Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds

REGULATION OF SYSTEMIC RISK THROUGH CONTRIBUTORY ENDOGENOUS AGENT-BASED MODELING

Aurora J. Bristor

Applied Mathematics, Statistics,
and Scientific Computation
University of Maryland
College Park, MD 20742, USA

Sean L. Barnes
Michael C. Fu

Robert H. Smith School of Business

University of Maryland
College Park, MD 20742, USA

ABSTRACT

The Financial Stability Oversight Council (FSOC) was created to identify and respond to emerging threats to the stability of the U.S. financial system. The research arm of the FSOC, the Office of Financial Research, has begun to explore agent-based models (ABMs) for measuring the emergent threat of systemic risk. We propose an ABM-based regulatory structure that incentivizes the honest participation and data contribution of regulated firms while providing clarity into the actions of the firms as endogenous to the market and driving emergent behavior. We build this scheme onto an existing ABM of a single-asset market to examine whether the structure of this scheme could provide its own benefits to market stabilization. We find that without regulatory intervention, markets acting within this proposed structure experience fewer bankruptcies and lower leverage buildup while returning larger profits for the same amount of risk.

1 INTRODUCTION

"Systemic risk refers to the risk of collapse of an entire complex system as a result of the actions taken by the individual component entities or agents that comprise the system" (Chen, Iyengar, and Moallemi 2013). More than five years after the financial crisis and more than three years after the establishment of the Financial Stability Oversight Council (FSOC) and the Office of Financial Research (OFR) to control systemic risk, it remains unregulated. However, "before we can hope to manage the risks of financial crises effectively, we must be able to define and measure those risks explicitly" (Lo 2008). Systemic risk is as multi-faceted and highly dynamic as the complex system it threatens, which challenges our ability to measure it. Further, the individual component entities collectively generating the risk have no incentive to pool their rich data resources as input to that measurement. We propose a regulatory structure designed to overcome these specific challenges of multi-dimensional dynamic measurement and individual incentive. We propose a cooperative structure between the FSOC and the participants of the financial system using an agent-based model (ABM). We argue that this structure would have the flexibility, precision, and depth of data to support the FSOC's mandate to measure and manage systemic risk. It is the aim of the remainder of this work to investigate whether the proposed structure results in a mitigation of the market risks.

The first challenge to the measurement of systemic risk is its multi-dimensional nature. In the first working paper of the OFR, the authors collected 31 academic models on systemic risk, each covering a different facet of financial stability. Still, the authors cautioned that they were not providing an exhaustive survey and that "even if an exhaustive overview of the systemic risk literature were possible, it would likely be out of date as soon as it was written" (Bisias et al. 2012). Systemic risk is also dynamic. New instruments for financial flexibility and access are constantly being invented and refined. We propose the use of an agent-based model to provide for the measurement of the multi-dimensional and dynamic aspects of systemic risk. As an ABM is executed, a record or calculation could be made from any of its variables, allowing it to report on multiple dimensions. Reporting across the time steps of the simulation would capture dynamic shifts in the model. This would keep regulators up to date and allow them to use flexible and process-based regulations. The repeatability of the outcomes of an ABM given the same inputs makes it a valuable tool for transparent counterfactuals, sensitivity analyses, and stress tests. Furthermore, possessing the flexibility to perform in nonlinear cases makes ABMs valuable during the fast-paced and shifting environment of an unfolding crisis.

The second challenge facing the measurement of systemic risk is the availability of input data. The "individual component entities or agents that comprise the system" (Chen, Iyengar, and Moallemi 2013) lack a collective incentive to share their deep and narrow market data to create the collaborative set required for the breadth and depth of data in a complex ABM of the market and its systemic risk. Despite a desire to minimize exposure to risk and support such a model, individual entities face strong data privacy concerns about working together. Government, on the other hand, is suited to deliver public goods not otherwise generated by the private market. We therefore propose a contributory model hosted by the government. This kind of data sharing exists already in the private sector. With privacy protections, individual firms deliver data to a collector, creating a large market database. The data collector compensates the firms for this data by providing access to the models and market analysis it generates. Putting this business model into practice for systemic risk measurement, the regulated market participants would provide data to the OFR and FSOC, which serve the role of data collector and model distributor in our proposed cooperative regulatory structure. Following the incentive structure in the private sector, the FSOC and OFR could incentivize proper data delivery with model access and systemic risk analysis generated from the ABM. With such access, participants could better estimate impacts to their risk profiles from their business proposals or market forecasts. Yet with limited data or model knowledge or control, firms would have to submit honest and timely data to best inform their decisions using the results of the regulator's model.

The FSOC should further invite market participants to engage in companion modeling to ensure that the model provides sufficient information to the regulators and the regulated without providing too much information to market competitors. This proposal then offers a concise package that incorporates incentives for honest data delivery and assistance in producing and maintaining a model that would be flexible, up to date, testable for counterfactuals, predictive, and available for use by regulators and the regulated alike. Such a measurement tool would be "multi-dimensional, adaptive, real-time, [and] able to incorporate illogic of human choice" (Lo 2008).

