

ABSTRACT

Title of Document: REGULATION OF SYSTEMIC RISK
THROUGH CONTRIBUTORY
ENDOGENOUS AGENT-BASED MODELING

Aurora Jean Cassells Bristol, Master of Science,
2013

Directed By: Professor Michael Fu, Department of Decision,
Operations and Information Technologies,
& Professor Sean Barnes, Department of
Decision, Operations and Information
Technologies

The Financial Stability Oversight Council (FSOC) was created to identify and respond to emerging threats to the stability of the United States financial system. The research arm of the FSOC, the Office of Financial Research (OFR), 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 the 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.

REGULATION OF SYSTEMIC RISK THROUGH CONTRIBUTORY
ENDOGENOUS AGENT-BASED MODELING

By

Aurora Jean Cassells Bristor

Thesis submitted to the Faculty of the Graduate School of the
University of Maryland, College Park, in partial fulfillment
of the requirements for the degree of
Master of Science
2013

Advisory Committee:
Professor Michael Fu, Chair
Professor Sean Barnes, Co-Chair
Professor David Levermore

© Copyright by
Aurora Jean Cassells Bristol
2013

Table of Contents

Table of Contents	ii
List of Tables	iii
List of Figures	iv
Chapter 1: Motivation & Proposal	1
Chapter 2: Methodology	13
Chapter 3: Verification, Validation & Calibration	28
Chapter 4: Experimental Results & Insights	43
Chapter 5: Conclusion and Recommendations	51
Appendix	55
Glossary	59
Bibliography	62

List of Tables

Table 1 Parameter & Variable List	16
Table 2 The share of choices made at each trial count level that were the least aggressive option available	39
Table 3 Effective market stability as measured in three variables averaged across the 30 markets of the two fund/choice scenarios of the nested model and their base model controls.....	45
Table 4 Aggregate market risk as measured in four variables across the two fund/choice scenarios of the nested model and their base model controls.....	46
Table 5 Aggregate wealth & returns across the two fund/choice scenarios of the nested model and their base model controls	47
Table 6 Fund-level wealth and bankruptcy statistics averaged across the 30 markets for the five- and ten-fund models.....	48
Table 7 Fund-level Sharpe ratios averaged across the 30 markets for the five- and ten-fund models.....	50

List of Figures

Figure 1 Base Model Control vs. Experimental Nested Model	14
Figure 2 Flow Chart of 1 Time Step	20
Figure 3 Outer-level Parameter Update Flow	23
Figure 4 Base Paper Results	41
Figure 5 Mock Base Paper Results on Time Step Limited Model	42
Figure 6 Nested Model Results	43
Figure 7 Flow Chart of 1 Time Step Example	55
Figure 8 Hedge Fund Time Step Flows	56
Figure 9 Nested Ten Fund Aggressiveness vs. Wealth.....	57
Figure 10 Nested Five Fund Aggressiveness vs. Wealth.....	58

Chapter 1: Motivation & Proposal

Five years after “the [financial] system broke down” and “repercussions of [Lehman’s] collapse ... reverberated across our economy” (Rep. Waxman, 2008), systemic risk remains unregulated and unmeasured. The *Dodd-Frank Wall Street Reform and Consumer Protection Act* (Dodd Frank Act), intended to address that risk, was passed in 2010 yet is less than halfway finalized. Perhaps fueled by the failures of other post-crisis efforts to calm the markets, one of the fundamental pieces of legislation that developed the momentum to pass was the very first section of the Dodd Frank Act in Title I: the establishment of the Financial Stability Oversight Council (FSOC) and the Office of Financial Research (OFR).

The OFR is created as a research group to support the FSOC by providing it with the depth of knowledge necessary to perform a rigorous task.

The FSOC has three broad mandates: (1) to identify risks to financial stability arising from events or activities of large financial firms or elsewhere; (2) to promote market discipline by eliminating participants’ expectations of possible government bailouts; and (3) to respond to emerging threats to the stability of the financial system. (Bisias, Flood, Lo, & Valavanis, 2012)

The FSOC was not tasked with a limited mission to address one facet of a financial meltdown such as providing liquidity¹ to a market, balancing inflation, or managing the employment rate. Its task is to regulate systemic risk more extensively.

¹ See Glossary for a full list of terms.

“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). This task is comprehensive and challenging. We classify two groups of complicating factors to the measurement of systemic risk, both of which stem from its general definition: “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, & Moallemi, 2013) First, systemic risk is as multi-faceted and highly dynamic as the complex system that it threatens. This is a complication of the theoretical mechanics of measurement. Second, the individual entities intrinsic to the system and collectively generating the risk have the richest and most detailed individual data on themselves, but no incentive to pool that data for the measurement of systemic risk. This presents a complication of the inputs to measurement. We propose a regulatory structure designed to overcome these specific challenges of multi-dimensional dynamic measurement and individual incentive restricting available inputs.

We propose a cooperative structure between the FSOC and the participants of the financial system it is designed to oversee using an agent-based model (ABM). Below, we detail the nature of the two factors complicating the measurement of systemic risk and how an ABM in use in a contributory structure would address those issues. 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. We further detail an intriguing potential consequence of the proposed structure. It is then the aim of the remainder of this thesis to investigate whether that consequence results

in the mitigation of some of those market threats that the proposed structure is designed to help monitor.

The first challenge to the theoretical mechanics of systemic risk measurement is its multi-dimensional aspect. Although there has been a large increase in the volume of research on systemic risk, this research has not been straightforward, cohesive, or even easily catalogued. To what breadth and depth of detail systemic risk should be measured is not clear. In the first working paper of the OFR (Bisias et al., 2012), 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.

While each of the approaches surveyed in this paper is meant to capture a specific challenge to financial stability, we remain agnostic about what is knowable. The system to be measured is highly complex, and so far, no systemic risk measure has been tested “out of sample”, i.e., outside the recent crisis. Indeed, some of the conceptual frame-works that we review are still in their infancy and have yet to be applied. Moreover, 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)

While each model offered a view on systemic risk, none of them completely incorporated any of the others or was itself a complete model of the risks involved. Even a catalogue of the models did not prove definitive, as the authors offered different taxonomies based on data requirement, supervisory scope, event or decision time horizon, and research method.

The second difficulty in the theory of systemic risk measurement is its dynamic nature. New instruments for financial flexibility and access are constantly being invented, altered, and refined. In the drive for profit and productivity, the envelope is continually pushed and occasionally reshaped: free checking, mortgage-backed securities, over-the-counter derivatives and credit-default swaps are all financial innovations. The risk is as shifting as the system to which it is intrinsically linked.

We propose the use of a tool as dynamic and multi-faceted as the systemic risk we need to define and measure: an agent-based model. As noted in the OFR's third working paper, most models do not address multiple dimensions and do not allow for emergent effects:

[A]bsent a model that can chart the course of events during financial disruptions, it is difficult to assess the value of these measures, especially given the many changes in the financial landscape that occur over time which lead to new vulnerabilities and paths through which shocks can propagate across the system. And existing models of the financial system are partial equilibrium models that are not built for this task. These models are designed to rebound back from shocks and return to the equilibrium state rather than running off the rails.

To make up for these weaknesses, a model of the financial system needs to be flexible, in order to provide results without extensive historical data either for parameterization or for testing, and in order to describe markets that fall into disarray.
(Bookstaber, 2012)

An ABM is one way to provide for the multi-dimensional and dynamic aspects of systemic risk in trying to measure it and to begin to meet the mandate for the OFR and FSOC.

In agent-based modeling (ABM), a system is modeled as a collection of autonomous decision-making entities called agents. Each agent individually assesses its situation and makes decisions on the basis of a set of rules. Agents may execute various behaviors appropriate for the system they represent – for example, producing, consuming, or selling. Repetitive competitive interactions between agents are a feature of agent-based modeling, which relies on the power of computers to explore the dynamics out of the reach of pure mathematical methods. At the simplest level, an agent-based model consists of a system of agents and the relationships between them. Even a simple agent-based model can exhibit complex behavior patterns and provide valuable information about the dynamics of the real-world system that it emulates. In addition, agents may be capable of evolving allowing unanticipated behaviors to emerge. (Bonabeau, 2002)

As an ABM is executed, a record or calculation could be made from any of its multiple variables, allowing it to report on multiple dimensions. This output would enable the examination of correlations between several different aspects of systemic risk and perhaps allow for the development of a single risk metric from any portion of the broad time series of data. Reporting on multiple variables allows the ABM to be multi-faceted, and reporting on those variables across the time steps of the simulation would capture dynamic shifts in the model. An ABM of the market could track different parameters of the regulated bodies and allow those parameters to be updated by the modelers as circumstances change. Further, the model could be written to be responsive to changes in input or even recoded to adapt to changes in the political landscape or in the world of financial innovation. Such a dynamic and responsive model 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 recreate and perform in

non-linear cases also makes ABMs valuable during the fast-paced and shifting environment of a crisis as it unfolds.

The second complication facing the measurement of systemic risk, the complication of the inputs to measurement, is less a matter of theoretical mechanics and more a matter of implementation. The incentive of the “individual component entities or agents that comprise the system” (Chen et al., 2013) has not been properly directed towards the measurement of systemic risk, ultimately limiting the available input data. The individual market participants would each like to measure and account for all of their risks. Yet due to their competitiveness as self-interested individuals, they do not have the collective incentive required to share their deep and narrow market data to create a collaborative set. Such collaboration would be required to provide the breadth and depth of data necessary for accuracy in a complex ABM of the market’s systemic risk. Our proposal guides individual incentive towards that collaboration by overcoming the roadblocks currently faced at an individual level. Despite a desire to minimize exposure to risk, individuals also face privacy concerns and competitive deterrents from working towards the public good of a stable marketplace.

The motivation of individual market participants to account for all possible risks is observed in their post-crisis behavior. Lehman Brothers would have had an interest in how its actions were impacting its risk profile long before it collapsed. Had the company known how to appropriately hedge or avoid those risks that bankrupted it or even had a better way to manage events as they unfolded, the course of the crash could have been very different. Yet even now, with Lehman and others

collapsed or bought out of the marketplace, the market participants have not demonstrably measured how they might be impacted by the risk of the system. That is, while the incentive is there, subsequent developments have not materialized. Instead, market participants have exhibited an attempt to control their exposure to systemic risk by taking what independent action they can in severely restricting other aspects of their individual risk profiles.

