only; the entire analysis works also for multi-assets situations.

[5]If agent j has probability measure $\mathbb{P}^j$, we could use the average of the $\mathbb{P}^j$ as a reference measure.

to the common-beliefs situation, the portfolios have non-zero quadratic variation, which we interpret as a proxy for the volume of trade, and we study this in Section 6.

Section 4 addresses the 'beauty contest' metaphor of Keynes [26]. In this Section, we consider whether the individual agents in the model would do better to *publicly profess beliefs they do not believe in.* The point of doing this is that their objective is defined in terms of their true beliefs, yet the equilibrium is characterised by the professed beliefs which guide their trading and consumption decisions. It may be (and it turns out to be) that they can improve their objective by adopting beliefs which are in some precise sense *a dynamic mixture of their true beliefs and a population-average of beliefs.* We contrast this with the recent study of Allen, Morris & Shin [1], where the asset prices are defined in terms of average expectation operators which do not compose in a time-consistent fashion. One consequence of this is that the prices are not derived from a state-price density, whereas in our situation they are. We believe that the time-inconsistency of their average-expectation operators depends strictly on the overlapping-generations structure assumed in their model, where each individual lives for just two periods. In such a story, an agent cannot *directly* compare consumption now and consumption five periods in the future, because five periods in the future he will not be consuming. The comparison can only be via the intermediate pricing achieved in markets at the intervening times. Indeed, in an overlapping-generations model with diverse beliefs but with agents who live for a random length of time which may be arbitrarily large, Brown & Rogers [5] find a state-price density which determines prices in the usual way.

In the next section, Section 5, we study the discrete-time analogue of the continous-time situation of Section 3, but with a difference. Starting from the observation that it is typically much easier to gather information on the stock price of a firm than on its dividend process, we imagine now that some agents think that *the stock price is a multiple of the dividend* (as it would be in a homogeneous market.) Otherwise, they believe that the changes in the log dividend are independent identially-distributed normal variables, whose variance they know, but whose mean has a normal prior, which they attempt to learn. Their beliefs are updated by the changes in *price*; but their beliefs enter into the calculation of the price also, so there is a natural feedback mechanism from beliefs into prices. It is possible to carry the analysis quite a long way, but the story is ultimately too complicated to study in general except by simulation. We present some simulation results which show how the mistaken belief that the stock is a multiple of the dividend can produce some very substantial and abrupt changes in price - bubbles and crashes. In general terms, having more diligent[6] agents in the economy reduces the frequency and severity of these big changes.

We place in an appendix a very simple-minded model-fitting exercise; this is not because the study is not of intrinsic interest, but rather because it differs in style from the mainly theoretical body of the paper. We take the diverse-beliefs model with log agents and try to fit it to various sample moments of the dataset of Shiller[7], as Kurz, Jin

[6]We shall refer to an agent as diligent if he actually uses the changes in log dividend - not the changes in log price - to update his beliefs.

[7]This dataset can be downloaded from `http://www.econ.yale.edu/~shiller/data.htm`

& Motolese [28] do. We find good agreement using a model with just three agents, and having reasonable parameter values. This supports the view that diverse beliefs may be able to resolve the equity premium puzzle, but the ability to match a few moments is not of course sufficient to justify a statistical model. Weizmann [42] and Jobert, Platania & Rogers [22] both analyze the equity premium puzzle from the point of view of a representative Bayesian agent, and find reasonable values for parameter estimates, but do not present evidence that the fitted models do any better than just fitting constants to the data.

Section 7 concludes and maps out directions for future research.

2 Diverse beliefs equilibria.

We are going to derive a general equilibrium for a dynamic economy with $J \geq 2$ agents, containing a single productive asset, whose output process $(\delta_t)_{t\geq 0}$ is observable to all agents. We shall suppose that time is continuous, and that δ is adapted to a filtration $(\mathcal{F}_t)_{t\geq 0}$ which is known to all agents. To cover various technical issues, we shall assume that the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}^0)$ satisfies the usual conditions; see [36] for definitions and further discussion.

Though the J agents all have the same information, they do not share the same beliefs about the distributions of the processes they observe. We suppose that agent j thinks that the true probability is $\mathbb{P}^j$, a measure locally equivalent to $\mathbb{P}^0$, with density process Λ^j

$$\Lambda_t^j = \frac{d\mathbb{P}^j}{d\mathbb{P}^0}\bigg|_{\mathcal{F}_t}, \tag{2.1}$$

which is a positive martingale.

The objective of agent j is to obtain

$$\sup E^j \int_0^\infty U_j(t, c_t^j)dt \tag{2.2}$$

where the supremum is over all consumption policies which keep the wealth of agent j positive. Here, U_j is some strictly increasing time-dependent utility, such that $U_j(t, \cdot)$ satisfies the Inada conditions. Notice that even if all agents have the same U_j, their objective is calculated taking expectations under their different $\mathbb{P}^j$, and so differences in beliefs will result in different optimal behaviour.

We use this to find the state-price density process ζ^j of agent j. Agent j's objective can be written in the equivalent forms

$$E^j \int_0^\infty U_j(t, c_t^j)dt = E^0 \int_0^\infty \Lambda_t^j U_j(t, c_t^j)dt. \tag{2.3}$$

Now consider the price that agent j is willing to pay at time s for a contingent claim which pays amount Y_t at time $t > s$. Denote this price by $\pi_s^j(Y_t)$ [8]. By considering

[8] Here, Y_t is some bounded $\mathcal{F}_t$-measurable random variable

the change in agent j's objective from buying this (marginal) contingent claim, the first order conditions give:

$$0 = \pi_s^j(Y_t) U_j'(s, c_s^j)\Lambda_s^j - E^0 \left[Y_t U_j'(t, c_t^j)\Lambda_t^j | \mathcal{F}_s \right]. \tag{2.4}$$

Rearrangement gives

$$\pi_s^j(Y_t) = E^0 \left[Y_t \frac{U_j'(t, c_t^j)\Lambda_t^j}{U_j'(s, c_s^j)\Lambda_s^j} \,\middle|\, \mathcal{F}_s \right] \tag{2.5}$$

So we see that agent j has state price density given by:

$$\zeta_t^j = U_j'(t, c_t^j)\Lambda_t^j \tag{2.6}$$

If we assume that the market is complete (or that we have a central planner equilibrium), then the agents must agree on the price of all contingent claims. So looking at the expression for $\pi_s^j(Y_t)$ and recalling that Y_t is arbitrary, we must have

$$\zeta_{t,s}^j = \frac{U_j'(t, c_t^j)\Lambda_t^j}{U_j'(s, c_s^j)\Lambda_s^j} \tag{2.7}$$

is the same for all j. Hence

$$\boxed{\zeta_t \nu_j = U_j'(t, c_t^j)\Lambda_t^j} \tag{2.8}$$

where ν_j is some $\mathcal{F}_s$ random variable. In particular, if there exists some value t_0 such that $\mathcal{F}_{t_0}$ is trivial[9] then we deduce that ν_j is in fact just a constant.

Now that we have (2.8), deriving equilibrium prices follows from market clearing in the usual way. In more detail, defining[10] the inverse marginal utilities I_j by

$$I_j(t, U_j'(t, y)) = y \tag{2.9}$$

for any $y > 0$, then

$$I_j(t, \zeta_t \nu_j / \Lambda_t^j) = c_t^j.$$

Summing on j and using market clearing gives

$$\sum_j I_j(t, \zeta_t \nu_j / \Lambda_t^j) = \delta_t. \tag{2.10}$$

This is an implicit equation for the unknown ζ in terms of the known quantities δ and Λ^j. In most examples, it will not be possible to invert this relationship explicitly, though under special assumptions, such as the assumption of log utilities made in the next section, we can do something.

[9] This will be the case in the example looked at in this paper

[10] The assumed properties of U_j ensure that I_j is well defined.

Thus starting from the dividend process, and the agents' beliefs, we have obtained an expression for the state-price density, which allows us to price contingent claims; for example, the time-t price of the stock is simply

$$S_t = E^0 \left[\int_t^\infty \frac{\zeta_u \delta_u}{\zeta_t} du \,\middle|\, \mathcal{F}_t \right]. \tag{2.11}$$

REMARKS. (i) In the case where all agents have the same beliefs (thus $\Lambda^J \equiv 1$ for all j), this reduces to the familiar expression for the state-price density as the marginal utility of optimal consumption.

(ii) Notice that the situation is completely general; there is no assumption about the nature of the stochastic processes, nor is there any assumption about the nature of the diverse beliefs. No such assumption is needed for (2.8).

(iii) **Rational overconfidence.** Kurz remarks that "a majority of people often expect to outperform the empirical frequency measured by the mean or median". In other words, each of the agents believes that they will usually do better than the average. In our setup, this result comes for free. If $\tilde{c}_t$ is any consumption stream and c_t^j is agent j's optimal consumption stream, then we have

$$E^j \int_0^\infty U_j(t, c_t^j) dt \geq E^j \int_0^\infty U_j(t, \tilde{c}_t) dt \tag{2.12}$$

This follows simply from the fact that c_t^j is agent j's optimal consumption stream. In general, different agents will choose a different consumption stream, *even if they have the same utility functions*; even if they do have the same utilities, each agent believes that he will do better (on average) than all the other agents.

3 Log agents.

Getting a reasonably explicit form for the state-price density process ζ is key to making progress, and for the rest of the paper unless explicitly stated to the contrary *we shall make the simplifying assumption*

$$\boxed{U_j(t, x) = e^{-\rho_j t} \log x} \tag{3.1}$$

for some positive ρ_j. This leads to an explicit form for the state-price density, and from that, expressions for the wealth processes of the individual agents, the equilibrium price of the stock, and the equilibrium dynamics of the riskless rate when we assume specific dynamics for the dividend process.

Under the assumed form (3.1) for the utility, the relation (2.8) for the state-price density simplifies to

$$\frac{e^{-\rho_j t} \Lambda_t^j}{c_t^j} = \nu_j \zeta_t. \tag{3.2}$$

The wealth process of agent j is thus

$$\begin{aligned}
w_t^j &= E^0\left[\int_t^\infty \frac{\zeta_u c_u^j}{\zeta_t}\,du \,\middle|\, \mathcal{F}_t\right] \\
&= E^0\left[\int_t^\infty \frac{e^{-\rho_j u}\Lambda_u^j/\nu_j}{\zeta_t}\,du \,\middle|\, \mathcal{F}_t\right] \\
&= \zeta_t^{-1} e^{-\rho_j t}\Lambda_t^j/\nu_j\rho_j \qquad (3.3)\\
&= c_t^j/\rho_j \qquad (3.4)
\end{aligned}$$

The derivation exploits the fact that Λ^j is a $\mathbb{P}^0$-martingale.

Using (3.2), market clearing gives

$$\delta_t = \sum_j c_t^j = \zeta_t^{-1}\sum_j \frac{e^{-\rho_j t}\Lambda_t^j}{\nu_j},$$

and hence by rearrangement

$$\zeta_t = \delta_t^{-1}\sum_j \frac{e^{-\rho_j t}\Lambda_t^j}{\nu_j}. \qquad (3.5)$$

Since the stock is in unit net supply, and the bank account in zero net supply, we can quickly identify the stock price, using (3.3):

$$S_t = \sum_j w_t^j = \zeta_t^{-1}\sum_j \frac{e^{-\rho_j t}\Lambda_t^j}{\rho_j\nu_j}. \qquad (3.6)$$

Substituting from ζ from (3.5) leads to

$$\boxed{S_t = \delta_t\,\frac{\sum_j e^{-\rho_j t}\Lambda_t^j/\rho_j\nu_j}{\sum_j e^{-\rho_j t}\Lambda_t^j/\nu_j}}$$

Notice that in this case the price-dividend ratio takes a particularly simple form:

$$\frac{S_t}{\delta_t} = \frac{\sum_j e^{-\rho_j t}\Lambda_t^j/\rho_j\nu_j}{\sum_j e^{-\rho_j t}\Lambda_t^j/\nu_j}, \qquad (3.7)$$

which we shall have need of later when it comes to fitting various moments to the Shiller dataset in Section A. If all the agents have the same beliefs, this is just a deterministic function of time, but with heterogeneous beliefs this becomes a random process. Notice also that the price-dividend ratio depends only on the likelihood-ratio martingales, and not on the underlying dividend process, though this property is special to the log case.