Another potential benefit of our proposed structure is a shift in perspective by market participants, who can generate systemic risk as an emergent effect from their combined actions. This kind of effect is natural in ABMs, which support the idea of agents as endogenous actors. In other types of models, each institution is instead exogenous to the market. These financial models, and implicitly those who use them, assume the market will continue to set prices without the participant. This supports the idea of systemic risk as a negative externality. Individual firms participating in our proposed regulatory structure, utilizing access to an ABM, would begin to implicitly consider themselves instead as endogenous to the market and systemic risk. We seek to discover if our proposed regulatory structure leads its participants toward a more stable market through their implicit shift in perspective from exogenous to endogenous actors.

2 METHODOLOGY

We wish to examine whether the construction of our proposed regulatory scheme can improve the stability of the marketplace. Access to an ABM of systemic risk in the market is provided as an incentive for market participants to provide accurate and timely data to the FSOC. The acknowledged benefit is in the additional information that could help firms improve their risk strategies and profit margins. With an

ABM's implicit view of agents as endogenous to the market, the firms using such a model could see their contribution towards market volatility. They could evaluate their short- and long-term risk and profit from working towards a more secure and stable market or a more volatile market (Smith, Suchanek, and Williams 1988). It is unclear whether there would be a sufficiently consistent motive for the actors to collectively move the market, and in which direction.

We test this question using an ABM to perform a repeatable experiment measuring market outcomes based on the collective acts of individual firms. We first model—as a control—a market outside our regulatory structure whose agents consider themselves exogenous to the market. We then model the market within the proposed regulatory scheme, in which the market participants can see that their actions endogenously influence their market. We refer to the ABM used for the control as the base model. The ABM of the proposed regulatory structure is modified from the base model by giving each market participant the option to alter its strategy based on simulations of its performance in the base model. We thus refer to this regulatory model as the nested model. The outer-level agents of the nested model are the market participants with model access in our regulatory structure, and the inner-level model represents the model they use to evaluate future decisions. We provide a visual depiction of the base and nested structures in Figure 1.

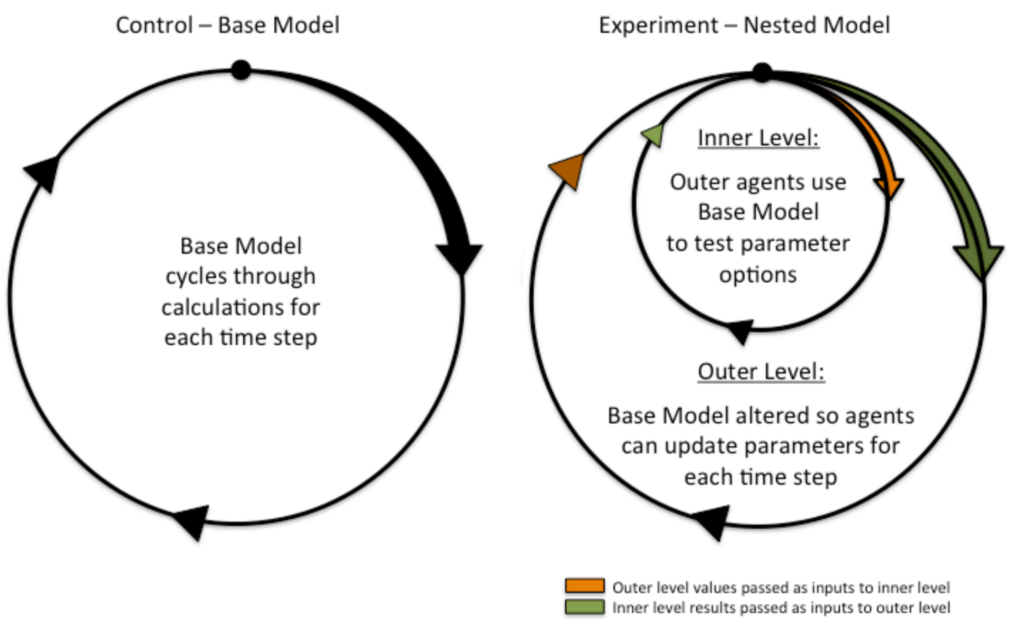

Figure 1: The unaltered base model is used as a control. An altered version of the base model is used as the outer level of the nested model to simulate firms' behavior when they are given access to an ABM of their market for testing hypotheticals. The unaltered base model is used in the inner level of the nested model as the ABM to which the outer level firms have access.

We build on the market ABM of Thurner, Farmer, and Geanakoplos (2012) as the base model for our experiment. Using a simple established model as the base minimizes new assumptions and computational error. The base model is still able to demonstrate sufficient market volatility to be able to show market stabilization. Other simple market ABMs model the stock market (Outkin 2012) or housing and mortgage-backed securities markets (Geanakoplos et al. 2012, Goldstein 2011), but this one in particular models the sort of market we are looking to regulate: the hedge funds and banking industry. We describe the base model, then detail the necessary adaptations for its use in the outer level of our experiment's nested model.

There are four types of agents in the base model: hedge funds, noise traders, banks, and investors. The hedge funds and noise traders generate the activity in the model as they buy and sell a single fixed-value no-dividend asset. The hedge funds only buy when the price is below the fixed value and hold the asset until forced to sell by a margin call from the banks. Modeling the funds to act based on knowledge of the true value reflects the assumption that the hedge funds are rational investors valuing the asset according to its fundamentals. The noise traders, on the other hand, represent traders who do not base buy and sell decisions on the stock's fundamental value. Their demand generates mispricings in the market by driving the price of the stock over or under the fixed fundamental value. This mispricing allows the hedge funds to take a long position (i.e., buying and holding) when the mispricing is advantageous. The banks provide loans to the hedge funds to buy more of the asset, but require the hedge funds to sell off if the asset value insufficiently supports the loan. The investors track the return on investment for each hedge fund and invest capital into those funds whose forecasted performance is better than a benchmark return.

