Agent-Based Analysis of Asset Pricing under Ambiguous Information

Ben-Alexander Cassell and Michael P. Wellman
Computer Science & Engineering
University of Michigan
Ann Arbor, MI 48109-2121 USA
{bcassell,wellman}@umich.edu

Keywords: agent-based finance, empirical game theory, equity premium

Abstract

In a representative agent model, the behavior of a social system is described in terms of a single aggregate decision maker. Such models are popular in economic and finance research, largely due to their analytic tractability, but fail to account for real-world agent heterogeneity. Agent-based simulation models naturally incorporate such heterogeneity, and we exploit this capability to investigate a recent model from the finance literature proposed by Epstein and Schneider (ES), and its ability to explain the classic equity premium puzzle in risky asset pricing. In addition to the *ambiguity-averse* trading strategy adopted by the representative agent in the ES model, we consider simple Bayesian strategies. Rather than impose a particular strategy profile, we employ an empirical game-theoretic approach to derive stable market compositions among the set of candidate strategies.

For most market configurations that we examined, ambiguity-averse pricing was not present in equilibrium support. We do, however, find ambiguity-averse pricing in equilibrium support for a market configuration analogous to an illiquid asset. For none of the market configurations that we examined were we able to find significant equity premia. Both our use of strategic equilibrium as a market composition concept, and the actions of our simulated market microstructure contribute to removing any equity premium. These findings underscore the need to verify that results from abstract representative-agent models are supportable in a higher-fidelity model where heterogeneity and strategic interactions are taken into account.

1. INTRODUCTION

As our computational powers increase, high-fidelity simulation becomes an increasingly attractive option for investigating implications of alternative agent behaviors. In simulation we can conduct controlled experiments, and modulate environmental assumptions in a fine-grained manner. Despite the many benefits of the *agent-based modeling* (ABM) approach, its adoption has been slow in the mainstream economic and finance research communities. Research problems in these fields are often amenable to agent-based simulation, since macroeconomic variables and asset prices are derived through repeated agent interaction, yet Leombruni and Richiardi (2005) estimate that ABM accounts for less than .03% of research published in the top 20 economic journals.

Leombruni and Richiardi attribute this to the prevalent belief that agent-based models are "(i) difficult to interpret and generalize, and (ii) difficult to estimate". The difficulty of generalization stems in part from the potential sensitivity of results to the particular configurations of agent types and strategies chosen by the modeler. ABM designs inherently allow more degrees of freedom than a representative agent model, and none produce the sort of crisp results obtainable from an analytical model; however, they do provide a natural way of aggregating the effects of heterogeneous agent behavior through simulating market microstructure, a facet often overlooked by analytical models.

We propose an ABM methodology that addresses concerns of generalization. First, rather than impose a profile of agent strategies, we employ game-theoretic analysis. For a market domain, we simulate the market of interest to construct a game model. Equilibria of this game are taken as the expected market composition with respect to the pool of candidate strategies. We can then estimate outcome variables of interest with respect to the expected market composition through repeated simulation and the application of variance reduction techniques.

Following some relevant background material and general description of our approach, we present a case study. We investigate a recent model of ambiguity aversion in asset pricing and its ability to resolve a famous problem in finance research: the equity premium puzzle. We find that when our chosen markets are simulated at equilibrium, neither ambiguity aversion, nor risk aversion, substantially resolve the equity premium puzzle.

2. FINANCE BACKGROUND

2.1. Representative Agent Models

Despite a burgeoning literature in agent-based financial modeling, it remains almost *de rigeur* in mainstream academic finance research to model aggregate behavior in financial markets analytically with a single representative agent. Such agents represent in aggregate form the preferences of all the agents in the market, and generally are assumed to hold all of the assets in a market in equilibrium. From these abstractions, pricing and other behavior is derived. Chapman and Polkovnichenko (2009) have recently shown that the lack of agent heterogeneity inherent in the representative agent approach may have important implications, particularly when the agent is not an expected utility maximizer. Specifically, these authors demonstrated that adding even one more agent to a market can qualitatively change the conclusions of an asset pricing study.

Although a two-agent model suffices to demonstrate the weakness of a representative agent model, it may also fail to support the agent heterogeneity required to explain market behavior. Consider a two-agent model, where one agent opts not to trade in equilibrium. Our estimated equity premium would then correspond to half of the pool of traders not participating in that market. This level of granularity may inaccurately estimate equity premium since no trades are conducted and only half of the market is participating in price formation. Additionally, if our two-agent model corresponds to two different pricing strategies it may not be reasonable to assume that the representative agent opting out in equilibrium implies that no agents of that strategy would participate in the market we are trying to model.

Because a single representative agent sets prices directly, such models need not consider how prices are actually determined by agent interaction through market mechanisms; however, this can lead to aggregate pricing that is inconsistent with real-world market operation. Consider the agent that has the lowest value among agents participating in a continuous double auction for a risky asset; if the market has a large number of orders outstanding then this agent's bids are never seen by the market. His buy price is too low to reach the top of the order book and his sell orders match at the higher market bid. We can thus see that in a thick market, the prices of agents with low values for an asset should not factor into calculations of equity premium, since their prices are never seen by the market. With as few as three agents, market microstructure can affect the allocation of assets, and therefore should not be ignored out of hand, as this, in turn, can affect market pricing, efficiency, and other outcomes.

2.2. Agent-Based Modeling

It might seem the way forward is to construct analytical models with greater numbers of agents; however, as the number of agents increases so does the complexity of analysis. Additionally, the incorporation of market microstructure can cause the complexity of the environment description to increase along with population granularity. Beyond a point, the only feasible computational approach is bottom-up simulation of agent behavior.

Despite limited recognition by mainstream economic journals, ABM approaches have been employed extensively in the social sciences, including finance. LeBaron (2006) surveys the agent-based finance literature, and discusses the motivations and limitations of the approach.

2.3. Ambiguity Aversion and the Equity Premium Puzzle

In standard models of asset pricing, investors demand a higher rate of return as the risk of an asset increases. This is a direct consequence of risk-averse utility: given a choice among two assets with the same expected value, the one with lower risk provides greater expected utility. This increase of rate of return on risky assets as opposed to risk-free alternatives is known as the equity premium. The *equity premium puzzle* refers to the apparent disparity between the observed equity premium and what would be predicted based on current models of investor risk preferences. If investors were not risk averse, classical economic theory would suggest that the price of stocks would rise until the point where expected return on stock was exactly equal to the expected return on risk-free assets. Since investing in stock has inherent risk, and traders are not in general risk neutral, stock prices should be lower than this point; however, even when accounting for risk aversion, the average return on stock is significantly higher than the return on treasury bills. Since this phenomenon was identified by Mehra and Prescott (1985), it has received a great deal of attention in finance research literature. DeLong and Magin (2009) provide an up-to-date survey characterizing the current state of knowledge and debate surrounding the U.S. equity premium.

One path taken by economists to explain the equity premium puzzle is to posit forms of non-standard preferences or decision rules. One example is *ambiguity aversion*, a cognitive phenomenon famously identified by Ellsberg, wherein decision makers prefer actions where the chance elements are objectively clear, even at significant sacrifice of expected utility (Halevy, 2007). Epstein and Schneider (2008) (ES) argue that this aversion can be justified in a dynamic context, when information quality is taken into account. The quality of information revealed under an ambiguous prospect may be less useful, in proportion to the degree of ambiguity. To incorporate this ambiguity, ES extend the work of Gilboa and Schmeidler (1989) who demonstrate that preferences for known risks can be captured by worst-case reasoning over non-unique priors. ES show that worst-case reasoning amounts to an asymmetric response to information depending on its content, since the worst case when receiving positive news is that it is not very informative about future dividend movement while worst case when receiving negative news is that it is very informative about future dividend movement. With this rationalization they develop a model of asset pricing with an ambiguity-averse representative agent, and demonstrate that this can explain why even a market of well-diversified investors may still demand compensation for the idiosyncratic risk associated with each asset they hold, since diversification does not reduce ambiguity in the same way that it mitigates risk. This in turn could explain an equity premium, even among savvy investors, and other phenomena of interest.

3. EMPIRICAL GAME MODEL OF ASSET PRICING

Taking the ES model as a starting point, we seek to address two questions. First, given the possibility of multiple strategies, is the ambiguity-averse strategy actually present in equilibrium? If traders gain no benefit from being averse to ambiguity, we would expect traders who are averse to ambiguity to be displaced by those who are not, calling into question the validity of modeling the whole market as a single ambiguity-averse trader. Second, in a model with agent heterogeneity and an active market mechanism, does pricing according to the ES model generate significant equity premium? To answer these questions, we constructed an ABM for asset pricing, and performed empirical game-theoretic analysis to evaluate strategy candidates. The full model includes elements that specify the market mechanism, asset definition, and agent strategies.

3.1. Market and Asset Models

Many representative agent models, including ES, give little consideration to market microstructure. When analyzing price formation via ABM, however, it becomes necessary to specify a mechanism by which the market operates. In many stock markets, agent interaction is mediated through a *continuous double auction* (CDA) (Friedman and Rust, 1993), and we adopt this mechanism in our investigation. In a CDA, both buyers and sellers submit prices and are matched continuously. Following ES, the model we examine consists of two types of assets, one offering a fixed return and the other offering a variable return. Agents are able to exchange these two assets through the CDA by specifying how much of the risk-free asset they are willing to offer or accept for a share of the risky asset.

3.2. Agent Strategy Composition

Given a market mechanism, we must also specify how agents will act within the mechanism to determine prices. The equilibrium bidding strategy for a CDA (or any dynamic market mechanism) in this context is unknown, thus we must evaluate a space of candidates to determine an appropriate composition of agent strategies in the model. We adopt a version of an approach called *empirical game-theoretic analysis* (EGTA). The EGTA framework (Wellman, 2006) performs agent-based simulation to generate sample payoffs for candidate strategy profiles, and from them induces a game form. The learned game model then serves as the basis for game-theoretic analysis, which identifies stable strategy profiles (e.g., Nash equilibria). In contrast to other popular equilibria discovery approaches, such as social learning (Ellison and Fudenberg, 1993), reinforcement learning (Hu and Wellman, 2003; Littman, 1994), or genetic programming (Chen, Duffy, and Yeh, 2002), finding Nash equilibria from an empirical game requires no assumptions about how agents arrive at equilibria, which by definition are outside the environment specification of the representative agent models that inform our study.

Our set of strategy candidates starts with the representative agent employed in the ES model. This strategy calculates an asset price using the ES formulation of ambiguity aversion (AA), given its own private information. The strategy then bids in a straightforward manner based on that price. We then add a second strategy candidate, based on a standard Bayesian (B) pricing model, that likewise bids straightforwardly given its price calculation. These two strategies are then parameterized by whether or not they incorporate risk aversion into their pricing strategy. We offer the following justification for this strategy parameterization: even if an agent is averse to the risk presented by short term proposition of whether to hold shares of the risky asset, pricing according to the agents true value may not be expected-utility maximizing. If, as a trader, I observe that there is a significant equity premium enforced by the market, my expected earnings are increased by ignoring my risk aversion, since I am always able to acquire the asset at a discount relative to its expected return.

There is a huge space of alternatives to these strategies, varying both the method of pricing, and mapping of prices to actual bids over time. Nevertheless, we focused on these strategies, given that our primary goal is to scrutinize the hypothesis that AA behavior is a plausible basis for modeling asset prices. Intuitively, for the AA strategy to be tenable, it should be minimally competitive with the natural alternative, B. Adding strategies to this mix only makes for a more stringent test.

3.3. Estimating the Empirical Game

Our construction of an empirical game through simulation follows the basic steps outlined by Jordan, Kiekintveld, and Wellman (2007) for EGTA:

1. Approximate the full game by reducing the effective number of players.

2. Run simulations covering all distinct strategy profiles.

3. Apply variance reduction techniques to estimate the outcomes for each strategy profile in the reduced game.

4. Search for equilibria in the resulting empirical game.

Our market simulation incorporates $N = 120$ agents. With the pool of four possible strategies, there are $\binom{123}{3} = 302,621$ distinct profiles, taking into account that agents' roles are symmetric. Sampling from all of these profiles would be unmanageable; therefore, to contain the profile space further, we employ the hierarchical reduction proposed by Wellman et al. (2005b), where multiple agents are constrained to adopt the same strategy. Specifically, we analyze a four-player version of the game, where each player selects a strategy to be played by 30 trading agents. With this reduction, there are only $\binom{7}{3} = 35$ distinct strategy profiles to sample. Although we have restricted the strategic degrees of freedom (thus sacrificing some fidelity), we retain agent heterogeneity in terms of beliefs and preferences at the full granularity of 120 agents.

Through averaging repeated simulations, we generate estimates of the expected payoff of playing each strategy profile in the underlying game. As in prior EGTA studies (Wellman et al., 2005a) as well as recent analyses of computer poker (White and Bowling, 2009), we apply variance reduction techniques in order to obtain statistically meaningful results with a feasible number of samples. For the experiments conducted in this paper, we employ the method of control variates (Lavenberg and Welch, 1981), though other traditional Monte Carlo tools may also provide significant variance reduction.

Given an empirical game model, we can apply standard game-theoretic solution concepts, such as Nash equilibrium. For these experiments we used *replicator dynamics* (Schuster and Sigmund, 1983) to derive symmetric mixed-strategy ϵ-Nash equilibria. Replicator dynamics considers a population playing strategies according to a distribution that is iteratively updated as a function of each strategy's performance against the current distribution. Friedman (1991) demonstrated that fixed points of this iterative process correspond to Nash equilibria with respect to the fitness function used to evaluate performance.

We examine markets where agents are risk averse to varying degrees. Agents in our simulations exhibit constant relative risk aversion (CRRA) according to the utility function used by Mehra and Prescott (1985), $u(c) = \frac{c^{1-\alpha}-1}{1-\alpha}$. For $\alpha = 0$ this utility function collapses to $c-1$, and for $\alpha = 1$ utility is defined by $\lim_{\alpha\to 1} u(c) = \log(c)$. Regardless of whether or not agents incorporate risk aversion into their pricing, when aggregating agent payoffs each agent's payoff is calculated with respect to their CRRA utility function. This allows us to examine whether risk-averse pricing is found in equilibrium when agents can strategically ignore their aversion, and determine what effect risk aversion should have on the measured equity premium.

4. ASSET PRICING UNDER AMBIGUOUS INFORMATION

Our case study addresses the scenario of asset pricing under ambiguous information investigated by Epstein and Schneider (2008). These authors model pricing of a risky asset (i.e., an equity security, or stock) by a representative agent that is averse to ambiguity. Though there are several available models of ambiguity aversion, we chose the ES formulation because it is prevalent in recent work in finance research.

In our simulations, in each trading period, the agents—*traders*—are given a piece of news that is partially informative about future dividend payments. Traders use these signals to update their belief about the current value of stocks. With updated beliefs, the traders submit some constant number of single-unit limit orders to the order book, where they are processed according to CDA market rules. At the end of each quarter, q, traders receive dividend payments on the stock they own, and interest payments on their cash balance. Each quarter is defined to be some constant number of trading periods. We concern ourselves with equity premium estimates on the length of a year for consistency with prior work, where a year is defined to be four quarters. The remainder of this section provides further detail and describes the important assumptions of the simulation model.

4.1. Market Conditions

Our market consists of two types of assets. A risk-free asset (cash) yields quarterly interest payments at a fixed rate, r. A risky asset (stock) yields quarterly dividend payments d_q, which fluctuate according to the mean-reverting process used by ES:

$$d_q = \kappa\bar{d} + (1-\kappa)d_{q-1} + \delta_q,$$

where δ_q is a shock to the dividend value at quarter q, $\delta_q \sim Normal(0, \sigma_\delta^2)$, and $\kappa \in (0,1)$ is the degree of reversion toward the mean dividend value, $\bar{d}$. If the dividend value drops below zero, traders are paid no dividend, though future dividend movement continues from the subzero value.

4.2. News

News is sent to all traders in every trading period, and every trader receives the same piece of news. A piece of news consists of a signal, s_t, representing intangible information available about the asset, and the most recent dividend payment, d_q. The signal is partially informative about the next dividend payment,

$$s_t = \delta_{q+1} + \epsilon_t,$$

where δ_{q+1} is the shock in the dividend that will be observed at the next dividend payment, and ϵ_t is the noise in the signal, $\epsilon_t \sim Normal(0, \sigma_\epsilon^2)$, $\sigma_\epsilon^2 \in [\underline{\sigma}_\epsilon^2, \bar{\sigma}_\epsilon^2]$. For our experiments,

$[\underline{\sigma}_\varepsilon^2, \bar{\sigma}_\varepsilon^2]$ constrain the space of possible beliefs about the news-generating process. As in ES, traders are assumed unable to learn the true news-generating process but can use news to improve their estimates of δ_{q+1} and update their prices.

4.3. Traders

Trader i is given two private values, $\underline{\sigma}_{\varepsilon,i}^2$ and $\bar{\sigma}_{\varepsilon,i}^2$, representing the trader's beliefs about the lower bound and upper bound of possible variance in the noise added to signals. $\underline{\sigma}_{\varepsilon,i}^2$ and $\bar{\sigma}_{\varepsilon,i}^2$ are drawn uniformly from $[\underline{\sigma}_\varepsilon^2, \bar{\sigma}_\varepsilon^2]$, $\underline{\sigma}_{\varepsilon,i}^2 < \bar{\sigma}_{\varepsilon,i}^2$. Each agent is also assigned a coefficient of risk aversion, α_i, and seeks to optimize its CRRA utility. Traders incorporate these private values into their valuation of asset, leading to a market that is heterogeneous in price.

All pricing strategies considered in our experiment calculate approximations of a *certainty equivalent price* (CEP), that is, the price at which the trader is indifferent to acquiring an additional share of the risky asset. To calculate the CEP, traders find the price, p_Δ such that

$$EU(h_{cash}, h_{shares}) = EU(h_{cash} - p_\Delta, h_{shares} + 1), \qquad (1)$$

where h_{cash} and h_{shares} are the trader's current holdings in cash and shares respectively. $EU(h_{cash}, h_{shares})$ is the expected utility function for the trader and is given by

$$\int_{-\infty}^{\infty} [u(h_{cash} + \frac{h_{shares}}{1+r}(p^* + \kappa\bar{d} + (1-\kappa)d_q + x))]dF_{\delta_q}(x),$$

where $dF_{\delta_q}(x)$ is the derivative of the believed cumulative distribution function for dividend shock and p^* is an estimate of next period's price for the risky asset.

Since δ_q is Gaussian, there is no closed-form solution to (1), and we must numerically integrate. For our experiments we set p^* to the traditional Capital Asset Pricing Model (CAPM) price, $\frac{\bar{d}}{r}$. The intuition for choosing this value is that the market determines the price of the asset in the future, and in the absence of knowledge about other market participants we should expect future prices to reflect the long term expected value of the asset, the CAPM price. We also bound the integral at the mean $\pm$ 4 standard deviations to capture 99.99% of possible outcomes. Without these bounds the integral is undefined since the CRRA utility function tends to negative infinity as $c \to 0$. Finally, we calculate p_Δ through a fixed point iteration. A similar approach is used to calculate sell prices, replacing $h_{shares}+1$ with $h_{shares}-1$. The next two subsections describe how the two families of pricing strategies utilize their private values to modify their CEP estimate.

For our experiments, each trader constructs a price schedule for buying and selling a constant number of units of the risky asset in each trading period. When generating prices for shares beyond the first that an agent wishes to acquire (or sell) in a trading period, they assume that the previous share was

Action	Price in Cash
Buy 3^{rd} share	28.80
Buy 2^{nd} share	29.50
Buy 1^{st} share	30.00
Sell 1^{st} share	30.40
Sell 2^{nd} share	30.70
Sell 3^{rd} share	30.95

Table 1. Example price schedule for a risk averse trader

acquired (or sold) at the price that they submitted to the market for that share. Sufficiently risk-averse traders thus generate a pricing schedule akin to that exhibited in Table 1. A bid to sell is valid only if the agent currently holds at least that quantity of the risky asset, but agents may borrow at the risk-free rate of interest in order to buy risky assets. Orders remain on the order book until they are filled or the agent submits a replacement bid. We simulate a limited form of asynchrony within a trading period by sequencing the orders randomly for submission to the CDA mechanism.

4.3.1. Ambiguity-Averse Pricing

AA traders are willing to accept that any variance in noise in the range $[\underline{\sigma}_{\varepsilon,i}^2, \bar{\sigma}_{\varepsilon,i}^2]$ is possible, and therefore maintain a set of priors. When ambiguity-averse traders receive a piece of news, they update their beliefs of the distribution of next period dividend shock, F_{δ_q}, using Bayesian updating over their set of priors. After updating their beliefs, AA traders price according to the posterior distribution that minimizes their expected value of holding the asset.

In addition to maintaining a set of priors, AA traders are also distinguished by enforcing a premium for enduring future ambiguous news. As in ES, AA traders discount their CEP by

$$[\gamma(\underline{\sigma}_{\varepsilon,i}^2) - \gamma(\bar{\sigma}_{\varepsilon,i}^2)]\frac{\sigma_\delta}{r\sqrt{2\pi\gamma(\bar{\sigma}_{\varepsilon,i}^2)}}.$$

This discount depends on on the trader's private values but is constant for the trader through a simulation run.

4.3.2. Bayesian Pricing

Bayesian traders treat all news as equally informative and thus respond symmetrically to news regardless of its content. For our experiments Bayesian traders use the average of $\bar{\sigma}_{\varepsilon,i}^2$ and $\underline{\sigma}_{\varepsilon,i}^2$ to form a single value for $\sigma_{\varepsilon,i}^2$. Thus Bayesian traders hold only a single prior for performing Bayesian updating and do not enforce a premium for enduring ambiguous news.

5. EXPERIMENTS

We examined several possible market parameterizations to investigate equity premium. Since the space of market parameterizations is continuous and factorial, such an investigation

Parameter	Setting
Q	40
κ	0.05
$[\underline{\sigma}_\varepsilon^2, \bar{\sigma}_\varepsilon^2]$	[0.0, 0.1]
σ_ε	0.01
σ_δ	0.01
$\bar{d}$	0.5
r	0.01
cash endowment	1000
stock endowment	20
trading periods per quarter	10
schedule size	5
α	$\sim Normal(1.0, 0.2)$

Table 2. Base market configuration for experiments.

Name	Modifications from BASE
1T1N	trading periods per quarter = 1, schedule size =1
HSHK	$\sigma_\delta = 0.05$, $\kappa = 0.2$
LAMB	$[\underline{\sigma}_\varepsilon^2, \bar{\sigma}_\varepsilon^2]$ = [0.0, 0.01]

Table 3. Other market configurations tested.

cannot hope to be exhaustive. Instead we aimed to consider a small set of reasonable and representative market configurations.

Some guidance for parameter setting can be taken from the literature. For example, Exley, Mehta, and Smith (2004) cover various measures of mean reversion fairly comprehensively. The distribution for α, the coefficient of risk aversion, was chosen to be consistent with the survey of behavioral studies in Mehra and Prescott (1985). The range of possible ambiguity in news, however, may be impossible to estimate from a finite sample, and thus there is little guidance available for setting this parameter. This difficulty in determining the news generating process is one of the underpinnings of the ES model, so the lack of established parameter settings is not surprising. In the absence of authoritative prescriptions for parameter settings, we set parameters to create potentially interesting scenarios. Table 2 presents the base market configuration that we considered, labeled as BASE throughout, while Table 3 summarizes the other configurations we examined, with respect to how they deviate from BASE.

To generate payoff matrices we used a four-player hierarchical reduction from a 120-player game. For each market configuration, we sampled 30 times from each of the 35 strategy profiles. Once the outcome samples are recorded, we applied the method of control variates to estimate the outcomes for each strategy profile. Payoff matrices are then constructed from the estimated outcomes.

	B:noRA	B:RA	AA:noRA	AA:RA
BASE	0.657	0.343	0	0
1T1N	0	0.134	0.345	0.520
HSHK	0	1.0	0	0
LAMB	0	1.0	0	0

Table 4. Equilibria found through replicator dynamics.

5.1. Feasibility

The first test of a representative agent pricing model is whether or not we would expect any agents to be adopting that strategy in equilibrium. For each market configuration we used replicator dynamics to find a stable equilibrium, using a population that is initially distributed uniformly among the possible strategies. When a stable equilibrium is observed that features only one strategy or one family of strategies, we consider it feasible to construct a representative agent model of the market.

Table 4 presents the symmetric Nash equilibrium that we found for each market configuration. In the table, "noRA" signifies that the strategy does not incorporate risk aversion in generating its prices. Every strategy considered had a non-negligible presence in equilibrium for at least one market configuration. For HSHK and LAMB, pure strategy equilibria were found for the strategy B:RA. Additionally, the equilibrium discovered for BASE is comprised solely of Bayesian pricing strategies. For these market configurations a representative agent model could be constructed.

We found no such equilibria where ambiguity aversion constituted a feasible pricing model when the market is in strategic equilibrium. The one configuration for which we found an equilibrium featuring ambiguity aversion is 1T1N. This market configuration both restricts the traders' ability to satisfy their preferences (through limiting them to one share bought and one share sold per quarter), and limits the traders' attempts to improve their estimates of the next period's dividend movement (by giving them only one piece of news per quarter). This scenario seems in some sense contrived, but may be analogous to markets that are very thin. For such markets, trades occur infrequently, and the news coverage of assets in these markets can be expected to be less than popular stocks. Longstaff (2005) argues that many important classes of assets are illiquid in this sense. Since the interaction between market configuration and equilibrium composition is difficult to discern a priori, having multiple reasonable strategies makes it hard to determine which pricing strategy should be the basis of a representative agent model without first running this simulation-based analysis.

The Bayesian pricing that incorporated risk aversion was present in all of the equilibria we found. Though our exploration of the strategy space was not, and could not have been, exhaustive, this presents support for truthfully reporting your value when submitting limit orders to the market under certain conditions.

5.2. Equity Premium Estimation

For each market configuration we constructed equity premium estimates through simulating those markets with the equilibrium composition. Since we are concerned with calculating equity premium in equilibrium conditions, the simulator is used in this phase to aggregate repeated market interactions over a variety of markets and trader states.

Equity premium per quarter, P_q, is calculated as the difference between the return on the risky asset and the return on the risk-free asset over the quarter

$$P_q = \frac{\bar{p}_q + d_q}{\bar{p}_{q-1}} - (1+r).$$

Essentially we are comparing the return of \$1 invested in the risky asset to \$1 invested in the risk-free asset over the period. In the rare event that no trades take place in a given quarter, we make no observations of equity premium. We then estimate the equity premium per year by multiplying the quarterly average by four.

Figure 1 presents our equity premia estimations at the discovered equilibria. The highest equity premium found for any equilibrium composition was less than 0.05%. Comparing these values to the equity premium estimates from U.S. stock market data presented by Fama and French (2002), between 2.55 and 7.43%, we see that none of these constitute a significant equity premium. This is not really surprising as our use of equilibrium as a basis for exploration implies that there are no beneficial strategic deviations. If there were significant a equity premium, then some agent can benefit from deviating. Additionally, the use of the CDA in a market with enough traders is going to restrict the ability of low prices to be seen by the market. Finally, for the trader behavior we simulated, traders with low values for the risky asset are eventually forced out of the market, since they will quickly sell off all of their assets and their buy prices are unlikely to be competitive. These factors, and possibly others, appear to have driven out even the equity premium that would have been predicted by a simple risk-averse representative agent model, making the equity premium puzzle yet more puzzling.

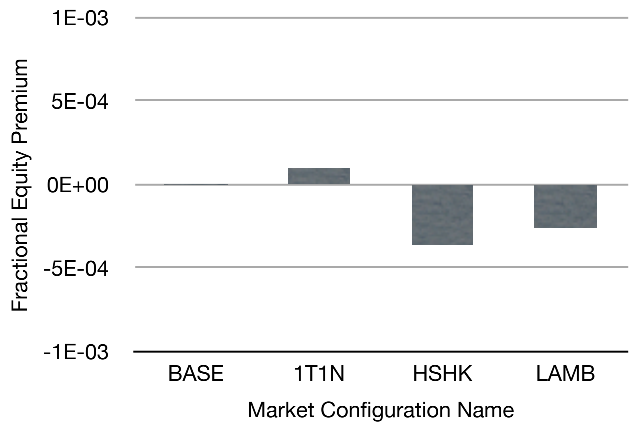

Figure 1. Calculated Equity Premium.

6. DISCUSSION

When market variables are derived through agent interaction, the estimation of these variables can be very complex. Heterogeneity of beliefs and pricing strategies among agents may impede meaningful aggregation of agent behavior. Realistic models of such variables should also take into account the relevant market microstructure. These aspects of real-world market interactions are often ignored by modelers in favor of analytical tractability. We address these issues through the use of an ABM methodology. This approach allows us to extend analytical pricing models to investigate the market dynamics of agents interacting through a real market mechanism. Additionally, through the use of equilibrium analysis, we address the contention that agent-based simulation is difficult to generalize. Game-theoretic concepts are familiar and accepted in mainstream economics, and should help ease concerns from that research community.

Employing empirical game-theoretic analysis allowed us to find market equilibria, abstracting from the methods by which individual agents may arrive at their respective strategies. Within this framework we suggested evaluating the feasibility of representative agent models through direct competition with other models of behavior. For our experiments we conducted this evaluation by adding pricing strategies that were obvious modifications to standard expected utility based pricing, which should constitute a litmus test for behavioral pricing models.

We applied this approach to investigate a model of ambiguity aversion presented by Epstein and Schneider. The empirical game-theoretic analysis suggested that their model of ambiguity-averse pricing can be an equilibrium strategy for certain market configurations, but it is not in general a feasible model of trader behavior in our simulated markets. Additionally, neither ambiguity-aversion nor moderate levels of risk-aversion were found to generate significant equity premia when the market is at a strategic equilibrium. The greatest equity premium we found for this pricing strategy was considerably less than even conservative estimates of the equity premium in the U.S. stock market over the period from 1951–2000. Our experiments lend further support to the argument presented by Chapman and Polkovnichenko (2009), as heterogeneity in agents, and market microstructure, lead to no equity premium being observed for markets where a simple risk-averse representative agent model would have predicted a premium.

REFERENCES

Chapman, D. A., and Polkovnichenko, V. 2009. First-order risk aversion, heterogeneity, and asset market outcomes. *Journal of Finance* 64(4):1863–1887.

Chen, S.-H.; Duffy, J.; and Yeh, C.-H. 2002. Equilibrium selection via adaptation: Using genetic programming to model learning in a coordination game. *The Electronic Journal of Evolutionary Modeling and Economic Dynamics*.

DeLong, J. B., and Magin, K. 2009. The U.S. equity return premium: Past, present, and future. *Journal of Economic Perspectives* 23(1):193–208.

Ellison, G., and Fudenberg, D. 1993. Rules of thumb for social learning. *Journal of Political Economy* 101(4):612–43.

Epstein, L. G., and Schneider, M. 2008. Ambiguity, information quality, and asset pricing. *Journal of Finance* 63(1):197–228.

Exley, J.; Mehta, S.; and Smith, A. 2004. Mean reversion. In *Finance and Investment Conference*.

Fama, E. F., and French, K. R. 2002. The equity premium. *Journal of Finance* 57:637–659.

Friedman, D., and Rust, J., eds. 1993. *The Double Auction Market: Institutions, Theories, and Evidence*. Addison-Wesley.

Friedman, D. 1991. Evolutionary games in economics. *Econometrica* 59(3):637–666.

Gilboa, I., and Schmeidler, D. 1989. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics* 18(2):141–153.

Halevy, Y. 2007. Ellsberg revisited: An experimental study. *Econometrica* 75:503–536.

Hu, J., and Wellman, M. P. 2003. Nash Q-learning for general-sum stochastic games. *Journal of Machine Learning Research* 4:1039–1069.

Jordan, P. R.; Kiekintveld, C.; and Wellman, M. P. 2007. Empirical game-theoretic analysis of the TAC supply chain game. In *Sixth International Joint Conference on Autonomous Agents and Multi-Agent Systems*, 1188–1195.

Lavenberg, S. S., and Welch, P. D. 1981. A perspective on the use of control variables to increase the efficiency of Monte Carlo simulations. *Management Science* 27(3):322–335.

LeBaron, B. 2006. Agent-based computational finance. In Tesfatsion, L., and Judd, K. L., eds., *Handbook of Agent-Based Computational Economics*. Elsevier.

Leombruni, R., and Richiardi, M. 2005. Why are economists sceptical about agent-based simulations? *Physica A: Statistical Mechanics and its Applications* 355(1):103–109.

Littman, M. L. 1994. Markov games as a framework for multi-agent reinforcement learning. In *Eleventh International Conference on Machine Learning*, 157–163.

Longstaff, F. A. 2005. Asset pricing in markets with illiquid assets. In *AFA 2006 Boston Meetings*.

Mehra, R., and Prescott, E. C. 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15(2):145–161.

Schuster, P., and Sigmund, K. 1983. Replicator dynamics. *Journal of Theoretical Biology* 101:19–38.

Wellman, M. P.; Estelle, J.; Singh, S.; Vorobeychik, Y.; Kiekintveld, C.; and Soni, V. 2005a. Strategic interactions in a supply chain game. *Computational Intelligence* 21:1–26.

Wellman, M. P.; Reeves, D. M.; Lochner, K. M.; Cheng, S.-F.; and Suri, R. 2005b. Approximate strategic reasoning through hierarchical reduction of large symmetric games. In *Twentieth National Conference on Artificial Intelligence*, 502–508.

Wellman, M. P. 2006. Methods for empirical game-theoretic analysis (abstract). In *Twenty-First National Conference on Artificial Intelligence*, 1552–1555.

White, M., and Bowling, M. 2009. Learning a value analysis tool for agent evaluation. In *Twenty-First International Joint Conference on Artificial Intelligence*, 1976–1981.