Rev. Econ. Design
DOI 10.1007/s10058-012-0124-8

ORIGINAL PAPER

A Bayesian approach to experimental analysis: trading in a laboratory financial market

Marco Cipriani · Riccardo Costantini · Antonio Guarino

Received: 28 December 2011 / Accepted: 27 April 2012

Abstract We employ a Bayesian approach to analyze financial markets experimental data. We estimate a structural model of sequential trading in which trading decisions are classified in five types: private-information based, noise, herd, contrarian and irresolute. Through Monte Carlo simulation, we estimate the posterior distributions of the structural parameters. This technique allows us to compare several non-nested models of trade arrival. We find that the model best fitting the data is that in which a proportion of trades stems from subjects who do not rely only on their private information once the difference between the number of previous buy and sell decisions is at least two. In this model, the majority of trades stem from subjects following their private information. There is also a large proportion of noise trading activity, which is biased towards buying the asset. We observe little herding and contrarianism, as theory suggests. Finally, we observe a significant proportion of (irresolute) subjects who follow their own private information when it agrees with public information, but abstain from trading when it does not.

Keywords Experimental economics · Herd behavior · Contrarian behavior · Bayesian methods

JEL Classification C92 · D8 · G14

M. Cipriani (✉)
Federal Reserve Bank of New York, New York, NY, USA
e-mail: marco.cipriani@ny.frb.org

R. Costantini · A. Guarino
University College London, London, UK
e-mail: riccardo.costantini.10@ucl.ac.uk

A. Guarino
e-mail: a.guarino@ucl.ac.uk

Published online: 06 June 2012

1 Introduction

In this paper we present a Bayesian approach to the analysis of experimental data. We consider a dataset coming from an experimental study (Cipriani and Guarino 2005a) in which subjects trade an asset in sequence with a market maker. To explain the order flow observed in the experiment, we build a structural model of sequential trading. Whereas in a Perfect Bayesian Equilibrium (PBE) all traders follow their private information, in the experiment we observe deviations from equilibrium behavior. To explain them, in the structural model that we estimate we allow trades to be the outcome of different trading strategies: (a) follow one's own private information; (b) act against one's own private information and herd on the predecessors' actions when sufficient evidence has built in favor of either buying or selling the asset; (c) act as a contrarian, going against one's own private information and the predecessors' majority action; (d) abstain from trading when private information contradicts the predecessors' majority action; (e) choose randomly between the available choices. Using a Bayesian approach, we estimate the posterior distribution of the structural parameters governing the arrival process of trades in the market.

While there is now a well established tradition of structural estimation in experimental economics (in particular, the Quantal Response Equilibrium), to the best of our knowledge, a Bayesian approach to statistical inference has never been used. We believe that such approach has interesting features. In particular, it allows us to compare in a natural way non-nested models, that is, models that are not a restricted version of a more general one. In this way, we can determine which specification fits the data best. In particular, our definitions of herding, contrarianism and "irresoluteness" depend on the choice of the price after which subjects do not use their own information; different choices lead to different estimates of the model's parameters. Bayesian techniques offer a natural way of comparing them. Moreover, Bayesian econometrics allows the experimenter to incorporate any prior information he may have on the parameters' distribution, and to estimate models through simulation when it would be difficult to do so with Maximum Likelihood.

We find that the model best fitting the data is that in which a proportion of trades comes from subjects who do not rely only on their private information, once the difference between the number of previous buy and sell decisions is at least two. In this model, the majority of trades stems from subjects following their private information. There is also a very large proportion of noise trading activity, which is biased towards buying the asset. Finally we observe very little herding and contrarianism, as theory suggests. We observe a significant proportion of subjects who follow their own private information when it agrees with public information, but abstain from trading when it does not.

The paper is organized as follows. Section 2 describes the experiment, Sect. 3 the empirical methodology, Sect. 4 the estimation strategy and Sect. 5 the results. Section 6 concludes. The Appendix contains the graphs of the parameters' posterior distributions.

2 The experiment

We use the data collected by Cipriani and Guarino (2005a) in an experimental study with undergraduate students of all disciplines at New York University. In particular, we use the data of the "Flexible Price" and "No History" treatments in Cipriani and Guarino (2005a). For each treatment, the experiment was repeated for four sessions. We, therefore, consider a total of eight sessions. In each session, there were 12 participants acting as traders. The procedures of the experiment replicate a simple model of sequential trading in a dealer's market. Here we only summarize them and refer the reader to Cipriani and Guarino (2005a) for a detailed illustration.

The experiment was paper and pencil. At the beginning of each session, the experimenters gave written instructions (available on request) to all subjects. The instructions were read aloud in an attempt to make the structure of the game common knowledge.

Each session consisted of ten rounds, in which subjects traded an asset in sequence. In each round, each subject was called to trade only once, with the sequence being randomly determined by the experimenter. Before each round, an experimenter determined the value of the asset, which could be high (100) or low (0) with equal probability.

An experimenter acted as market maker, setting the price at which subjects could trade. Before trading, each subject received a binary signal (a bad signal 0 or a good signal 1) with precision 0.7 on the value of the asset. After observing the signal, the subject decided whether he wanted to buy, to sell or not to trade at the price set by the market maker. In the "Flexible Price" Treatment, previous subjects' decisions were public information, along with the prices at which they had been taken. In the "No History" Treatment, previous subjects' decisions and prices were not public information. Therefore, a subject knew only his signal and the price at which he could trade.[1] At the end of each round, that is, after all 12 participants had traded, the realization of the asset value was revealed.

In the experiment, the numeraire was a fictitious currency called lira. Payoffs were computed as follows: if a subject bought, he earned (100 + Value − Price) liras. If he sold, he earned (100 + Price − Value) liras. If he decided not to trade, he earned 100 liras. That is, subject were endowed with 100 liras at the beginning of each round in order to avoid negative payoffs.