A set of simultaneous equations defines the relationships between these four agents. All market activity is performed in a series of T time steps and is based on finding the market-clearing price that satisfies this set of equations for each new noise trader expenditure at time t. The noise trader demand is inversely proportional to the market-clearing price, as the expenditure of the noise traders is generated without regard to the price of the asset. As the price drops, the same investment can buy more shares, increasing the noise traders' demand. The hedge funds' demand values are also interdependent with the market-clearing price, but through a more complex set of linked equations outlined in Figure 2.

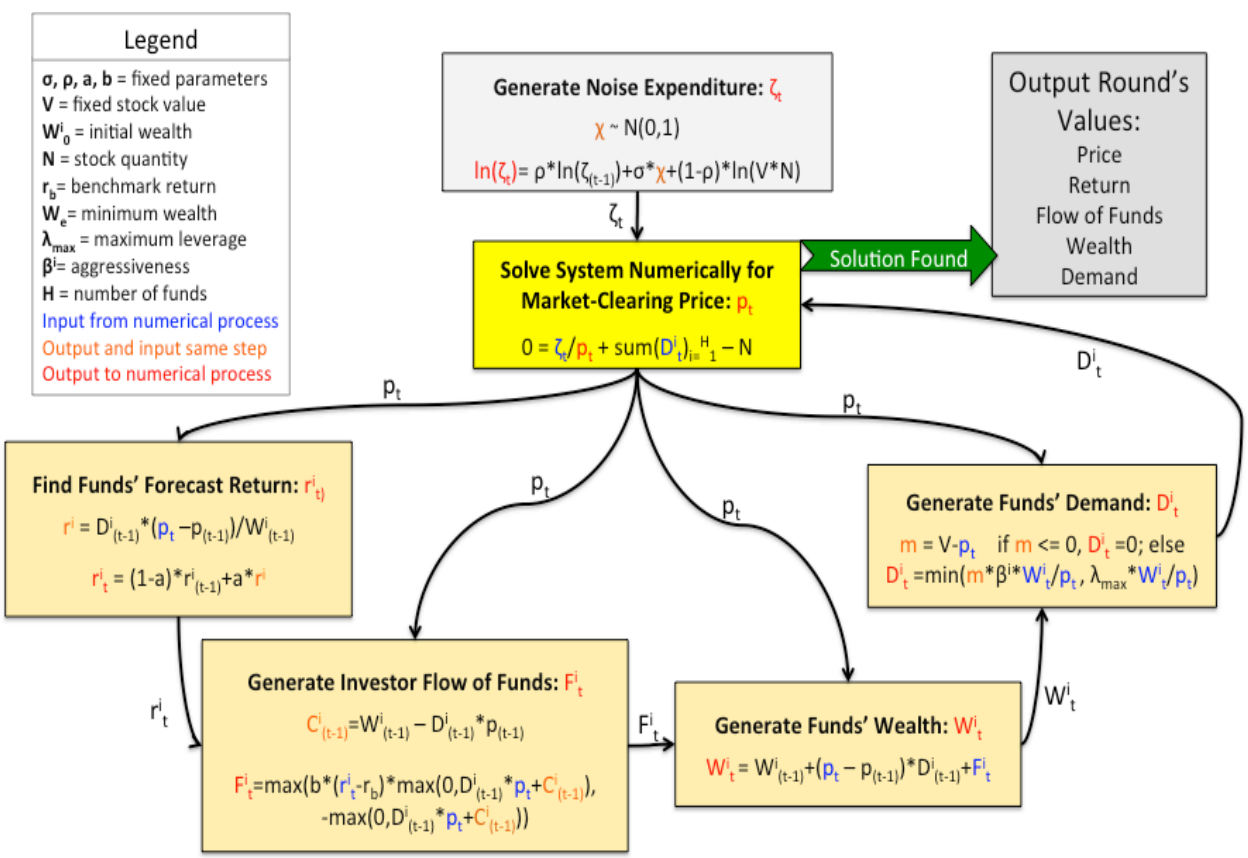

Figure 2: The only value calculated independently in each time step is the noise expenditure. All other values are historical or rely on the market-clearing price, which in turn depends on the demand generated by the sequence of equations beginning with last period's returns.

These equations cover investor returns and flow of capital, hedge fund wealth, and the hedge fund demand limited by the banks' leverage limit. The new market price determines the return on investment over the last time step, which in turn influences the investors' flow of capital to the hedge fund. The current retail value of the held assets and the investor flow of capital determine the wealth of each hedge fund. The available wealth of each hedge fund and the new price relative to the fixed fundamental value V are what ultimately set the demand for each fund. Thus, the market-clearing price equation is the avenue through which the confluence of individual requirements for return, invested capital, and wealth yields collective outcomes as it balances the demands of the noise traders and all hedge funds. The nested model operates on the same market and set of agents as the base model. The difference is that the nested model provides the hedge fund agents in the outer level with the ability to change the one parameter within their control: their aggressiveness β in response to a mispricing m.