This is about as far as we can get without some more specific assumptions on the nature of the dividend process. We shall now assume that the dividend process satisfies

$$d\delta_t = \delta_t \sigma_t (dX_t + \alpha_t^* \, dt) \tag{3.8}$$

where X is an $(\mathcal{F}_t)$-Brownian motion under $\mathbb{P}^0$, and σ is some strictly positive bounded previsible process. The process α^* is of course unknown to the agents. Concerning the agents' beliefs, we suppose that

$$d\Lambda_t^j = \Lambda_t^j \alpha_t^j dX_t \tag{3.9}$$

where the α^j are previsible processes. Thus under the measure $\mathbb{P}^j$ the process X becomes a Brownian motion with drift α^j (by the Cameron-Martin-Girsanov Theorem; see [36], IV.38 for an account).

The equation (3.5) for the state-price density gives

$$\zeta_t \delta_t = \sum_j \frac{e^{-\rho_j t} \Lambda_t^j}{\nu_j} \equiv L_t, \tag{3.10}$$

say. A little Itô calculus gives us

$$dL_t = L_t(\bar{\alpha}_t \, dX_t - \bar{\rho}_t \, dt) \tag{3.11}$$

where

$$\bar{\alpha}_t \equiv \sum_j q_t^j \alpha_t^j, \quad \bar{\rho}_t \equiv \sum_j q_t^j \rho_j, \tag{3.12}$$

and where

$$q_t^j \equiv \frac{e^{-\rho_j t} \Lambda_t^j / \nu_j}{\sum_i e^{-\rho_i t} \Lambda_t^i / \nu_i}. \tag{3.13}$$

The dynamics of the riskless rate follow easily from (3.10), (3.11). We have

$$d\zeta_t = \zeta_t(-r_t dt - \kappa_t dX_t)$$

where

$$r_t = \bar{\rho}_t + \sigma_t(\alpha_t^* + \bar{\alpha}_t) - \sigma_t^2, \tag{3.14}$$

$$\kappa_t = \sigma_t - \bar{\alpha}_t. \tag{3.15}$$

We can also derive the dynamics of the stock price. After some routine calculations, we arrive at

$$dS_t = S_t \big\{ (\kappa_t + a_t)(dX_t + \kappa_t dt) + r dt) \big\} - \delta_t dt, \tag{3.16}$$

where

$$a_t \equiv \frac{\sum \alpha_t^j e^{-\rho_j t} \Lambda_t^j / \nu_j \rho_j}{\sum e^{-\rho_j t} \Lambda_t^j / \nu_j \rho_j}$$

is an average of the α_t^j using weights different from the q_t^j. This allows us to identify the volatility σ^S of the equilibrium stock price, namely

$$\sigma_t^S = \kappa_t + a_t = \sigma_t - \bar{\alpha}_t + a_t. \tag{3.17}$$

In general, this is different from the volatility σ_t of the dividend process, even if that volatility is constant. If all agents agreed, it is immediate from (3.7) that the volatility of the stock is the same as the volatility of the dividend process; this illustrates again the general principle that heterogeneous beliefs will generate fluctuations which would be absent in a model where all agents agree. Observe also that if $\rho_j = \rho$ is the same for all j, then $a_t = \bar{\alpha}_t$, and hence (using (3.17)) $\sigma_t^S = \sigma_t$. This checks out with what we would get from (3.7), which implies that $\delta_t = \rho S_t$ when all the impatience parameters are the same.

4 Diverse beliefs and beauty contests.

The model assumptions are the same as in Section 3; we will additionally assume that $\rho_j = \rho$ for all j, so that the only differences between agents are their beliefs (represented by the likelihood-ratio martingales Λ^j) and their initial wealths. For simplicity, we will assume that agents are aware of the beliefs of others insofar as these are expressed through a 'market average' belief, to be defined later. This is consistent with the 'beauty contest' metaphor of Keynes [26], where people adjust their views in the direction of what is perceived to be the general view.

Would it benefit agents to publicly pretend that their beliefs are different, given by likelihood-ratio martingales $\tilde{\Lambda}^j$ instead of Λ^j? By doing this, the agents affect the equilibrium state-price density, which is now determined by

$$\delta_t \tilde{\zeta}_t = e^{-\rho t} \sum_j \nu_j^{-1} \tilde{\Lambda}_t^j \equiv e^{-\rho t} \bar{\Lambda}_t, \tag{4.1}$$

and the consumption they achieve in equilibrium, which is now given as

$$c_t^j = e^{-\rho t} \tilde{\Lambda}_t^j / \nu_j \tilde{\zeta}_t. \tag{4.2}$$

From the calculation

$$w_0^j = \tilde{\zeta}_0^{-1} E \int_0^\infty \tilde{\zeta}_s c_s^j \, ds = (\tilde{\zeta}_0 \rho \nu_j)^{-1}$$

we see that $\nu_j^{-1} \propto w_0^j$, and so the mixing weights which produce $\bar{\Lambda}$ from the $\tilde{\Lambda}^j$ are always the same, whatever $\tilde{\Lambda}^j$ the agents select. Notice that

$$\begin{aligned} d\bar{\Lambda}_t &= \bar{\Lambda}_t \frac{\sum \nu_j^{-1} \tilde{\Lambda}_t^j \tilde{\alpha}_t^j dX_t}{\bar{\Lambda}_t} \\ &\equiv \bar{\Lambda}_t \left(\sum p_t^j \tilde{\alpha}_t^j \right) dX_t \\ &\equiv \bar{\Lambda}_t \bar{\alpha}_t \, dX_t \end{aligned} \tag{4.3}$$

where $p^j_t \equiv \nu_j^{-1}\tilde\Lambda^j_t/\bar\Lambda_t$, a probability distribution, and $\bar\alpha$ is a suitably-chosen time-varying convex combination[11] of the individual $\tilde\alpha^j$. We can think of $\bar\Lambda$ as defining a new probability $\bar P$ defined by the usual likelihood-ratio recipe

$$\left.\frac{d\bar P}{dP}\right|_{\mathcal{F}_t} = \bar\Lambda_t, \tag{4.4}$$

and by the Cameron-Martin-Girsanov Theorem, we have that

$$\bar X_t \equiv X_t - \int_0^t \bar\alpha_s \, ds \tag{4.5}$$

is a $\bar P$-Brownian motion.

Agent j's objective is to maximise

$$\begin{aligned}
E^0\int_0^\infty e^{-\rho t}\Lambda^j_t \log(c^j_t)\,dt &= E^0\int_0^\infty e^{-\rho t}\Lambda^j_t \log(e^{-\rho t}\tilde\Lambda^j_t/\nu_j\tilde\zeta_t)\,dt \\
&= E^0\int_0^\infty e^{-\rho t}\Lambda^j_t \log(\tilde\Lambda^j_t/\bar\Lambda_t)\,dt + \kappa \\
&= \bar E\int_0^\infty e^{-\rho t}\,\frac{\Lambda^j_t}{\bar\Lambda_t}\,\log(\tilde\Lambda^j_t/\bar\Lambda_t)\,dt + \kappa \\
&\equiv \Phi + \kappa,
\end{aligned}$$

say, where κ is an unimportant constant. To develop this expression further, we note that

$$\frac{\Lambda^j_t}{\bar\Lambda_t} = \mathcal{E}((\alpha^j - \bar\alpha)\cdot\bar X)_t, \tag{4.6}$$

where $\mathcal{E}$ denotes the Doleans exponential. Thus agent j tries to maximize

$$\begin{aligned}
\Phi &= \bar E\int_0^\infty e^{-\rho t}\,\mathcal{E}((\alpha^j-\bar\alpha)\cdot\bar X)_t\ \log(\mathcal{E}((\tilde\alpha^j-\bar\alpha)\cdot\bar X)_t)\,dt \\
&= \bar E\int_0^\infty e^{-\rho t}\,\mathcal{E}((\alpha^j-\bar\alpha)\cdot\bar X)_t \left\{ \int_0^t (\tilde\alpha^j_s - \bar\alpha_s)\,d\bar X_s - \tfrac12\int_0^t (\tilde\alpha^j_s-\bar\alpha_s)^2\,ds \right\} dt \\
&= \bar E\int_0^\infty e^{-\rho t}\int_0^t \mathcal{E}((\alpha^j-\bar\alpha)\cdot\bar X)_s \left\{(\alpha^j_s-\bar\alpha_s)(\tilde\alpha^j_s-\bar\alpha_s) - \tfrac12(\tilde\alpha^j_s-\bar\alpha_s)^2\right\} ds\,dt \\
&= \bar E\int_0^\infty \rho^{-1}e^{-\rho s}\,\mathcal{E}((\alpha^j-\bar\alpha)\cdot\bar X)_s\,(\tilde\alpha^j_s-\bar\alpha_s)(\alpha^j_s - \tfrac12(\bar\alpha_s+\tilde\alpha^j_s))\,ds \\
&= E^j\int_0^\infty \rho^{-1}e^{-\rho s}\,(\tilde\alpha^j_s-\bar\alpha_s)(\alpha^j_s - \tfrac12(\bar\alpha_s+\tilde\alpha^j_s))\,ds.
\end{aligned}$$

This expression is to be maximised over the professed beliefs $\tilde\alpha^j$, which it must be remembered are also present in $\bar\alpha$. To make this dependence explicit, let us temporarily drop the time subscript, and write

$$\bar\alpha = p^j\tilde\alpha^j + b_j \equiv p^j\tilde\alpha^j + \sum_{i\neq j} p^i\tilde\alpha^i.$$

[11]... different, of course, from the $\bar\alpha_t$ defined at (3.12) ..

The task is therefore to maximise the quadratic

$$(\tilde{\alpha}^j(1-p^j)-b_j)(\alpha^j-\tfrac{1}{2}b_j-\tfrac{1}{2}(1+p^j)\tilde{\alpha}^j)$$

over $\tilde{\alpha}^j$. Routine calculus yields

$$\tilde{\alpha}^j = \frac{p^j\, b_j}{1-(p^j)^2} + \frac{\alpha^j}{1+p^j}$$

Cross-multiply this equation to learn

$$\frac{1-(p^j)^2}{p_j}\,\tilde{\alpha}^j = b_j + \frac{\alpha^j(1-p^j)}{p^j},$$

and now add $p^j\tilde{\alpha}^j$ to both sides to discover

$$\left(\frac{1-(p^j)^2}{p_j}+p^j\right)\tilde{\alpha}^j = \bar{\alpha} + \frac{\alpha^j(1-p^j)}{p^j}.$$

From this now we have

$$p^j\tilde{\alpha}^j = (p^j)^2\,\bar{\alpha} + \alpha^j p^j(1-p^j), \tag{4.7}$$

and summing on j gives us

$$(1-\sum(p^j)^2)\bar{\alpha} = \sum \alpha^j p^j(1-p^j).$$

This gives us an expression for $\bar{\alpha}$ in terms of the known quantities p^j and α^j, namely

$$\bar{\alpha} = \frac{\sum \alpha^j p^j(1-p^j)}{\sum p^j(1-p^j)}. \tag{4.8}$$

Combining with (4.7) allows us to express agent j's optimal professed belief as

$$\tilde{\alpha}^j_t = p^j_t\,\bar{\alpha}_t + (1-p^j_t)\alpha^j_t, \tag{4.9}$$

which bears this appealing interpretation: *agent j will do best to profess beliefs which are a convex combination of his own true beliefs and the 'average' belief of the population.* Notice also that the weights p^j used in (4.9) are *the same as the weights used to form the population average $\bar{\alpha}$ from the individual α^j*. Thus if agent j's beliefs have little weight in the population average, such as would happen if ν_j^{-1} were small - equivalently, w_0^j is small - then agent j gives less weight to the population average, trusting his own beliefs more. Again, agent j's weight is raised when Λ^j_t is relatively large - that is to say, if events have proved agent j correct; then the population average is drawn towards agent j's beliefs, an entirely plausible effect.

