Beliefs, Shocks, and the Emergence of Roles in Asset Markets: An Agent-Based Modeling Approach

Evan Albers*
Khoury College of Computer Sciences
Northeastern University
Boston, MA, United States
albers.e@northeastern.edu

Mohammad T. Irfan
Department of Computer Science
Bowdoin College
Brunswick, ME, United States
mirfan@bowdoin.edu

Matthew J. Botsch
Department of Economics
Bowdoin College
Brunswick, ME, United States
mbotsch@bowdoin.edu

ABSTRACT

Although predictive AI models have grown to dominate computational finance, they are often limited in their applications when it comes to studying interventions and explaining behavioral outcomes. Financial economics, on the other hand, has a rich history of analytical approaches to asset-pricing theory, often requiring sweeping assumptions. In this paper, we construct an agent-based model of asset markets that is able to dispense with onerous restrictions on agent behaviors and beliefs, while having analytical validity and providing insights into the functioning of asset markets. In particular, we evaluate our models with respect to several traditional financial economic theories like Tobin's separation theorem and the capital asset pricing model (CAPM). We devise a network representing trades to show the emergence of different roles played by the agents. We study interventions, such as shocks, and explain the outcomes using our model. Finally, we investigate the effects of noise trading and show that noisy agents converge to different equilibrium points due to their differences in beliefs. Put together, this paper presents an agent-based model that can be used to study the effects of heterogeneous beliefs and risks of the agents and shocks to assets at a systemic level, thereby connecting localized agent and asset characteristics to global or collective outcomes.

KEYWORDS

Agent-Based Modeling, Computational Finance, Multi-Agent Simulation, Asset Market, Risk, Portfolio, Shocks, Noise Trading

ACM Reference Format:
Evan Albers, Mohammad T. Irfan, and Matthew J. Botsch. 2024. Beliefs, Shocks, and the Emergence of Roles in Asset Markets: An Agent-Based Modeling Approach. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 9 pages.

1 INTRODUCTION

We present an agent-based model (ABM) of asset markets grounded in financial economic theory with the goals of studying agent behaviors—e.g., the emergence of behavioral roles—and examining and explaining the effects of interventions. An ABM is chosen as the modeling paradigm because it can support heterogeneity to a degree that is typically precluded by analytical approaches, while also incorporating the theoretical foundation of modern portfolio theory. Although predictive models have been effective for tasks like price prediction in computational finance, our key focus here is performing interventions, such as introducing shocks and evaluating their effects. Additionally, we want the model to engage in a dialogue with the existing economic theory. These objectives not only justify our modeling approach but also set this study apart from prior studies.

Our study is motivated by Robert Lucas's famous critique of macroeconomic models [23]. Lucas makes the case that when policymakers change the economic environment, they need to consider how agents will respond to their changes. For example, if we make policies based on some statistical relationship we observe within an economic environment, the relationship itself will be affected by our policies due to the reaction by the agents. Our study captures the dynamics between the economic environment and agent behaviors in the context of asset markets by going beyond traditional restrictive assumptions like belief homogeneity and the absence of noise traders [13].

ABMs have a long history in economics (see [16, 22, 28, 34] for surveys). Despite this, Farmer and Axtell [14] state that "at present ABM[s] are not widely used in mainstream economics." It is hard to pinpoint exactly why this is the case, but we believe there is ample opportunity at the intersection of economics and multiagent systems.

Much of the relevant work in the literature is either too broad (e.g., modeling European economy [12], ecological econmies [17], human systems [6]), or focuses on specific components (e.g., learning and estimation [1, 21, 27, 38], momentum and reversal trading [36, 37], behavioral bias and strategies [4, 10]), or present different solution concepts altogether (e.g., out-of-equilibrium approach [2]).

Closest to ours is the Santa Fe artificial market study [21], whose designers investigate the effect of multiple trading strategies on price but include only one risky asset, thereby limiting any inference with respect to inter-asset pricing dynamics. Other works focus on the effect of specific trading strategies through a game-theoretic analysis of Nash equilibria within a market [33, 35]. Nash equilibria being computationally hard in general, these studies come at great computational cost. Additionally, the scope of these investigations is limited largely to a single asset and singular specific instances of agent behavior, with limited consideration for the more general properties that drive asset pricing.

There has also been significant use of stylized ABMs, in applications as varied as electricity markets, Treasury markets, and

*EA worked on this research as an undergraduate student at Bowdoin College.

Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), N. Alechina, V. Dignum, M. Dastani, J.S. Sichman (eds.), May 6 – 10, 2024, Auckland, New Zealand.

even local fish markets [19, 20, 31]. While these may be compelling applications, the models are too stylized to adapt here.

We demonstrate the utility of ABMs in the study of finance by reexamining standard financial economic theory using an ABM. The aim is to achieve a more comprehensive understanding of the implications of economic theory than could be achieved by analytic approaches alone. The capital asset pricing model (CAPM) is a prime candidate to demonstrate the power of ABMs in analyzing theoretical conclusions. Not only is it amenable at a large scale due to modern computational power, but it also has significant potential for increased insight. Representing an equilibrium with equations, for example, obscures any inter-agent interactions that might occur in arriving at that equilibrium. These interactions could be revealed and analyzed with an ABM.

Beyond connecting with the modern financial economic theory, we perform simulations to study what roles agents play in an asset market. For this, we devise a network representation of trades, which we call asset-flow networks. We show interesting outcomes, such as the emergence of the role of "dealers" due to agent characteristics. To our knowledge, prior ABMs on asset markets did not investigate the emergence of roles.

Furthermore, we investigate and explain the effects of interventions, such as shocks to an asset. Interestingly, we show that a large enough shock to an asset can cause a ripple effect that moves the other assets in the *same* direction as the shocked asset. This may be counter-intuitive but can be reasoned using our ABM.

We also study the effect of heterogeneous beliefs in light of the literature on noise trading [5, 11, 29]. We show that noise trading may lead to multiplicity of equilibria.