Each hedge fund in the outer level is given access to the inner-level model to test their aggressiveness parameter β for the next time step t. They run the inner-level model on different values of β and choose the one that returns the best outcome as diagramed in Figure 3. The inner-level model is structured like the base model, running for T time steps on the four types of agents. The base model as a control is given the same inputs at the start of each market simulation as the outer-level market simulation runs. To better inform the outer-level funds, however, the inner-level model runs T time steps with initial values that correspond to the outer-level funds at time t. In other words, any hedge fund in the outer-level runs its inner-level model using data from its own market environment. Thus, a hedge fund experiments with its own β, but the β's for the other funds in the inner-level model are fixed as those of the other outer-level funds at time t. We take on the role of the FSOC in providing all updated market data to the ABM for each fund's simulation. Note that not even the regulator needs to know these values. The regulatory system could utilize a fully-homomorphic encryption scheme, designed to provide encrypted results calculated on encrypted data (Gentry 2009).

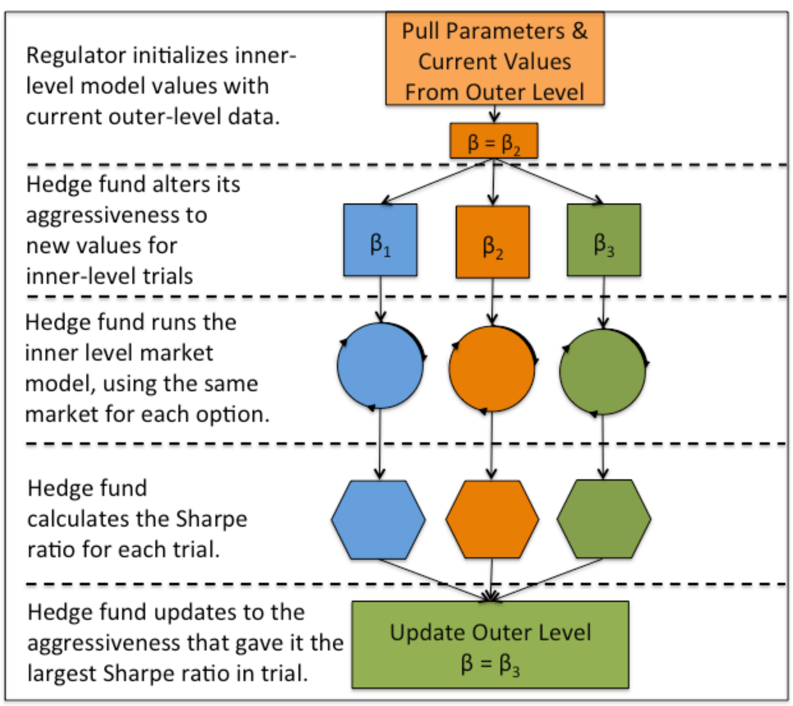

Figure 3: The flow for one of the outer-level hedge funds in using the inner-level model to update its aggressiveness β. The regulator initializes the inner-level simulations run by the hedge funds individually.

To facilitate comparison of different strategies, the variance reduction technique of common random numbers is used to make the inner-level model used by the outer-level funds the same for each trial of β. These markets are generated independently from the market in which the outer-level funds operate to simulate the firms' lack of information about the future of their own market. Finally, these markets must be independent across the trial sets of the hedge funds to simulate their independent market forecasts. These requirements on the trial markets can be met by controlling the random numbers that generate noise trader demand. As shown above in Figure 2, noise trader demand is the only equation exogenous to the hedge funds. That demand is driven by a recursive function that calls the random variate χ. If we control χ, we control the exogenous markets of all simulations.

Once all of their options are run through the inner model, each outer-level fund must choose the value of β they wish to implement in the next time step of the outer-level model. Wealth, return, and bankruptcy events are possible data points to use in optimizing the utility of the hedge funds. However, these outputs may all return different values depending on which random market the outer-level fund is running in the inner-level model. Rather than look at any one metric of profit or loss, we measure utility using the Sharpe ratio, a measure of risk-weighted return: E(Return) / E(Risk). It serves as a consistent measure across different risk and profit environments of the expected return for a given amount of risk. A higher Sharpe ratio indicates that in the same risk environment, more profit is realized. Alternatively, for the same amount of profit, there is less associated risk. Each outer-level fund updates its β in the next round to the value that produces the highest Sharpe ratio in simulated inner-level markets.

Providing outer-level hedge funds with choice requires a design for the funds' options as well as their decision process. We chose a simple set of options for each fund: increase the aggressiveness, decrease the aggressiveness, or maintain the same aggressiveness. In the original base model, there are ten hedge funds with initial β's in intervals of five. So that each fund can test whether it would be better off with one of its competitors' values, each fund can change its aggressiveness by +/- 5, or not at all. We also tested a percent-change-based approach, but this led to exponential growth in aggressiveness with no measurable market impact. We constrain the aggressiveness to be at least five because a value of zero or a negative value is without meaning in this model. We did not set a cap on the funds' possible choices, however, to see if and how they would choose to become more aggressive than their initial values. We also decided to experiment with the aggressiveness parameter as a set of values assigned to different class variables. In this scenario, rather than allowing the funds to alter their aggressiveness only marginally from their current level, the choice is based on which aggressiveness class yields the best results. We limit this second scenario to half the number of funds in the original base model to limit computational run times. The original ten funds with three options require 30 trials, whereas five funds each selecting from the others' aggressiveness classes require 25 trials. Ten funds in the latter scenario would require 100 trials.