Notice also that the weighted average of the professed beliefs of the various agents is the weighted average of their actual beliefs, $\bar{\alpha}$: from (4.9) and (4.8),

$$\sum p^j_t\tilde{\alpha}^j_t = \sum (p^j_t)^2\bar{\alpha}_t + \sum p^j_t(1-p^j_t)\alpha^j_t = \bar{\alpha}_t.$$

Of course, this discussion is somewhat idealized, in that the agents would need to know what the market average belief $\bar{\alpha}$ was in order to know what beliefs to profess. Nevertheless, the point is clear: by professing beliefs more consistent with the overall beliefs of others, agents can improve their objective, and even if the agents do not know *exactly* what the average belief $\bar{\alpha}$ may be, by moving towards that, they will do better. Thus even if an agent does not know for sure where he should shift his professed beliefs to, he will gain even by moving partially in the right direction. Thus we have substantiated Keynes' verbal metaphor in a well-specified neoclassical financial model; the only unconventional ingredient is diversity of beliefs, modelled through different probability measures.

5 Diverse mistaken beliefs

We have seen in Section 4 how it may be advantageous to agents to pretend to believe something which they do not actually believe; the agents know what is going on, but they consciously act differently. In this Section, we shall study what is in some sense the opposite situation, where the agents do not completely understand the market around them, but nevertheless act in accordance with the analysis of Section 3. Once again, we restrict the discussion to agents with log utilities, and for reasons which will become apparent we shall work in discrete time.

In practice, it may be very hard to learn about the dividends of an asset; dividend payments are infrequent, and are often smoothed in various ways which limit their usefulness as indicators of the state of a firm. On the other hand, the stock price is usually easy to get hold of; it is available daily or more frequently; and it provides what is arguably a more sensitive indicator of the state of the firm. In a market of log agents with common beliefs and common impatience parameter ρ, the stock price is simply a multiple of the dividend process, $\delta_t = \rho S_t$; see (3.7). So we shall consider a situation where *agents observe the stock price, and assume that it is a constant multiple of the dividend process.* This introduces a natural and simple feedback mechanism from prices to beliefs. The agents assume that the log returns of the observed stock prices are actually the changes in $\log \delta$, and they modify their beliefs in the light of this knowledge - but those modified beliefs then feed back into the stock prices.

To carry this analysis further, we record the following result, whose proof is a straightforward exercise.

Proposition 1. *Suppose that $X_1, X_2, \ldots$ are independent $N(\mu, \tau^{-1})$ random variables, where τ is known, but μ is not known. Starting with a $N(\hat{\mu}_0, (K_0\tau)^{-1})$ prior for μ, the posterior mean $\hat{\mu}_t$ for μ, and the posterior precision τ_t given $\mathcal{Y}_t \equiv \sigma(X_1, \ldots, X_t)$, satisfy*

$$\tau_t = K_t\tau \equiv (t + K_0)\tau, \tag{5.1}$$

$$K_t\hat{\mu}_t = K_0\hat{\mu}_0 + \sum_{i=1}^{t} X_i. \tag{5.2}$$

The joint density of $(X_1, \dots, X_t)$ *is*

$$\lambda_t \equiv \exp\left\{-\frac{\tau}{2}\sum_1^t X_i^2 + \frac{\tau}{2}(K_t\hat{\mu}_t^2 - K_0\hat{\mu}_0^2)\right\}\left(\frac{\tau}{2\pi}\right)^{t/2}\sqrt{\frac{K_0}{K_t}}. \tag{5.3}$$

Remarks. (i) Notice that the joint density of $(X_1, \dots, X_t)$ under the assumption that these are independent gaussians with zero mean and variance τ^{-1} will be

$$\lambda_t^0 \equiv \exp\left\{-\frac{\tau}{2}\sum_1^t X_i^2\right\}\left(\frac{\tau}{2\pi}\right)^{t/2}.$$

Thus if we take this as the reference measure, the likelihood-ratio martingale takes the simple form

$$\Lambda_t = \lambda_t/\lambda_t^0 = \exp\left\{\frac{\tau}{2}(K_t\hat{\mu}_t^2 - K_0\hat{\mu}_0^2)\right\}\sqrt{\frac{K_0}{K_t}}. \tag{5.4}$$

(ii) How does λ_t change to λ_{t+1} when the new observation X_{t+1} is seen? If we write

$$X_{t+1} = \hat{\mu}_t + \varepsilon, \tag{5.5}$$

then some simple calculations from (5.2) give us the updating

$$\hat{\mu}_{t+1} = \hat{\mu}_t + \frac{\varepsilon}{K_{t+1}}. \tag{5.6}$$

Using this and (5.3) we are able to derive the updating

$$2\log(\lambda_{t+1}/\lambda_t) = -\tau\varepsilon^2\frac{K_t}{K_{t+1}} + \log\left(\frac{K_t}{K_{t+1}}\right) + \log\tau \tag{5.7}$$

for λ.

Working in discrete time, the arguments of Sections 2 and 3 go through with minor change, giving us

$$\zeta_t\delta_t = \sum_j e^{-\rho_j t}\Lambda_t^j/\nu_j \tag{5.8}$$

exactly as before (3.5), and the analogue

$$\begin{aligned} \zeta_t S_t &= \sum_j \frac{e^{-\rho_j t}\Lambda_t^j}{\nu_j(e^{\rho_j}-1)} && (5.9)\\ &\equiv \sum_j \frac{e^{-\rho_j t}\Lambda_t^j}{\tilde{\nu}_j} && (5.10) \end{aligned}$$

of (3.6) for the ex-dividend stock price S_t at time t.

As we remarked earlier, the agents are supposed to see the stock price and assume that it is a multiple of the dividend process. The discrete-time analogue of the dynamics (3.8) assumed previously for δ is to suppose that the random variables $X_t \equiv \log(\delta_t/\delta_{t-1})$ are independent $N(\mu, \tau^{-1})$. *Thus the agents will assume that the random variables* $\log(S_t/S_{t-1})$ *are independent gaussians with common (unknown) mean and (known) precision*[12]. If we have determined the λ_n^j and S_n for $n \leq t$, we use the price/dividend ratio from (5.8) and (5.9) to determine the value of $\xi \equiv \log(S_{t+1}/S_t)$:

$$\begin{aligned} \frac{S_{t+1}}{\delta_{t+1}} &= \frac{S_t e^{\xi}}{\delta_{t+1}} \\ &= \frac{\sum_j e^{-\rho_j(t+1)} \lambda_{t+1}^j / \tilde{\nu}_j}{\sum_j e^{-\rho_j(t+1)} \lambda_{t+1}^j / \nu_j} \\ &= \frac{\sum_j e^{-\rho_j(t+1)} (\lambda_{t+1}^j/\lambda_t^j) \lambda_t^j / \tilde{\nu}_j}{\sum_j e^{-\rho_j(t+1)} (\lambda_{t+1}^j/\lambda_t^j) \lambda_t^j / \nu_j}. \end{aligned} \tag{5.11}$$

In the expression (5.11), everything is known except the ratios $\lambda_{t+1}^j/\lambda_t^j$; and these are related (via (5.7) and (5.5)) to the unknown value $\xi = \log(S_{t+1}/S_t)$. Hence we are able to find (numerically) the value of ξ which solves the updating equation, and from this work out how the price of the asset evolves. To make a meaningful comparison, we consider the ratio of the price S_t (which arises under the mistaken belief that the price is a multiple of the dividend) to the price S_t^* which arises if the agents are able to observe the dividend process exactly. If this ratio is close to one, then the effects of the mistaken assumption is small.

The combined effects of all these assumptions are too complicated to be analyzed except numerically, so we have carried out a number of simulations. Throughout, we supposed that the annualised volatility of the dividend process is 0.25, the actual annualised growth rate is 1.5%, and the time between observations is one day (thus the moments of each log price change are those implied by the annualised figures).

The characteristics of the agents are generated randomly. One feature which we took care to build in is that if we perform a simulation with n_1 agents, and then repeat with the same random seed but with $n_2 > n_1$ agents, then the first n_1 agents in the second simulation are identical to the n_1 agents used in the first. The distributions of the different characteristics are as follows. The ρ_j are supposed to be drawn uniformly from $[0.04, 0.33]$, corresponding to mean look-ahead times ranging from 3 to 25 years. The assumed values of τ for the agents are drawn uniformly from $[0.4\tau^*, 1.05\tau^*]$, where τ^* is the true value used for the simulations. The prior means for the annualised growth rate were drawn uniformly from $[-0.05, 0.15]$, and all the ν_j are assumed to be equal to 1.

We performed a number of runs with the same random seed (and therefore the same realised sample path of δ) for 30 agents, and for 50 agents. The different runs were also

[12] We move to discrete time because in continuous time the quadratic variation of the price process would immediately tell the agents that this hypothesis is false.

distinguished by the different numbers of agents who are assumed to be diligent, in the sense that some of the agents might update their posteriors *seeing the true values of δ, and believing that they are correct.* Thus if all the agents were diligent, then the prices observed are formed exactly as described in Section 3; the ratio of the ideal stock price S to the dividend process is given by (3.7).We denote this ideal stock price by S^* for the purposes of the discussion of this section, to distinguish it from the price S actually computed at (5.11). The various figures shown come as two panels, the upper showing the log of the ratio S^*/δ, and the lower showing the log of the ratio S/S^*. The different figures differ in the number of agents assumed to be diligent; for the same total number of agents, the upper panel should be the same, and visual inspection shows that this is the case.

For 30 agents, we show in Figure 1 the behaviour of the price when no agent is diligent; the repeated ramping up followed by sharp falls is the most obvious feature[13], and the range of values covered is quite high, from about -0.2 to nearly 0.4. Changing one agent to diligent, we still see a choppy price path, Figure 3, though the ramp-ups are less pronounced, and the overall range of the trajectory is smaller. The overall level however is quite different. Increasing the number of diligent agents to 5, Figure 3 largely eliminates the peaky behaviour of the previous two plots, and it would be natural to conjecture that this more orderly behaviour becomes more prevalent as the number of diligent agents rises, but the plot Figure 4with 10 diligent agents suggests otherwise. The final plot Figure 5 in the series, with 25 diligent agents, still shows quite a wide range of variation of S from S^*; only one in six of the agents is mistakenly interpreting the price as a multiple of the dividend, and yet the log of the price ratio ranges from below -0.2 to over 0.1.

The Figures 6, 7, 8, 9, 10, 11, show the corresponding results for 50 agents, with similar qualitative features; notice particularly the dramatic crash when no agent is diligent!

What we see in the various results are qualitative features which in other contexts might be described as herding, or attributed to behavioural effects. The present framework is able to generate such qualitative features strictly within the neoclassical framework of finance; all agents are behaving rationally, the only point is that they have misinterpreted what the market prices actually are. This market is definitely *not* always right.

6 Volume of trade.

Having derived an expression (3.3) for the wealth, taking an Itô expansion gives

$$dw_t^j = w_t^j\{-\rho_j dt + (\alpha_t^j + \kappa_t)dX_t + (r_t + \kappa_t^2 + \alpha_t^j\kappa_t)dt\} \tag{6.1}$$

[13]Other simulations generate ramping down followed by sharp rises.