The need for data privacy and security by financial market participants is the clearest roadblock to data collaboration. In a fast-paced world of financial innovation, trade secrecy is an important protection. One firm cannot obtain proprietary data from any other, nor will it divulge its own market positions, profiles, or networked relationships to its competitors. A complete model of systemic risk would have to rely on data from the whole system, but broad coverage of proprietary market data is not something to which any single market participant has access. Individual participants and researchers do not have access to sufficiently granular data to incorporate the details of the market into a model of risk.

In addition to privacy and secrecy challenges, there is a game-theoretic roadblock to an individual market participant's incentive to define and measure systemic risk. As a public good, security from systemic risk is not something any one individual can benefit from promoting for its own sake. The efforts of any one firm cannot bring it about, and in the meantime those same efforts may make the firm less profitable and drive it from the marketplace.

Government, on the other hand, is best suited to deliver public goods not otherwise generated by the private market. The passage of the Dodd Frank Act and

the creation of the FSOC and OFR, like the Federal Reserve, Federal Deposit Insurance Corporation, and other government-sponsored enterprises before them, were attempts to fill this role. The market as a whole and the individuals and governments impacted by it have the incentive and can together develop the means to measure systemic risk even though the individual market participants may not be sufficiently incentivized or able. This can take the form of the means to write and enforce rules designed to pull proprietary data from market participants to force their participation in the government's enactment of the public good, but this merely poses another negative externality onto the market. Our proposal is designed to take advantage of the strength of government as an impartial partner who can provide privacy protection and explicit incentive to participate in the regulatory structure. Rather than treat the government as a mere tool for enforcement, we see the opportunity for ubiquitous market engagement that can be used to help shift the incentive of individual firms.

To address the issue of harnessing private incentive to gain access to proprietary data delivered in good faith, we propose a contributory model hosted by the government. An example of this mutually beneficial data sharing exists already in the private sector. With privacy protections in place, individual firms deliver their data to a corporate data collector, creating a large market database. The collector uses that database for modeling and profit-seeking enterprises. Though the data providers cannot access the collective proprietary data of the market to build their own models, the data collector pays for the supplied data with access to its own models and that comparative market analysis made possible by the proprietary data collection itself.

The regulated market participants are the data providers while the OFR and FSOC act in the position of data collector and model distributor in our proposed cooperative regulatory structure. The model made possible by this collection of contributed proprietary data is the ABM of the market that measures systemic risk. Following the incentive structure in the private sector, the FSOC and OFR could incentivize accurate and timely data delivery from contributing market participants by offering model access and systemic risk analysis generated from the ABM. With access to the ABM of systemic risk, participants could query the model in advance of a particular set of actions to estimate the impact of their choices on their own risk profiles. As in the private sector arrangement, a private party wishing to design such a model for its own use would not have access to all of the data provided into the contributory model. Therefore, the firms' best approach would be to honestly participate in the delivery of data to the regulator and make decisions based on the results of the model generated by the FSOC.

For this model generation, the FSOC should 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. Each participant would be incentivized to provide a substantial amount of institutional insight and manpower to the model creation process if they believed they would benefit from a more valid model. The impact of each institution looking to promote its own benefit would in some way balance the potential biases of such modeling assistance. If sufficiently balanced, this help could remain a net benefit to the modeling process required of the FSOC.

The companion modeling approach is an iterative participatory process involving stakeholders, regulatory agencies, and researchers from multiple disciplines in a repeated looping through a three-stage cycle: field work and data analysis, model design, and computational experiments. Agent-based modeling and role-playing games constitute important aspects of this process. The objective is the management of complex problems through a continuous learning process rather than the delivery of definitive problem solutions. (Tsfatsion, 2005)

The contributory nature of our proposal addresses some of the most pressing practical problems facing implementation of an ABM of systemic risk. It incentivizes accurate delivery of the required proprietary data and simultaneously provides assistance in working through the complexity of that data and the ABM it supports.

The OFR proposed the use of an ABM to address its multi-dimensional and dynamic modeling needs (Bookstaber, 2012). We believe this approach is best, yet recognize that working with private corporations is a matter of addressing incentives. Thus, we suggest the OFR develop for use by the FSOC an ABM of the market and its participants to measure systemic risk. These market participants would constantly update the model with new transaction and valuation information. In turn, they would be provided access to the model and its results. The proposed contributory ABM regulatory structure 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 the regulators and regulated. Such a measurement tool would be “multi-dimensional, adaptive, real-time, [and] able to incorporate illogic of human choice” (Lo, 2008).

A potential consequence of our proposed structure is a shift in perspective by the market participants. Systemic risk as defined above is a phenomenon to which the market participants are endogenously related. That is, systemic risk is something they can bring about as an emergent effect from their combined individual actions. This kind of emergent effect is a natural outcome for ABMs, which support the idea of its agents as endogenous actors. Many other models, for computational simplicity, presume each institution is instead exogenous to the market. These financial models, and implicitly the market participants who use them, assume that the market will continue to set and derive prices with or without the modeled participant present. Such models do not support engaging in private actions to mitigate public risk, but support the idea of systemic risk as a negative externality. Individual firms participating in our proposed regulatory structure and utilizing access to an ABM would begin to implicitly consider themselves instead as endogenous to the market. They would see how their choices alter the market and impact their own risk profile.

This implicit shift in perspective on their place in the market opens the question of whether firms would be individually incentivized to move towards a more stable market. They would be able to see that from the collection of individual actions a stable market could be created in the same manner as a market collapse. 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.

In Chapter 2 we describe the experiment and the model we designed for its implementation. Chapter 3 describes the methods used to verify, validate, and

calibrate the model outlined in Chapter 2. In Chapter 4, we discuss the results of the experiment and what they can tell us about the potential benefits of our proposal. Finally, we make closing remarks in Chapter 5, including proposals for further study based on our results and modeling scheme.

Chapter 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 incentive for market participants to provide accurate and timely data to the FSOC. The acknowledged benefit is in the additional information that should help firms improve their risk strategies and profit margins. That same model access may also alter the firms' assumptions about their level of influence over the market. With an ABM's implicit view of agents as endogenous to the market, the firms using such a model may see the possibilities of greater security in working towards a stable market or greater profitability in working towards a more volatile market (Smith, Suchanek, & Williams, 1988). It is unclear whether there would be a sufficiently consistent motive for enough of 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 individuals. 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 model 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 an ABM. We refer to this regulatory model as the

nested model. The outer-level agents of the nested model are the market participants with choice in our regulatory structure, and the inner-level model represents each participant's simulation of its potential strategies. We provide a visual depiction of the base and nested structures in Figure 1.

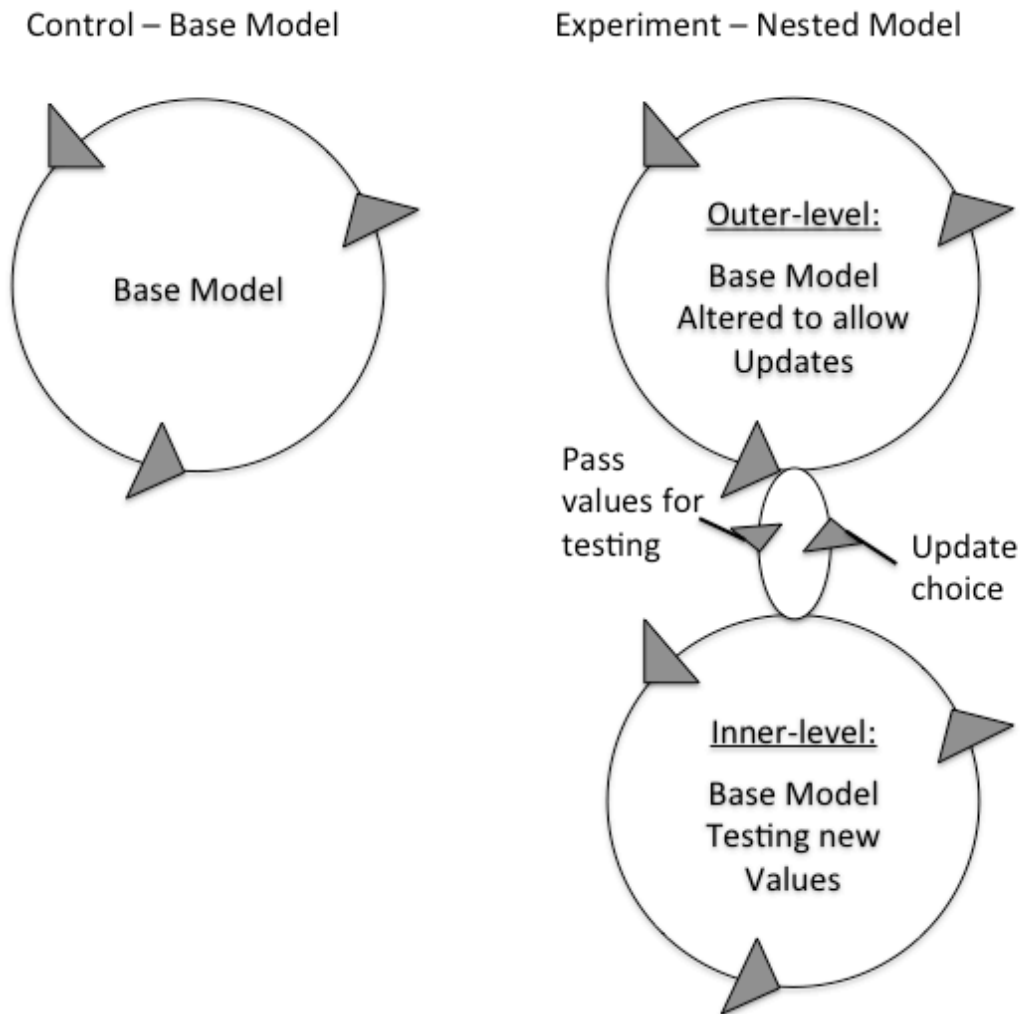


Figure 1 Base Model Control vs. Experimental Nested Model
 The unaltered base model is used as a control. As an agent-based model of the market, it is also the inner level of the nested model that the outer-level firms use to decide how to make their update. An altered version of the base model is used in the outer level of the nested model to simulate the firms' behavior in markets when those firms are given the capacity to update.

We build on the simple market ABM of Thurner, Farmer, and Geanakoplos (2012) as the base model for our experiment. Using a simple pre-developed model as the base minimizes assumptions and computational error. Within that simplicity, the base model is still able to demonstrate appropriate market volatility such that the proposed regulatory scheme would be able to show an improvement in 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 what changes were necessary to adapt the base model for 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 cannot short the stock or employ any other sort of optionality. They 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 based on its fundamentals. The noise traders, on the other hand, are called such because they represent traders who do not base buy and sell decisions on the fundamentals of the stock's value. Thus, they both make decisions based on 'noise' and generate 'noise' in the data of supply and demand for the stock. Their uneducated demand generates a mispricing 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 leverage to the hedge funds to buy more of the asset, but require the hedge funds to meet a margin call if they exceed their maximum leverage. That is, the hedge funds must sell some of the asset to maintain a required maximum loan-to-value ratio on the loan from the banks to the value of the stock held as collateral. The investors track the return on investment for each hedge fund and invest capital that can be leveraged with the banks into those funds whose forecasted performance is better than a set benchmark return. The parameters and variables used to calculate and control these activities are listed in Table 1.

Table 1 Parameter & Variable List

Fixed		Time Dependent		Hedge Fund Dependent	
a	Parameter for calculating Returns	m_t	Mispricing	C_t^i	Cash
b	Investor Sensitivity to Return	p_t	Price	D_t^i	Demand
H	Total Hedge Funds	ζ_t	Noise Trader Expenditure	F_t^i	Investor Funds
N	Supply	χ_t	Random Variate	r_t^i	Raw Returns
r_b	Benchmark Return			r_t^i	Forecast Returns
T	Total Time Steps			W_t^i	Wealth
V	Fixed Value of the Asset			β^i	Aggressiveness
W_e	Minimum Wealth				
W_0	Initial Wealth				
λ_{max}	Leverage Limit				
ρ	Parameter for calculating Noise Trader Expenditure				
σ	Random Walk Parameter for Noise Trader Expenditure				

All market activity is performed in a series of T time steps. These time steps are iterations of the simultaneous equations that define relationships between the four types of agents. At each time step t , the model solves for a market-clearing price p_t , which is the numerical solution to a system of simultaneous equations. To be a market-clearing price, p_t must generate sufficient demand such that the total sum of the demand for the stock from the hedge funds and noise traders is equivalent to the total supply N .

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 to represent the lack of fundamental analysis performed by such traders. Their demand is calculated in the equation for the market-clearing price as the amount of demand supported at a price by the given expenditure. As the price drops, the same expense can buy more stock, increasing the noise traders' demand. The hedge funds' demand values are also interdependent with the market-clearing price, but through a more complicated set of nested equations. These equations are designed to incorporate the incentives of the hedge funds and the other two agents in the model, the bankers and investors.

The nested equations cover investor returns and flow of capital, hedge fund wealth, and the hedge fund demand as 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 change in price, the associated 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.

Independently of the market-clearing price, the noise traders generate an amount ζ_t that they will spend to purchase the asset. The expenditure is solved for from the equation

$$\ln(\zeta_t) = \rho * \ln(\zeta_{t-1}) + \sigma * \chi_t + (1 - \rho) * \ln(V * N) \quad (1)$$

where χ_t is a random variate drawn from a normal distribution with mean 0 and standard deviation 1, V is the value of the asset, and N is the asset supply. In the absence of the hedge funds and with $\rho < 1$, this expenditure would generate a price which is mean reverting around the fixed value V . With χ_t distributed $N(0,1)$, the product $\sigma\chi_t$ is distributed $N(0, \sigma^2)$ and $\ln(\zeta_t)$ is normally distributed around a weighting of its prior value and $\ln(VN)$. Being the only time-variant value independent of the price, this expenditure determines the nature of the market exogenous to the hedge funds.

With this value for expenditure fixed, the total noise demand given by

$$\frac{\zeta_t}{p_t} \quad (2)$$

is inversely proportional to the price p_t . The market-clearing equation must find this price at time step t such that the sum of the noise trader demand and fund demands D_t^i exactly equals the asset supply N .

$$y = \frac{\zeta_t}{p_t} + \sum_{i=1}^H D_t^i - N = 0 \quad (3)$$

With the hedge funds only allowed to take a simple long position in the market, hedge fund demand is only generated when the price p_t is less than the fixed value V , generating a mispricing $m_t = V - p_t > 0$. Then, an individual hedge fund i has its demand given in each time step t by

$$D_t^i = \min\left(\frac{\beta^i * m_t * W_t^i}{p_t}, \frac{\lambda_{max} * W_t^i}{p_t}\right) \quad (4)$$

where β^i is the hedge fund aggressiveness parameter, W_t^i is the fund's wealth at time step t , and λ_{max} is the maximum leverage allowed for all funds by the banks. Thus, the demand is proportional to the wealth, mispricing on the asset and aggressiveness of the hedge fund, but limited by the maximum leverage the banks allow for the funds.

It can be seen that the demand (4) is a function of the price and that the price (3) is a function of the demand. The demand also depends on all of the other functions, and so acts as the conduit for their influence over the price. This is shown in Figure 2 where we depict the logical layout of the set of simultaneous equations that define the relationships between the four types of agents.

The wealth at time t is also directly impacted by the price. Wealth is defined as the value of the held assets plus cash, where cash will have a negative value if the fund is borrowing from the banks (5). We take the sum of the previous period's wealth plus the increased value of held assets and this round's flow of cash from investors, F_t^i , to determine the amount of wealth available to each hedge fund i for purchasing stock in the current round t at price p_t (6).

$$W_{t-1}^i = C_{t-1}^i + D_{t-1}^i * p_{t-1} \quad (5)$$

$$W_t^i = W_{t-1}^i + (p_t - p_{t-1}) * D_{t-1}^i + F_t^i \quad (6)$$

Substituting, we see that (6) can also be written as the sum of the previous period's cash with the revised retail value of the currently held assets under price p_t and the current influx of cash from investors F_t^i (7). This leaves only two avenues to grow or lose wealth in a new round: the change in price and the influence of the investor's capital.

$$W_t^i = C_{t-1}^i + D_{t-1}^i * p_t + F_t^i \quad (7)$$

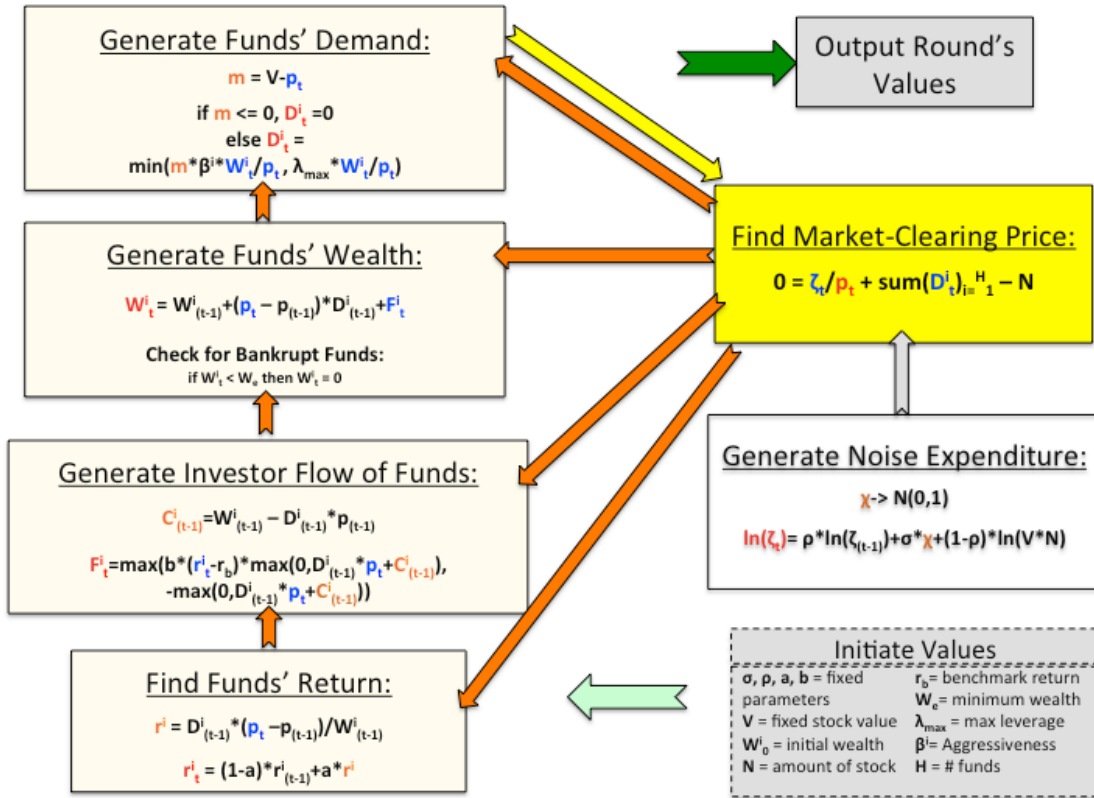


Figure 2 Flow Chart of 1 Time Step

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. A numerical example of this flow for a 3-fund market is provided in the appendix.

The flow of capital from investors F_t^i , meanwhile, is calculated as

$$F_t^i = b * (r_t^i - r_b) * (D_{t-1}^i * p_t + C_{t-1}^i) \quad (8)$$

The investors will invest cash into a hedge fund if they expect that fund to generate returns on their capital larger than a benchmark return r_b that they could expect to receive on average in some other market. In the event that the expected return r_t^i is greater than the benchmark return, this equation returns the positive amount the investors put into the hedge fund proportional to the parameter b that determines the investors' sensitivity to return, the margin of their success, and the fund's pre-investor wealth. In the event that the expected return r_t^i is less than the benchmark return, this value will be negative and the investors will attempt to withdraw that amount. This withdrawal is capped such that the maximum amount of wealth investors can take from a hedge fund is its total wealth.

The expected return r_t^i used by the investors to determine their flow of funds is a moving average on the actual return r_t^i . This is a “well-documented” behavior according to the authors, citing Busse (2001), Chevalier and Ellison (1997), and others. This behavior “guarantees steady-state behavior with well-defined long-term statistical averages of the wealth of the hedge funds” (Turner, Farmer, & Geanakoplos, 2012).

$$r_t^i = (1 - a) * r_{t-1}^i + a * r_t^i \quad (9)$$

With $a < 1$, this weighted average sets the expected return to be an adjustment on the prior time step's expectation in the direction of the actual return. As a approaches 1, more value is placed on the actual return. In the base model, we set a to 0.99. The actual return r_t^i is the product of the demand and the change in price as a share of the

wealth in the previous time step. This is the final equation in the set of nested simultaneous equations generating each fund's demand interdependent on price.

$$r_t^i = \frac{D_{t-1}^i * (p_t - p_{t-1})}{W_{t-1}^i} \quad (10)$$

The base model also places a lower bound on the wealth for each hedge fund. If any fund falls below 10% of the initial wealth value W_0 , the hedge fund is assumed to be bankrupt. The model sets the fund's wealth and demand to 0, and after a set number of time steps, replaces the bankrupt fund in the market with a new hedge fund having the same initial parameters as the one that exited the market.

The nested model operates on the same market structure and agents as the base model. The only difference is that the nested model gives the hedge fund agents in the outer level the ability to change one parameter in the model that is descriptive of a characteristic within their control: their aggressiveness β in response to a market mispricing m . We modified the base model to include this option in the outer-level model.

Each hedge fund in the outer level is given access to the inner-level model to select their aggressiveness parameter for the next round. They run the inner-level model for different trial variations on the aggressiveness parameter and choose the β that returns the best outcome from those trials as diagramed in Figure 3. The inner-level model is structured exactly like the base model, running for T time steps on the four types of agents. While the base model is given the same inputs at the start of each market simulation as the outer-level market simulation runs, the inner-level model is given different inputs. To better inform the outer-level funds, the inner-level

model runs T time steps on the four types of agents whose initial values are the values of the outer-level simulation at time t . That is, any one hedge fund in the outer-level runs its trials of the inner-level model using data from its current market environment. For each of the trials, a hedge fund experiments with a different β for itself, but the β 's for the other funds in the inner-level model are fixed at those of the other outer-level funds at time t . We take on the role of the FSOC in providing this updated data to the ABM for each fund to use for simulation.

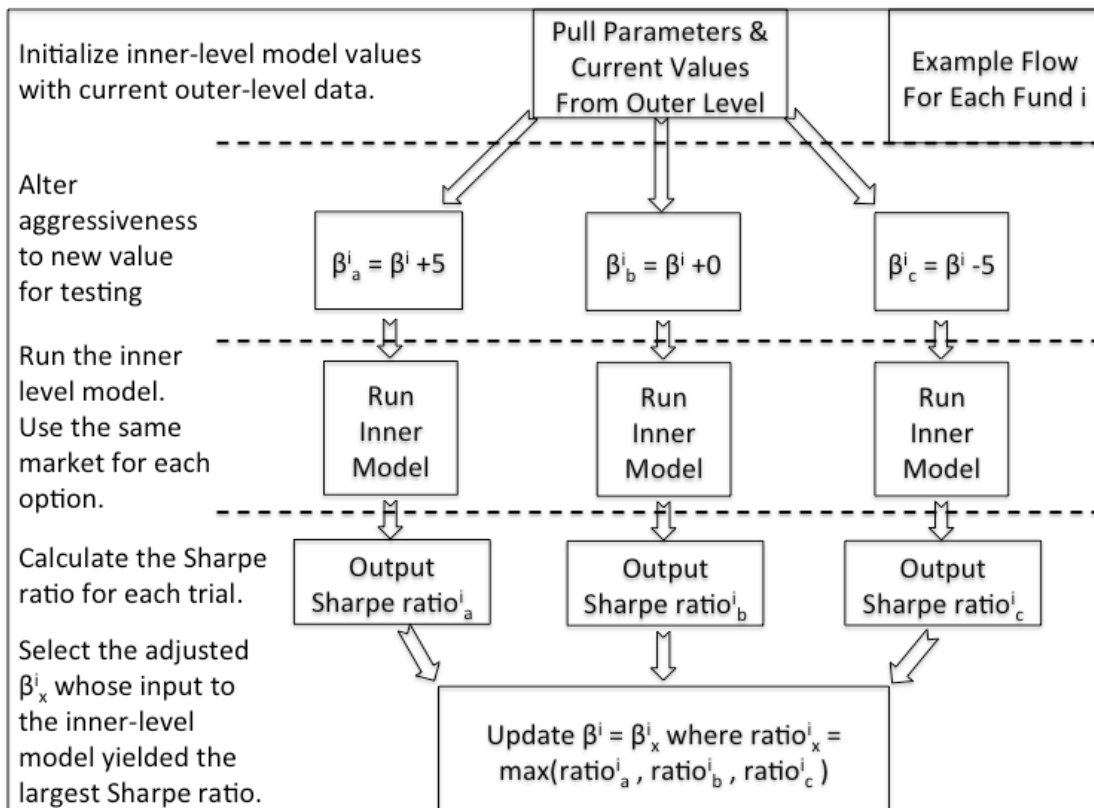


Figure 3 Outer-level Parameter Update Flow

The flow of the inner-level model in updating the β for an outer-level hedge fund i . The regulator prepopulates the outer-level values as the initial values for the inner-level simulation. Each hedge fund then alters its current β and runs trials of the inner-level model. The outer-level fund chooses to update its β to the one that yielded the greatest amount of return given risk as measured by the Sharpe ratio.

The markets used by the outer-level funds in the inner-level model must be the same for each trial of β to generate a fair comparison. These markets must also be different from the market in which the outer-level funds operate to simulate the firms' lack of information about the future of their own market, and different across the trials of the different hedge funds to simulate lack of collaboration or different market forecasts produced by different firms. These requirements on the trial markets can be met by controlling the random numbers that generate the noise trader demand. As shown above in Figure 2, the noise trader demand is the only equation exogenous to the hedge funds. From equation (1), 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.

Computerized random number generators generate random numbers from a seed value and a deterministic algorithm. Given the same seed value, the random number generator runs its algorithm and generates the exact same sequence of random numbers. We only need to ensure that the random number generators for the inner-level trials of an individual hedge fund are all given the same seed value in one time step in order to generate the same χ and thus the same exogenous markets for an even comparison of the fund's test β 's. To meet the other requirements on the markets, we need to ensure that this seed value is different from the seed values of the other hedge funds' trials at that same time step, and also different from the seed value for the outer-level market's χ .

This type of control is possible with a process known as common random numbers. This process controls the seed values to deliver repeatable random values

for each trial simulation. Not only does it allow direct comparison of trial β 's for each outer-level hedge fund as required for the experiment, this same method allows direct market comparison of the results for the base and nested models as detailed in Chapter 4. Clark (1990) discusses the use of common random numbers for such output comparison.

[Assume] that the simulation output is stochastic and has a distribution that is unknown to the analyst. This paper presents methods for making comparisons by either ranking alternatives or selecting the best alternative with respect to a single output performance measure. To increase the effectiveness of the comparisons, these methods are able to use the variance reduction technique of common random numbers. [...] Variance reduction techniques offer the potential for reducing the number of replications required for an experimental objective. [...] The ability to use the same random numbers for each alternative offers the potential for sharpening the comparisons. This is an example of a variance reduction technique and is called *common random numbers*. It is easier to apply than many of the other variance reduction methods. (Clark, 1990)

Although the influence of the noise traders is determined by controlled random values, the inner-level model run by the funds is otherwise prepopulated with the outer-level parameters at time t . This includes the wealth and aggressiveness of the other funds, which is proprietary information each fund requires to be kept unknown. It is not necessary to assume that the hedge funds in the outer level each have access to these updated proprietary values, however. While they are allowed access to the inner-level model to make assessments for their choice of β , the details of the model and its inputs are not made known to them – it is a black box operation. The regulator is the owner of the model in this scenario and prepopulates the market parameters of the inner-level model with the current details of the outer-level funds

without revealing that data to the outer-level funds. Not even the regulator needs to know the raw values. The regulatory system could operate in a fully-homomorphic encryption scheme, which is a scheme designed to provide encrypted results calculated on encrypted data (Gentry, 2009).

Once all options are run through the inner model, the outer-level funds must choose the value for β they wish to implement in the next time step of the outer-level model. There are many possible data points to use in optimizing the utility of the hedge funds. Wealth, return, and bankruptcy events are all model outputs from the agent-based model. However, these outputs may all return greatly different values depending on the random market on which the outer-level fund is running the inner-level model. Rather than look at any one metric of profit or loss, we measure utility using the Sharpe ratio.

$$\text{Sharpe ratio} = \frac{E[\text{Return}]}{E[\text{Risk}]} \quad (11)$$

The Sharpe ratio is a measure of risk-weighted return. It serves as a consistent measure across different risk and profit environments of the expected return for the same amount of risk. A higher Sharpe ratio would indicate that in the same risk environment, more profit could be made. Alternatively, for the same amount of profit, there is less accompanying risk. Each fund updates its β to the tested value of aggressiveness that produces the highest Sharpe ratio in simulated markets.

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 newly updated parameters. The noise traders' expenditure ζ_i for the outer-level market is applied and the outer-level model solves for the market-clearing price

given the equations for wealth, investment, demand and return as in the base model.

Though the outer-level agents implicitly develop an endogenous view of themselves in the market by running simulations of their own positions, the simulations they run are of agents still considering themselves exogenous to the market. However, the only way to model perceived endogeneity throughout all levels of simulation is to generate an infinite cascade of nested models. Therefore, we provide the outer level with an ABM of funds considering themselves as exogenous to the market for simplicity. This simplicity at the cost of accuracy may limit the capacity of the outer-level funds to appropriately judge the value of an update to their strategy and also dampen the potential market impact of the shift to an endogenous perspective by market participants. Nonetheless, models should test results from their simplest formulations before implementing more complex practice, so we retained this modeling assumption and proceeded with the validation process.

Chapter 3: Verification, Validation & Calibration

There is a continually developing process of validation, verification, and calibration (VVC) for agent-based models to 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. Many articles have been written on the verification and validation process, often aimed at specific computer programs or markets, or at larger industry. Even general principles are still being developed and debated (Matteo Richiardi, 2006). Although the concept is easy to describe, the execution is a set of processes as complex as the models it seeks to cover. VVC provides the surety of use required for a computer-programmed ABM that would be derived in mathematical models from a set of proofs and axioms.

Verification ensures that there are no ‘bugs’ in the computer program: that no error prevented the program from running to completion, and no quirk of syntax allowed the code to run to completion while performing unintended computations. Validation is the process to ensure that the right model was selected to address the given problem. Calibration is used to adjust the value of unmeasured or underived parameters.

Model Verification is substantiating that the model is transformed from one form into another, as intended, with sufficient accuracy. Model verification deals with building the model right. The accuracy of transforming a problem formulation into a model specification or the accuracy of converting a model representation from a micro flowchart form into an executable computer program is evaluated in model verification.

Model Validation is substantiating that the model, within its domain of applicability, behaves with satisfactory accuracy consistent with the [modeling and simulation] objectives. Model validation deals with building the right model.
(Balci, 1997)

Specific verification steps depend on the computer program being used, but involve checking output at successive phases of calculation. According to Sargent (2007), “To develop a valid simulation model, several versions of a model are usually developed prior to obtaining a satisfactory valid simulation model. During every model iteration, model verification and validation are performed.” Specific validation steps depend on the context of the model and the available parties for development. Calibration is an iterative and often numerical process. Sargent goes on to provide an alphabetical list of verification, validation, and calibration techniques that can be used subjectively or objectively, and usually in combination:

1. Animation
2. Comparison to Other Models
3. Degenerate Tests
4. Event Validity
5. Extreme Condition Tests
6. Face Validity
7. Historical Data Validation
8. Historical Methods
9. Internal Validity
10. Multi-stage Validation
11. Operational Graphics
12. Parameter Variability-Sensitivity Analysis
13. Predictive Validation
14. Traces
15. Turing Tests

For the base and nested model, twelve of the fifteen techniques were used as described in his work. Similar to the original base model authors, we did not include animation, predictive validation or a Turing test. Animation is a technique for

exhibiting the behavior of a system graphically over time such as the movement of parts through a factory, and there is no spatial component to the models to view for accuracy. Also, neither the base nor nested models are sufficiently detailed to make falsifiable predictive outcomes or to fool an expert into thinking its outcomes are true values for a Turing test.

The other twelve tests break down into two subgroups in our testing: those that are designed to demonstrate the plausibility of the model's initial design and those that are designed to check numeric outcomes or set numeric inputs to ensure the model behaves according to expectation. Face validity, the historical methods technique and multi-stage validation ensure plausibility. The rest are different angles on numeric output checks.

Face validity is having the support of subject matter experts (SMEs). SMEs can validate the model assumptions whether those assumptions are made about the structure of the model, its equations or parameters, or the data on which the model is built. Sargent's historical methods are rationalism, in which it is assumed that everyone knows whether underlying assumptions are true and reasonable; empiricism, which requires every assumption and outcome be empirically validated; and positive economics. "Positive economics requires only that the model be able to predict the future and is not concerned with a model's assumptions or structure" (Sargent, 2007). What he calls multi-stage validation is a combination of these three historical methods where possible. For our models, the outcomes cannot be tested empirically nor the future predicted, but the model assumptions are common enough to be supported by rationalism.

The second subgroup of validation techniques is intrinsically tied to the verification and calibration processes as all are meant to check the correctness of numeric values. These techniques can also overlap each other in scope, but for the sake of implementation in different environments, it is still good to have them listed separately. Comparison to other models requires a sufficiently comparable market model to test correctness of overall output. Degenerate tests look for appropriate outcomes as inputs approach the limits, while extreme condition tests require plausible outcomes from any combination of extreme or unlikely factors in the system. Also encroaching on the degenerate test is sensitivity analysis, which looks for appropriate model behavior in response to changing values of the inputs and internal parameters. Event validity asks if the occurrences of events in the model compare appropriately to real systems, overlapping the scope of historical data validation that tests if the model behaves in the same manner as available historical data. Internal validity involves performing several runs of the model to test for consistency across multiple outputs. Operational graphics and trace testing both involve tracking data points over time to ensure the data maintains the proper values. Operational graphics charts various performance measures, while trace testing follows the logic of particular agents or entities. Assuming that model and structure plausibility were obtained at the outset by the first subgroup, this second subgroup of VVC techniques ensures that the assumptions and logic considered plausible are implemented correctly and yield expected or reasonable values.

We first outline the VVC behind the base model, which is both the control for the experiment and also the core of the nested model. Most of its validation,

verification and calibration can be drawn from the original paper, but we add a description of the VVC for the extra steps necessary to make the base model available for use by the nested model. We then discuss the VVC of the additional components brought on by the creation of the nested model.

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 established through the credentials of the authors of the original model and the model's subsequent academic support. The authors are established SMEs in finance, economics, and complexity theory. Other subject matter experts have added face validity to the model by publishing the paper in a peer-reviewed journal and using it as a citation in more than fifty other published papers. The OFR added institutional face validity by naming it as the example model in the area of banks and asset managers in a discussion on ABM (Bookstaber, 2012).

The rationality of the base model is tied to its simplicity. Simplicity in modeling helps minimize the need for more complex VVC operations.

The goal of agent-based modeling is to enrich our understanding of fundamental processes that may appear in a variety of applications. It is important to keep the model as simple as possible. When a surprising result occurs, it is very helpful to be confident that one can understand everything that contributed to the model. Simplicity is also helpful in giving other researchers a realistic chance of extending one's model in new directions. While the topic being investigated may be complicated, the assumptions underlying the agent-based model should be simple. The complexity of agent-based modeling should be in the simulated results, not in the assumptions of the model.
(Axelrod, 1997)

The assumptions of the model are relatively few. There is no input data except for the data encoded at the beginning of the model and explicitly provided. Thus, it

cannot be missing, extreme-valued, or of a different distribution than the model specifications expect. The structural assumptions are also few, as the authors maintained a simple model. This simple rationality can be seen in the following areas:

- 1) The total wealth of a hedge fund is only composed of cash and one asset.
- 2) There are only two types of traders, neither of which is illogical.
- 3) The value traders can only take a long position on the asset, with no optionality.
- 4) Funds leverage from the banks at 0% interest.
- 5) Banks and investors are ‘deep-pocketed’, so there is no tightening of liquidity.
- 6) No accounting is necessary for how much stock was bought at which price for gains/loss analysis.
- 7) Investors, noise traders, and banks are each represented by one equation accounting for their influence as a representative aggregate without attempting to model individuals’ market choices.

While the limitations on the model keep it simple, so do the underlying structure, function, and parameter choices as described in Chapter 2. 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.”

The numeric validation of the original base model was not described explicitly, but was made evident throughout the discussion in the paper. Many graphs and tables were provided, part discovery and part numeric confirmation of the simple

rational design. 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. An example is the value of the benchmark return, r_b , which was determined through sensitivity analysis and discussed in general terms.

If the benchmark return is set very low then funds will become very wealthy and will buy a large quantity of the asset under even small mispricings, preventing the mispricing from ever growing large. This effectively induces a hard floor on prices. If the benchmark return is set very high, funds accumulate little wealth and play a negligible role in price formation. The interesting behavior is observed at intermediate values of r_b where the funds' demand is comparable to that of the noise traders. (Turner et al., 2012)

Accepting the VVC of the base model as presented in the original paper, we still needed to retest the model with the numeric subgroup of tests. We received the validated original base model as a set of MATLAB scripts but needed to translate that into a working SAS model. We planned to build the nested model in SAS once the base model was established, and so we needed the base model to be implemented in SAS to use it as a control with common random numbers. We performed iterative rounds of verification and validation as we converted the base model into SAS.

The most heavily used VVC techniques in this part of the process were trace testing, operational graphics, and the comparison to other models. We repeatedly used trace testing to make sure that the agents in the SAS base model were following

the equations that defined their behavior in the proper sequence, as well as operational graphics to check that the patterns claimed in the original base paper were generated by the translated model. With the SAS model running as expected, we compared it to the validated MATLAB version on specified inputs. The equivalence of the SAS and MATLAB base models anchors the new version to the VVC from the original paper.

The only structural difference between the outer level of the nested program and the base model is that the funds in that outer level are given the power of choice, whose power and effect we wish to measure. Plausibility for the structure of this choice is based on rationalism. The numeric VVC we performed prior to generating experimental results checks numeric outcomes and sets numeric inputs to ensure the model behaves according to expectations. It will be in the examination of the results that the final validation is implicitly performed: checking for reasonable values returned from changes to the inputs, event validity, or extreme conditions.

Giving the outer-level hedge funds the power of choice at each time step 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 could test whether it would be better off with one of its competitors' values, we give each fund the option to change its aggressiveness by +/- 5, or not at all.² We constrain the aggressiveness to be at least five because a value of zero or a negative value is

² We also tested different percentages by which to let the funds change their β . This led to exponential growth in aggressiveness with no measurable market impact.

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. Operational graphics showed that this occasionally generated values in the hundreds. We also decided to experiment with the aggressiveness parameter as a set of values assigned to different class variables, one for each possible level of aggressiveness. In this scenario, rather than allowing the funds to increase or decrease their aggressiveness in the manner of addition, the choice would be on which aggressiveness class yielded the best results. We limited this second scenario to half the number of funds in the original base model to limit the computational complexity. The original ten funds with three options require thirty trials, and five funds each selecting from five aggressiveness classes require twenty-five. One hundred inner-level trials would have been required in each time step t had we given all ten funds the choice of any of ten class values.

During each time step, the hedge funds calculate the Sharpe ratio (11) for each simulated aggressiveness option and select the β that delivers the greatest return given the risk. The Sharpe ratio is a simple business finance tool that is useful for comparing performance across risk profiles. We chose this to be the measure for our outer-level hedge funds because we anticipate that both the profitability and risk measures will change in the output of the inner-level model under different β 's.

Numeric VVC for the nested model began with a comparison to the SAS base model using common random numbers. Before performing our experiment looking for variation due to choice, we ensured that without choice the nested and base models 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. Sensitivity analysis was performed 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 without choice. 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 as if from the point of view of an outer-level hedge fund running its options in both the ten-fund three-choice (10x3) and five-fund five-choice (5x5) scenarios.

To test the impact of lowering the information available to the hedge funds on the consistency of choice, we recalculated the classic Sharpe ratio on subgroups of 100 markets. By using the same 100 markets throughout the test we were able to compare directly how the same information viewed differently might alter the pattern of choice. We started with the most informed choice based on all 100 markets. For five of the six tested markets, the aggressiveness that returned the highest Sharpe ratio was the least aggressive option available to the funds. With this option as the most selected, we then tested how often this same choice was made as the 100 markets were broken up into smaller subgroups providing less information per choice. We found that even looking at only one run of the base model, the funds chose the same option in the majority of cases as they had in the test using 100 runs. Further, fund 9 of the ten-fund market, which had initially chosen the middle of its three aggressiveness options when given 100 trials on which to base one choice, began choosing the least aggressive option in line with the other funds with greater frequency as it was given less information by running fewer trials per choice. This brought it into greater consistency with the other funds' selections as shown in Table 2.

We decided to run 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 reduce the computational complexity and maximize 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 accuracy of judgment.

Table 2 The share of choices made at each trial count level that were the least aggressive option available

		5 Fund Market, t = 0			10 Fund Market, t = 0		
		Fund 1	Fund 3	Fund 5	Fund 2	Fund 5	Fund 9
Aggressiveness β_0	5	5	15	25	10	25	45
	Trials/ Choice	# Choices			# Choices		
100	1	100%	100%	100%	100%	100%	0%
50	2	100%	100%	100%	100%	100%	50%
25	4	100%	100%	100%	100%	100%	25%
10	10	100%	100%	100%	90%	90%	20%
5	20	95%	100%	80%	90%	75%	20%
2	50	70%	74%	64%	68%	66%	36%
1	100	56%	65%	52%	54%	60%	40%

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. Figures 4 and 5 show this comparison

for the original base model at 50,000 time steps and the SAS 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 6). 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 travel 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 were still valid in the nested model. These results affirm that the nested model maintained the validity established in the base model.

We performed 12 of 15 of Sargent's VVC processes to ensure the surety of use of the nested model. We started by returning to the source and reiterating the VVC of the base model, and validated the proper transition of the base model into the SAS programming environment. The validated base model provided the platform against which to validate the nested model. We then verified, validated, and tested all aspects of the method and performance of choice in the nested model, from the rationale of the design to its appropriate implementation. These tests properly completed, we were then able to match the output patterns of the nested model with the claims of the base model, completing the VVC.

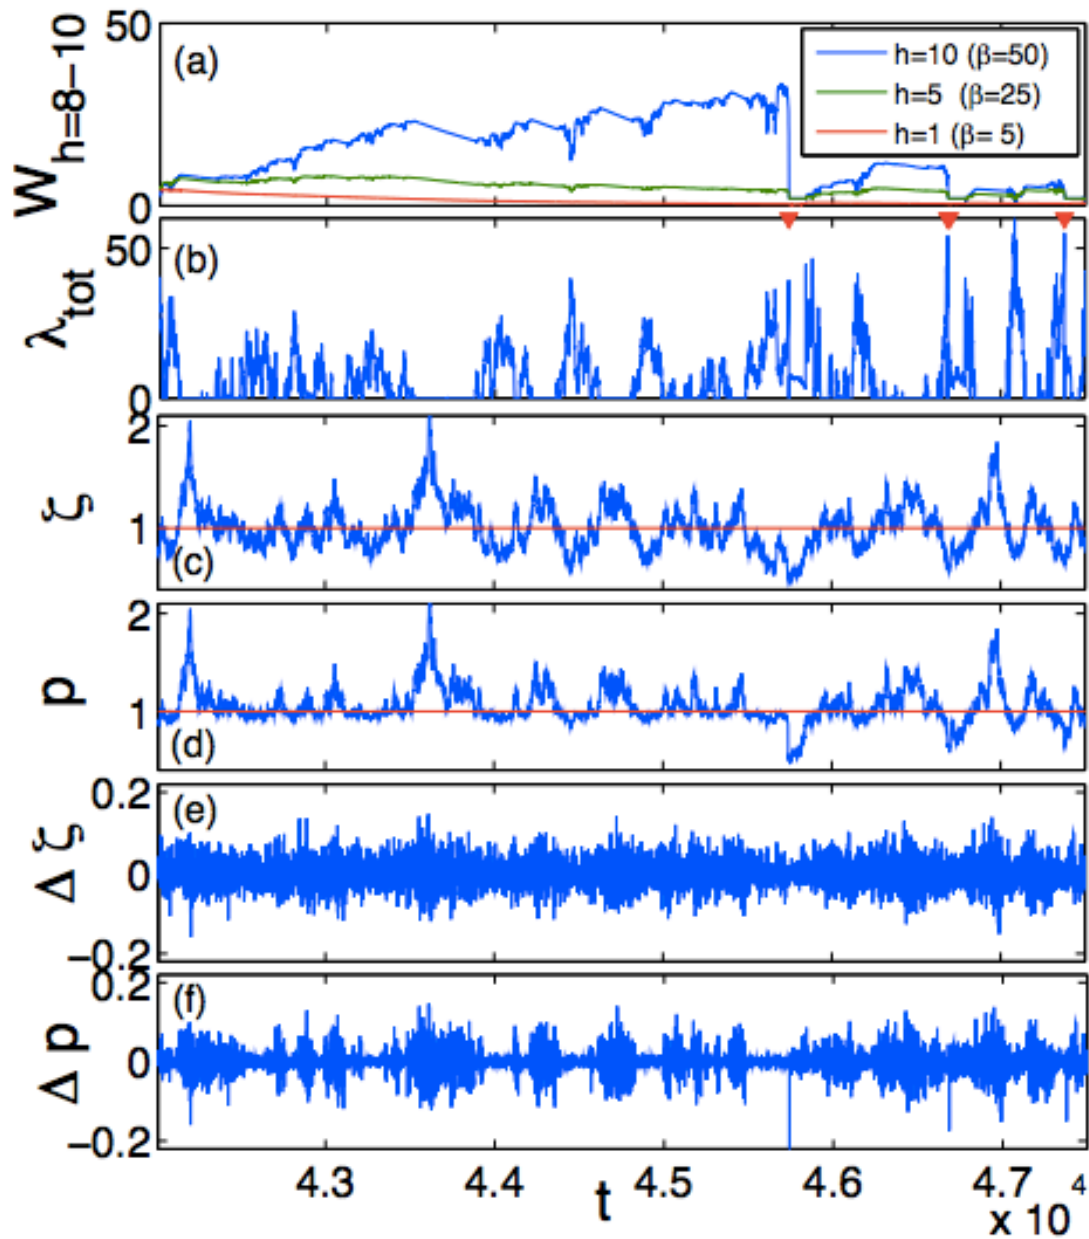


Figure 4 Base Paper Results

This figure taken from the base paper, pg. 14, with the most representative data for the performance of a market, using 50,000 time steps




Figure 5 Mock Base Paper Results on Time Step Limited Model
 One of several figures generated by our SAS base model in finding the minimum of 5,000 time steps required to generate comparable market behavior to the validated base model




Figure 6 Nested Model Results

A market's worth of results of the 5-fund 5-change complete nested model set in the manner of the base paper demonstrate the preserved relationships of the validated base model into the nested models.

Chapter 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 for this experiment. Price volatility, market crashes, and bankruptcy behavior provide an account of the effective market stability. These measures provide data on what took place in the market. The amount of leverage in the marketplace and its dynamics provide information on the market risk underlying this model's market. As the base paper showed, there is a relationship between increased leverage and market crashes. Finally, the individual and overall market wealth and investor returns address the issue of incentivized participation and whether the benefit of additional knowledge enhances performance for the funds. We compared the results of the base and nested models across these different metrics generated 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, we report the more conservative test.

We look first to the effective market stability measures as the most direct answer to the question of market impact under the regulatory proposal. The results, as shown in Table 3, were mixed by the different variables measured but overall show a slightly more stable market when the agents consider themselves endogenous to the market. 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 almost half the number of bankruptcies as its comparable base

model with only 45 bankruptcies over the 30 markets, and the ten-fund three-choice (10x3) nested model also came in with fewer at 441. The drop in bankruptcy rate of the 10x3 model is not significant, but the drop in bankruptcies for the five-fund nested model is significant at the $\alpha = .05$ level. There were more market crashes – defined as a collapse in total market wealth of 50% or more – in both versions of the nested model than in their correspondent base models. Neither difference was statistically significant ($\alpha = .05$). In many cases, multiple market crashes were counted in less than fifteen time steps of the same market. This relative closeness in time makes these market crashes effectively part of the same market downturn, weakening the impact of the metric. Finally, comparing the average standard deviation of the asset price across the 30 markets did not turn up any statistical difference ($\alpha = .05$) between the base and nested model. The 5x5 and 10x3 models demonstrate greater or equal effective market stability than their base market controls in the aggregate, accounting for the markets where the nested models counted multiple market crashes in the same downturn.

Table 3 Effective market stability as measured in three variables averaged across the 30 markets of the two fund/choice scenarios of the nested model and their base model controls

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 $\alpha = .05$ level. Standard errors are reported in parentheses.

The model of endogenous market participants yielded increased effective market stability, but another important factor in true market stability 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 et al., 2012). In the nested models, this risk is reduced: the average leverage is lower and the standard deviation of leverage is smaller. These metrics are significantly more stable ($\alpha = .05$) in the 5x5 model than the five-fund base model. The 10x3 and 5x5 nested models had greater values for the maximum overall and individual fund leverage in any time step, but these outliers did not overshadow the statistically lower leverage of the endogenous agents in the five-fund model (Table 4).

Table 4 Aggregate market risk as measured in four variables across the two fund/choice scenarios of the nested model and their base model controls.

Model	Average Across Markets		Max Across Time Steps	
	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 $\alpha = .05$ level. Standard errors are reported in parentheses.

The greater effective market stability shown in the decreased bankruptcies and the lower underlying risk measured in decreased leverage were both generated through the collective self-interested actions of the endogenously acting hedge funds.

Looking at the funds' results, we measure the success of their attempts to improve their market position.

The values of wealth and return given in Table 5 show that the 5x5 nested model is less profitable than the base five-fund model when measuring the average wealth and return in the market place across all hedge fund positions, regardless of any one fund's price of bankruptcy. The lower average wealth of the 5x5 and 10x3 models is significant at the $\alpha = .05$ level. In both the 5x5 and 10x3 scenarios, funds do not have an incentive to participate in such a structure if the success of the nested funds is measured by wealth or return alone.

Table 5 Aggregate wealth & returns across the two fund/choice scenarios of the nested model and their base model controls

Model	Average Across Markets		Extremes Across Time Steps			
	Average Wealth	Average Return	Max Overall Return	Min Overall Return	Max Fund Return	Min Fund Return
Base 5	18.23 (4.5)	0.09% (0.02%)	8.7%	-17.2%	14.6%	-28.6%
5x5	13.84 (5.42)	0.09% (0.02%)	32.2%	-14.2%	58.8%	-20.4%
Base 10	8.7 (1.89)	0.04% (0.03%)	30.2%	-33.8%	100.9%	-61.5%
10x3	7.41 (1.68)	0.03% (0.03%)	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.

Whereas any one fund would like to generate the most wealth, it is also true that those with the greatest profits are the most prone to bankruptcy. This effect is observable in Table 6 in both the five- and ten-fund base models, where there is a

statistically significant difference ($\alpha = .05$) between the average wealth of the funds with the fewest bankruptcies and the funds with the most.

Table 6 Fund-level wealth and bankruptcy statistics averaged across the 30 markets for the five- and ten-fund models.

		Model			
		Base 5		5x5	
		Avg. Wealth	Avg. # Bankruptcies	Avg. Wealth	Avg. # Bankruptcies
Fund 1	$\beta_0 = 5$	4.93 (0.49)	0.00 (0.00)	13.70 (1.34)	0.33 (0.12)
Fund 2	$\beta_0 = 10$	13.51 (1.61)	0.13 (0.08)	14.08 (1.14)	0.37 (0.11)
Fund 3	$\beta_0 = 15$	17.65 (1.80)	0.53 (0.12)	13.44 (1.12)	0.30 (0.12)
Fund 4	$\beta_0 = 20$	21.06 (1.88)	1.00 (0.19)	13.42 (1.29)	0.30 (0.11)
Fund 5	$\beta_0 = 25$	34.00 (3.32)	1.27 (0.24)	14.55 (1.26)	0.23 (0.09)
Overall		18.23 (0.82)	2.93 (0.57)	13.84 (0.99)	1.53 (0.48)
Range		29.07	1.27	1.13	0.13

		Model			
		Base 10		10x3	
		Avg. Wealth	Avg. # Bankruptcies	Avg. Wealth	Avg. # Bankruptcies
Fund 1	$\beta_0 = 5$	3.57 (0.42)	0.00 (0.00)	7.53 (1.00)	1.17 (0.33)
Fund 2	$\beta_0 = 10$	6.73 (0.90)	0.10 (0.07)	5.22 (0.42)	1.9 (0.33)
Fund 3	$\beta_0 = 15$	6.81 (0.88)	0.63 (0.18)	6.97 (0.69)	1.17 (0.23)
Fund 4	$\beta_0 = 20$	6.42 (0.67)	1.03 (0.21)	5.95 (0.76)	1.23 (0.27)
Fund 5	$\beta_0 = 25$	5.97 (0.50)	1.80 (0.35)	7.47 (0.79)	1.23 (0.29)
Fund 6	$\beta_0 = 30$	7.10 (0.63)	2.10 (0.40)	6.81 (0.62)	1.90 (0.38)
Fund 7	$\beta_0 = 35$	8.12 (0.50)	2.50 (0.45)	8.18 (1.03)	1.23 (0.38)
Fund 8	$\beta_0 = 40$	10.39 (0.77)	2.73 (0.47)	7.70 (0.91)	1.70 (0.34)
Fund 9	$\beta_0 = 45$	13.78 (1.09)	2.77 (0.48)	10.90 (1.65)	1.97 (0.45)
Fund 10	$\beta_0 = 50$	18.12 (1.67)	2.80 (0.48)	7.36 (0.88)	1.20 (0.28)
Overall		8.70 (0.34)	16.47 (2.95)	7.41 (0.31)	14.70 (2.01)
Range		14.55	2.80	5.68	0.80

NOTE: Bold values are significantly different between base and nested versions at the $\alpha = .05$ level. Base model wealth and bankruptcy values are significantly different between the smallest β_0 and largest β_0 funds 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 outer-level funds based their choice on the Sharpe ratio (11) to account for the variation across markets demonstrated by such measures as wealth, bankruptcy, and returns. This measure will also highlight the most profitable choice given the observed correlation between increased risk and increased profit. The Sharpe ratio is designed to maximize the benefit for a given value of risk. The average of the final Sharpe ratios are shown in Table 7 for the base and nested versions to compare the risk-weighted profitability. The nested models had both decreased risk and decreased profit across most funds. Also across most individual funds the Sharpe ratio increased. Those funds that instead showed a significant decrease ($\alpha = .05$) were three of the four funds fixed at the least aggressive β in the ten-fund base model and whose incidence of bankruptcy increased in the nested version. These are imperfect correlations, however. Fund 1 of the five-fund model was fixed at the least aggressive β in the base model and experienced a significant increase ($\alpha = .05$) in the number of bankruptcies in the 5x5 model, but still increased its Sharpe ratio. 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 independent incentive for firms to participate in a contributory data-sharing plan that would grant them access to an agent-based model of their market.

Table 7 Fund-level Sharpe ratios averaged across the 30 markets for the five- and ten-fund models.

Sharpe Ratios		Model			
		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.

The results of the experiment across effective market stability and underlying market risk demonstrate moderate improvement over the base model. Some of the weakness of impact was likely from the inability of the endogenously acting agents to sufficiently predict their own outcomes from a single trial of each option of the base model. Running more trials per option would mitigate the uncertainty demonstrated in the sensitivity analysis. Even with limited information, individual funds were able to use it 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.

Chapter 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. As discussed, an agent-based model may provide the best solution to the problem of measuring outcome paths in the non-linear environment of the financial markets. There would be unprecedented awareness 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.

A model of such a scale is not likely to be ready soon, however. Even the cryptographic tools necessary for housing and computing the sensitive data that would be necessary are still under construction and discussion (Flood, Katz, Ong, & Smith, 2013). Andrew Lo suggested in his 2008 testimony before Congress that we allow market participants to calculate their own risk metrics to be delivered to the regulating agency to preserve privacy. Yet this practice invites fraud and error. Allowing the FSOC to calculate the measure of systemic risk based on submitted data instead promotes consistency and would allow the oversight body further calculations to provide it with a broader look at the market. Our proposal relies on such data submission and regulatory calculations to be done in privacy. As the OFR pointed

out in their first working paper and most recently developed in their eleventh, there is new and developing work in cryptography

that enable[s] individuals to maintain the privacy of their data through encryption algorithms that allow third parties to compute aggregate statistics across multiple individuals while preserving the privacy of each individual. [...] Although still in experimental stages of development, these so-called “secure multi-party computational” and “fully homomorphic encryption” algorithms will likely revolutionize the way in which systemic risk is measured and managed.

(Bisias et al., 2012)

For the purposes of this proposal, fully homomorphic encryption (FHE) would serve to let the participants send encrypted data that would not be decrypted in transmission or in computation. Yet at the same time, this data could serve as useful input to the FSOC’s calculations as if it had been unencrypted. The FSOC would need to be able to measure the contributors’ impacts on systemic risk as well as the overall network effects without seeing the specifics of a company’s data or business processes. Insofar as the data is housed, it would need to be maintained with a level of security no less than the most secure measures taken by the institutions themselves. Safeguards would need to be taken, but the protection of corporate secrets is already a technology in use and could be extended to the government’s access.

The aim of this thesis 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. The agent-based model run in the experiment of this thesis was allowed to use recent information from the other funds on the assumption that the market regulator would be the model facilitator, maintaining the fund-level data in a secure and encrypted format. A non-agent based model used for forecasting instead would be more linear than an ABM and would only have available the data from each fund's own history and general market information. Nonetheless, funds in the present real markets use many forecasting methods of such linear and information-limited scope to great advantage. 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 thesis invites and supports such further inquiry into the use of agent-based models as excellent non-linear 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.

Appendix

Base Model Time Step Walk-Through

The initialized parameters in the bottom right of Figures 2 and 7 are fixed at the beginning of each program. The only value generated independently of the others in every time step is the noise expenditure. All other values are generated by finding the market-clearing price. This price feeds directly into all other equations that are connected sequentially from the returns generated in the last round to the return's implications for the demand generated in the current round. This demand feeds back into the market-clearing equation for the price.

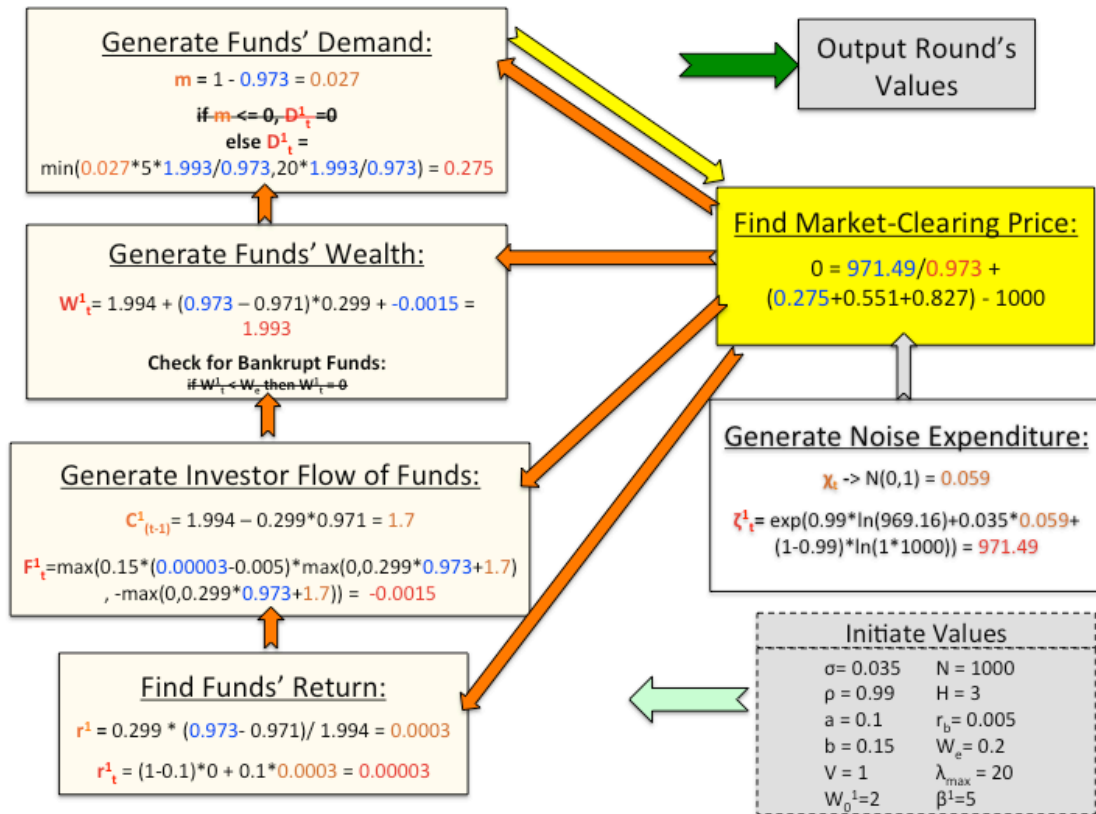


Figure 7 Flow Chart of 1 Time Step Example

In Figures 2 and 7, blue numbers are used to represent inputs generated elsewhere in the flow, orange numbers are generated and used within the same group of equations, and red numbers are the outputs from each equation set that serve as one of the following equation set's inputs. Black numbers are fixed inputs, either from the fixed parameters set at the beginning of the program or from values generated in the prior time step.

The numbers for the example in Figure 7 were generated on a 3-fund market. While the diagram depicts the flow for fund 1, there are two other funds dependent on and influencing the market-clearing price. All three funds are represented in the market-clearing equation by the contribution of their demands. The demand for fund 1 is the blue 0.275 summed together with the demands of funds 2 and 3 in the price block. These contribute to the red 0.973, which is the numerical solution to the simultaneous equation problem – the market-clearing price. The market-clearing price equation in a 3-fund market is the intersection of three otherwise independent flow sequences for each fund, visually depicted in Figure 8.




Figure 8 Hedge Fund Time Step Flows
The time step flows of all hedge funds interact through the market-clearing price equation.

10x3 Model Aggressiveness

The aggressiveness in the 10-fund nested model was allowed to grow without bound as funds chose to add 5, 0, or -5 to their parameter β . Though the average market aggressiveness showed dramatic shifts, this did not seem to have any direct impact on the other variables under observation. A graph showing a comparison of the wealth of funds in one ten-fund market along with the average β is shown in Figure 9. Figure 10 shows the same data for a five-fund market.




Figure 9 Nested Ten Fund Aggressiveness vs. Wealth

The wealth of five funds of the ten-fund nested model plotted along with the average market aggressiveness shows no signs of being impacted by the dramatic parameter shifts.

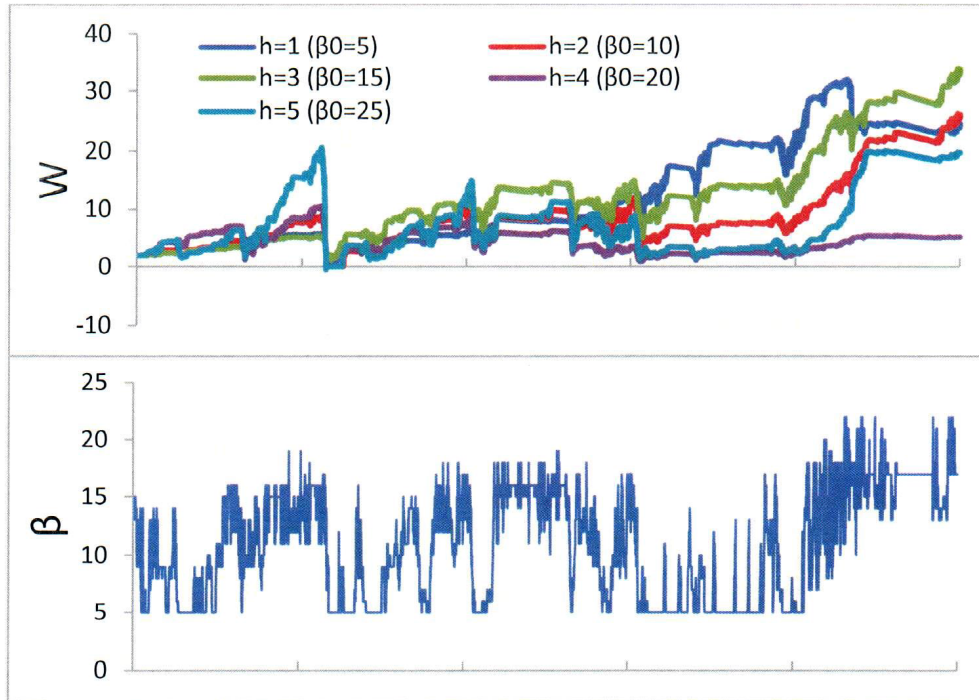


Figure 10 Nested Five Fund Aggressiveness vs. Wealth

The wealth of all funds in the five-fund nested model plotted along with the average market aggressiveness shows no signs of being impacted by the parameter, but the parameter is also much more stable than in the ten-fund model.

Though the rest of the output data from the 10x3 model is in line with expectation, the unbounded growth of the 10x3 model β 's does not line up with the understanding of the variable. If a normally diverse market contained funds with appetites ranging from a conservative 5 to an aggressive fund 50, what would market aggressiveness be when over 100? The aggressiveness represents the fund's demand for underpriced assets – the demand and price they are willing to pay. With numbers that far exceed the initial parameters, these values become meaningless. On the other hand, the aggressiveness parameter for the 5x5 model is a class variable, taking values from conservative to aggressive but not beyond that in either direction. This configuration seems much more plausible as a real-world situation.

Glossary

capital: cash-equivalent used to buy assets. If a firm is borrowing, the capital is used as collateral against the loan.

credit-default swap: an insurance policy against the default on credit (failure to pay debt) of a third party. The purchaser pays the seller for the guaranty of debt payment if the third party defaults on the loan.

default: failure to make payment on a loan

derivative: a type of financial instrument whose value is derived by algorithm or schedule from an underlying entity or asset

dividend: payment made as a distribution of profits to shareholders

emergent: developing, or generated by a collective process

externality: an impact from an activity on a third party who did not choose the activity or its effects

Federal Deposit Insurance Corporation: a United States government corporation created by the Banking Act of 1933 to guaranty deposits at participating banks

Federal Reserve: the central bank of the United States, created by the Federal Reserve Act in 1913 to maximize employment, stabilize prices, and moderate interest rates

financial instrument: any tradeable asset

fully homomorphic encryption: a cryptographic system which supports addition and multiplication on encrypted inputs and returns encrypted outputs.

government sponsored enterprise: financial services corporations created by the government to enhance credit for borrowers and increase market strength by reducing risk to investors

hedge fund: initially designed to hedge or limit the risk to a portfolio of a market crash, hedge funds are aggressively managed investment portfolios limited to accredited investors

hedge: an investment position reducing the risk of an asset by taking an offsetting risk position on a related asset

leverage: debt used to finance asset purchases

liquidity: the degree to which an asset can be bought or sold in the market without affecting the asset's price. Liquidity is characterized by a high level of trading activity. Assets that can be easily bought or sold are known as liquid assets.

long position: buying an asset when it is perceived to be undervalued, with the expectation that it will rise in value

margin call: a demand by a lender that an investor deposit further collateral to cover possible losses

market-clearing: equality of supply and demand. A market-clearing price is thus one at which there is the same amount of demand for an asset as there is supply.

MATLAB: computer program for numerical computing developed by MathWorks. Specially suited to working with matrices.

mortgage-backed securities: a type of asset-backed security that is secured by a pool of mortgages

noise traders: investors whose decisions are not based on the fundamentals of the stock's value, even if some other technical analysis is involved

optionality: having an embedded financial option, such as the right to repay before the end of the loan term

over-the-counter: negotiated privately rather than on a regulated exchange with standardized terms

risk profile: the threats to which an entity is exposed as well as the entity's appetite for and management of those risks

SAS: statistical analytics software developed by SAS Institute. Specially suited to processing large data line by line.

scripts: self-contained MATLAB programs which require no input to run

security: a negotiable financial instrument representing value dependent on some asset, then said to be 'secured'

short: selling a borrowed security into the market with the promise to purchase from the market to repay the debt in the future. Usually performed when the seller expects the price to decline.

spread: difference between two prices or values

stress-test: a simulation performed on a portfolio to determine financial solvency in the event of a crisis

utility: an economic term referring to total benefit gained from a certain good or service

valuation: determination of economic value

Bibliography

- Axelrod, R. (1997). Advancing the Art of Simulation in the Social Sciences. *Complexity*, 3(2), 16–22.
- Balci, O. (1997). Verification, Validation and Accreditation of Simulation Models. *Proceedings of the 1997 Winter Simulation Conference*.
- Bisias, D., Flood, M., Lo, A. W., & Valavanis, S. (2012). A Survey of Systemic Risk Analytics. *Office of Financial Research, Working Pa.*
- Bonabeau, E. (2002). Agent-based modeling: methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences of the United States of America*, 99 Suppl 3, 7280–7287. doi:10.1073/pnas.082080899
- Bookstaber, R. (2012). Using Agent-Based Models for Analyzing Threats to Financial Stability. *Office of Financial Research, Working Pa.*
- Chen, C., Iyengar, G., & Moallemi, C. C. (2013). An Axiomatic Approach to Systemic Risk. *Management Science*, 59(6), 1373–1388. doi:10.1287/mnsc.1120.1631
- Clark, G. M. (1990). Use of Common Random Numbers in Comparing Alternatives. In *Proceedings of the 1990 Winter Simulation Conference* (pp. 367–371). Retrieved from http://informs-sim.org/wsc90papers/1990_0066.pdf
- Flood, M., Katz, J., Ong, S., & Smith, A. (2013). Cryptography and the Economics of Supervisory Information: Balancing Transparency and Confidentiality. *Office of Financial Research, Working Pa.*
- Geanakoplos, J., Axtell, R., Farmer, D. J., Howitt, P., Conlee, B., Goldstein, J., ... 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. *Proceedings of the 41st annual ACM symposium on Symposium on theory of computing STOC 09*, 19, 169. doi:10.1145/1536414.1536440
- Goldstein, J. (2011). An Agent-Based Model of the Interaction between the Housing and RMBS Markets.

- Lo, A. W. (2008). *Hedge Funds , Systemic Risk , and the Financial Crisis of 2007 – 2008*.
- Matteo Richiardi, R. L. N. S. and M. S. (2006, January 31). A Common Protocol for Agent-Based Social Simulation. JASSS. Retrieved from <http://jasss.soc.surrey.ac.uk/9/1/15.html>
- Outkin, A. V. (2012). An Agent-based Model of the Nasdaq Stock Market : Historic Validation and Future Directions .
- Sargent, R. G. (2007). Verification and Validation of Simulation Models. *Proceedings of the 2007 Winter Simulation Conference*, 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. doi:10.2307/1911361
- Tesfatsion, L. (2005). AGENT-BASED COMPUTATIONAL ECONOMICS : A CONSTRUCTIVE APPROACH TO ECONOMIC THEORY, (October 2003).
- Turner, S., Farmer, J. D., & Geanakoplos, J. (2012). Leverage Causes Fat Tails and Clustered Volatility. *Quantitative Finance*, 12(5), 695–707.
- U.S. House. 111th Congress. "H.R. 4173:Dodd-Frank Wall Street Reform and Consumer Protection Act."