During the experiment, the market maker updated the price in a Bayesian way, assuming that all subjects behaved rationally. Economic theory predicts that rational subjects should always follow their signal, that is, they should buy after seeing a good signal and sell after seeing a bad one. No one should decide not to trade, as private information allows subjects to earn money in expected value. Therefore, when a subject bought the asset, the price was updated assuming that he had seen a good signal. Similarly, when a subject sold the asset, the price was updated assuming that he had observed a bad signal. Finally, in the case of a no trade, the price was kept constant.

[1] Cipriani and Guarino (2005a) report that subjects' behavior was not significantly different in the two treatments.

3 Empirical methodology

In the PBE of the financial economy implemented in the laboratory, agents always follow their private information. Cipriani and Guarino (2005a) show that while the PBE has some predictive power, there are deviations from equilibrium in the experimental data. Our purpose is to explain the order flow observed in the laboratory by allowing subjects to have different trading strategies. In particular, based on the discussion of the experimental data by Cipriani and Guarino (2005a), we postulate that trades can be the expression of five different strategies:

1. With probability p, a trade stems from a subject who simply relies on private information, buying with a good signal and selling with a bad signal. We refer to this type of trade as a "private-information based trade."
2. With probability n, a trade stems from a subject acting as noise trader, whose actions are altogether random. In particular, a noise trader buys, sells or does not trade the asset with probabilities ε_b , ε_s and $1 - \varepsilon_b - \varepsilon_s$.
3. With probability h, a trade stems from a subject acting as a "herder." Specifically, let us denote by p_t^d the price faced by the subject playing at time t of round d.[2] We consider two price thresholds, $T^L \leq 50 \leq T^H$. A herder follows his signal (behaving like private-information based traders, described under point 1 above) whenever $T^L \leq p_t^d \leq T^H$; he buys independently of his signal for $p_t^d > T^H$ and sells independently of his signal for $p_t^d < T^L$. Since the price is higher than the threshold T^H when there have been more buy than sell orders, and lower than the threshold T^L when there have been more sell than buy orders, this trading strategy consists in following the majority trading action chosen by the predecessors.
4. With probability c a trade stems from a subject acting as a "contrarian." Exactly as herders, contrarians follow their signal whenever $T^L \leq p_t^d \leq T^H$. In contrast to herders, however, they go *against* the pattern of past trading activity whenever the price is outside the interval defined by the two thresholds. In particular, they sell for a high price $p_t^d > T^H$, and buy for a low price $p_t^d < T^L$.
5. Finally, with probability i, a decision stems from a subject acting as an "irresolute." Like herders and contrarians, irresolutes follow their signal when $T^L \leq p_t^d \leq T^H$. However, if $p_t^d > T^H$ they only buy with a good signal; with a bad signal they abstain from trading. Conversely, if $p_t^d < T^L$ they only sell with a bad signal; with a good signal they abstain from trading. In other words, when the consensus that has built in the market goes against their private signal, these traders behave in an "irresolute" manner and decide not to take a position.

A few considerations are in order. First, a no trade can only come from a noise trader or from an irresolute trader. All the other strategies described above only entail either selling or buying the asset.[3]

[2] We counted rounds sequentially, by session and treatment. Round 1 was the first round of the first session of the Flexible Price Treatment, round 11 was the first round of the second session of the Flexible Price Treatment, round 41 was the first round of the first session of the No History Treatment, and so on.

[3] Cipriani and Guarino (2005b) estimate the proportion of noise versus information based traders (i.e., traders of type 1 and 2 only), using a frequentist approach.

Second, we will estimate our model for different levels of the thresholds and then compare the fit. In particular, we will compare the fit of the model estimated with three different sets of thresholds: $\{T^L = 50, T^H = 50\}$, $\{T^L = 30, T^H = 70\}$ and $\{T^L = 15, T^H = 85\}$. Note that, in our laboratory experiment, there is a one-to-one relationship between prices and trade imbalance (i.e., the difference between the number of buy orders and sell orders from the beginning of the round until a subject trades). As a result, the price thresholds correspond to different levels of the trade imbalance. In particular, a threshold of 50 corresponds to a trade imbalance of zero (equal number of buys and sells in the past history of trades). The 30–70 thresholds correspond to a trade imbalance of -1 and 1 (that is, the number of sells exceeds the number of buys by one, or vice versa). Finally, the 15–85 thresholds correspond to a trade imbalance of -2 and 2 (that is, the number of sells exceeds the number of buys by two, or vice versa). The thresholds (either directly defined in terms of the price or of the trade imbalance) define what we mean by herding behavior or by contrarianism. This means that, e.g., a trade stems from a subject acting as a herder if he follows the crowd when the absolute value of the trade imbalance is at least 1 (for $\{T^L = 50, T^H = 50\}$), at least 2 (for $\{T^L = 30, T^H = 70\}$), or at least 3 (for $\{T^L = 15, T^H = 85\}$).

Third, $p+c+i$ represents the probability that a trade stems from a trader following a strategy which, although not in line with the PBE, is not random. In particular, it can be rationalized in terms of following the crowd (herders), going against it (contrarians), or being willing to trade only when all private and public information point to the same direction (irresolute).

We want estimate the vector of parameters $\Phi = \{p, n, h, c, i, \varepsilon_b, \varepsilon_s\}$, which defines the arrival of different types of decisions in the laboratory. To this purpose, we construct the likelihood function of the data. Let us denote the history of trades for round d by $h^d = \{x_t^d : x_t^d \in \{Buy, Sell, NoTrade\}, t = 1, 2, \ldots, 12\}$. In the same fashion, let us denote the history of signals and prices in round d as $s^d = \{s_t^d : s_t^d \in \{0, 1\}, t = 1, 2, \ldots, 12\}$ and $p^d = \{p_t^d : p_t^d \in [0, 100], t = 1, 2, \ldots, 12\}$. The likelihood function can then be written as

$$\mathcal{L}\left(\Phi; \{h^d\}_{d=1}^D, \{s^d\}_{d=1}^D, \{p^d\}_{d=1}^D\right) = \Pr\left(\{h^d\}_{d=1}^D | \Phi, \{s^d\}_{d=1}^D, \{p^d\}_{d=1}^D\right),$$

where $D = 80$ is the total number of rounds in our dataset. Since rounds are independent,[4] the likelihood function can be expressed as the product of the likelihood of trading decisions in each round d, that is,

$$\mathcal{L}\left(\Phi; \{h^d\}_{d=1}^D, \{s^d\}_{d=1}^D, \{p^d\}_{d=1}^D\right) = \prod_{d=1}^{D} \Pr\left[h^d | \Phi, s^d, p^d\right].$$

4 The independence across rounds stems from the maintained hypothesis in our structural estimation that the behavioral model we presented above is the "true" data generating process; that is, that all heterogeneity in subjects' decisions (and all possible time-dependence within rounds) is captured by the parameters that we estimate. Note that our assumption that the parameters governing subjects behavior are constant across sessions is consistent with the fact that we do not observe much cross-session heterogeneity in subjects' behavior.