However, the wealth dynamics of agent j can be expressed in terms of the portfolio process π^j as

$$dw_t^j = \pi_t^j(dS_t + \delta_t dt) - c_t^j dt + (w_t^j - \pi_t^j S_t) r_t dt. \tag{6.2}$$

Comparing coefficients and using (3.17) leads to the identification

$$\boxed{\pi_t^j = \frac{w_t^j(\alpha_t^j + \kappa_t)}{\sum_i w_t^i(\alpha_t^i + \kappa_t)} = \frac{w_t^j(\alpha_t^j + \kappa_t)}{S_t(a_t + \kappa_t)}.} \tag{6.3}$$

Equation (6.3) gives us an expression for the amount of the risky asset held by agent j. In the case where all the agents have the same belief, we have that:

$$\pi_t^j = \frac{e^{-\rho_j t}/\nu_j \rho_j}{\sum_i e^{-\rho_i t}/\nu_i \rho_i}$$

hence there is no volatility in the evolution of π_t^j. However, when agents do disagree, then there will be a lot of volatility in π_t^j. To show what can happen, we will suppose that σ is constant, all the α_j are constant, and that $\rho_j = \rho$ for all j. The expression (6.3) for the proportion held by agent j is now simply

$$\begin{aligned} \pi_t^j &= \frac{w_t^j(\alpha^j + \sigma - \bar{\alpha}_t)}{\sigma S_t} \\ &= \frac{\zeta_t w_t^j(\alpha^j + \sigma - \bar{\alpha}_t)}{\zeta_t \sigma S_t} \\ &= \frac{\Lambda_t^j(\alpha^j + \sigma - \bar{\alpha}_t)}{\sigma \nu_j(\sum \Lambda_t^i/\nu_i)} \end{aligned} \tag{6.4}$$

The defining expression for $\bar{\alpha}_t$, simplified in this situation to

$$\bar{\alpha}_t = \frac{\sum \alpha^j \Lambda_t^j/\nu_j}{\sum \Lambda_t^j/\nu_j}\ , \tag{6.5}$$

leads after some calculations to

$$\begin{aligned} d\bar{\alpha}_t &= -\bar{\alpha}_t^2\, dX_t + \frac{\sum(\alpha^i)^2 \Lambda_t^i \nu_i}{\sum \Lambda_t^j/\nu_j}\, dX_t + \text{finite-variation terms} \\ &= \frac{\sum(\alpha^i - \bar{\alpha}_t)^2 \Lambda_t^i/\nu_i}{\sum \Lambda_t^j/\nu_j}\, dX_t + \text{finite-variation terms} \\ &\equiv v_t\, dX_t + \text{finite-variation terms,} \end{aligned}$$

say. Suppose that $d\pi_t^j = \theta_t^j dX_t +$ finite-variation terms. Multiplying (6.4) throughout by $\sum \Lambda_t^i/\nu_i$, and expanding gives

$$\left\{\, \theta_t^j + \pi_t^j \bar{\alpha}_t \,\right\}\left(\sum \Lambda_t^i/\nu_i\right) = \frac{\Lambda_t^j}{\sigma \nu_j}\left\{\, \alpha^j(\sigma + \alpha^j - \bar{\alpha}_t) - v_t \,\right\}$$

after some calculations. Rearranging, and recalling (3.13), we obtain the expression

$$\begin{aligned}\theta_t^j &= -\pi_t^j \bar{\alpha}_t + q_t^j \{\, \alpha^j(\sigma + \alpha^j - \bar{\alpha}_t) - v_t \,\}/\sigma & (6.6)\\ &= q_t^j \left[\frac{(\alpha^j - \bar{\alpha})^2}{\sigma} - \frac{v_t}{\sigma} + \alpha^j - \bar{\alpha} \right] & (6.7)\end{aligned}$$

with some calculation.

Notice that the sum of the θ^j is zero, as it must be. The absolute value of θ_t^j can be interpreted as the volume of trade in the risky stock by agent j. Hence the square root of the quadratic variation of the vector θ can be interpreted as the *total* volume of trade. The representation (6.7) shows that in general terms the volume of trade gets bigger with greater diversity of beliefs.

7 Conclusions

This paper has shown how to deal with diverse beliefs of agents in a completely general manner; the key observation is that we should model agents' beliefs as probability measures, whose likelihood-ratio martingales enter naturally into the optimality criterion, and thence into equilibrium prices. Abstract expressions for the state-price density and for the equilibrium stock price arise simply from the analysis, and are visibly analogous to (but extensions of) the corresponding expressions with no diversity of belief. An immediate first result is an explanation of the phenomenon of rational overconfidence.

By specializing to the case of log agents, the equilibrium can be computed quite explicitly, and its properties studied. We find quite simple and explicit expressions for the riskless rate, the stock price, the risk premium and the volatility of the stock price, in terms of the fundamentals of the problem, namely, the dynamics of the dividend process and the beliefs of the agents, expressed as likelihood-ratio martingales. Diversity of belief generates an active market, and we are able to find an expression for the volatility of the agents' holdings of the stock, which we interpret as a proxy for volume of trade. Moreover, we are able to show that under the assumption of diverse beliefs, there is benefit to individual agents to act as if their beliefs were different from what they truly believe; such actions modify the equilibrium in such a way that the agents true objectives are improved. This is therefore an analysis which explains the 'beauty contest' phenomenon commented on and postulated by Keynes, using no modelling elements other than rational expectations equilibrium and diverse beliefs. In particular, it is not necessary to introduce any 'behavioural' concepts, nor are the agents' objectives in any way unconventional.

Staying within this strictly neoclassical financial framework, we find a mechanism to generate bubbles and crashes, by supposing that some agents assume that the observed stock prices are actually constant multiples of the dividend process (as would be the case in a homogeneous market). We do not need to resort to the paraphernalia of behavioural finance - the bubble is generated by entirely rational agents, some of whom happen to be rational and mistaken.

There remain many interesting questions to be studied in this area. For example, can diverse beliefs create an economic rôle for money, by (say) imposing leverage constraints which more money will ease? The paper [30] is a first step down this road. Are there tractable examples where the agents have utilities different from log, and if so, what do the solutions look like? These and other questions are in principle amenable to a correctly-formulated modelling of diverse beliefs, which this paper has attempted to present.

A Fitting annual return and consumption data.

Kurz [27] uses his model of diverse beliefs to fit various sample moments of the Shiller data set, and we perform a similar study here.

We take a very simple version of the model, with just three agents who never change their beliefs, so we assume that the α^j are constant. We also take σ_t to be constant.

The quantities of interest are shown in the table below; we list both the empirical value[14] and the values as produced by fitting our model.

Table 1: Simulation Results

	Fitted	Empirical
Mean price/dividend ratio	25.74	25
Standard deviation of price/dividend ratio	3.67	7.1
Mean return on equity	0.072	0.07
Standard deviation of return on equity	0.108	0.18
Mean riskless rate	0.010	0.01
Standard deviation of riskless rate	0.057	0.057
Equity Premium	0.0615	0.06
Sharpe Ratio	0.330	0.33

The results shown were generated by choosing $\sigma = 0.543, \alpha^* = -0.01, \alpha^1 = 0.209, \alpha^2 = 0.711, \alpha^3 = -0.05, \rho_1 = 0.158, \rho_2 = 0.01, \rho_3 = 0.680, \nu_1 = 7.88, \nu_2 = 3.39, \nu_3 = 1$.

From the table above, we see that the diverse beliefs model with these parameter values gives quite a good fit to the sample moments considered by Kurz *et al.*. Only the standard deviation of the price/dividend ratio is substantially off the empirical value, a sample moment which we note was not fitted very closely by Kurz either, probably because the volatility of recorded annual consumption is in general too small to explain the observed volatility in stock returns. Nevertheless, the model seems to be doing a reasonable job explaining these figures given the very specific assumptions made.

B Bayesian learning.

The case in which all the α are constant corresponds to that in which the agents all start with a belief about the behaviour of the dividend process and stick with this forever. Such a setup is in some senses unsatisfactory, because even if the agents were to observe that the behaviour of the dividend were very different to their initial beliefs about it, they would still keep with these initial beliefs.

[14] These empirical values are calculated by Kurz and are based on the Shiller data set. They are based on monthly data from the S&P 500 between 1871 and 1998. See [27] and [28] for further details.

We therefore consider the case of Bayesian agents, who learn as they observe data. Bayesian learning is a huge topic which has been studied by [3], [4], [16], [9], [35] among others. For example, Guidolin and Timmermann [16] look at a discrete time case in which the dividend process can have one of two different growth rates over each time period and the probability of each growth rate is unknown to the agents. The agents are learning, so this affects the way that the stock price is calculated and hence the dynamics of the stock and options prices. Again, David and Veronesi [9] look at a continuous time model in which at any given time, the economy can be in one of two states; boom and recession. The agents do not observe this state directly, but instead must infer it from their observations of the dividend process.

We take a very unsophisticated model of Bayesian learning, which for completeness summarises a story told before else where; see, for example, Brown, Bawa & Klein [3], Brennan & Xia [4], or Rogers [35] for much the same material.

An agent observes a Brownian motion with drift:

$$Y_t = X_t + bt$$

where X is a $\mathbb{P}$-Brownian motion and b is some unknown constant. Instead of making an initial guess at the value of b and sticking with it, the agent gives a prior distribution to the unknown parameter b and then updates this prior distribution as time progresses. If the agent was sure about b, then he would have:

$$\left.\frac{d\mathbb{P}}{d\mathbb{P}^0}\right|_{\mathcal{F}_t} \equiv \Lambda_t = \exp\{bX_t - \frac{1}{2}b^2 t\}$$

However, the agent gives b a normal prior distribution with mean β and precision ϵ. [15] It follows that the change of measure the agent works with is given by:

$$\begin{aligned}\Lambda_t &= \int_{-\infty}^{\infty} \sqrt{\frac{\epsilon}{2\pi}} \exp\{-\frac{\epsilon}{2}(b' - \beta)^2 + b' X_t - \frac{1}{2}(b')^2 t\} db' \\ &= \sqrt{\frac{\epsilon}{\epsilon + t}} \ \exp\left\{\frac{X_t^2 + 2\beta\epsilon X_t - \epsilon(\beta)^2 t}{2(\epsilon + t)}\right\}\end{aligned}$$

This gives:

$$\Lambda_t = \Lambda_t \alpha_t dX_t$$

where:

$$\alpha_t = \frac{X_t + \beta\epsilon}{\epsilon + t} \tag{B.1}$$

This is of the form described in Section 2, but the α_t are now adapted processes rather than constants. Thus, our model can deal with intelligent agents who update their beliefs, as well as the simple agents who always hold the same beliefs.

[15] This is equivalent to having variance ϵ^{-1}

References

[1] Allen, F., Morris, S., and Shin, H. Beauty contests and iterated expectations in asset markets. *Review of Financial Studies 19* (2006), 719–752.

[2] Angeletos, G. M., and Pavan, A. Efficient use of information and social value of information. Tech. rep., MIT, 2006.

[3] Bawa, V. S., Brown, S. J., and Klein, R. W. *Estimation risk and optimal portfolio choice*. North-Holland, Amsterdam, 1979.

[4] Brennan, M. J., and Xia, Y. Stock price volatility and equity premium. *Journal of Monetary Economics 7* (2001), 265–296.

[5] Brown, A. A., and Rogers, L. C. G. Diverse beliefs in a single-asset economy with parameter uncertainty. Preprint, Statistical Laboratory, University of Cambridge, 2008.

[6] Brown, D., and Jennings, R. On technical analysis. *Review of Financial Studies 2* (1989), 527–551.