Broadly speaking, our study is an attempt to return to the first principles that drive asset pricing. We exclusively implement the principles of agent behavior that underpin modern portfolio theory [24–26] and CAPM [13] to establish a dialogue between our models and modern economic theories and to enhance our understanding of what exactly happens when one relaxes certain underlying assumptions in a systematic manner.

2 PRELIMINARIES

We will frequently use the terms *return* and *risk*. Broadly speaking, return represents the percentage gain or loss on an investment. It is treated as a random variable. Risk, usually measured as the standard deviation of returns, represents the uncertainty associated with future returns.

2.1 The Capital Asset Pricing Model (CAPM)

The CAPM is grounded in a series of assumptions regarding investors (agents), and their relationship to financial capital [13]. Agents trade assets (stocks and bonds). Their decisions are guided by two fundamental asset attributes: risk and return. Asset returns are stochastic in nature, with an expected return, and an associated standard deviation of return that represents their risk. Agents wish to maximize the amount of return they achieve, while minimizing the amount of associated risk.

The real world contains many assets. The question agents must answer is what combination of assets will maximize their utility (i.e., provide more return at less risk). Consider a two-asset world

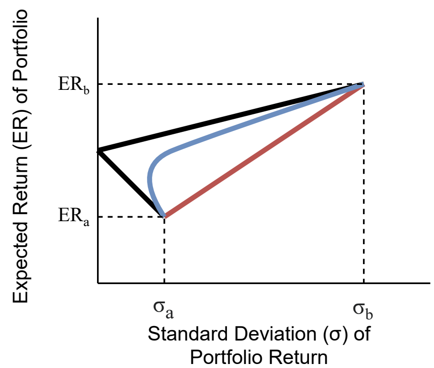

Figure 1: Three different worlds in which assets a and b are in a portfolio. Red line: a and b are perfectly positively correlated; blue: imperfectly correlated (positively or negatively); black: perfectly negatively correlated. In any world, given the proportions of a and b in a portfolio, we can compute the σ of the portfolio return, which gives us the ER of the portfolio. Notably, on the blue and black curves, it is possible to increase the ER while reducing σ (by going left).

in which agents wish to determine a set of portfolio weights that will maximize their utility. Figure 1 illustrates how portfolio risk and expected return depend on the correlation between the assets: the two assets can be perfectly positively correlated, perfectly negatively correlated, or imperfectly correlated. Perfect positive correlation (i.e., prices of the two assets change identically) results in risk and expected return increasing one-for-one, represented by the straight red line. Perfect negative correlation (i.e., a decrease in the price of one asset is perfectly matched by an increase in the other) makes it possible to eliminate all risk by choosing the apporpriate portfolio weights. This is signified by the piece-wise linear black line. Finally, imperfect correlation results in something akin to the curved blue line: portfolio risk can be reduced but not eliminated via diversification.

Asset returns being stochastic, a portfolio of assets is a weighted function of stochastic variables [13]. To determine the standard deviation of a portfolio of assets, we must know the covariance between each pair of assets. Suppose σ_a and σ_b are the standard deviations in return of assets a and b, respectively. The standard deviation σ_P in return of a portfolio P containing assets a and b can be represented by the following equation:

$$\sigma_P = \sqrt{w^2\sigma_a^2 + (1-w)^2\sigma_b^2 + 2w(1-w)Cov(a,b)}. \quad (1)$$

Here, $0 \leq w \leq 1$ is the weight parameter dictating the convex combination of assets a and b in portfolio P (i.e., w fraction of P is a and $1-w$ fraction b), and $Cov(a,b)$ represents the covariance in return of assets a and b.

Additionally, we know that covariance can be represented in terms of the correlation between two variables and their respective standard deviations as follows, where $\rho_{a,b}$ is the correlation between two assets a and b:

$$Cov(a,b) = \sigma_a\sigma_b\rho_{a,b}.$$

We can then define the standard deviation, σ_P, of the return of a given portfolio in terms of the standard deviation of and correlation between its two constituent assets, a and b.

$$\sigma_P = \sqrt{w^2\sigma_a^2 + (1-w)^2\sigma_b^2 + 2w(1-w)\sigma_a\sigma_b\rho_{a,b}}.$$

Economists [7, 15] as well as AI and multiagent systems researchers [18], [30, p. 26] are fond of saying that there is no such thing as a free lunch. However, the nature of the standard-deviation equation above suggests otherwise. Notably, the standard deviation of a portfolio is not linearly related to the standard deviations of its constituent assets. If we were to find assets that are anything but perfectly positively correlated, we can achieve a greater expected return while reducing the total quantity of risk we "pay" in terms of the standard deviation for such return (please refer to the blue and black curves in Figure 1). Therefore, diversifying non-perfectly positively correlated assets is indeed a free lunch [9].

Agents can achieve a portfolio with higher return at a lower level of risk simply by diversifying their portfolio. Therefore, in a multi-asset world, there must be some optimal portfolio. This is the portfolio that maximizes expected return relative to risk. Put another way, this is the most "cost effective" portfolio, in the sense that risk is the cost that investors pay for return, and this portfolio maximizes the return-risk ratio. At this point, it is worth considering another type of asset that is present within the market: the risk-free asset.

The risk-free asset occupies a special role in the theory of financial economics. The standard deviation of the risk-free return is zero (hence the name risk-free)—therefore it is fixed. It helps reduce investor optimization to a much simpler exercise, as we will see next.

2.2 Tobin's Separation Theorem

Tobin's separation theorem tells us how agents build their optimal portfolio [32]. Each agent first determines an optimal *risky* portfolio (i.e., optimal weights on all risky assets). They then optimize their own utility function by dividing their holdings into two potentially unequal parts: one part is invested in the optimal portfolio and the other part in the risk-free asset. Individuals with varying risk preferences will vary in their proportion of risky versus risk-free asset holdings. For example, risk-averse agents will allocate more towards the risk-free asset than risk-tolerant agents will. However, *the relative proportions of individual risky assets will be the same for all agents*, even if different agents invest different amounts in risky assets overall. The optimal proportions of risky assets is known as the *market portfolio* because it is the same for all agents.