In the same fashion, conditional on the signal s_t^d and the price p_t^d, an action x_t^d is independent of the other actions, and, therefore, we can write

$$\Pr\left[h^d|\Phi, s^d, p^d\right] = \prod_{t=1}^{12} \Pr\left[x_t^d|\Phi, s_t^d, p_t^d\right].$$

Therefore, the probability of observing a subject buying the asset at time t of round d is

$$\begin{aligned}\Pr\left[x_t^d = Buy|\Phi, s_t^d, p_t^d\right] &= p\mathbf{1}_{\{s_t^d=1\}} + h\mathbf{1}_{\{p_t^d>T^H\}\cup\{s_t^d=1,T^L\leq p_t^d\leq T^H\}} \\ &\quad + c\mathbf{1}_{\{p_t^d<T^L\}\cup\{s_t^d=1,T^L\leq p_t^d\leq T^H\}} + i\mathbf{1}_{\{s_t^d=1,T^L\leq p_t^d\}} + n\varepsilon_b,\end{aligned}$$

where $\mathbf{1}_{\{E\}}$is the indicator function taking value 1 if the event E occurs and 0 if not.

Similarly, the probability of observing a subject selling the asset at time t of round d is

$$\begin{aligned}\Pr\left[x_t^d = Sell|\Phi, s_t^d, p_t^d\right] &= p\mathbf{1}_{\{s_t^d=0\}} + h\mathbf{1}_{\{p_t^d<T^L\}\cup\{s_t^d=0,T^L\leq p_t^d\leq T^H\}} \\ &\quad + c\mathbf{1}_{\{p_t^d>T^H\}\cup\{s_t^d=0,T^L\leq p_t^d\leq T^H\}} + i\mathbf{1}_{\{s_t^d=0,\, p_t^d\leq T^H\}} + n\varepsilon_s.\end{aligned}$$

Finally the probability of observing a no trade is given by

$$\Pr\left[x_t^d = No\ Trade|\Phi, s_t^d, p_t^d\right] = n(1-\varepsilon_b-\varepsilon_s) + i\left[\mathbf{1}_{\{s_t^d=1,\, p_t^d<T^L\}} + \mathbf{1}_{\{s_t^d=0,\, p_t^d>T^H\}}\right].$$

4 Estimation strategy

We estimate the posterior distribution of the parameters governing the arrival process of trades in the market through a Bayesian approach. We compute the parameters' posteriors under two different assumptions on the priors: (1) that all parameters are uniformly distributed in the [0, 1] interval; (2) that they are all distributed as a beta distribution, $B(0.5, 0.5)$. Both priors are usually used in Bayesian econometrics as "uninformative priors." Our results are largely unaffected by the choice of the prior, therefore we only report those for the uniform distribution (those for the beta are available upon request).

In order to compute any posterior moment or statistics, we draw a sample (a chain) from the parameters' posterior distribution, with a methodology called Markov Chain Monte Carlo (MCMC). We start from an initial parameter vector Φ_0, drawn from the prior distribution. We then draw another proposed parameter vector Φ^* drawn from a proposal distribution $J(\Phi)$ (in our case, a uniform distribution). We set

$$\Phi^1 = \left\{\begin{array}{ll} \Phi^* & \text{with probability} \quad \min\left(\dfrac{f(\Phi^*|data)/J(\Phi^*)}{f(\Phi^0|data)/J(\Phi^0)}, 1\right), \\ \Phi^0 & \text{otherwise,} \end{array}\right\}$$

where $f(\Phi^*|data)$ is the conditional density distribution of the parameter vector. Once we compute Φ^1, we draw another proposed parameter vector Φ^{**}, and set Φ^2 equal to either Φ^{**} or Φ^1. We continue the procedure until we have drawn a chain of the desired length.

Let us provide some intuition for the algorithm, and, to this purpose, let us focus on the case in which the proposal distribution is also uniform, in which case $\frac{f(\Phi^*|data)/J(\Phi^*)}{f(\Phi^0|data)/J(\Phi^0)} = \frac{f(\Phi^*|data)}{f(\Phi^0|data)}$. At each stage t of the chain, the algorithm sets Φ_t equal to the proposed parameter vector Φ^* if the posterior distribution for Φ^* is higher than that of Φ_{t-1}; if instead the posterior distribution of Φ^* is lower than that of Φ_{t-1}, the new element of the chain Φ_t is set equal to Φ^* with a probability that is higher the higher the ratio between the two distributions $f(\Phi^*|data)$ and $f(\Phi^0|data)$.

It is relatively straightforward to show that the algorithm (which is called a Metropolis-Hasting algorithm with independent sampling) produces a sample of parameter vectors $\Phi^0, \Phi^1 \dots \Phi^n$ that converges to the parameters' posterior distribution.[5] For each estimated model, we drew a chain of length 500, 000; moreover, in order to mitigate the effect of the starting distribution, we discarded the first 150, 000 elements of the chain.[6]

The Bayesian approach allows us to test easily between models that are not nested. In particular, as we mentioned in the previous section, herding, contrarianism and irresoluteness can be defined with respect to different levels of price thresholds T^L and T^H. For each two models H_1 and H_2 that we want to compare, we can compute their Bayes Factor $B_{12} = \frac{\Pr(data|H_1)}{\Pr(data|H_2)}$; the higher the ratio (that is, the higher the probability that the data is generated by model H_1 rather than by model H_2) the stronger the evidence in favor of model H_1 vis à vis model H_2. Following Kass and Raftery (1995) we interpret any value of $2\log B_{12}$ above 6 as strong evidence against H_2, and any value between 2 and 6 as positive evidence against the H_2.[7] In order to compute the Bayesian factor we use the harmonic mean identity method described in Raftery et al. (2006).