[7] Buraschi, A., and Jiltsov, A. Model uncertainty and option markets with heterogeneous beliefs. *Journal of Finance 61*, 6 (2006), 2841–2897.

[8] Cabrales, A., and Hoshi, T. Heterogeneous beliefs, wealth accumulation and asset price dynamics. *Journal of Economic Dynamics and Control 20* (1996), 1073–1100.

[9] David, A., and Veronesi, P. Option Prices with Uncertain Fundamentals: Theory and Evidence on the Dynamics of Implied Volatilities. *Working paper* (2000).

[10] Detemple, J., and Murthy, S. Intertemporal asset pricing with heterogeneous beliefs. *Journal of Economic Theory 62* (1994), 294–320.

[11] Diamond, D., and Verrecchia, R. Information aggregation in a noisy rational expectations equilibrium. *Journal of Financial Economics 9* (1981), 221–235.

[12] Fan, M. Heterogeneous beliefs, the term structure and time-varying risk premia. *Annals of Finance 2* (2006), 259–285.

[13] Frankel, J. A., and Rose, A. K. A survey of empirical research on nominal exchange rates. In *Handbook of International Economics III*. North Holland, Amsterdam, 1995, ch. 33, pp. 1689–1729.

[14] Grossman, S. J., and Stiglitz, J. On the impossibility of informationally efficient markets. *American Economic Review 70* (1980), 393–408.

[15] Grundy, B., and McNichols, M. Trade and revelation of information through prices and direct disclosure. *Review of Financial Studies 2* (1989), 495–526.

[16] Guidolin, M., and Timmermann, A. Option prices under bayesian learning: Implied volatility dynamics and predictive densities. *Working paper* (2001).

[17] Harris, M., and Raviv, A. Differences of opinion make a horse race. *Review of Financial Studies 6* (1993), 473–506.

[18] Harrison, J. M., and Kreps, D. Speculative investor behavior in a stock market with heterogeneous expectations. *Quarterly Journal of Economics 92* (1978), 323–336.

[19] He, H., and Wang, J. Differential information and dynamic behaviour of stock trading. *Review of Financial Studies 8* (1995), 914–972.

[20] Hellwig, C. Public announcements, adjustment delays and the business cycle. Tech. rep., Department of Economics, UCLA, 2002.

[21] Hellwig, C. Heterogeneous information and welfare effects of public information. Tech. rep., Department of Economics, UCLA, 2005.

[22] Jobert, A., Platania, A., and Rogers, L. C. G. A Bayesian solution to the equity premium puzzle. Preprint, Statistical Laboratory, University of Cambridge, 2006.

[23] Jouini, E., and Napp, C. Consensus consumer and intertemporal asset pricing with heterogeneous beliefs. *Review of Economic Studies 74* (2007), 1149–1174.

[24] Judd, K. L., and Bernardo, A. E. Asset market equilibrium with general tastes, returns, and informational asymmetries. *Journal of Financial Markets 1* (2000), 17–43.

[25] Kandel, E., and Pearson, N. D. Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy 4* (1995), 831–872.

[26] Keynes, J. *The General Theory of Employment, Interest and Money.* Macmillan, 1936.

[27] Kurz, M. Rational Diverse Beliefs and Economic Volatility. *Prepared for the Handbook of Finance Series Volume Entitled: Handbook of Financial Markets: Dynamics and Evolution* (2008).

[28] Kurz, M., Jin, H., and Motolese, M. Determinants of stock market volatility and risk premia. *Annals of Finance 1*, 2 (2005), 109–147.

[29] Leland, H. E. Who should buy portfolio insurance? *Journal of Finance 35*, 2 (1980), 581–94.

[30] Li, T. M., and Rogers, L. C. G. Lucas economy with trading constraints. Preprint, Statistical Laboratory, University of Cambridge, 2009.

[31] Lucas, R. E. Expectations and the neutrality of money. *Journal of Economic Theory 4* (1972), 103–124.

[32] Morris, S., and Shin, H. S. Social value of public information. *American Economic Review 92* (2002), 1521–1534.

[33] Morris, S., and Shin, H. S. Central bank transparency and the signal value of prices. *Brookings Papers on Econonic Activity 2* (2005), 1–66.

[34] Nakata, H. A model of financial markets with endogenously correlated rational beliefs. *Economic Theory 30*, 3 (2007), 431–452.

[35] Rogers, L. C. G. The relaxed investor and parameter uncertainty. *Finance and Stochastics 5* (2001), 131–154.

[36] Rogers, L. C. G., and Williams, D. *Diffusions, Markov Processes and Martingales*. Cambridge University Press, 2000.

[37] Singleton, K. Asset prices in a time-series model with disparately informed, competitive traders. In *New Approaches to Monetary Economics: Proceedings of the Second International Symposium in Economic Theory and Econometrics* (1987), W. Barnet and K. Singleton, Eds., Cambridge University Press.

[38] Townsend, R. Market anticipations, rational expectations and Bayesian analysis. *International Economic Review 19* (1978), 481–494.

[39] Varian, H. R. Divergence of opinion in complete markets: a note. *Journal of Finance 40* (1985), 309–317.

[40] Varian, H. R. Differences of opinion in financial markets. In *Financial Risk: Theory, Evidence and Implications* (Boston, 1989), C. C. Stone, Ed., Kluwer Academic Publishers.

[41] Wang, J. A model of competitive stock trading volume. *Journal of Political Economy 102* (1994), 127–167.

[42] Weizmann, M. A unified Bayesian theory of equity puzzles. Preprint, Harvard University, 2005.

[43] Wu, H. M., and Guo, W. C. Speculative trading with rational beliefs and endogenous uncertainty. *Economic Theory 21* (2003), 263–292.

[44] Wu, H. M., and Guo, W. C. Asset price volatility and trading volume with rational beliefs. *Economic Theory 23* (2004), 795–829.

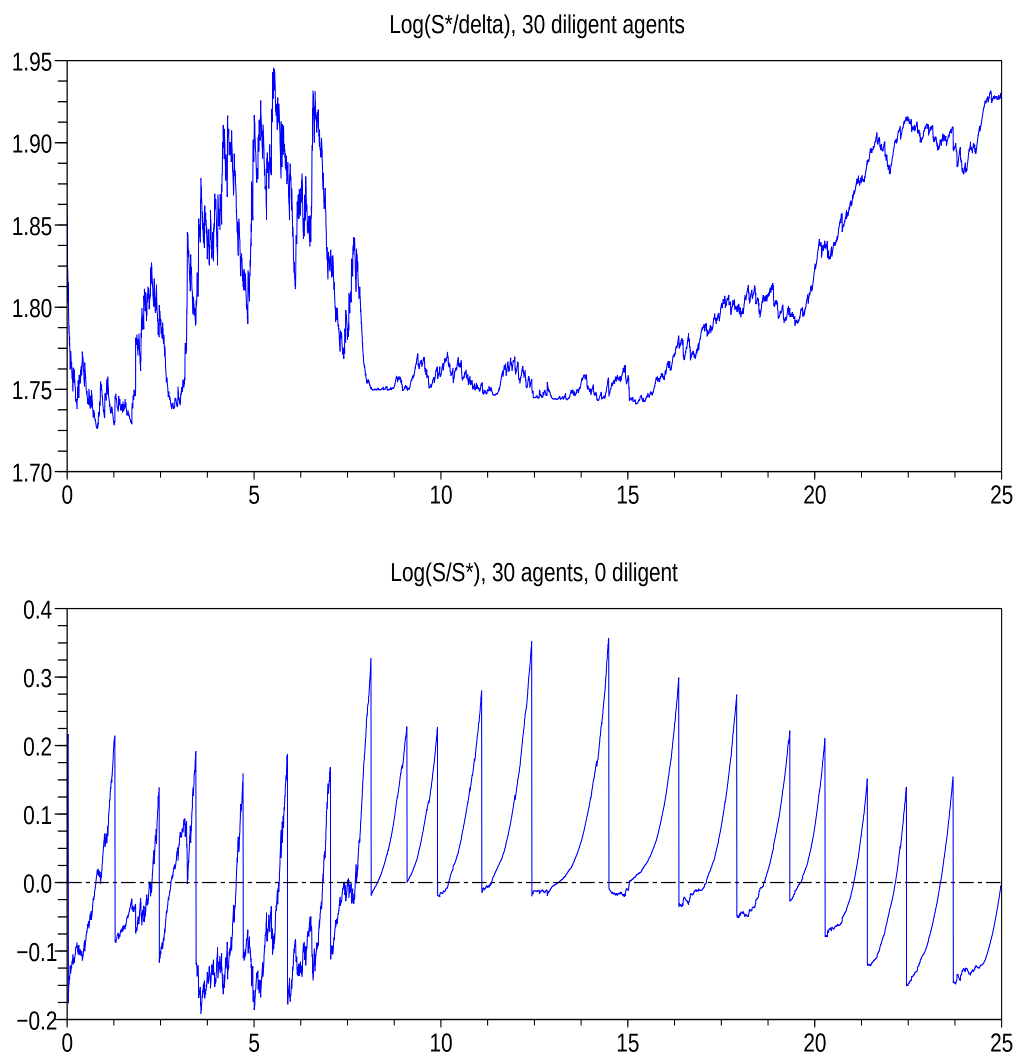

Log(S*/delta), 30 diligent agents
1.95
1.90
1.85
1.80
1.75
1.70
0
5
10
15
20
25
Log(S/S*), 30 agents, 0 diligent
0.4
0.3
0.2
0.1
0.0
−0.1
−0.2
0
5
10
15
20
25

Figure 1:

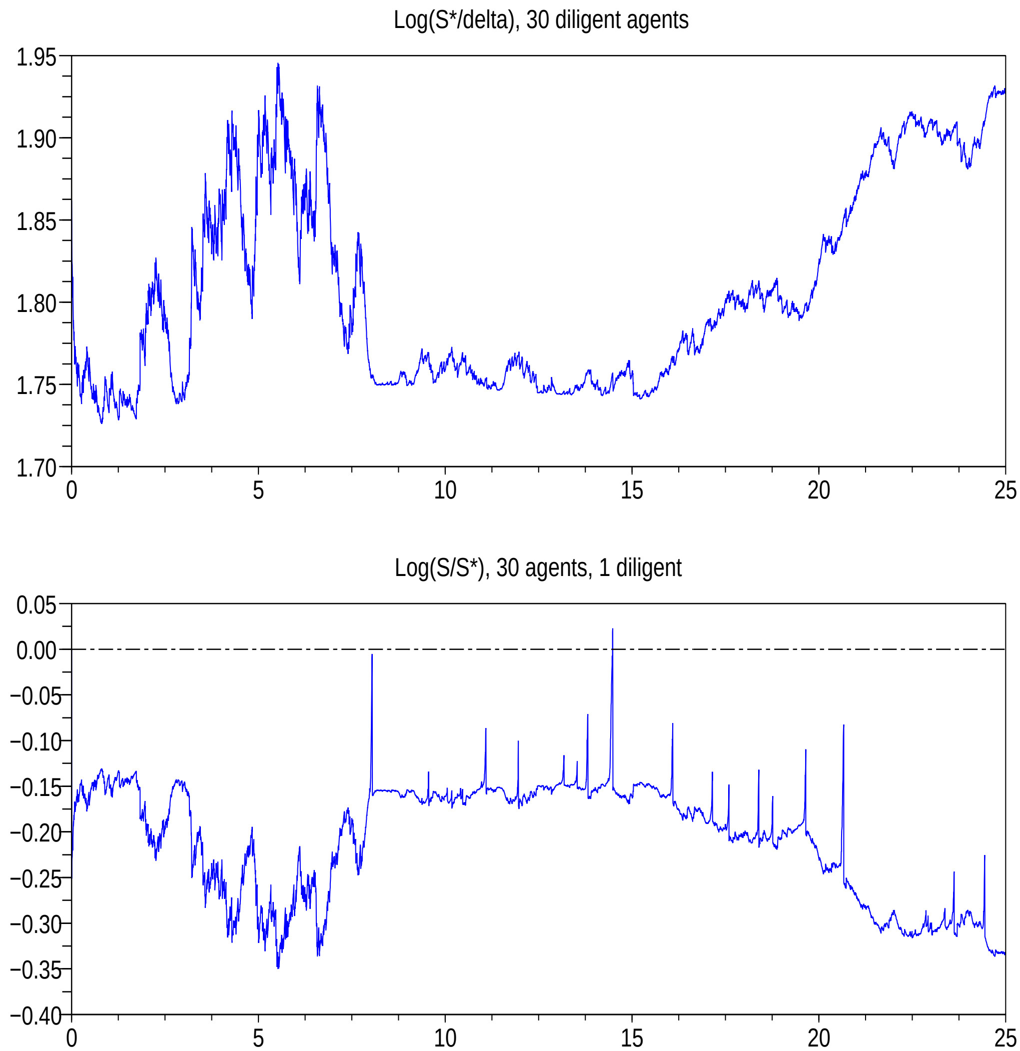
Log(S*/delta), 30 diligent agents
1.95
1.90
1.85
1.80
1.75
1.70
0
5
10
15
20
25
Log(S/S*), 30 agents, 1 diligent
0.05
0.00
−0.05
−0.10
−0.15
−0.20
−0.25
−0.30
−0.35
−0.40
0
5
10
15
20
25