Tobin's separation theorem is interesting for a handful of reasons. The first is that agent preferences about risk have no bearing on the composition of the optimal portfolio; individual stock picking has no place in a world in which agents are optimizing according to the theorem. In fact, to stock-pick would be to adopt a higher-risk and lower-return portfolio, thereby leaving money on the table. Another interesting conclusion is that independent of an asset's particular characteristics, it might still be worth including in one's portfolio for the purpose of diversification.

2.3 Connecting Tobin to CAPM

Tobin's separation theorem brings us back to CAPM. Consider a world in which all agents have the same beliefs about the assets. Such information symmetry is a common presumption in many economic models. If each agent has the same belief about the risk and return characteristics of an asset, then for each agent, the optimal risky-asset portfolio will be the same. If each agent holds the same risky-asset portfolio, then the market capitalization of the assets will correspond to their proportions within the optimal portfolio. We can use this to derive the CAPM formula for the expected return of an asset i [13]:

$$ER_i = R_f + \beta_i(ER_M - R_f). \tag{2}$$

Here, ER_i is the expected return on asset i, R_f is the risk-free rate of return, ER_M is the expected return of the market portfolio, and β_i, typically known as the "beta" of asset i, is defined as

$$\beta_i = \frac{Cov(R_i, R_M)}{\sigma_M^2} = \rho(R_i, R_M)\frac{\sigma_i}{\sigma_M}. \tag{3}$$

Above, σ_i and σ_M denote the standard deviation of the return of i and the market portfolio M, respectively, and $\rho(R_i, R_M)$ denotes the correlation between i and M. That is, β_i represents the volatility of i with respect to M. Positive β_i means i's expected return moves in tandem with M's, whereas negative means it goes in the opposite direction. The magnitude of β_i signifies how much i's return amplifies M's.

Looking back at Equation 2, the expected return of i is composed of two components: the risk-free return R_f and the market risk premium $(ER_M - R_f)$ multiplied by β_i.

3 OUR MARKET MODEL

Our model consists of agents whose interactions comprise the market. There are no "market makers" or other explicit roles. In our market model, *agents* trade *assets*.

3.1 Assets

Our model consists of a set of agents that trade assets. Our assets have two fundamental characteristics: payoff and variance.

P_i : Expected payoff of asset i in dollars.

σ_i^2 : Variance of payoff of asset i in dollars squared.

Payoff is distinct from the expected return of an asset. Payoff denotes the final dollar value of the asset at the end of the trading period. The expected return of the asset at any one point is the payoff of the asset divided by the current market price. Payoff is intended to represent the value of the given asset at the end of the trading period. Included in this value is any sort of temporal discounting or dividends paid to the asset.

Our conception of payoff and variance is integral to examining the functioning of our market simulation. By denoting these characteristics in dollar values, we allow for the expected return on assets to be determined entirely by the market; return will adjust based on the prices that agents are willing to pay for assets. If prices increase, because payoff is held constant, expected return will decrease, or vice-versa.

3.2 From Agents to a Market

Our model of agent behavior entails very few presumptions about the nature of agents. A brief thought experiment demonstrates the fundamental thinking that undergirds the agents in our model.

Consider a world that has no active market: there are merely two agents who might wish to trade with each other. Agent A has a portfolio that consists exclusively of the risky asset, and agent B has lots of cash but no risky assets whatsoever. They both are subject to some utility function for which risk and return are inputs. A wants to receive a high price for their risky asset. B wishes to pay a low price, to maximize expected return. Each makes an offer. If the prices overlap, they trade; otherwise they recalculate their utilities with new prices and make new offers. If no trades occur at an initial ask price, agent A will revise down their price, recalculate utility, and determine whether they still wish to trade. A similar process occurs for agent B; if no trades occur at a given bid, B will revise up their bid and recalculate their utility.

Our model makes no presumptions about whether agents will trade, nor does it characterize orders as a distribution of some kind. It simply confers upon agents the ability to evaluate assets, and trade according to their preferences. We define (investor) agents as agents who are subject to a utility function of the following form.

$$U(\sigma, \mu) = \mu - \frac{1}{2} r\sigma^2.$$

Above, μ represents the expected return of the agent's portfolio, r represents the risk coefficient, and σ represents the standard deviation of the portfolio's return. Utility is positively related to expected return, and negatively related to risk. The risk coefficient r represents an agent's relationship with risk. Increasing values of r make agents more risk averse, while lower values make them more risk tolerant. We presume in our market that no agent is "risk-seeking" (i.e., having a negative value of r).

We can use Tobin's separation theorem and agent utility functions to determine the optimal allocation between the market portfolio and the risk-free asset. We know that by Tobin's separation theorem, any utility maximizing bundle U will consist of some combination of the risk-free asset and the market portfolio, in which the market portfolio will have weight w, and the risk-free asset will have weight $1 - w$. We can then characterize U as follows.

$$ER_U = wER_M + (1 - w)R_f, \text{ and}$$
$$\sigma_U = w\sigma_M.$$

Recall that R_f is risk-free, so $\sigma_f = 0$. We can plug these constraints into the agent utility function, take the derivative with respect to w to maximize, and solve for the optimal weight of the market portfolio as a function of ER_M and σ_M:

$$U(\sigma, \mu) = wER_m + (1 - w)R_f - \frac{1}{2} r(w\sigma_M)^2.$$
$$\frac{\partial U}{\partial w} = ER_M - R_f - rw\sigma_M^2.$$

Setting $\frac{\partial U}{\partial w} = 0$, we get

$$w = \frac{ER_M - R_f}{r\sigma_M^2}.$$

Here, w represents the fraction of an agent's total holdings (the sum of the value of their risky asset holdings, and the risk-free asset) that they ought to allocate towards the market portfolio.

Consider again our single asset world with agents A and B. Each agent calculates their optimal value of w, and then compares w to the current fraction of their holdings comprised by the risk-free asset. For agent A, w is less than 1; therefore, to achieve their utility-maximizing bundle U_A they must sell off some of the risky asset. Conversely, agent B determines by a similar process that achieving their optimal bundle U_B requires purchasing the risky asset. They both submit limit orders at the prices that each used in their calculation of w. If their limit orders do not overlap, then they must adjust their prices and reevaluate whether, at the adjusted prices, the optimal w still differs from the current allocation within their portfolio. If it does, then they will continue to adjust limit order prices until either trades occur or the expected return changes as a function of the price offered to satisfy the current weight of the asset within their portfolio.

Having outlined the basic asset evaluation process, we may now offer a comprehensive picture of all of the attributes that comprise an agent in our model. Each agent has the following attributes, denoted here as attributes of agent A:

- $\Theta_A = \{a_0, a_1, a_2, ...a_N\}$, where Θ_A signifies agent A's portfolio, and contains integers representing the number of shares held by the agent of assets 0 through N.
- $\phi_A = \{p_0, p_1, p_2...p_N\}$, where ϕ_A represents agent A's current beliefs about asset prices. For any asset i, $\phi_A(i)$ reflects A's belief about the price of i. After i is traded, $\phi_A(i)$ reflects the market price of i.[1]
- r_A, the risk coefficient of agent A.
- τ_A, the quantity of risk-free asset held by agent A.

There are some minor details regarding the simulation. We define the "refresh rate" of agents to be the interval of time they wait to reevaluate their trades. Additionally, we set a "step size" for the agents. The step size determines the amount by which each agent increments their limit orders if they do not execute a trade. It primarily influences the rate of convergence and how stable the eventual market equilibrium is.

Next, we give a high-level pseudocode of our simulation scheme in Algorithms 1, 2, and 3. Please note that the pseudocode is meant for an overall understanding only. We omit many details, including the multi-threaded implementation of our simulation.

Algorithm 1 Outline of the overall simulation

```
while Simulation running do
    for each agent A do
        Opt_A ← Optimal portfolio as calculated in Section 4.1
        if Θ_A ≠ Opt_A then
            Trade according to Algorithm 2
            Update φ_A according to Algorithm 3
```

[1]Note that prices in our model emerge endogenously through trading.

Algorithm 2 An arbitrary agent A's Trading Behavior

for each asset i in Θ_A **do**
    if $\Theta_A(i) < Opt_A(i)$ **then** ▷ Buy because $Opt_A(i)$ is greater
        Buy asset i
    else if $\Theta_A(i) > Opt_A(i)$ **then**
        Sell asset i

Algorithm 3 Updating ϕ_A

if Trade of asset i occurs **then**
    $\phi_A(i)$ = trade price
else if Order times out **then**
    if Order was a buy order **then**
        Increase $\phi_A(i)$ according to Section 3.2
        cancel order for asset i ▷ Cancel order, not relevant
    else ▷ Must be a sell order
        Decrease $\phi_A(i)$ according to Section 3.2
        cancel order for asset i

4 MULTIAGENT SIMULATION

We run a series of simulations to study interventions, noise, and importantly, to establish a dialogue with standard financial economic theory. We programmed our simulations on MAXE, a message-based multi-agent simulation platform designed for designing and implementing agent-based models [3]. It supports a variety of agent classes (as in object-oriented programming) relevant to financial modeling. More generally, it is a system that allows one to specify agents that exchange messages with each other.

MAXE stands in contrast to ABIDES, an alternative agent-based modeling system with the goal of modeling the NASDAQ exchange as closely as possible [8]. We found MAXE's framework to be more general purpose, which suited our goals. Additionally, MAXE's implementation in C++ allows for greater speed at the agent messaging level [3]. We used the *ProRata exchange algorithm* of MAXE and implemented our agents in Python.

We simulate a market at two different scales: 30 agents trading three assets and 500 agents trading 10 assets. The larger scale is particularly useful for validating the model. The smaller scale is useful for visualization. The simulations were performed on a high-performance computing (HPC) grid.

4.1 The CAPM World: Homogeneous Beliefs

It might be the case that we inadvertently dictate the outcome of the simulation. Therefore, some method of verification is necessary.

We adopt a two-pronged approach to validating the results of the model. We know that CAPM and Tobin's separation theorem are beyond analytic reproach. We argue that if we construct a simulation that adheres to the fundamental presumptions of Tobin's separation theorem and CAPM, and the data resulting from running the simulation confirms these theorems, then our choice of mechanism does not influence the final equilibrium price (at least in a manner unaccounted for by these well-understood theories of asset-pricing), and can be used to relax presumptions and move beyond the basic world in which these presumptions apply.

We follow a standard procedure for computing CAPM efficient portfolios [39][Ch 12.5].

Convergence of Mean-Variance Characteristics
Recall Tobin's separation theorem: There is a single optimal portfolio, and agents ought to act on their beliefs about asset returns to determine the exact weights of this portfolio (the market portfolio), and allocate their total holdings between this market portfolio and the risk free asset.

One presumption of CAPM is that each agent holds the same beliefs. Given homogenous belief and Tobin's separation theorem, agents will converge to the same market portfolio.

We can use this conclusion to verify the validity of our simulation, given a market in which agents have homogeneous beliefs. If we run the simulation and all agents converge on the same portfolio, then we can conclude that our market design is in fact a mere extension of analytic presumption.

To confirm this, we recorded the mean-variance (return-risk) characteristics of each agent's current portfolio immediately after they complete a transaction over the course of the simulation. We then color each data-point to indicate its temporal position: blue for earlier and red for later.

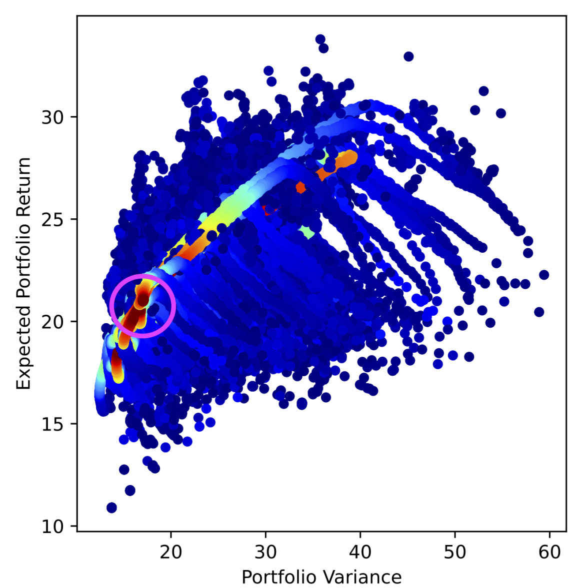

Figure 2: The mean-variance characteristics of 500 different agents' portfolios over the course of a trading period. Color represents timeline: blue to red. Agents converge on roughly the same market portfolio (circled in magenta), as predicted by CAPM and Tobin's separation theorem.

The mean-variance characteristics shown in Figure 2 demonstrates that the agents do in fact converge to a single portfolio. Additionally, because we have specified the homogeneous beliefs of each agent, and because we know that each agent is optimizing their portfolio in the same manner as prescribed by Tobin's

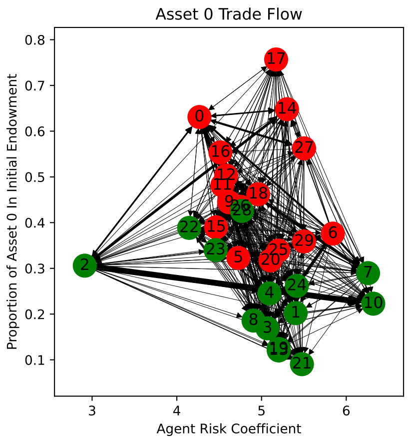

Figure 3: Effect of the proportion of initial endowment and risk coefficient on trade flow. Higher risk-coefficient indicates less risk tolerance and more risk averseness. Green nodes are net buyers and red net sellers.

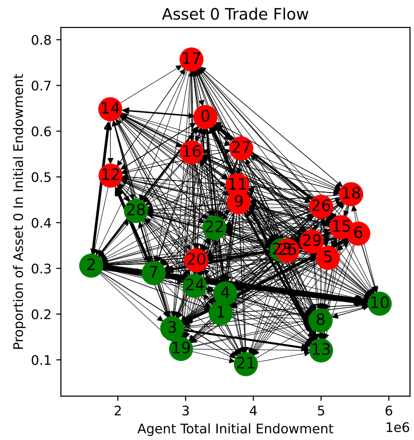

Figure 4: Effect of Asset 0's proportion in the initial endowment and the total initial endowment on trade flow. The edges and node colors have the same meaning as in Figure 3.

theorem, we can conclude that these agents have converged upon the optimal market portfolio. This is a strong piece of evidence supporting the validity of our simulation.

To evaluate whether the equilibrium achieved by our simulation is the same equilibrium predicted by CAPM, we first calculate what CAPM would predict and compare it with the expected returns at an equilibrium in our simulation. We have found that *these two values are approximately equal,* with any variation easily attributable to frictions stemming from the lack of fractional share ownership. This is significant additional evidence of our model's validity.

Asset-Flow Networks: The Visible Hand

We devise *an asset-flow network*, where each node represents an agent in the market, and each edge represents a trade that occurs between two agents; edge width is determined by the volume of trade. Asset-flow networks show the exchanges between agents of a single asset within the multi-asset market. Visualizations are done using the 30-agent model.

Figure 3 shows an asset-flow network for asset 0. We have similar networks for assets 1 and 2 that are not shown. Agents with higher proportions of a given risky asset tend to sell the risky asset to other agents within the market. Some notable outliers include agent 2 in the market for asset 2, and agents 22 and 12 in the market for asset 1 (not shown). Upon investigation, the cause of these outliers is initialization. For example, asset 2 was initialized with a higher price than assets 0 and 1, and also appears to have comprised a significant proportion of agent 2's portfolio. This gives agent 2 an apparently high initial endowment of asset 2, but it is somewhat artificial as this is not reflective of the asset's final equilibrium price.

We also generate asset-flow networks to check whether the volumes of agents' initial endowments affected the flow of assets in the market. Figure 4 does not show any strong evidence for this. The proportion of an asset in the initial endowment seems to be the factor that matters most.

Emergence of Roles

These networks offer interesting insights on the emergence of roles. One example is the role of agent 2 within the market for asset 0. We can see that agent 2 sells an incredible amount of asset 0 to agent 10. Note that agent 2 has a high tolerance for risk but little wealth, while agent 10 has a comparatively low tolerance for risk but higher initial wealth. This might indicate the role of a "dealer" for agent 2 because agent 10 must purchase a massive volume of asset 0 as agent 10 is quite wealthy but has a very low proportion of asset 0 in their initial portfolio. Furthermore, consider the scenario of agents with high proportions of the risky asset in their portfolios submitting ask limit orders to the market. Because agent 2 has a higher risk tolerance, they will accept a lower return (i.e., higher price) than agent 10. Therefore, agent 2 will purchase shares of asset 0 before agent 10 will. Agent 2 ends up selling said shares, as agent 10 eventually bids the price high enough to justify the trade-off, and the comparative attractiveness of other assets compels agent 2 to decrease the proportion of asset 0 in their portfolio.

Once again, we see the value of ABM because it would be impossible to model such asset flows from data generated by a predictive model that does not model agent behavior or market mechanisms.

Price Convergence

We observe significant variations in price at the beginning of the trading period, which generally settles into an equilibrium path as time goes on. As shown in Figure 5, the price for asset 1 exhibits a sort of "bouncing" about a range of prices (asset 0 and 2 show similar behavior). This is a reflection of the step size and the inability to trade fractional shares of assets.

4.2 Interventions

So far we have proceeded with a static conception of asset returns. However, real markets are dynamic and intended to reflect changes in the risk or return of assets. We rerun previous experiments to establish the behavior of agents in response to shocks, and to

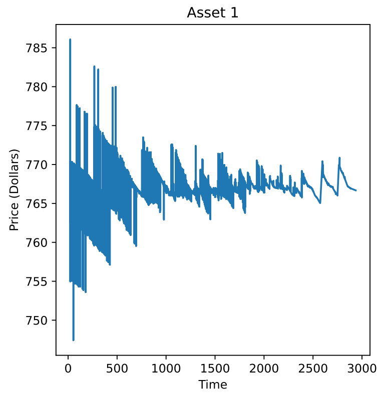

Figure 5: Convergence of asset 1 price. There is some “bouncing,” especially at the beginning, as agents search within a range to make trades. It later converges to an equilibrium.

determine the effect that noise traders have in changing the market equilibrium in response to shocks. We use the 30-agent model.

In this experiment, we implemented a negative payoff shock of 10% to asset 2 at timestamp 5000, and a positive payoff shock of 11.11% at timestamp 10000. Figure 6 presents interesting results on cross-asset price dynamics. We can see the returns of assets 0 and 1 undergo an increase in expected return (i.e., their price falls) timed exactly to the negative shock to asset 2 (the opposite happens for positive shocks). We observe across other simulations a coordination of returns across assets in the market in response to the payoff shocks to a single asset, much as we can still observe in the three asset case. In this case, notably, the returns of all three assets move in the *same direction*, which may seem counter-intuitive.

Using our model, we can explain the effects of interventions. Intuitively, one would hypothesize an increase in the return of the asset whose payoff is negatively shocked (e.g., price falls due to bad news) and a decrease in the return of the other assets in the market (i.e., price increases). This hypothesis is driven by the intuition that assets are fundamentally competing against other assets in the market for buyers. By shocking the payoff of one asset in the market, we made one asset less competitive, thereby lowering the competitive pressure on the other assets and making them more attractive. Another intuition about market behavior also supports this hypothesis. Prices in the market are a function of agents' relationship with risk; if agents are extremely risk tolerant, prices will be high and returns low, and vice versa if they are extremely risk-averse. Because we see no shift in the agent composition of the market, we would expect no fundamental change in the risk premium present in the market after shocking the payoff of an asset in the market. To hold the risk premium constant, the returns of the other assets in the market would have to fall to make up for the rise in the return of shocked asset. In short, this argument would also lead us to believe that the returns of other assets in the market should fall in response to a negative shock to the payoff of one asset. However, this is not what we observe.

To explain this, we have to resort to agents' risk-return optimization in our simulation. Agents ought to purchase the optimal portfolio. However, changes to payoff characteristics also change the market portfolio. Large enough shocks will also affect the *fraction* of their total holdings that agents wish to allocate to the market portfolio overall. Consider what happens when we substantially decrease the payoff of an asset. First, the total amount of possible payoff decreases. This makes the market less attractive relative to the risk-free rate. If the fraction of our total holdings that we ought to devote to the market portfolio decreases, then agents will sell off *all assets*. The effect of this slight decrease might reasonably be enough to counteract the increase in price we would otherwise observe, holding the fraction of the market portfolio within our total holdings constant.

In short, when we negatively shock the payoff of an asset, it decreases the size of the “pie” we call market portfolio. Then, even if some other asset's slice of the pie (i.e., its payoff) increases as a result, due to the size of the whole pie decreasing, the newly increased slice may still be smaller than than the previous slice. For small enough shocks, the overall size of the pie may not decrease so much, and the opposite might be observed. In either case, we can do model-based calculations to explain the outcome. Explanations like this are yet another reason to use agent-based models to interrogate underlying modeling components and assumptions.

4.3 Noise Trading: Heterogeneous Beliefs

A *noise trader* is an individual with imperfect information about the true value of the assets they are trading [5, 11, 29]. Shleifer and Summers [29] discuss noise traders and offer a very general definition of noise traders. They define a noise trader as one who deviates from so-called “rational” trading behavior in a systematic way (not randomly). We use this definition to implement such traders and observe their effect on the market in a variety of ways, from portfolio convergence to the asset-flow patterns.

Using the 30-agent model, we first generate an alternative set of beliefs in which asset 0 has a payoff that is 10% higher than in the original (”rational”) asset characteristic set. We then set the beliefs of five agents within the market to this alternative set and rerun the simulation, generating similar graphs in an effort to tease out what effect noise traders have on asset pricing. Let us first examine the mean-variance plot.

As we can see in Figure 7, the addition of optimistic noise traders results in two equilibrium market portfolios. This violates the presumption that undergirds CAPM, as we would expect it to, given that we have relaxed the premise that agents have homogeneous beliefs about asset characteristics. We can see that there is a smaller equilibrium cluster with higher expected return representing the portfolio that the noise traders converge on. In addition to observing the higher-return portfolio of the noise traders, we also observe the agents allocate a higher proportion of their holdings to the asset for which they had optimistic beliefs.

We next examine the asset-flow networks to evaluate how noise traders affect asset movement. Agents 0-4 being the designated noise traders, Figure 8 demonstrates the asset-flow networks for asset 0 (the subject of the noise traders' optimism) within the same market as before. Comparison with the previous asset-flow networks for asset 0 in the non-noisy simulation shows that there is a significant shift in the relationship between specific agents in the

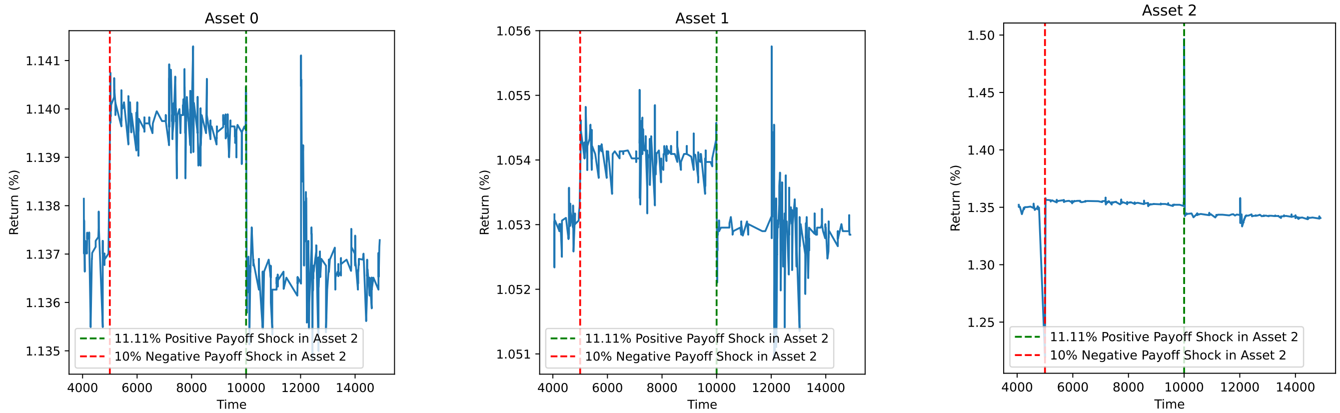

(a) Price changes in asset 0 **(b) Price changes in asset 1** **(c) Shocks on asset 2**

Figure 6: The ripple effect of shocks on asset 2. The return on asset 2 over the course of the simulation is shown in (c). We can observe significant and immediate changes in return in response to payoff shocks, which fade quickly in response to trading. As shown in (a) and (b), the returns on assets 0 and 1 interestingly move in the *same direction* as asset 2 (after shocks).

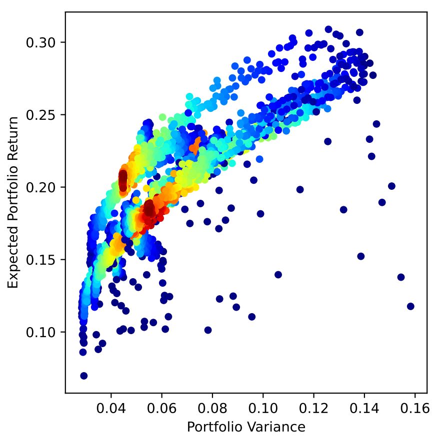

Figure 7: The mean-variance characteristics of portfolios of 30 agents, five of whom are noise traders. There are two distinct equilibria.

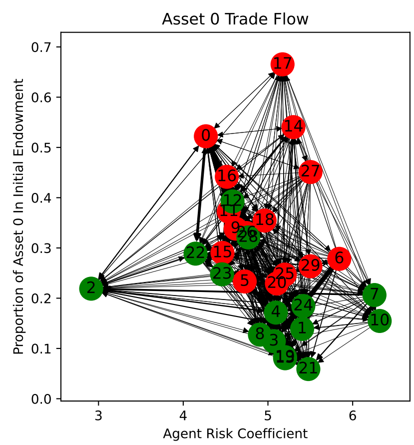

Figure 8: Asset-flow network of asset 0 in the noisy simulation. Node 2 is no longer a dealer as they were in Figure 3.

simulation as a result of adding noise traders. In particular, agent 2 appears to have lost their previous dealer role (shown in Figure 3), becoming rather an unremarkable net buyer of risky asset 0. Agent 0 also has a shift in its role, appearing to play a significant role for asset movement across the market. Interestingly, the proportion in the initial endowment is still an overriding determinant of agent behavior, even compared with a shift in payoff belief of 10%.

This experiment provides some insight into how slight changes to the presumptions of CAPM can have drastic and immediate effects not only on market equilibria but also on agent roles.

5 CONCLUSION

Proceeding from the predominant analytic theory of asset pricing, we have been able to achieve two goals: to verify that our model is able to engage in a dialogue with the existing economic theory and to expand beyond the original set of assumptions that undergird the predominant asset-pricing theory. We demonstrated a variety of ways of understanding the functioning of the market and inter-agent interaction. There are many promising avenues for future research, including a systematic investigation of noise-trading strategies, examining the change in agent-owner composition due to shocks, and more broadly, accounting for myriad types of assets.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their helpful suggestions and many kind words, which made this joint CS-Economics research and all the hard work worthwhile. MTI is grateful to the NSF for support from Award IIS-1910203.

REFERENCES

[1] Simone Alfarano, Thomas Lux, and Friedrich Wagner. 2006. Estimation of a simple agent-based model of financial markets: An application to australian stock and foreign exchange data. *Physica A: Statistical Mechanics and its Applications* 370, 1 (2006), 38–42.

[2] W Brian Arthur. 2006. Out-of-equilibrium economics and agent-based modeling. *Handbook of computational economics* 2 (2006), 1551–1564.

[3] Peter Belcak, Jan-Peter Calliess, and Stefan Zohren. 2022. Fast Agent-Based Simulation Framework with Applications to Reinforcement Learning and the Study of Trading Latency Effects. http://arxiv.org/abs/2008.07871 arXiv:2008.07871 [cs, q-fin].

[4] Mario A Bertella, Felipe R Pires, Ling Feng, and Harry Eugene Stanley. 2014. Confidence and the stock market: An agent-based approach. *PloS one* 9, 1 (2014), e83488.

[5] Fischer Black. 1986. Noise. *The Journal of Finance* 41, 3 (1986), 528–543. https://doi.org/10.1111/j.1540-6261.1986.tb04513.x _eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1986.tb04513.x.

[6] Eric Bonabeau. 2002. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the national academy of sciences* 99, suppl_3 (2002), 7280–7287.

[7] Carlos Bousoño-Calzón, Josué Bustarviejo-Muñoz, Pedro Aceituno-Aceituno, and José Joaquín Escudero-Garzás. 2019. On the economic significance of stock market prediction and the no free lunch theorem. *IEEE Access* 7 (2019), 75177–75188.

[8] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. 2020. ABIDES: Towards high-fidelity multi-agent market simulation. In *Proceedings of the 2020 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*. ACM, Miami, FL, USA, 11–22.

[9] John Y Campbell. 2000. Diversification: A bigger free lunch. *Canadaian Investment Review* 13 (2000), 14–15.

[10] Filippo Castiglione. 2000. Diffusion and aggregation in an agent based model of stock market fluctuations. *International Journal of Modern Physics C* 11, 05 (2000), 865–879.

[11] J. Bradford De Long, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann. 1990. Noise Trader Risk in Financial Markets. *Journal of Political Economy* 98, 4 (1990), 703–738. https://www.jstor.org/stable/2937765 Publisher: University of Chicago Press.

[12] Christophe Deissenberg, Sander Van Der Hoog, and Herbert Dawid. 2008. EURACE: A massively parallel agent-based model of the European economy. *Applied mathematics and computation* 204, 2 (2008), 541–552.

[13] Eugene F. Fama and Kenneth R. French. 2004. The Capital Asset Pricing Model: Theory and Evidence. *Journal of Economic Perspectives* 18, 3 (Sept. 2004), 25–46. https://doi.org/10.1257/0895330042162430

[14] J. Doyne Farmer and Robert L. Axtell. 2022. *Agent-Based Modeling in Economics and Finance: Past, Present, and Future.* INET Oxford Working Papers 2022-10. Institute for New Economic Thinking at the Oxford Martin School, University of Oxford. https://ideas.repec.org/p/amz/wpaper/2022-10.html

[15] Kristin J Forbes. 2007. The microeconomic evidence on capital controls: no free lunch. In *Capital Controls and Capital Flows in Emerging Economies: Policies, Practices, and Consequences*. University of Chicago Press, Chicago, IL, USA, 171–202.

[16] Brian Heath, Raymond Hill, and Frank Ciarallo. 2009. A survey of agent-based modeling practices (January 1998 to July 2008). *Journal of Artificial Societies and Social Simulation* 12, 4 (2009), 9.

[17] Scott Heckbert, Tim Baynes, and Andrew Reeson. 2010. Agent-based modeling in ecological economics. *Annals of the New York Academy of Sciences* 1185, 1 (2010), 39–53.

[18] Stefan J Johansson. 2002. No free lunches in multi-agent systems. *International Journal of Computational Intelligence and Applications* 2, 04 (2002), 411–421.

[19] Alan P. Kirman and Nicolaas J. Vriend. 2001. Evolving market structure: An ACE model of price dispersion and loyalty. *Journal of Economic Dynamics and Control* 25, 3-4 (March 2001), 459–502. https://ideas.repec.org/a/eee/dyncon/v25y2001i3-4p459-502.html

[20] Deddy P. Koesrindartoto, Junjie Sun, and Leigh Tesfatsion. 2005. *An Agent-Based Computational Laboratory for Testing the Economic Reliability of Wholesale Power Market Designs.* Staff General Research Papers Archive 12388. Iowa State University, Department of Economics. https://ideas.repec.org/p/isu/genres/12388.html

[21] Blake LeBaron. 2002. Building the Santa Fe artificial stock market. *Physica A* 1 (2002), 20.

[22] Blake LeBaron. 2006. Agent-based computational finance. *Handbook of computational economics* 2 (2006), 1187–1233.

[23] Robert E. Lucas. 1976. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy* 1 (1976), 19–46. https://doi.org/10.1016/S0167-2231(76)80003-6

[24] Myles E Mangram. 2013. A simplified perspective of the Markowitz portfolio theory. *Global journal of business research* 7, 1 (2013), 59–70.

[25] Harry markowitz. 1952. Portfolio Selection. *The Journal of Finance* 7, 1 (1952), 77–91.

[26] W Megginson. 1996. A historical overview of research in finance. *Journal of finance* 39, 2 (1996), 323–346.

[27] Takanobu Mizuta. 2020. An agent-based model for designing a financial market that works well. In *2020 IEEE symposium series on computational intelligence (SSCI)*. IEEE, IEEE, Canberra, Australia, 400–406.

[28] Egle Samanidou, Elmar Zschischang, Dietrich Stauffer, and Thomas Lux. 2007. Agent-based models of financial markets. *Reports on Progress in Physics* 70, 3 (2007), 409.

[29] Andrei Shleifer and Lawrence H Summers. 1990. The Noise Trader Approach to Finance. *Journal of Economic Perspectives* 4, 2 (May 1990), 19–33. https://doi.org/10.1257/jep.4.2.19

[30] Yoav Shoham and Kevin Leyton-Brown. 2008. *Multiagent systems: Algorithmic, game-theoretic, and logical foundations.* Cambridge University Press, New York, NY, USA.

[31] Leigh Tesfatsion, Valentin Petrov, and James Nicolaisen. 2001. *MARKET POWER AND EFFICIENCY IN A COMPUTATIONAL ELECTRICITY MARKET WITH DISCRIMINATORY DOUBLE-AUCTION PRICING.* Economic Reports 18195. Iowa State University, Department of Economics. https://EconPapers.repec.org/RePEc:ags:iowaer:18195

[32] James Tobin. 1958. Liquidity Preference as Behavior Towards Risk. *Review of Economic Studies* 25, 2 (1958), 65–86. https://EconPapers.repec.org/RePEc:oup:restud:v:25:y:1958:i:2:p:65-86.

[33] Elaine Wah and Michael P. Wellman. 2017. Latency arbitrage in fragmented markets: A strategic agent-based analysis. *Algorithmic Finance* 5, 3-4 (April 2017), 69–93. https://doi.org/10.3233/AF-160060

[34] L Wang, Kwangwon Ahn, C Kim, and C Ha. 2018. Agent-based models in financial market studies. In *Journal of Physics: Conference Series*, Vol. 1039. Institute of Physics Publishing, Institute of Physics Publishing, Phuket, Thailand, 012022.

[35] Xintong Wang, Christopher Hoang, Yevgeniy Vorobeychik, and Michael P. Wellman. 2021. Spoofing the Limit Order Book: A Strategic Agent-Based Analysis. *Games* 12, 2 (May 2021), 46. https://doi.org/10.3390/g12020046

[36] JR Wei, JP Huang, and PM Hui. 2013. An agent-based model of stock markets incorporating momentum investors. *Physica A: Statistical Mechanics and its Applications* 392, 12 (2013), 2728–2735.

[37] Chun-Xia Yang, Rui Wang, and Sen Hu. 2013. Modeling and analysis of an agent-based model for Chinese stock market. *Physics Letters A* 377, 34-36 (2013), 2041–2046.

[38] Yi Zhang, Zhe Li, Yongchao Zhang, et al. 2020. Validation and calibration of an agent-based model: A surrogate approach. *Discrete Dynamics in Nature and Society* 2020 (2020).

[39] Eric Zivot. 2021. *Introduction to Computational Finance and Financial Econometrics with R.* Bookdown. https://bookdown.org/compfinezbook/introcompfinr/