5 Results

In Table 1, we report the modes, means and standard deviations of the parameters' posterior distributions for the three models we estimated: $\{T^L = 50, T^H = 50\}$, $\{T^L = 30, T^H = 70\}$ and $\{T^L = 15, T^H = 85\}$. Table 2 reports the 0.05, 0.10, 0.90 and 0.95 quantiles for each parameter. In the rest of the section, we will comment on the results in terms of posterior means, pointing out when the mode gives a slightly different picture. Note that all the parameters are quite precisely estimated: the stan-

[5] Note that we implemented the algorithm parameter by parameter sequentially (and not for the whole parameter vector at the same time). For a rigorous explanation of the MCMC methodology see, for example, Gelman et al. (2003), Koop (2003) and Bolstad (2010).

[6] Acceptance rates across parameters (that is, the probability that a new proposed parameter vector Φ^* is accepted as part of the Markov chain) and model specifications were on average 24 % for the uniform prior, and 17 % for the beta prior (low acceptance rates slow down convergence to the posterior distribution).

[7] We report the thresholds in terms of twice the natural logarithm of the Bayes Factor as it is of the same scale as the likelihood ratio test statistics.

Table 1 Posterior mode, mean and standard deviation (uniform prior)

	$\{T^L = 50, T^H = 50\}$			$\{T^L = 30, T^H = 70\}$			$\{T^L = 15, T^H = 85\}$		
	Mode	Mean	Std.	Mode	Mean	Std.	Mode	Mean	Std.
p	0.39	0.36	0.04	0.25	0.29	0.05	0.21	0.26	0.06
n	0.30	0.39	0.05	0.45	0.41	0.04	0.44	0.45	0.03
i	0.19	0.14	0.04	0.20	0.17	0.04	0.24	0.20	0.05
h	0.13	0.08	0.04	0.07	0.07	0.04	0.07	0.05	0.04
c	0.11	0.04	0.03	0.07	0.05	0.03	0.09	0.04	0.03
ε_B	0.31	0.28	0.05	0.30	0.30	0.04	0.31	0.31	0.04
ε_S	0.27	0.22	0.04	0.24	0.21	0.04	0.23	0.23	0.04

Table 2 Posterior percentiles (uniform prior)

	$\{T^L = 50, T^H = 50\}$				$\{T^L = 30, T^H = 70\}$				$\{T^L = 15, T^H = 85\}$			
	<0.05	<0.1	<0.90	<0.95	<0.05	<0.1	<0.90	<0.95	<0.05	<0.1	<0.90	<0.95
p	0.29	0.31	0.41	0.42	0.21	0.23	0.36	0.37	0.16	0.18	0.33	0.35
n	0.30	0.32	0.45	0.47	0.35	0.36	0.46	0.48	0.40	0.41	0.49	0.51
i	0.08	0.09	0.19	0.20	0.10	0.11	0.22	0.24	0.11	0.13	0.27	0.29
h	0.01	0.03	0.12	0.14	0.01	0.02	0.12	0.14	0.00	0.01	0.10	0.11
c	0.00	0.01	0.07	0.09	0.01	0.01	0.08	0.09	0.00	0.01	0.08	0.10
ε_B	0.19	0.21	0.34	0.36	0.24	0.25	0.35	0.37	0.25	0.26	0.35	0.37
ε_S	0.14	0.16	0.27	0.29	0.15	0.16	0.26	0.28	0.18	0.19	0.28	0.29

dard deviations of all the posterior distributions (in the three models that we estimate) range between 0.03 and 0.06. Moreover, although the priors were flat, the distribution of most parameters were bell-shaped, and relatively symmetric around the mean (see the Appendix).

Under all parametrizations for the thresholds, between one third and one fourth of all trades can be classified as coming from traders who followed their own private information ("private-information based traders"). The proportion of trades coming from subjects following their private information decreases, however, as the thresholds become more extreme (from 0.36 with a 50–50 threshold, to 0.26 with a 15–85 threshold).

In contrast, pure random trading (n) accounts for 40 % of all trading activity, and it is biased toward buying (with ε_B ranging from 0.28 to 0.31), rather then selling (ε_S between 0.21 and 0.23). The behavior of noise subjects shows little variation across models.

What about herding and contrarianism? As predicted by sequential trading models with asymmetric information (see, e.g., Avery and Zemsky 1998), there seems to be very little of it. In particular, only between 5 and 8 % of trades can be attributed to herding, whereas contrarianism ranges between 4 and 5 % (according to the chosen level of the thresholds). Note that these numbers are somewhat higher if we look at the mode of the posterior distributions, with herding and contrarianism reaching 13 and 11 % under the 50–50 threshold specification. Moreover, under most specifications, looking at the 0.1–0.9 percentiles, the level of herding ranges between 1 and 12 %, whereas that of contrarianism between 0 and 8 %. That is, a relatively high mass of probability is attached to levels of herding and contrarianism above 5 %. Indeed, in contrast to all other parameters, the distribution for h and c are both quite strongly skewed to the right, with posterior distributions having relatively high mass on the right tail (see Appendix).

In contrast to what theory predicts, however, there is a significant number of trades, on average between 15 and 20 % of the sample, stemming from subjects who act as "irresolute" traders. That is, stemming from subjects who abstain from trading when private (i.e., the signal) and public (i.e., the price) information do not agree with each other, but follow both when they agree. The presence of this type of subjects in the laboratory is an interesting departure from the theory, as abstention from trading is one of the main departure from equilibrium behavior (see Cipriani and Guarino 2005a).