Figure 2:

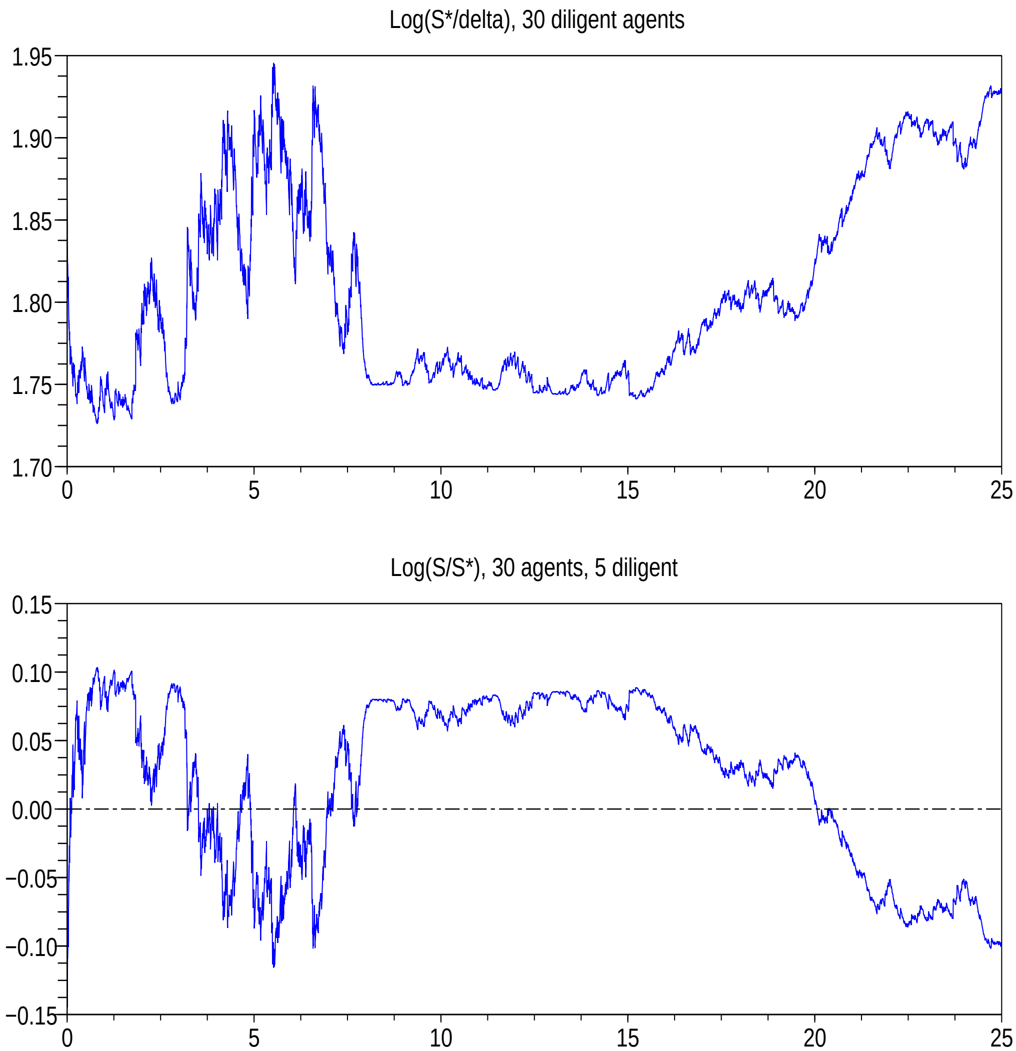
Log(S*/delta), 30 diligent agents
1.95
1.90
1.85
1.80
1.75
1.70
0
5
10
15
20
25
Log(S/S*), 30 agents, 5 diligent
0.15
0.10
0.05
0.00
−0.05
−0.10
−0.15
0
5
10
15
20
25

Figure 3:

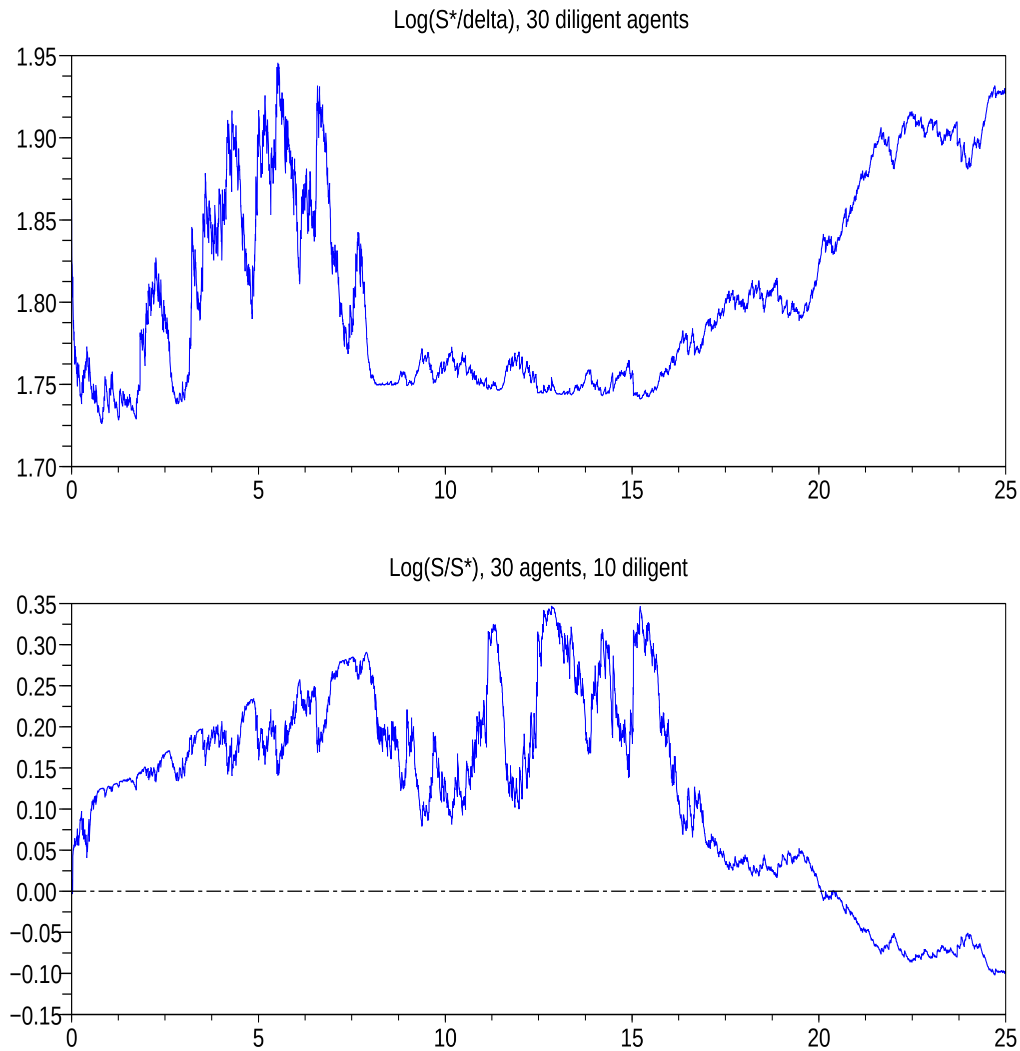
Log(S*/delta), 30 diligent agents
1.95
1.90
1.85
1.80
1.75
1.70
0
5
10
15
20
25
Log(S/S*), 30 agents, 10 diligent
0.35
0.30
0.25
0.20
0.15
0.10
0.05
0.00
−0.05
−0.10
−0.15
0
5
10
15
20
25

Figure 4:

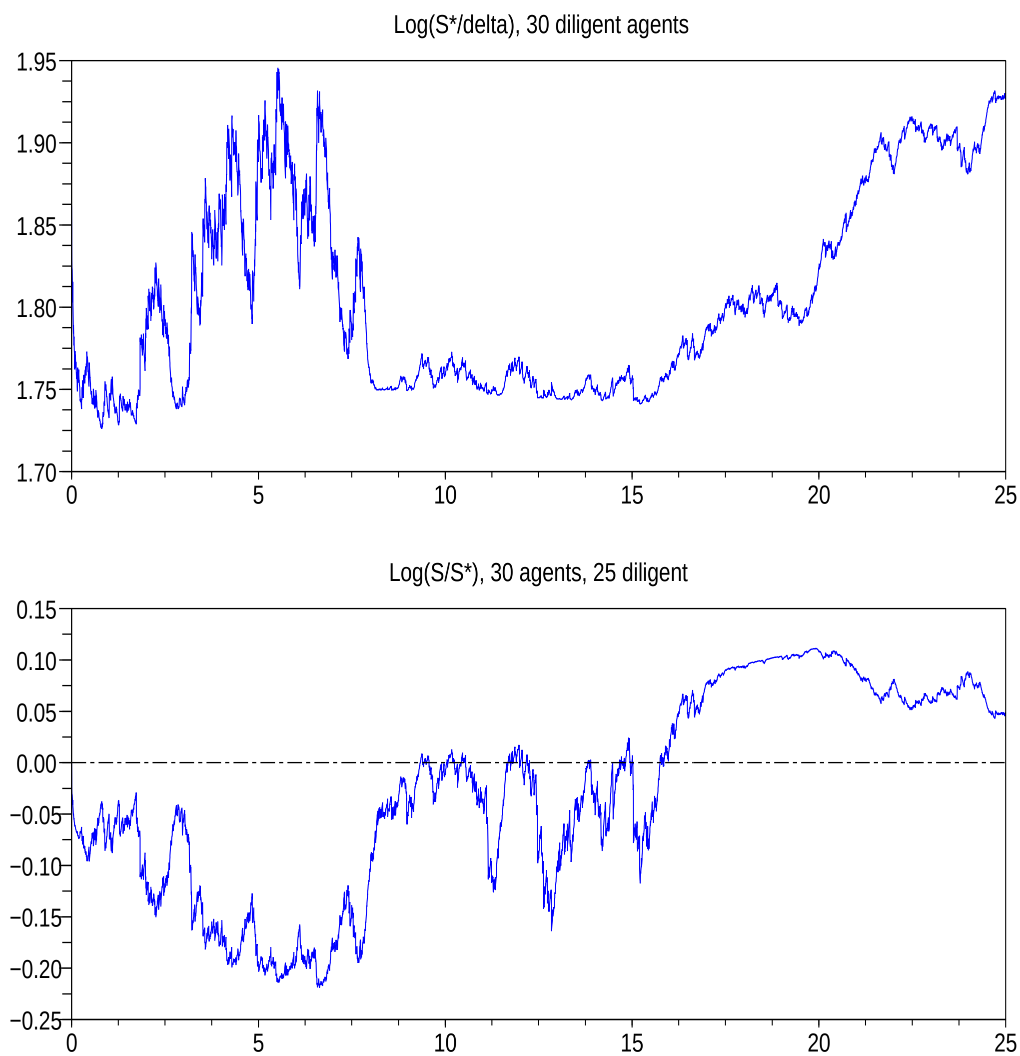
Log(S*/delta), 30 diligent agents
1.95
1.90
1.85
1.80
1.75
1.70
0
5
10
15
20
25
Log(S/S*), 30 agents, 25 diligent
0.15
0.10
0.05
0.00
−0.05
−0.10
−0.15
−0.20
−0.25
0
5
10
15
20
25