Once each hedge fund selects its aggressiveness based on results from the inner-level simulations, the outer level of the nested model simulates one time step with these updated parameters. The noise traders' expenditure for the outer-level market is applied and the outer-level model then solves for the market-clearing price given the equations for wealth, investment, demand and return as in the base model.

3 VERIFICATION, VALIDATION & CALIBRATION

Verification, validation, and calibration (VVC) for agent-based models ensure that the results accurately represent the intention of the model, the right model has been chosen for the task, and the parameters are properly set. Verification ensures that no error prevented the program from running to completion or performed unintended computations. Validation is the process to ensure that the right model was selected for the given problem. Calibration is used to adjust the value of parameters. We first outline the VVC behind the base model; most of this can be drawn from the original paper (Thurner, Farmer, and Geanakoplos 2012). We then discuss VVC for the steps necessary to make the base model available to the nested model, and for the nested model itself. Twelve of Sargent's 15 techniques for VVC (Sargent 2007) were used as described in his work. We did not include animation, predictive validation or a Turing test. The other twelve tests break down into two subgroups: those that ensure plausibility of the model's initial design, and those that check numeric outcomes or inputs.

The plausibility provided by the face validity and rationale of the base model is what sets the stage for the rest of the VVC process for both the base and nested models. Face validity of the base model is first established through the credentials of the authors of the original model, established subject matter experts (SMEs) in finance, economics, and complexity theory. Further face validity is provided by the model's academic support from the journal that published the work and the 50+ papers that cite it, including the OFR's reference to the model as their example in the area of banks and asset managers in a discussion on ABM (Bookstaber 2012). The rationale of the base model is what satisfies Sargent's 'historical methods' and 'multi-stage validation' criteria. While the limitations of the model keep it simple, so do the underlying structure, function, and parameter choices. Most conform to common economic reason, and the more complex equations are drawn from supported literature. The authors, in their own words, "build the simplest model possible."

Once the model and structure plausibility are obtained, the second subgroup of VVC techniques ensures that the assumptions and logic considered plausible are implemented correctly and yield expected or reasonable values. The numeric validation of the original base model was made evident throughout the discussion in the paper. Degenerate tests of mispricings at different levels produced expected demand values, extreme price conditions were checked and found to generate outcomes that could be seen as reasonable risk reduction strategies, and price movements and distributions were seen to align with historic data trends. Many runs of the base model demonstrated internal validity by telling the same economic story across various outputs. Operational graphics were used throughout the paper, and some of the sensitivity analysis was discussed.

Numeric VVC for the nested model began with a comparison to the base model using common random numbers. We ensured that the base model and the nested model without choice return the same output. We then used trace testing to validate the option of choice itself by ensuring the outer-level hedge funds were generating their trial options successfully and correctly interpreting and implementing the results. We also compared the inner-level model to the base model using set inputs and common random numbers. This ensured that the outer-level funds' tool for testing their options and observing their endogeneity in the system was operating correctly.

Beyond testing the validity of the structure of choice for the outer level of the nested model, we tested the amount of information available to make that choice. We performed sensitivity analysis in two areas: the number of inner-level model iterations required per trial of β to return consistent results across different random market environments, and the required length in time steps of each inner-level model iteration. These two aspects of the information provided to the outer-level hedge funds are intertwined. We needed to ensure that the amount of information provided to the hedge fund about the success of each trial of β would yield a noticeable trend above the noise of the random market simulations. If an outer-level fund could not make a consistent choice in its trials of options, then we could not expect the power of choice to produce meaningful deviation from the base model. On the other hand, providing each hedge fund less information for each trial would reduce computational complexity and duration, allowing us to perform more validation and experimental runs of the nested model. These sensitivity tests were performed from the perspective of an outer-level hedge fund running its options in both the ten-fund three-choice (10x3) and five-fund five-choice (5x5) scenarios.

We ran the nested model using a single inner-level model run to inform the choice of the hedge funds based on this sensitivity analysis. Each fund should make the same choice on average if given only one inner-level trial run as they would make with 100 trials, so an effect should still be observable. Furthermore, limiting the number of trial runs required for each of 25 or 30 choices at each time step would drastically increase the time available to run multiple replications of the nested model. However, in implementation, when the outer-level funds are running the inner-level simulations in real time and are not themselves the subject of experiment, more trials per choice would be used to ensure greater precision in judgment.

The final analysis on quantity of available information for the hedge funds' decision is on the length in time steps of each of the market simulations. The base, outer-level, and inner-level models each iterate across T time steps. Since the outer-level model is running the inner-level model at each of the outer-level time steps, the influence of T on the running time is quadratic in the nested model. Consequently, we performed an internal validation using operational graphics and ran several tests of the base model to find the minimum number of time steps required to demonstrate the results validated in the original paper. We ran base models for values of T ranging from the original value of 50,000 down to 500 time steps. Figures 4a and 4b show this comparison for the original base model at 50,000 time steps and the base model at 5,000 time steps, which was the smallest value at which the base model still demonstrated its validated relationships and sufficient market dynamics.

We used the same approach to validate the nested model against the base model (Figure 4c). The validity of the base paper's conclusions should not alter based on the modeled perceptions of the market participants: the change in price should still be more clustered than the change in noise demand, and both should vary randomly around their means. In addition, a buildup of leverage should correlate with a market crash. Similar to how we ensured these patterns were visible in the time-step-limited version of the base model, we verified these claims in the nested model.

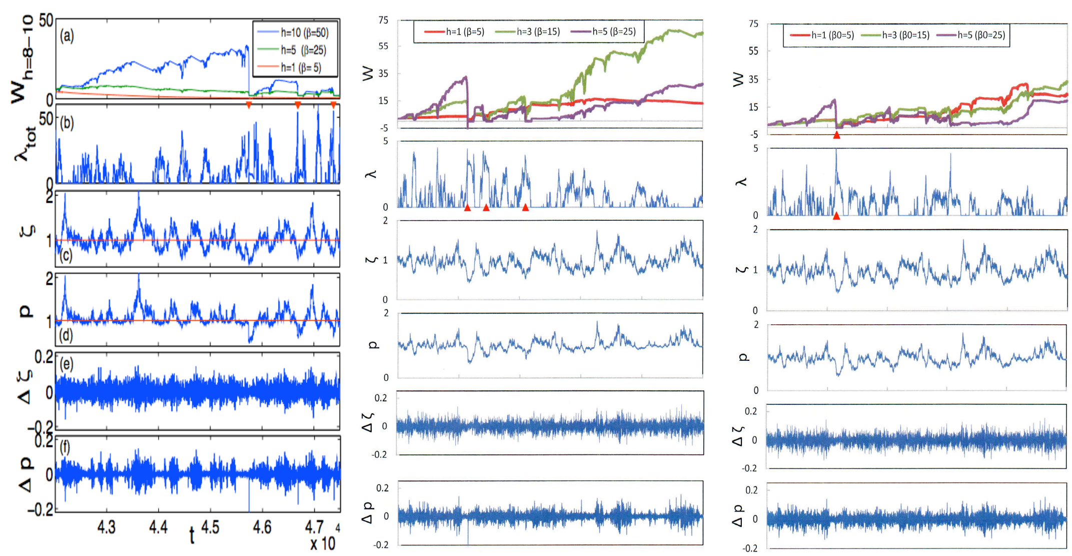

(a) (b) (c)

Figures 4a-c: Original 50k / Base 5k / Nested Comparison. Figure 4a is taken from the base paper, pg. 14, using 50,000 time steps. Figure 4b was generated by our base model in finding the minimum of 5,000 time steps required for comparable market behavior to the validated base model. Figure 4c is a market's worth of results of the 5-fund 5-change nested model set in the manner of the base paper to demonstrate the preserved relationships of the validated base model.

4 EXPERIMENTAL RESULTS & INSIGHTS

We developed and validated the nested ABM of the proposed regulatory scheme to determine whether its implicitly endogenous structure would lead market participants towards a calmer market. A few key metrics are of particular interest. Price volatility, market crashes, and bankruptcy behavior provide an account of effective market stability. Measures of leverage in the marketplace provide information on the risk underlying this model's market. Finally, individual and overall market wealth and returns address the issue of incentivized participation. We compared the base and nested models in 30 simulated markets of 5,000 time steps each using an unpooled two-sample t-test, assuming unequal variances. We also ran a paired t-test with comparable results.

The results were mixed by the variables measuring effective market stability, but overall show a slightly more stable market when the agents operate under our proposed regulatory structure (Table 1). Over the 30 markets, the base model with five hedge funds generated 87 bankruptcies, whereas the base model of ten funds generated 495. The five-fund five-choice (5x5) nested model generated approximately half the number of bankruptcies as its comparable base model with only 45 bankruptcies, and the ten-fund three-choice (10x3) nested model also came in with fewer at 441.

Table 1: Effective market stability as measured in three variables averaged across the 30 markets of the two fund/choice scenarios of the nested and base control models

Model	Total Bankrupt Funds	Total Market Crashes	Price Std. Dev.
Base 5	**2.9 (0.6)**	0.6 (0.1)	0.21 (0.01)
5x5	**1.5 (0.5)**	0.7 (0.2)	0.23 (0.01)
Base 10	16.5 (3.0)	2.0 (0.3)	0.21 (0.01)
10x3	14.7 (2.0)	2.5 (0.4)	0.21 (0.01)

NOTE: Bold values are significantly different between base and nested versions at the α = .05 level. Standard errors are reported in parentheses.

Another important factor is the underlying tension that drives and heightens the effective market volatility. Overleveraged markets are at a greater risk of experiencing a downturn as banks' margin calls put the price into a downward spiral (Thurner, Farmer, and Geanakoplos 2012). In the nested models, this risk is reduced (Table 2).

Table 2: Aggregate market risk as measured in four variables across the two fund/choice scenarios of the nested model and their base model controls.

	Average Across Markets		Max Across Time Steps	
Model	Average Leverage	Std. Dev. Leverage	Overall Leverage	Fund-level Leverage
Base 5	**1.19 (0.04)**	**0.89 (0.02)**	5.7	10.8
5x5	**0.84 (0.04)**	**0.61 (0.03)**	11.5	15.3
Base 10	1.41 (0.04)	1.03 (0.03)	15.9	20
10x3	1.32 (0.06)	0.99 (0.05)	16.6	20

NOTE: Bold values are significantly different between base and nested versions at the α = .05 level. Standard errors are reported in parentheses.

The greater effective market stability and lower underlying risk were generated through the collective self-interest of endogenously acting hedge funds. We then measure the success of their attempts to improve their market position. The values of wealth and return given in Table 3 show that the 5x5 nested model is less profitable than the base five-fund model by the average wealth and return in the market place across all hedge fund positions.

Table 3: Aggregate wealth & returns across the two fund/choice scenarios of the nested model and their base model controls

	Average Across Markets		Extremes Across Time Steps			
Model	Average Wealth	Average Return	Max Overall Return	Min Overall Return	Max Fund Return	Min Fund Return
Base 5	**18.23 (0.82)**	0.09% (0.004%)	8.7%	-17.2%	14.6%	-28.6%
5x5	**13.84 (0.99)**	0.09% (0.003%)	32.2%	-14.2%	58.8%	-20.4%
Base 10	**8.7 (0.34)**	0.04% (0.006%)	30.2%	-33.8%	100.9%	-61.5%
10x3	**7.41 (0.31)**	0.03% (0.005%)	16.6%	-29.0%	122.7%	-84.4%

NOTE: Returns measured include returns of 0 from bankrupt funds. Bold values are significantly different between base and nested versions at the $\alpha = .05$ level. Standard errors are reported in parentheses.

In both base models there is a strong positive correlation between the number of bankruptcies and average wealth achieved by each fund: 0.94 for the five-fund base and 0.73 for the ten-fund base model. That correlation is lost in the nested models, dropping to -0.38 for the five-fund model and 0.18 in the ten-fund version. In the nested models, the risk and rewards are more evenly distributed as funds drop their fixed distinction by β to compete on improved Sharpe ratio. The Sharpe ratio will highlight the most profitable choice given the observed correlation between risk and increased profit. The nested models had both decreased risk and decreased profit across most funds, as well as an increased Sharpe ratio (Table 4).

Table 4: Fund-level Sharpe ratios averaged across the 30 markets for the five- and ten-fund models.

		Base 5	5x5	Base 10	10x3
Fund 1	$\beta_0 = 5$	19.7% (0.8%)	20.1% (1.4%)	**17.2% (1.2%)**	**9.4% (1.7%)**
Fund 2	$\beta_0 = 10$	18.9% (1.1%)	20.2% (1.6%)	**16.7% (1.3%)**	**6.0% (1.6%)**
Fund 3	$\beta_0 = 15$	17.5% (1.2%)	20.8% (1.5%)	**14.1% (1.8%)**	**8.2% (1.8%)**
Fund 4	$\beta_0 = 20$	**15.9% (1.3%)**	**20.7% (1.5%)**	12.9% (1.8%)	8.8% (1.8%)
Fund 5	$\beta_0 = 25$	**15.3% (1.4%)**	**21.7% (1.4%)**	10.0% (2.1%)	8.8% (1.8%)
Fund 6	$\beta_0 = 30$			9.2% (2.1%)	8.9% (1.9%)
Fund 7	$\beta_0 = 35$			7.6% (2.2%)	11.7% (1.8%)
Fund 8	$\beta_0 = 40$			7.4% (2.2%)	8.9% (1.8%)
Fund 9	$\beta_0 = 45$			7.8% (2.1%)	7.8% (2.1%)
Fund 10	$\beta_0 = 50$			8.1% (2.1%)	9.2% (1.8%)
	Overall	17.7% (1.1%)	21.2% (1.6%)	10.8% (1.9%)	6.4% (1.3%)
	Range	4.4%	1.7%	9.8%	5.6%

NOTE: Bold values are significantly different between the base and nested models at the $\alpha = .05$ level. Standard errors are reported in parentheses.

By introducing another level of competition, the nested model returned more evenly distributed Sharpe ratios across funds just as it had with the average wealth and bankruptcy measures. This evidence of increased competitiveness and the positive impact on the Sharpe ratios of most funds provides some incentive for firms to participate in a contributory data-sharing plan that would grant them access to an agent-based model of their market.

The results of the experiment across effective market stability and underlying market risk demonstrate moderate improvement over the base model. Letting the outer-level funds run more than one trial per examined β would improve the judgment of the funds and strength of the results. Even with limited trials, individual funds were able to use that data to their advantage. While improving the stability of the market, the market participants were able to use the new information provided in the proposed structure to compete for individual benefit.

5 CONCLUSION AND RECOMMENDATIONS

Some saw the market crash of 2007 as a failure of the market to regulate itself, whereas some saw it as a failure of government regulation. Whether private or government control is favored or feasible, greater tools of understanding must be in the hands of the market participants and market facilitators, the regulated and the regulating. An agent-based model may provide the best solution to the problem of measuring outcome paths in the nonlinear environment of financial markets. There would be increased awareness of endogeneity and interconnected risk in the market if the OFR could compose such a model to be put into the hands of the regulated market participants and their regulators. New avenues for inquiry and stress testing would be available for governance.

The aim of this work has been to begin the investigation into whether such a scheme has the potential value to make the effort involved in creating an agent-based model of the market worthwhile. Taking a very simple but validated agent-based model of the hedge-fund market as a stand-in for what the OFR might make, we developed a nested agent-based model of our proposal in which the hedge-fund market uses an ABM of itself to direct its own progress. We sought to discover if a market informed by an agent-based model of itself is more stable, and whether any loss in profit would be offset by the potential for greater stability.

Market stability was most significantly improved in the decline in bankruptcies and leverage build-up in the five-fund model. While other metrics were more muted or unchanged, by no measure was the more informed market less stable overall than the model of uninformed agents. Bankruptcy declines came with wealth declines as well, but by factoring in the improved risk profile to the portrait of profit, we find that most funds able to run a forecasting ABM had a Sharpe ratio better than in the base model. This increase in market stability and return, given the risk, was particularly notable for having been generated by a very dynamic market of agents using a parameter-static market of agents as their forecasting tool.

We would recommend an inquiry into what would happen if we replaced the parameter-static agent-based model the outer-level funds are using as a forecasting tool with some other known forecasting methods currently in use. Having a comparison of the market results of such methods when used in the outer level with the results of this first experiment would better inform the value of the ABM as a strategic market tool.

This work invites and supports such further inquiry into the use of agent-based models as excellent nonlinear financial modeling tools. To obtain the necessary expertise, data, and participation for the development and use of such a model, we recommend companion modeling between the OFR and financial firms to generate a government-facilitated agent-based model with secured contributed proprietary information. This government-maintained model would help the market participants see themselves as endogenous to the market while providing transparency to the regulator, opening the door for greater regulation of the markets both by the government and by the market participants themselves.

REFERENCES

Bisias, D., Flood, M., Lo, A. W., & Valavanis, S. 2012. "A Survey of Systemic Risk Analytics," Office of Financial Research Working Paper No.0001

Bookstaber, R. 2012. "Using Agent-Based Models for Analyzing Threats to Financial Stability," Office of Financial Research Working Paper No. 0003.

Chen, C., Iyengar, G., & Moallemi, C. C. 2013. "An Axiomatic Approach to Systemic Risk." *Management Science*, *59*(6), 1373–1388.

Geanakoplos, J., Axtell, R., Farmer, D. J., Howitt, P., Conlee, B., Goldstein, J., Hendrey, M., Palmer, N. M., & Yang, C. (2012). "Getting at Systemic Risk via an Agent-Based Model of the Housing Market." *American Economic Review*, *102*(3), 53–58.

Gentry, C. 2009. "Fully Homomorphic Encryption Using Ideal Lattices." In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing STOC 09*, *19*, 169-178.

Goldstein, J. 2011. "An Agent-Based Model of the Interaction between the Housing and RMBS Markets." Department of Computational Social Science, George Mason University, Fairfax, Virginia. http://www.css.gmu.edu/images/GoldsteinCEF2011.pdf

Lo, A. W. 2008. "Hedge Funds , Systemic Risk , and the Financial Crisis of 2007 – 2008," testimony, November 13, 2008, before the U.S. House, Committee on Oversight and Government Reform.

Outkin, A. V. 2012. "An Agent-based Model of the Nasdaq Stock Market: Historic Validation and Future Directions." In Proceedings for the Conference of The Computational Social Science of the Americas 2012. http://computationalsocialscience.org/wp-content/uploads/2012/09/Outkin2012.pdf

Sargent, R. G. 2007. "Verification and Validation of Simulation Models," In *Proceedings of the 2007 Winter Simulation Conference*, edited by S.G. Henderson, B.Biller, M.-H. Hsieh, J. Shortle, J.D. Tew, and R.R. Barton, 124–137.

Smith, V. L., Suchanek, G. L., & Williams, A. W. 1988. "Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets." *Econometrica*, *56*, 1119–1151.

Thurner, S., Farmer, J. D., & Geanakoplos, J. 2012. "Leverage Causes Fat Tails and Clustered Volatility." *Quantitative Finance*, *12*(5), 695–707.

AUTHOR BIOGRAPHIES

AURORA BRISTOR earned an M.S. in Applied Mathematics in 2013 from the University of Maryland, College Park. She has done work for CoreLogic in academic outreach and mortgage risk modeling. Her current interest is in the application of agent-based models to financial policy analysis. She works with Fannie Mae assessing credit risk on mortgage acquisitions. Her email is wsc14bristor@gmail.com.

SEAN BARNES is an Assistant Professor of Operations Management in the Robert H. Smith School of Business at the University of Maryland. He received his Ph.D. in Scientific Computation from the University of Maryland in 2012. His current research interests are modeling the transmission of infectious diseases, healthcare analytics, simulation, agent-based modeling, and network analysis. His work has been published in the INFORMS Journal of Computing, IIE Transactions on Healthcare Systems Engineering, and Infection Control and Hospital Epidemiology. His email address is sbarnes3@umd.edu.

MICHAEL C. FU is the Ralph J. Tyser Professor of Management Science in the Robert H. Smith School of Business, with a joint appointment in the Institute for Systems Research and affiliate faculty appointment in the Department of Electrical and Computer Engineering, all at the University of Maryland. His research interests include simulation optimization and applied probability, with applications in supply chain management and financial engineering. He has a Ph.D. in applied math from Harvard and degrees in math and EECS from MIT. He served as WSC2011 Program Chair, and is a Fellow of INFORMS and IEEE. His email address is mfu@umd.edu.