Finally, we want to test which specification, in terms of thresholds, explains our data better. In order to do so, we compute the implied Bayes factors for the three models. When we compare $\{T^L = 50, T^H = 50\}$ with $\{T^L = 30, T^H = 70\}$ the computed Bayes factor 7.64 provides positive evidence in favor of the $30 - 70$ thresholds. In the same fashion, when $\{T^L = 50, T^H = 50\}$ is compared to $\{T^L = 15, T^H = 85\}$, we obtain positive evidence in favor of $\{T^L = 50, T^H = 50\}$, with a Bayes factor of 3.43. It is not surprising then to find out (strong) evidence in favour of $\{T^L = 30, T^H = 70\}$ vis à vis $\{T^L = 15, T^H = 85\}$. Therefore, the data support the inference that the level of prices that makes agents switch from rationality to herding, contrarianism and irresoluteness is 30–70, which corresponds to a trade imbalance of at least 2 in absolute value.

6 Conclusions

We have studied the data of an experimental financial market through a Bayesian approach. In our model of sequential trading, decisions are classified in five types: private-information based, noise, herd, contrarian and irresolute. Through MCMC simulation, we have estimated the posterior distributions of the structural parameters. This technique has allowed us to evaluate how different, non-nested models fit the data. While structural estimation is well established in experimental economics (an example is the widespread use of parameter estimation for the Quantal Response Equilibrium), the Bayesian approach to statistical inference has not been exploited so far. We believe that the use of more complicated structural models and the need of evaluating the fit of non-nested models will likely lead to a diffusion of this approach.

Acknowledgments Antonio Guarino gratefully acknowledges the financial support of the ERC and of the INET foundation. The views expressed herein are those of the authors and should not be attributed to the IMF, its Executive Board, or its management; the Federal Reserve Bank of New York, or the Federal Reserve System. All errors are ours.

Appendix Posterior distributions

See Figs. 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, and 15.

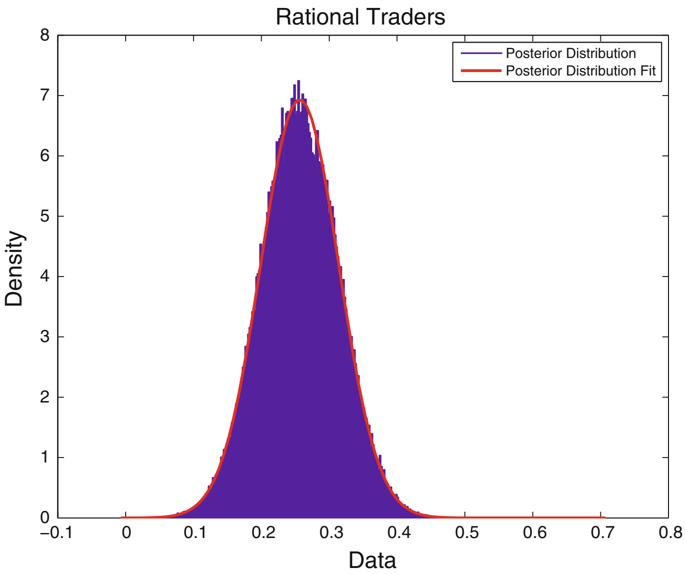

Fig. 1 Posterior distribution for the parameter p for the model with $\{T^L = 15, T^H = 85\}$

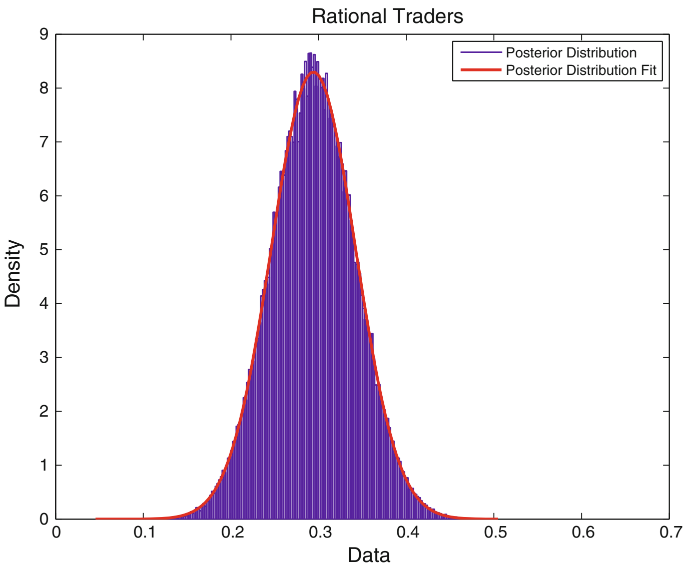

Fig. 2 Posterior distribution for the parameter p for the model with $\{T^L = 30, T^H = 70\}$

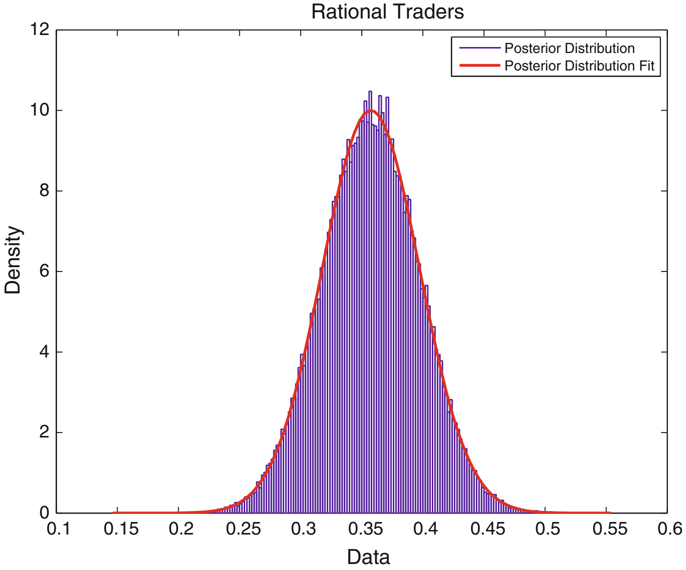

Fig. 3 Posterior distribution for the parameter p for the model with $\{T^L = 50, T^H = 50\}$

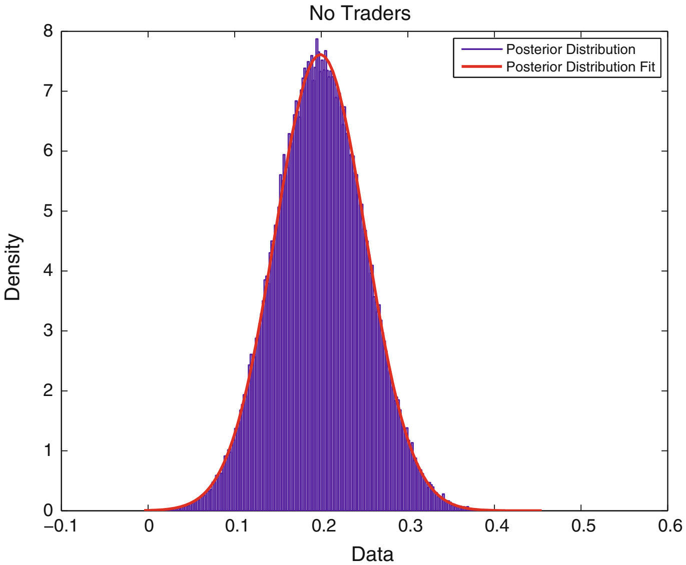

Fig. 4 Posterior distribution for the parameter i for the model with $\{T^L = 15, T^H = 85\}$

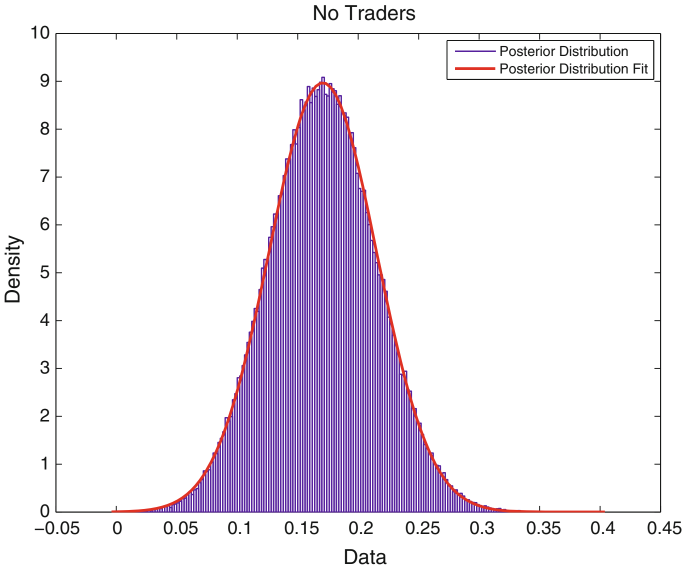

Fig. 5 Posterior distribution for the parameter i for the model with $\{T^L = 30, T^H = 70\}$

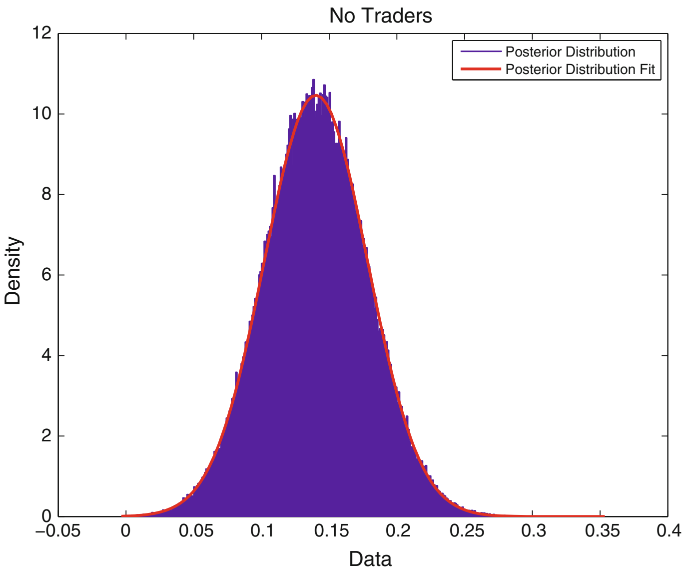

Fig. 6 Posterior distribution for the parameter i for the model with $\{T^L = 50, T^H = 50\}$

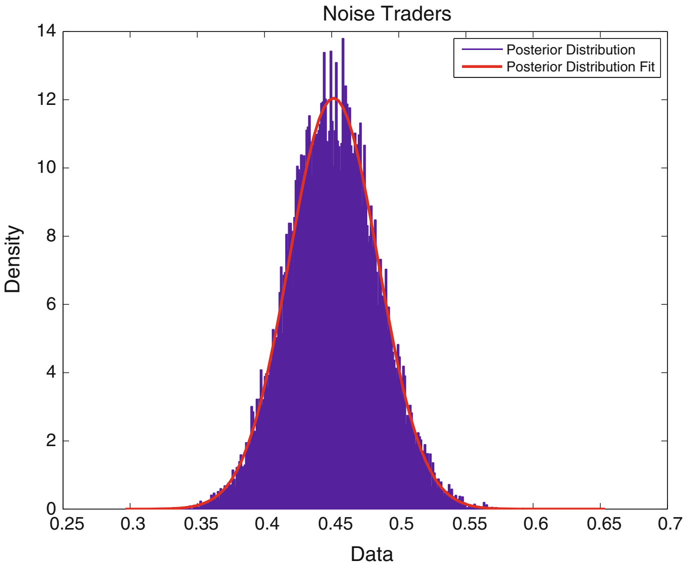

Fig. 7 Posterior distribution for the parameter n for the model with $\{T^L = 15, T^H = 85\}$

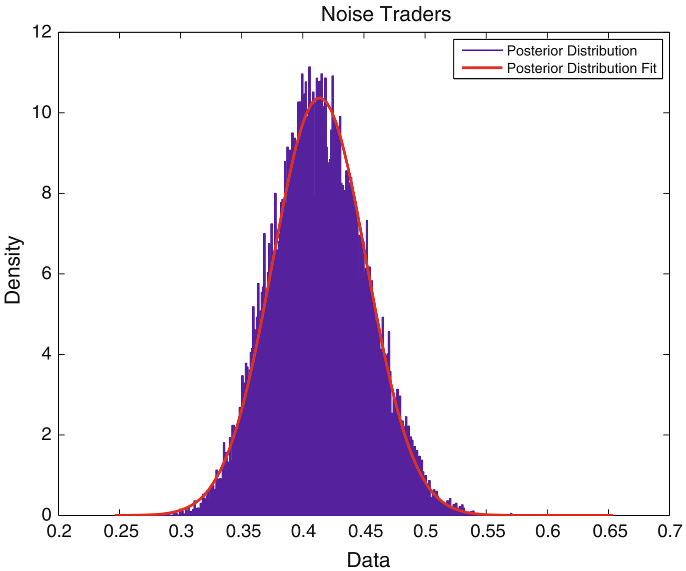

Fig. 8 Posterior distribution for the parameter n for the model with $\{T^L = 30, T^H = 70\}$

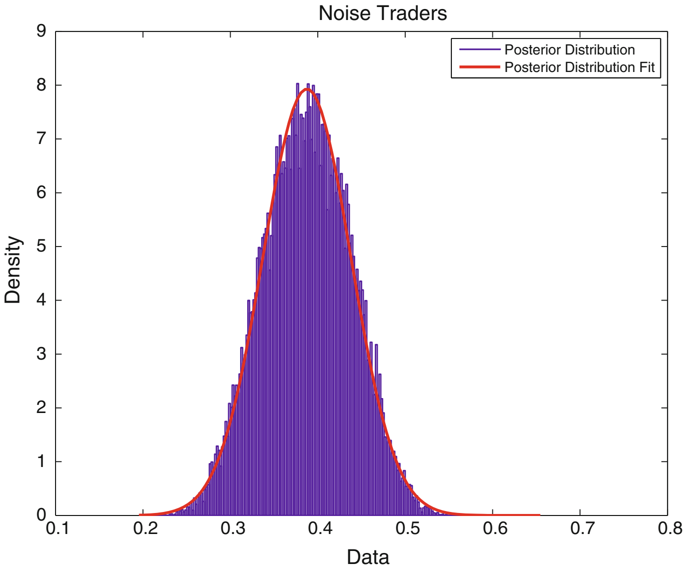

Fig. 9 Posterior distribution for the parameter n for the model with $\{T^L = 50, T^H = 50\}$

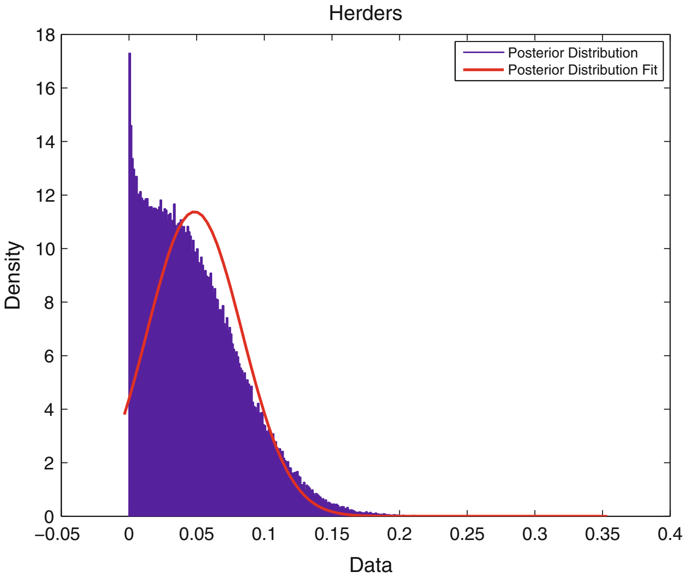

Fig. 10 Posterior distribution for the parameter h for the model with $\{T^L = 15, T^H = 85\}$

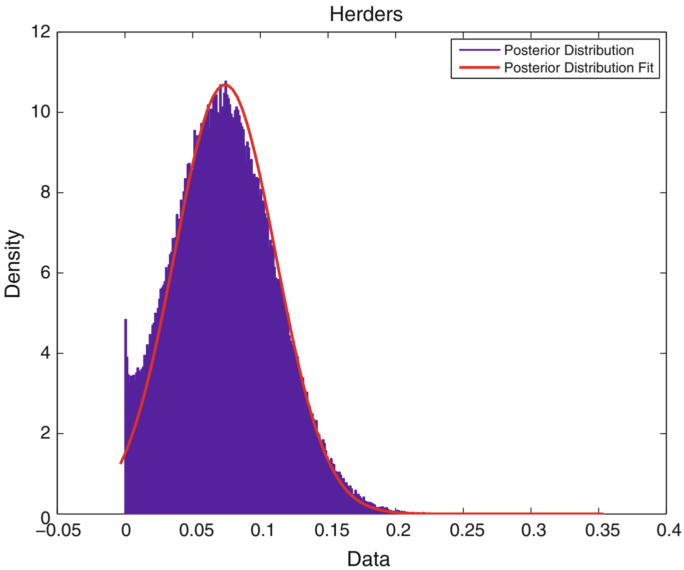

Fig. 11 Posterior distribution for the parameter h for the model with $\{T^L = 30, T^H = 70\}$

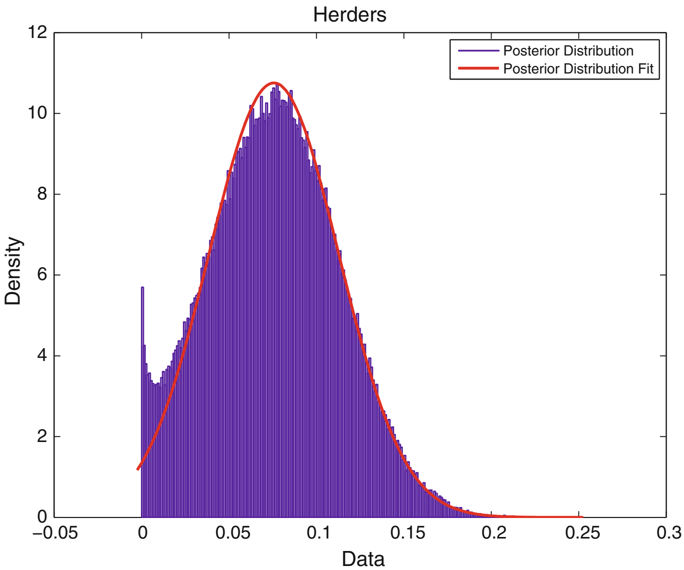

Fig. 12 Posterior distribution for the parameter h for the model with $\{T^L = 50, T^H = 50\}$

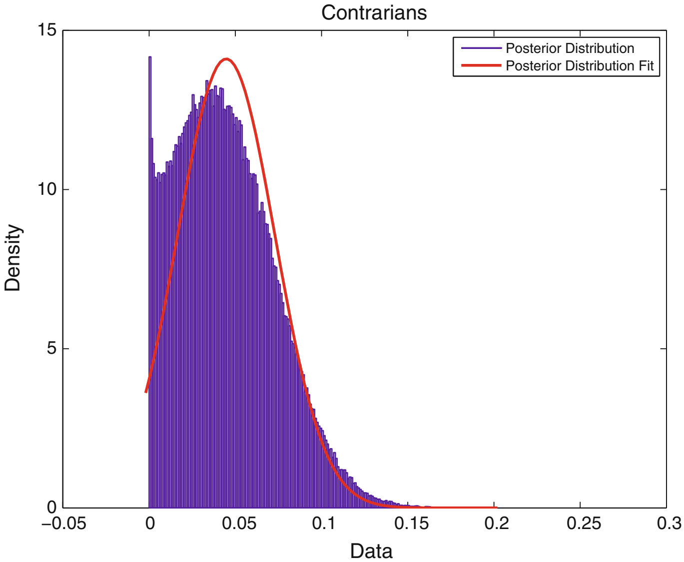

Fig. 13 Posterior distribution for the parameter c for the model with $\{T^L = 15, T^H = 85\}$

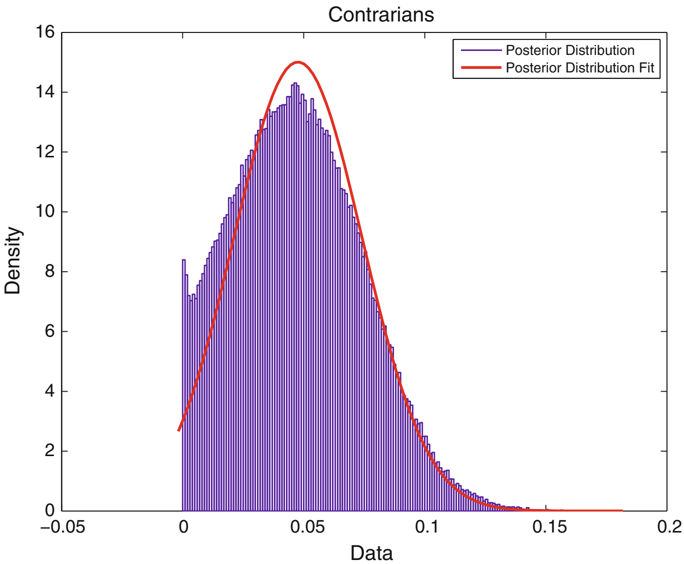

Fig. 14 Posterior distribution for the parameter c for the model with $\{T^L = 30, T^H = 70\}$

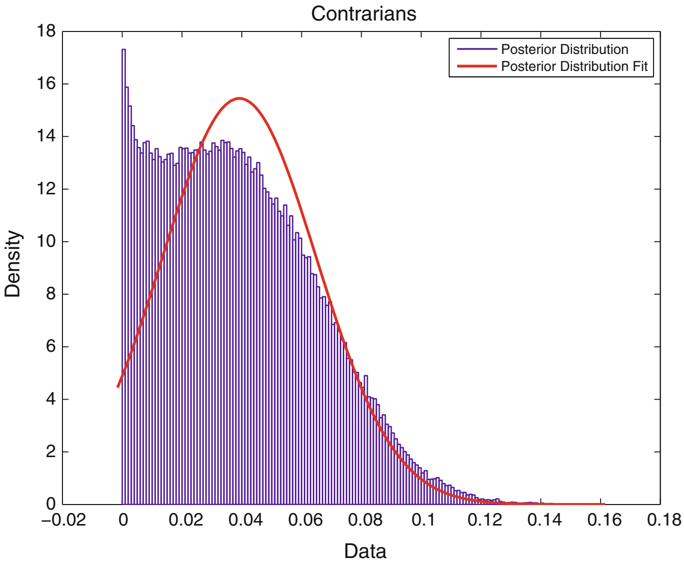

Fig. 15 Posterior distribution for the parameter c for the model with $\{T^L = 50, T^H = 50\}$

References

Avery C, Zemsky P (1998) Multidimensional uncertainty and herd behavior in financial markets. Am Econ Rev 88(4):724–748

Bolstad W (2010) Understanding computational Bayesian statistics. Wiley, New York

Cipriani M, Guarino A (2005a) Herd behavior in a laboratory financial market. Am Econ Rev 95(5): 1427–1442

Cipriani M, Guarino A (2005b) Noise trading in a laboratory financial market: a maximum likelihood approach. J Eur Econ Assoc 3(2–3):315–320

Gelman A, Carlin JB, Stern HS, Rubin DB (2003) Bayesian data analysis, 2nd edn. CRC Press, London

Kass R, Raftery A (1995) Bayes factors. J Am Stat Assoc 90(430):773–795

Koop G (2003) Bayesian econometrics. Wiley, New York

Raftery A, Newton M, Satagopan J, Krivitsky P (2006) Estimating the integrated likelihood via posterior simulation using the harmonic mean identity. Memorial Sloan-Kettering Cancer Center, Department of Epidemiology and Biostatistics WP Series, p 6