Figure 5:

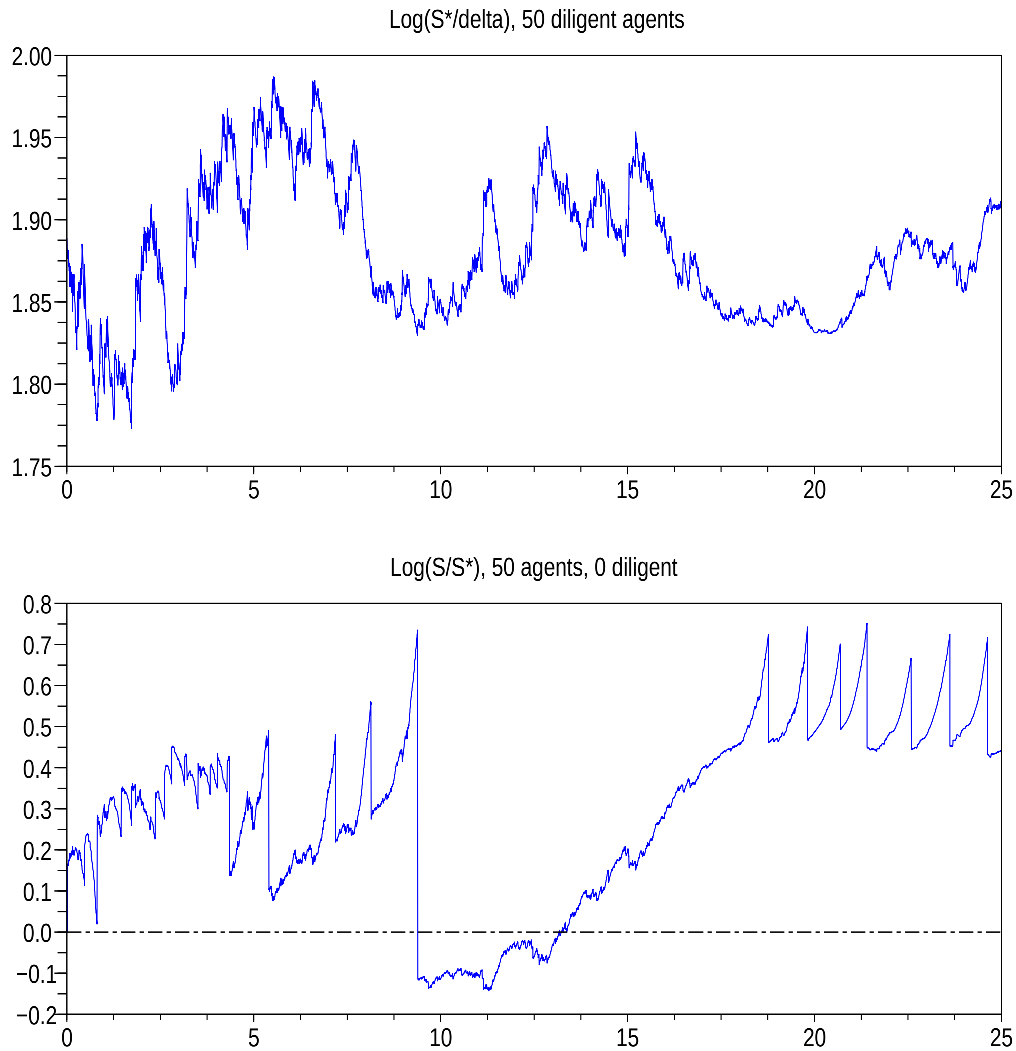
Log(S*/delta), 50 diligent agents
2.00
1.95
1.90
1.85
1.80
1.75
0
5
10
15
20
25
Log(S/S*), 50 agents, 0 diligent
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0.0
−0.1
−0.2
0
5
10
15
20
25

Figure 6:

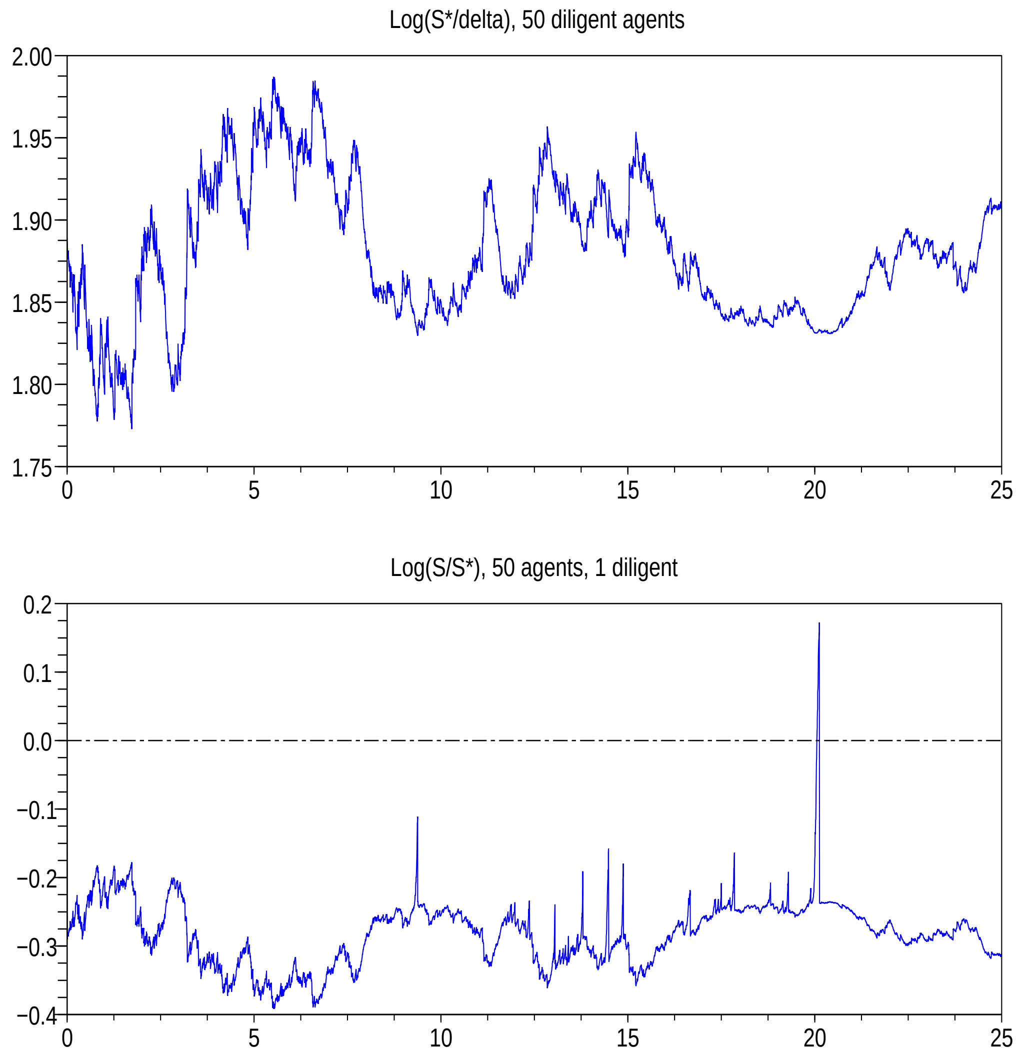
Log(S*/delta), 50 diligent agents
2.00
1.95
1.90
1.85
1.80
1.75
0
5
10
15
20
25
Log(S/S*), 50 agents, 1 diligent
0.2
0.1
0.0
−0.1
−0.2
−0.3
−0.4
0
5
10
15
20
25

Figure 7:

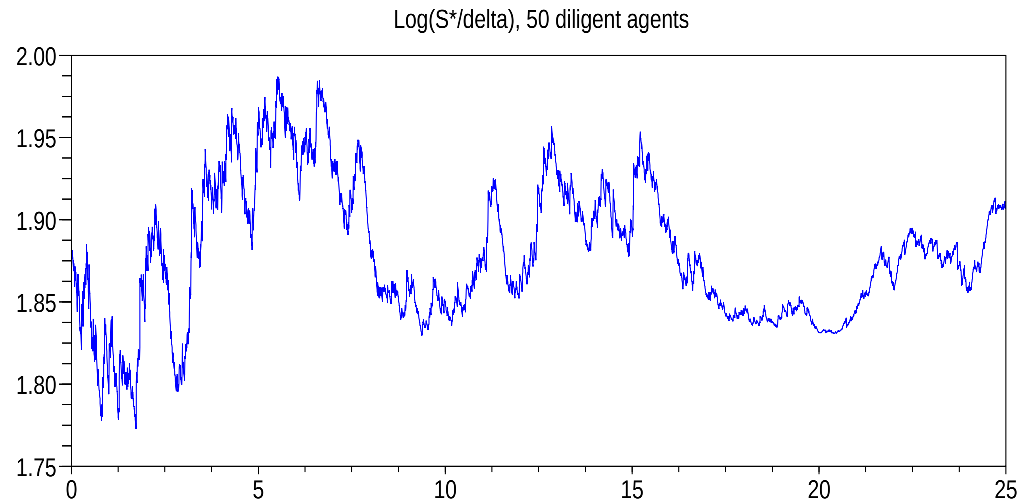

Log(S*/delta), 50 diligent agents
2.00
1.95
1.90
1.85
1.80
1.75
0
5
10
15
20
25

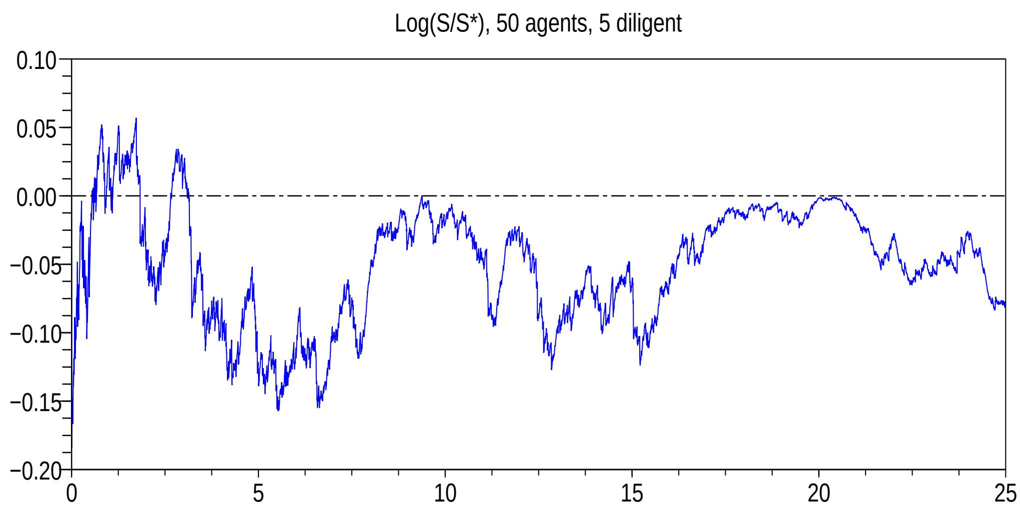

Log(S/S*), 50 agents, 5 diligent
0.10
0.05
0.00
−0.05
−0.10
−0.15
−0.20
0
5
10
15
20
25

Figure 8:

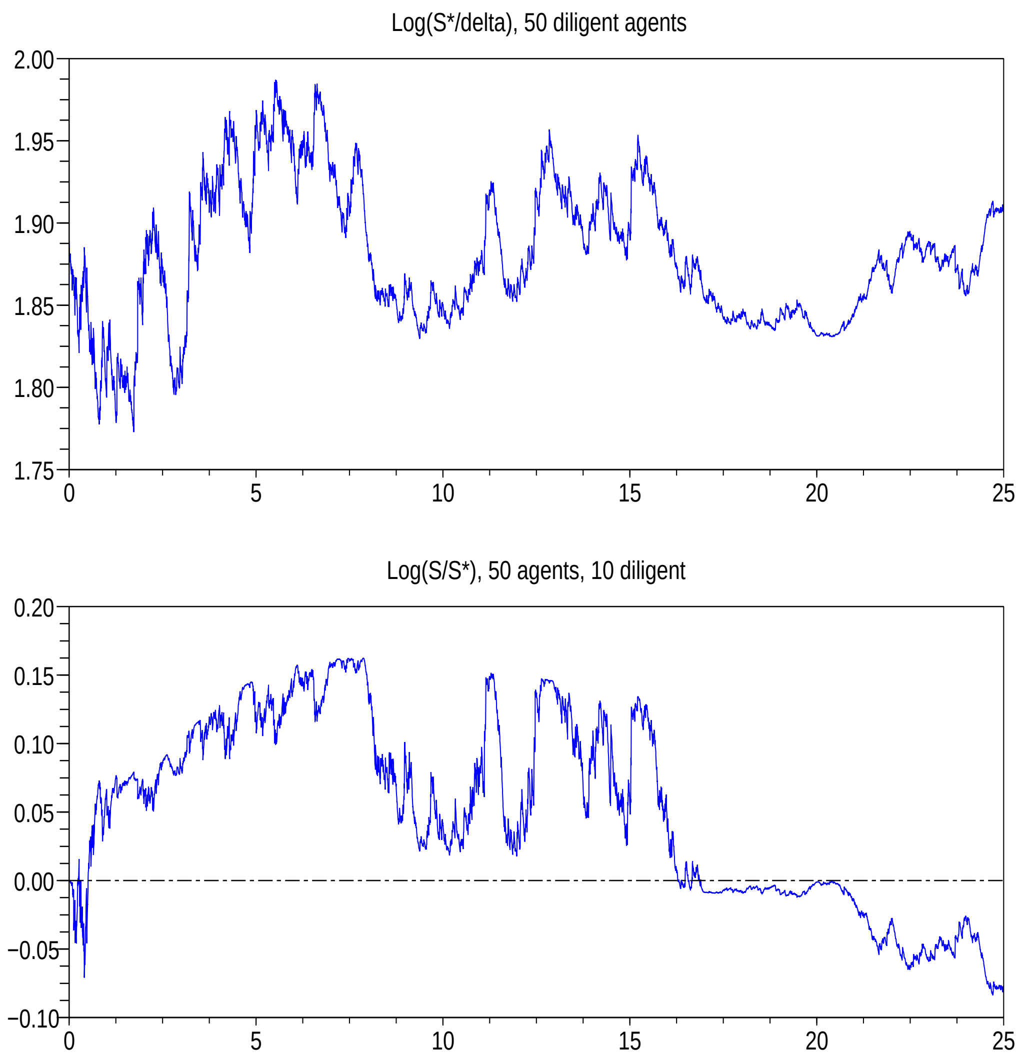
Log(S*/delta), 50 diligent agents
2.00
1.95
1.90
1.85
1.80
1.75
0
5
10
15
20
25
Log(S/S*), 50 agents, 10 diligent
0.20
0.15
0.10
0.05
0.00
−0.05
−0.10
0
5
10
15
20
25

Figure 9:

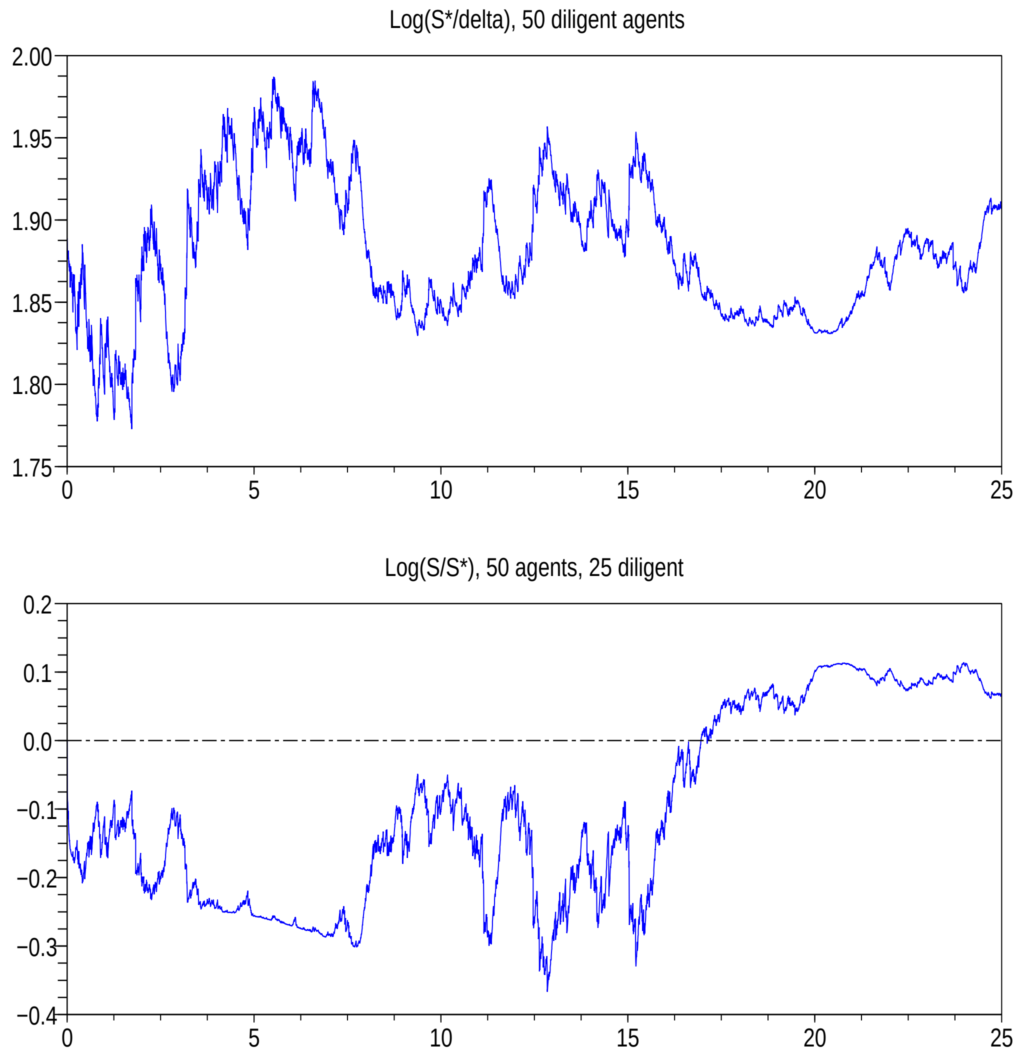
Log(S*/delta), 50 diligent agents
2.00
1.95
1.90
1.85
1.80
1.75
0
5
10
15
20
25
Log(S/S*), 50 agents, 25 diligent
0.2
0.1
0.0
−0.1
−0.2
−0.3
−0.4
0
5
10
15
20
25

Figure 10:

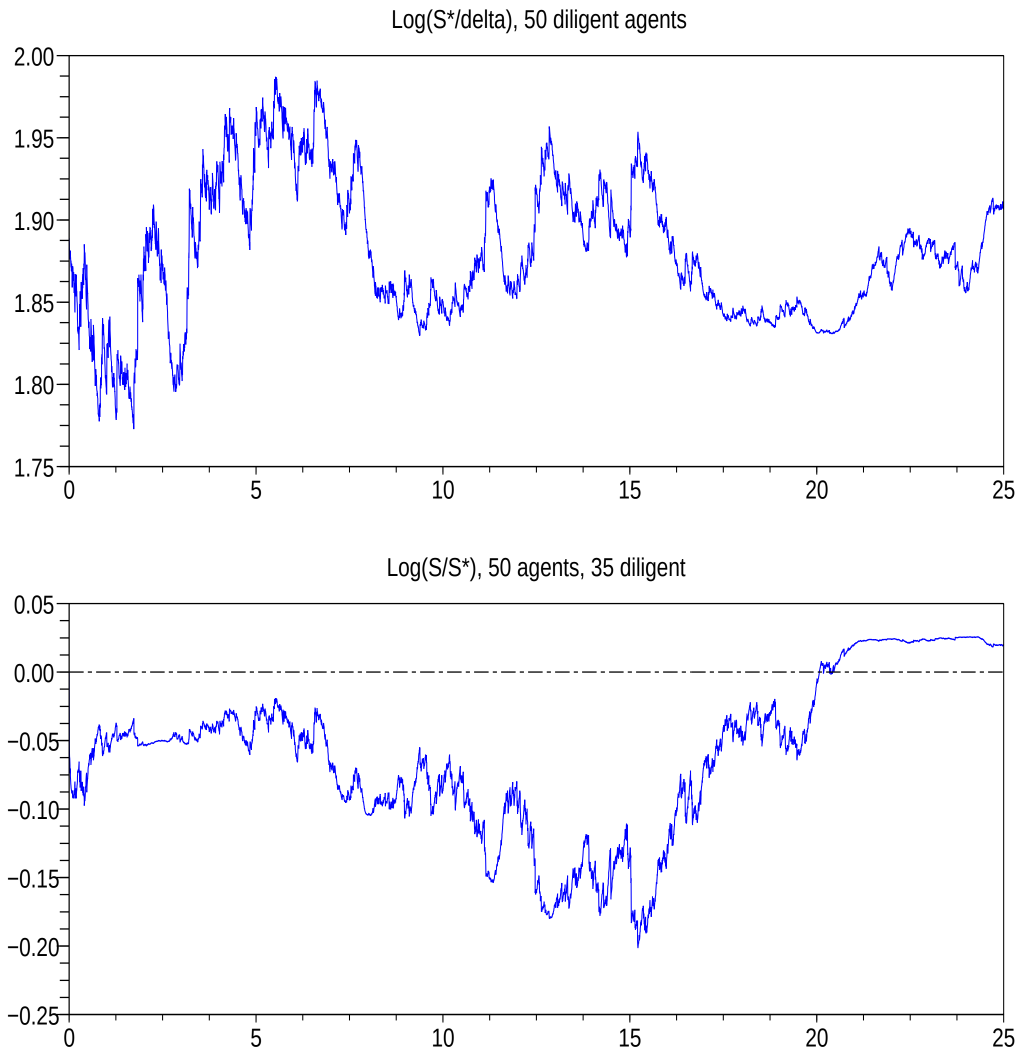
Log(S*/delta), 50 diligent agents
2.00
1.95
1.90
1.85
1.80
1.75
0
5
10
15
20
25
Log(S/S*), 50 agents, 35 diligent
0.05
0.00
−0.05
−0.10
−0.15
−0.20
−0.25
0
5
10
15
20
25

Figure 11: