

Simulating Heterogeneous Portfolio Choices and Financial Market Outcomes

Sebastian Benthall¹, Christopher D. Carroll², Zachary David³, John Liechty⁴, Alan Lujan⁵, Christopher McComb⁶ and Nicholas Skar-Gislinge⁷

¹New York University School of Law, 40 Washington Square South, New York, NY 10012, USA

²Department of Economics, 590 Wyman Hall, Johns Hopkins University, Baltimore, MD 21218, USA

³Urvin Finance, Washington, DC 20015, USA

⁴Smeal College of Business, Penn State University, University Park, PA 16802, USA

⁵Department of Economics, The Ohio State University, 410 Arps Hall, 1945 N. High St., Columbus, OH 43210, USA

⁶Department of Mechanical Engineering, Carnegie Mellon University, 5000 Forbes Ave, Pittsburgh, PA 15213, USA

⁷Copenhagen Center for Social Data Science, University of Copenhagen, Øster Farimagsgade 5 1353 Copenhagen K, Denmark

Abstract

Explaining the causes of financial crises requires relaxing some of the assumptions traditionally made by macroeconomic theory. We are able to simulate market crashes when consumer households have heterogeneous expectations of asset prices. These consumers earn, consume, and allocate their savings to either a risky or risk-free asset. As a departure from most macroeconomic models, the price of the risky asset is determined by a realistic agent-based simulation of a stock market, rather than through an idealized frictionless market model. Only a subset of agents in the market, the institutional investors, price the asset according to a rational expectations strategy based on an observed signal of the asset's fundamental value. The consumers try to infer this fundamental value from the noisy price signal under bounded rationality conditions that limit their memory of and access to the price process. The consumers pass all of their buy and sell orders to the market in aggregate through a broker. This dynamic process of strategic interaction between heterogeneous macroeconomic agents and the agent-based financial market determines a new equilibrium price.

We find that when the consumers have enough market power to overwhelm the institutional investors, the market will fail by driving the price of the risky asset to zero. This shows how one form of bounded rationality – ignorance about fundamental asset value – can contribute to financial crises. We also find that increasing the consumer market power leads to an increase in volatility (and a decrease in the price) of the risky asset. This suggests that there is an asymmetry whereby consumption dampens asset bubbles while accelerating price collapses.

Keywords

Heterogeneous Agent Modeling, Agent-Based Models, Financial Markets, Market Microstructure

AMPM 2022: 2nd Workshop in Agent-based Modeling & Policy-Making, December 14th, 2022, Saarbrücken, Germany


✉ spb413@nyu.edu (S. Benthall)

🌐 <https://sbenthall.net/> (S. Benthall)

🆔 0000-0002-1789-5109 (S. Benthall)



© 2023 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

1. Introduction

The dynamics of the macroeconomy are of nearly universal concern. But despite its importance, the science used to study the macroeconomy is still unable to predict and adequately respond to large financial market shocks, such as the Financial Crisis of 2007-2008. This is, in part, because economists traditionally have preferred parsimonious models of the economy with several well-motivated simplifying assumptions designed to characterize the economy's aggregate and general equilibrium behavior. This traditional theory of the macroeconomy, however, left the United States ill-prepared for the 2008 Financial Crisis, which involved dynamics that these simplified models ruled out or assumed away. Since that crisis, macroeconomics and adjacent financial researchers have explored models that relax these traditional assumptions. This paper presents one such modeling effort.

One obstacle to modeling financial crises is the rational expectations assumption, according to which all agents in the model know the model's information. Typically, this allows agents to converge on a price and prevents market failures. In our model, consumer households have heterogeneous expectations based on boundedly rational access to price history. Prices are determined by the interactions, via a broker, of the consumers' investing decisions with an agent-based model (ABM) of a financial market, populated by realistic financial actors such as institutional investors, market makers, and trend followers. Thus, we build on heterogeneous agent modeling (HAM) techniques to design an experiment in which consumers have heterogeneous expectations, are a departure from traditional representative agent assumptions. We find that the more market power that is wielded by boundedly rational consumer households relative to informed, fundamentalist institutional investors, the more likely the market is to crash to a price of \$0.

The Frontier of Macroeconomic and Financial Research

A traditional assumption of macroeconomics has been *rational expectations* [1, 2]. In a rational expectations model, agents' beliefs about uncertain outcomes are governed by equations that ensure that those beliefs are accurate with respect to the predictions of a model used by the agent. In a rational expectations model, an agent's beliefs about uncertain outcomes reflect the model's true environment such that the agents' predictions about future outcomes are accurate on average. This is a powerful constraint that in many cases ensures an equilibrium by inducing a feedback loop where agents can anticipate the environment and subsequently the agents' actions are predictable. However, rational expectations models assume away many bounded rationality conditions, which can arise because humans have limited cognitive capacity—e.g., limited ability to pay attention to information. These conditions are known to be both realistic and salient to the interactions between the macroeconomy and the financial system and leaving them out can lead to dramatically different dynamics compared to real-world phenomena. New methods of Agent-based Computational Economics have emerged to study macroeconomics and financial markets as dynamic systems of many interacting agents without the rational expectations constraint [3, 4].

In our model, we relax the rational expectations assumption for the consumer households in the model. Instead, households employ a boundedly rational learning process that uses historical

market pricing data to inform present beliefs about expected future rates of return. Agents in our model have a *memory function* that determines their weighting of historical prices. This function is parameterized to account for the relative importance of recent and past prices. These agents also have limits on their ability to attend to price information from the market.

Instead of using an unrealistic frictionless market clearing mechanism typical of rational expectations models, we instead draw on financial research methods and use an agent-based model (ABM) to simulate the financial market's role in changes to asset prices. The use of ABMs for modeling the financial market is yet another methodological shift in financial research made in response to the 2008 financial crisis [5, 6]. We build on a detailed agent-based simulated market, mimicking many of the features of a modern electronic market, and connect it programmatically to the households in the macroeconomy. The financial ABM is calibrated to several stylized facts about the financial system but is not itself amenable to a closed-form description. Hence, the use of the ABM reinforces the need for agents to have boundedly rational expectations of the market's performance.

A major development in macroeconomics in the last few decades has been the shift from Representative Agent (RA) modeling to Heterogeneous Agent Modeling (HAM). Earlier RA approaches modeled an entire sector of an economy as a single "representative agent". This approach, equivalent to a *mean-field theory* approach in physics, is appropriate when the aggregate dynamics of a system can be characterized by the central limit theorem, but otherwise breaks down [7]. The advancement of computational technology along with innovations in numerical methods has allowed economists to develop models that incorporate important differences across agents. For example, recent work in agent-based macroeconomics has developed models of *heterogeneous expectations* [8]. HAM builds on earlier work that shows how heterogeneity in the population can explain wealth inequality [9, 10, 11] as well as allowing economists to explicitly study the consequences of heterogeneity in, for example, time preference rate, age, and education on the marginal propensity to consume [12]. This shift towards HAM has contributed to better modeling of responses to crises from exogenous shocks, such as the COVID-19 pandemic, and the tailoring of policy responses [13].

In this paper, we present an integrated HAM-ABM approach that builds on [12] by considering heterogeneity not only on discount factor but also on the coefficient of risk aversion. Our toy economy consists of six sub-populations each characterized by a discount factor β (reflecting the agent's view of the time value of money) and a coefficient of constant relative risk aversion ρ (meaning that the proportion of their wealth that the agents will be willing to hold in the risky asset will decrease as their wealth increases, but at different rates depending on the value of ρ). These different subpopulations have predictable differences in their target levels of wealth (again based on the subpopulation's value of ρ), which are reflected in our model results. These different subpopulations may also engage the financial system differently and in a way that contributes to the overall price dynamics of the risky asset. We employ HARK (Heterogeneous Agent Resources and toolKit) [14], an open-source software toolkit for heterogeneous agent modeling, to build this part of the simulation. In our model, these households play the role of both consumers and retail investors.

This combination of boundedly rational expectations of asset prices, an ABM financial model, and heterogeneous agent modeling puts our model at the frontier of macroeconomic and financial research.

Findings

We find that the boundedly rational and heterogenous retail investing households (identified below as HARK agents) ‘break’ the simulation (force prices to \$0 or over \$5,000, with a starting value of \$100) as their relative market power increases relative to the institutional investors¹. This market power increases as more HARK agents pay attention to the risky asset (and consequently invest) and as the conversion rate between the macroeconomy (HARK) currency and the financial market (identified below as AMMPS (Agent-Market Modeling Platform and Simulator) currency increases. Interestingly a stronger dividend rate moderates the negative impact of their market power. One interpretation of this result is that the more that capital is unhinged from direct signals of fundamental asset value (in our case, the most recent dividend information), the more likely markets are to thrash and crash.

We also find that an increase in the HARK agent’s market power leads to an increase in volatility (and a decrease in the price) of the risky asset. Our results suggest that consumer behavior dampens asset bubbles while accelerating price collapses. The HARK agents’ behaviors are asymmetric, in that they tend to ‘deflate’ asset bubbles. It is currently unknown whether this is due to their increasing consumption as their wealth increases, to investing less of their income in the risky asset, or actually selling some of their risky asset holdings in order to supplement their income as they consume more than they make. On the other hand, when the risky asset has a negative price run this can evolve into a price collapse, as the desire of the HARK agents to consume is not diminished sufficiently to stop the selloff. As a result, increased volatility appears to be related more to persistent negative price runs as compared to persistent asset bubbles. We note that the lack of clarity on this point, suggests ways in which our simulation platform could be improved, such as by devising summaries of the risky asset’s price process, for each simulation run, that can reasonably distinguish between asset bubbles and price collapses.

Organization

The paper is organized as follows. In Section 2, we give an overview of the integrated model that combines both the HAM, HARK, model and the ABM financial market model. , In Section 3, we describe the design of our simulation study its parameters. Section 4 reports the results of the simulation. Section 5 discusses ways in which the simulation could be improved to make it both more robust as well as more realistic. There is also a Technical Appendix which describes more details of the model.

2. The Integrated Model Simulation Platform

The core innovation of this paper is the integration of a heterogeneous agent model (HAM) of the macroeconomy (the HARK model) with an agent-based model (ABM) of a financial market

¹We have identified these boundaries of ‘breaking’ the simulation because it is not reasonable to have a collapse in the price to 0 when the stock is still providing a dividend and it is correspondingly not reasonable to have a price that is more than a 50 fold increase of the starting value, over no more than 8 quarters, when the average rate of growth is around 10% per year. We acknowledge that a factor of 50x for the upper bound is arbitrary, but it seemed a reasonable guess to start.

	Information	Strategy	Heterogeneity
Institutional Investor	Daily dividend rate and intraday prices	Lucas asset price	Assumptions about consumer risk aversion and discount factors
Broker	Aggregate consumer orders	Executes all orders across the trading day	
Market Maker	Intraday prices	Place bid and ask quotes around current mid price to minimize inventory	
Technical Traders	Intraday prices	Trend following and mean reversion	Different technical indicators, memory and thresholds for entry and exit
0-Information Traders	None	Executes random orders	
Consumers	Daily prices	Invest based on inferred asset value	Risk aversion and discount factor

Table 1
Summary of agent strategies and sources of information.

for exchanging (and setting the price of) a risky asset (the Financial Market model). The assumptions of the model are summarized in Table 2. A graphical overview of the integrated system is given in Figure 1. In this section, we give an overview of both the HARK macroeconomy and the AMMPS ABM financial market components of our integrated simulation framework. The sources of information and strategies of the agents in the combined model is summarized in table 1.

2.1. The HARK Macroeconomy Component

Agents in the HARK model are endowed with wealth and income and have different risk preferences and discount factors. The HARK agents experience random shocks to their wealth and income during the process of the simulation (mimicking changes in employment, inheritances, disabilities, and so on). The primary tasks of the HARK agents are to determine the amount of wealth to consume each period, and then how much wealth to invest in a risky asset (akin to a market index like the S&P 500) and a risk-less asset (akin to a government bond).

Typically the allocation is done once a quarter, based on the HARK agent's view of the mean

Assumptions	
Assets	
Risky Asset	Yields daily dividends modelled as an exogenous random walk, subject to a lognormal shock each day. Dividends are assumed to be reinvested in the market for all agents except for consumers.
Riskless Asset	Yields a zero rate of return.
Agents	
Institutions	Employ the fundamental Lucas asset pricing model and consider the price of the risky asset to be determined solely by its dividends. These agents have limited liquidity and exhibit heterogeneous views regarding the relationship between dividends and asset prices, representing different views of the discount factor and risk aversion of the economy's consumers.
Zero Information Traders	Trade for reasons unrelated to the model, adding a flow of orders that induce fluctuations in supply and demand. Have a long term net zero position in the market.
Technical Traders	Form trading signals based solely on past price history and do not consider fundamental prices for the risky asset.
Market Makers	Supply liquidity by submitting limit orders on both sides of the market. They adjust their prices in order to maintain an average net zero position in the risky asset and adjust the width of their bid-ask spread in response to volatility.
Brokers	Execute the trading volume of consumers across the trading day, but do not take into account the asset price.
Consumers	Consume out of wealth and save in risky and riskless asset. Receive uncertain income with transitory and permanent shocks. Use previous day's closing market price to determine their demand for the risky asset, which is executed by brokers. They have limited attention and only participate in the market sporadically. They price the risky asset using the Lucas asset pricing model and infer the dividends process from the asset price. Vary by discount factor and relative risk aversion.

Table 2
Summary of assumptions made in this model.

and variance of the risky asset and the equity premium (the difference between the mean return of the risky asset and the risk-free return of the safe asset). We expand the HARK framework to mimic more realistic asset allocation behavior by agents (consumers) in the macroeconomy. These agents can be viewed as retail investors. A random subset of these agents make their allocation decisions each day, as opposed to all of the agents making their allocation decisions on the same day, once a quarter. Each day the total quantities of the risky asset that these agents




Figure 1: A schematic diagram of the components of the simulation and how they interact. HARK Agents engage in transactions with the risky asset, which get pooled by the *Buy Broker* and *Sell Broker* into aggregate buy and sell limit orders. These are executed in batches on the *Market*, which results in a changing asset price and consequent capital gains rate of return. This rate of return informs a *Financial Model* which is used to compute the new expected rate of return and market variance anticipated by the HARK agents.

want to either buy or sell are aggregated and sent to a *Buy Broker* and *Sell Broker*, respectively. These brokers then proceed to buy or sell on behalf of the agents in the financial market, where they interact with and compete with other financial market agents within the AMMPS ABM component of the simulation. The price for the risky asset is then determined by simulating one day of trading using the AMMPS financial market ABM, which includes the *Buy Broker* and *Sell Broker*, market makers, and institutional investors that are choosing a target price based on an exogenous dividend process. At the end of the trading day, the wealths of the macro agents are adjusted based on the closing price and dividend.

One of the key elements of the HARK simulation that we explore, in addition to the proportion of agents that pay attention to the market during a day, is how the agents make inferences about (or guess) the mean and variance of the risky asset—a key component to their calculation regarding how much of the risky asset to hold. In particular, we introduce a scheme for weighting the impact of the historic mean and variance of the price process along with the observed price process generated during the simulation. For some parameter values in this scheme, the HARK agents essentially always use the historical mean and variance of the risky asset, for other parameter values, they have a ‘relatively’ short memory with regards to the price of the risky asset and essentially disregard all but the last few weeks of price movements. The other major consideration that we parameterize and vary over the course of the simulation is the relative economic size of the HARK agents (in terms of the average number of shares of the risky asset) relative to the ABM agents in the AMMPS component. Due to page restrictions, a detailed discussion of the HARK algorithm and how the parameter values impact the weighting scheme are included in the accompanying Technical Appendix B.

2.2. The AMMPS ABM Financial Market Component

The AMMPS ABM simulated financial market, includes Institutional Investors, Market Makers, Buy and Sell Brokers, and a variety of noise traders. Below we give a sense of how the first three of these agents operate (with full details, along with details about the noise traders, given in the Technical Appendix C). A key feature of this simulation platform is that the market features a price time priority matching engine processing agent messages and sending market data in a FIX like format.

Institutional Investors

The institutional agent represents so-called ‘buy side’ market participants that take large positions in the market, which are executed across the trading day using a constant participation execution algorithm in order to minimize their market impact. The agents use a Lucas asset price model [15] where their current estimate of the value of the dividend is used in order to determine a fundamental value for the risky asset. Deviations in the price from this fundamental value trigger large-scale buying or selling, which is halted when the target position has been reached with regards to the agents’ fundamental price.

Market Makers

The market maker agents are the primary source of liquidity in the financial market simulation ecosystem. Market makers want to profit by facilitating trades to both buyers and sellers, typically placing passive orders at many different price levels simultaneously so that their orders are at the front of the queue for each price level in the market’s central limit order book. This gives them favorable time priority compared to signal-based agents. To profit from liquidity provision rather than short-term or long-term price trends, the agents frequently adjust their passive orders to try to trade equally on both sides of the market while achieving an average sell price that is higher than the average buy price. The particular pricing logic used by the Market Makers is detailed in Technical Appendix C.

Buy and Sell Brokers

The broker agents form the link between the macroeconomy, the HARK agents, and the AMMPS ABM financial markets by buying and selling on behalf of the HARK agents. Each time the market opens the *Buy* and *Sell Brokers* receive the buy and sell volume targets from the subset of the HARK agents who are paying attention to the risky asset that day. The buy and sell volumes are executed by the *Buy* or *Sell Broker* respectively, evenly across the day—but with an acceleration in trading towards the end of the day, if needed, in order execute all of the orders from the HARK agents. At the end of the day the *Buy* and *Sell Brokers* send the closing price back to the macro agents.

3. Simulation Study

Using the integrated model simulation platform, we conducted a simulation study where we sought to determine the impact of: 1) attention to whether to trade, 2) attention to price history,

Parameter	Description	Values
<i>Percent Attention</i>	Probability that a HARK agent pays attention and re-balances their portfolio.	(0.05, 0.35, 0.70, 0.95)
p_1	Decay rate governing how HARK agents weight the price generated during the simulation.	(0.01, 0.1, 0.5, 0.99)
p_2	Decay rate governing how HARK agents weight simulation data vs. historical mean and variance.	(0.01, 0.1, 0.5, 0.99)
<i>Dividend Growth Rate</i>	The mean of the lognormal shock to the dividend of the risky asset.	(1.00, 1.0005, 1.0015)
<i>Dollars per HARK Money Unit</i>	Scaling factor DPH between units of wealth used by the HARK agents and ‘dollars’ used to buy and sell the risky asset.	(1, 000, 5, 000, 10, 000, 15, 000)

Table 3

Simulation parameter description and values used in the simulation study. See the Technical Appendix for more details about these parameters and how they impact the simulation study.

3) the level of dividends, and 4) the amount of the money scaling factor, DPH , on the mean and volatility of the price process (as a way of measuring asset bubbles and asset collapses) and the overall level and distribution of wealth among the HARK agents.

The HARK part of the simulation is written in Python and is available under an open-source license. Further the simulated market for the risky asset (the financial market) is built in an agent-based simulation framework called AMMPS, written in C#. Future versions of this ABM platform are intended to be made available under an open-source license. The HARK Python code communicated with the AMMPS C# code via RabbitMQ, which was deployed in batch manner using containers in the Azure cloud system.

3.1. Structure of Simulation Study

We varied five parameters over a grid. All of the remaining parameters were set to values that gave us reasonable behavior from the integrated system. For each combination of parameters, we ran 15 versions of the integrated model (each with 8 quarters of 60 simulated 6.5 hour trading days).

The five simulation parameters (and their values)—which vary the HARK agents’ attention and memory, the rate at which the dividend of the risky asset grows, and the relative size of the HARK agents impact on the financial market—are given in Table 3.

3.2. Breaking the Simulation

We purposely used simulation parameter values that would ‘break’ the AMMPS financial market portion of the simulation platform, where ‘breaking’ meant when the market makers set prices at \$0 (or less) or \$5,000 (or more), corresponding to a 5,000% increase in the price of the risky asset over the simulation. We have identified these boundaries of ‘breaking’ the simulation because it is not reasonable to have a collapse in the price to 0 when the stock is still providing a dividend and it is correspondingly not reasonable to have a price that is more than a 50 fold

increase of the starting value, over no more than 8 quarters, when the average rate of growth is around 10% per year. We acknowledge that a factor of 50x for the upper bound is arbitrary, but it seemed a reasonable guess to start.

The simulation ‘breaks’ when the net buy or sell volume of the HARK agents exceeds the liquidity that can be provided by the other traders in the market, i.e. the institutional agents. In this case, the market makers are left with the excess volume and will adjust their prices, in an attempt to make the other agents take their inventory, but as these other agents have exhausted their liquidity the market makers are stuck with the excess volume causing further adjustment to their price. This starts a feedback loop between the HARK agents and the market makers, where the HARK agents’ desire for buying (selling) is increased further by the market makers’ continuous upwards (downwards) price adjustments. Essentially we test when the influence of the HARK agent’s demand for buying or selling shares in the risky asset overwhelmed the capacity of the non-broker agents in the AMMPS financial market. As not all simulation parameter values broke the simulation, we ran an analysis (a binary logit) where we determined the influence of the simulation parameters on the probability of ‘breaking’ the market.

Then, removing the ‘broken’ simulation from the data, we analyzed data (using regressions) from each of the remaining runs of the integrated model, capturing the mean and standard deviation of the percent change of the ‘daily’ price of the risky asset (*Market Mean* and *Market Standard Deviation*) during the 8 quarters of trading within each run of the integrated model, and the mean and standard deviation of the distribution of wealth among the HARK agents at the end of the simulation (*Average Wealth* and *Standard Deviation of Wealth*). We also summarized the distribution of wealth (mean and standard deviation) for all of the HARK agents and then within each heterogeneous subgroup of HARK agents that made up the universe of HARK agents—noting this distinction as *Across all Groups* and *Within Risk and Discount Groups*.

We ran the integrated model 15 times for each combination of the simulation parameters, and for each run the HARK agent composition was the same: there were 25 HARK agents in 6 different risk and discount groups (a 2×3 design) with two levels of discounting (0.95 and 0.97) and three levels of risk aversion (using a Constant Rate of Risk Aversion) (3.3, 6 and 8.67).

4. Results

After reporting on the conditions that lead to failures in the market (the market breaking), the two primary dynamics that we investigate are Market Mean and Volatility (of the risky asset) and Distribution of Wealth (across the HARK agents). The Market Mean and Volatility are based on the ‘daily’ price movement of the risky asset across the 8 quarters within each complete run of the integrated model. The dependent variable *Market Mean* is the average return of the price of the risky asset at the end of each complete run of the integrated model, and the dependent variable $\log(\text{Volatility})$ of the Market is the log of the standard deviation of the daily percent change of the risky asset over the 8 quarters.

When does the Market ‘Break’?

As reported in Table 4, an increasing number of HARK agents paying attention to the market (increased *Percent Attention*), results in an increased probability of an AMMPS market failure (or

	AMMPS Market Failure
Intercept	680.36 (71.61)***
Percent Attention	4.22 (0.18)***
Dividend Growth Rate	−694.85 (71.64)***
Dollars per HARK Money Unit	$1.1e - 03$ ($2.7e - 05$)***
AIC	3499.53
BIC	3528.92
Log Likelihood	−1745.77
Deviance	3491.53
Num. obs.	11470

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 4

Binary logistic: simulation parameters' impact on the probability that the AMMPS (financial market) portion of the simulation fails ('breaks') or is overwhelmed by the HARK agent's orders.

having the simulation 'break'); this is similar for an increase in the *Dollars per HARK Money Unit*. Both make sense in that they increase the relative market size of the *Buy/Sell Brokers* relative to the other AMMPS agents—which can be viewed as an increase in the market power of the retail investors versus the institutional investors and market makers—and can push intra-day prices to extremes. Alternatively, an increase in the *Dividend Growth Rate* decreases the probability of an AMMPS market failure, which is intuitive given that it appears that most failures occur when negative feedback forces prices to zero and a stronger *Dividend Growth Rate* generally keeps the price further from zero over the history of a simulation. Net of the influence of these parameters, neither of the memory parameters (p_1 and p_2) were statistically significant with regards to whether the market failed (or 'broke').

Next, we removed the simulations where the market 'broke' and then used the remaining simulation results to understand how the price dynamics of the risky asset and the distribution of wealth among the HARK agents varied as a function of the simulation parameters.

Market Mean and Volatility

We used the natural log of the standard deviation of the volatility of the rate of (percentage) return for the daily price of the risky asset over each simulation run because the original (non-log) values have significant positive skewness and large outliers.

As reported in Table 5, an increasing number of HARK agents paying attention to the market (increased *Percent Attention*), results in increased market volatility; this is similar to an increase in the *Dollars per Hark Money Unit*. Again, both make sense in that they increase the relative market size of the *Buy/Sell Brokers* relative to the other AMMPS agents allowing their utility functions to become more dominant and produce reinforcing, path-dependent behavior, and price movements. Interestingly, an increase in the *Dividend Growth Rate* decreases market volatility. This is likely due to the fact that in markets where there is strong upward pressure on the price, negative transient price movements are 'dampened' and are less likely to lead to feedback that collapses the price down toward zero.

	Mean of Market	Log(Volatility) of Market
Intercept	-0.92 (0.01)***	40.97 (1.60)***
Percent Attention	-1.4 - 04 (1.8e - 05)***	0.07 (2.9e - 03)***
Dividend Growth Rate	0.92 (0.01)***	-45.39 (1.59)***
Dollars per HARK Money Unit	-1.5e - 08 (1.6e - 09)***	7.7e - 06 (2.6e - 07)***
R ²	0.52	0.19
Adj. R ²	0.52	0.19
Num. obs.	8264	8264

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 5

Simulation parameters' impact on the mean and the Log(Volatility) of the daily (percent) return of the risky asset for each simulation run. The memory parameters, p_1 and p_2 are not statistically significant and, hence, are not included in the analysis.

The impact of the simulation parameters reverses with regard to the mean market price. As the *Dividend Growth Rate* increases, the market price tends to increase at the end of each simulation; this makes sense as an increased dividend is theoretically linked with a higher value for the risky asset. As both the *Percent Attention* and the *Dollars per HARK Money Unit* increase, the mean market price decreases. This is most likely due to higher values of these simulation parameters leading to increased market volatility. In those environments, a feedback cycle for negative price movements tends to arise, driving the price toward zero.

Both of these patterns suggest that the HARK agents' behaviors are asymmetric, in that they tend to 'deflate' asset bubbles by increasing their consumption as their wealth increases, by either investing less of their income in the risky asset or actually selling some of their risky asset holdings in order to supplement their income as they consume more than they make. Alternatively, when the risky asset is involved in a price collapse, the desire of the HARK agents to consume is not diminished sufficiently to stop the selloff, and as a result, increased volatility appears to be related more to persistent negative price runs as compared to persistent asset bubbles. This lack of clarity points to one way in which our simulation platform could be improved, which is by devising summaries of the risky asset's price process, for each simulation run, that can reasonably distinguish between asset bubbles and price collapses.

Wealth Distribution

Each complete run of the integrated model has 150 HARK agents, consisting of 6 subgroups of 25 agents, where the agents in each subgroup have the same risk preference and use the same discount factor. The initial distribution of wealth across the 150 HARK agents is random. The Wealth Distribution that we analyze is based on the wealth of the HARK agents at the end of the 8 quarters for a complete run of the integrated model. The dependent variable *Average Wealth* is the mean of the HARK agents' final wealth; the dependent variable *Spread of Wealth* is the standard deviation of the HARK agents' final wealth and can be viewed as a measure of wealth inequality. As previously noted, we calculate these dependent variables using all of the HARK agents (the *Across all Groups* case) and for the HARK agents within a specific risk and discount rate subgroup (the *Within Risk and Discount Group* case).

	Average Wealth	Spread of Wealth
Intercept ($p_2 = 0.01$)	-74.96 (2.56)***	-145.12 (21.37)***
Dummy Variable: $p_2 = 0.1$	0.08 (0.00)***	0.05 (0.04)
Dummy Variable: $p_2 = 0.5$	0.05 (0.01)***	0.03 (0.04)
Dummy Variable: $p_2 = 0.99$	0.05 (0.01)***	0.03 (0.04)
Dividend Growth Rate	76.25 (2.56)***	148.66 (21.36)***
Dollars per HARK Money Unit	$1.6e - 06$ ($4.1e - 07$)***	$1.4e - 05$ ($3.5e - 06$)***
R^2	0.13	0.01
Adj. R^2	0.13	0.01
Num. obs.	8264	8264

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 6

Simulation parameters' impact on Average Wealth and Spread of wealth among HARK agents for each simulation run (the *Across all Groups* case). The Dummy variables, for p_2 , show how the intercept changes when p_2 takes values that are different from the 'base case' of $p_2 = 0.01$. The level of attention and the memory parameter p_1 are not statistically significant and, hence, not included in this analysis.

When we summarize what drives differences in wealth across all HARK agents in a simulation, as reported in Table 6; we find that both the level of attention of the HARK agents as well as one of the market price memory parameters (p_1) had no impact on the distribution of wealth. Alternatively, the *Dividend Growth Rate* and *Dollars per HARK Money Unit* both had positive correlations. The first is not surprising as it suggests a log-normal type effect on the distribution of wealth, the larger the average wealth, the broader the spread of wealth. The second, suggests that the more active the HARK agents are, the greater their collective wealth—which supports one of the underlying assumptions of this simulation framework, which is that having HARK agents update their asset allocation quarterly will give very different results when compared with a more realistic random attention model. The other memory parameter (p_2) reflects how HARK agents weigh the historic mean and variance vs the price generated during the simulation, with low values corresponding to 'essentially' ignoring the historic mean. It turns out that the lowest of these values corresponds to an average decrease in the mean and spread of wealth. Perhaps this is due to the fact that when agents pay attention to the market prices, as opposed to the supposed long-term history, there are more feedback-driven price runs which tends to result in negative price runs heading toward zero.

When we summarize what drives differences in wealth within each HARK agent group, as reported in Table 7; we see the same patterns with respect to the HARK agent memory parameters as well as with respect to the *Dividend Growth Rate* and *Dollars per HARK Money Unit*, which we posit are for similar reasons. As for the preference parameters, a higher coefficient of relative risk aversion (*CRRRA*) has a positive effect and a higher *Discount Factor* has a negative effect on average wealth. The effect of *CRRRA* can be understood as a moderation on risk taking; agents with lower *CRRRA* parameters on average are able to enjoy some higher returns without exposing themselves to too much risk, ending the simulation with higher average wealth. The effect of the *Discount Factor* parameter is similarly expected and explained; agents with a high *Discount Factor* put an increasingly similar weight on today as they do tomorrow, leading to a strong

	Average Wealth	Spread of Wealth
Intercept ($p_2 = 0.01$)	-72.99 (2.76)***	-122.31 (15.03)***
Dummy Variable: $p_2 = 0.1$	0.08 (0.01)***	0.22 (0.03)***
Dummy Variable: $p_2 = 0.5$	0.05 (0.01)***	0.14 (0.03)***
Dummy Variable: $p_2 = 0.99$	0.05 (0.01)***	0.15 (0.03)***
Dividend Growth Rate	76.25 (2.76)***	133.22 (14.99)***
Dollars per HARK Money Unit	$1.6e - 06$ ($4.5e - 07$)***	$7.9e - 06$ ($2.4e - 06$)***
CRRA	0.11 (0.00)***	0.36 (0.00)***
Discount Factor	-2.74 (0.16)***	-10.56 (0.90)***
R ²	0.29	0.13
Adj. R ²	0.29	0.13
Num. obs.	49584	49584

*** $p < 0.001$; ** $p < 0.01$; * $p < 0.05$

Table 7

Simulation parameters' impact on Average Wealth and Spread of wealth among HARK agents for each simulation run (the *Within each Risk and Discount Groups* case). The Dummy variables, for p_2 , show how the intercept changes when p_2 takes values that are different from the 'base case' of $p_2 = 0.01$. The level of attention and the memory parameter p_1 are not statistically significant and, hence, not included in this analysis.

incentive to smooth consumption and save in a safe asset, which could prevent agents from participating in the market and benefit from higher returns to wealth.

5. Future Improvements

Clearly, there is scope for adding to (and refining) the grid points used in order to better understand how the current simulation system performs. The most likely candidates are exploring smaller values of the HARK agent's memory parameters (p_1 and p_2). Another simple improvement will be to make more detailed summaries of the simulation to use as additional dependent variables; natural candidates to consider include summaries of stylized facts of the risky asset price dynamics (e.g., AR parameter estimates, GARCH/ARCH parameter estimates) and functional data summaries of the price dynamics to extract functional trends, hopefully allowing us to identify and differentiate between asset bubbles and price collapses; they could also include summaries of the HARK agents utility functions.

The fact that demand from the HARK agents can overwhelm the capacity of the AMMPS market is, in some ways, not surprising, but we intend to change both the HARK agents and the AMMPS market in ways that reflect more realism and require the AMMPS market to 'break' for reasons more likely to arise in a real financial market. For the HARK agents, we intend to build in a scheme for identifying when a market has failed (e.g. when it has deviated sufficiently from the historic mean and variance) which would then justify the HARK agents changing their behavior (e.g. flight to quality) based on a more complex belief-based utility function. Under 'normal' conditions, HARK agents may behave more like the institutional investors in that they price the asset according to the Lucas asset pricing model, when evidence suggests they should

abandon this model then alternative actions would be taken.

For the AMMPS market, clearly, the Institutional Investors can and will need to become more realistic: using a Bayesian learning model to estimate the dividend process; having more realistic and flexible risk preferences as well as having flexible access to capital markets allowing their leverage to change. Both modifications will affect how much they will participate in the market, and hence their power in the market; similarly for market makers, having a dynamic access to credit would greatly affect their ability to provide liquidity at different points in the simulation.

Finally, adding a circuit breaker feature in the simulated exchange in AMMPS, would add a more realistic handling of events where the market for some reason moves very drastically in one direction when the liquidity of the market actors is exhausted. Instead of simply 'breaking' the market, a circuit breaker type of event could offer a path for the market to recover.

Acknowledgements

This material is based upon work supported by the Defense Advanced Research Projects Agency through cooperative agreement No. HR001 12220029 and the National Science Foundation under Grant No. 2105301. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the sponsors. One of the authors of this article was supported by the New York University Information Law Institute's Fellows program, which is funded in part by Microsoft Corporation.

References

- [1] J. F. Muth, Rational Expectations and the Theory of Price Movements, *Econometrica* 29 (1961) 315.
- [2] S. M. Sheffrin, *Rational Expectations*, Cambridge University Press, 1996.
- [3] B. LeBaron, L. Tesfatsion, Modeling Macroeconomies as Open-Ended Dynamic Systems of Interacting Agents, *American Economic Review* 98 (2008) 246–250.
- [4] L. Tesfatsion, Chapter 16 Agent-Based Computational Economics: A Constructive Approach to Economic Theory, in: *Handbook of Computational Economics*, volume 2, Elsevier, 2006, pp. 831–880.
- [5] J. D. Farmer, D. Foley, The economy needs agent-based modelling, *Nature* 460 (2009) 685–686.
- [6] R. Bookstaber, *The End of Theory*, Princeton University Press, 2017.
- [7] A. F. Siegenfeld, Y. Bar-Yam, An Introduction to Complex Systems Science and Its Applications, *Complexity* 2020 (2020) 1–16.
- [8] W. A. Branch, B. McGough, Heterogeneous Expectations and Micro-Foundations in Macroeconomics, in: *Handbook of Computational Economics*, volume 4, Elsevier, 2018, pp. 3–62.
- [9] M. Huggett, The risk-free rate in heterogeneous-agent incomplete-insurance economies, *Journal of Economic Dynamics and Control* 17 (1993) 953–969.
- [10] S. R. Aiyagari, Uninsured Idiosyncratic Risk and Aggregate Saving, *The Quarterly Journal of Economics* 109 (1994) 659–684.
- [11] P. Krusell, A. S. Jr., A. A. Smith, Jr., Income and Wealth Heterogeneity in the Macroeconomy, *Journal of Political Economy* 106 (1998) 867–896.
- [12] C. Carroll, J. Slacalek, K. Tokuoka, M. N. White, The distribution of wealth and the marginal propensity to consume, *Quantitative Economics* 8 (2017) 977–1020.
- [13] C. D. Carroll, E. Crawley, J. Slacalek, M. N. White, Modeling the Consumption Response to the CARES Act, *Finance and Economics Discussion Series* 2020 (2020) 107–141.
- [14] C. Carroll, A. Kaufman, J. Kazil, N. Palmer, M. White, The Econ-ARK and HARK: Open Source Tools for Computational Economics, in: *Proceedings of the 17th Python in Science Conference*, 2018, pp. 25–30.
- [15] R. E. Lucas Jr, Asset prices in an exchange economy, *Econometrica: journal of the Econometric Society* (1978) 1429–1445.
- [16] R. C. Merton, Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case, *The Review of Economics and Statistics* 51 (1969) 247.
- [17] P. A. Samuelson, Lifetime Portfolio Selection By Dynamic Stochastic Programming, *The Review of Economics and Statistics* 51 (1969) 239.
- [18] J. Y. Campbell, L. M. Viceira, *Strategic Asset Allocation*, Oxford University Press, 2002.
- [19] L. M. Viceira, Optimal Portfolio Choice for Long-Horizon Investors with Nontradable Labor Income, *The Journal of Finance* 56 (2001) 433–470.
- [20] J. Sabelhaus, J. Song, The great moderation in micro labor earnings, *Journal of Monetary Economics* 57 (2010) 391–403.

A. Technical Appendices: Introduction

In these Technical Appendices, we present details about the computational models used in the simulation platform, in the companion paper but which are not needed to gain a general understanding of the work.

B. The macroeconomic (HARK) agents

The macroeconomic agents in our simulation are both consumer households and retail investors.

B.1. Portfolio Choice Model

The macroeconomic HARK agents in our simulation are obtained from the Econ-ARK Heterogeneous Agents Resources and toolKit (HARK) library [14]. When a HARK agent pays attention, the allocation problem solved by these agents is of the kind first explored by Merton [16] and Samuelson [17] and further developed by Campbell and Viceira [18] [19]. These agents choose consumption and saving to maximize their present discounted expected utility over a finite (or infinite) life-cycle given by:

$$\max \mathbb{E}_t \left[\sum_{t=0}^T [\beta(1 - D_t)]^t u(c_t) \right]. \quad (1)$$

The agents have an isoelastic, or *constant relative risk aversion* (CRRA), utility function with parameter ρ

$$u(c) = \frac{c^{1-\rho} - 1}{1 - \rho} \quad (2)$$

Additionally, these agents experience life-cycle uncertainty such as employment, income, and mortality risk. When saving, these agents can choose between a safe asset with a low but certain return and a risky asset with a higher expected but uncertain return. The proportion of total assets that the agent chooses to invest in a risky asset is called the risky portfolio share, or risky share for short, and is denoted by ς_t . We can represent the overall return on the agent's portfolio \mathbb{R}_{t+1} as:

$$\mathbb{R}_{t+1} = R(1 - \varsigma_t) + \mathbf{R}_{t+1}\varsigma_t \quad (3)$$

$$= R + (\mathbf{R}_{t+1} - R)\varsigma_t \quad (4)$$

where the return of the risky asset is distributed $\log \mathbf{R} \sim \mathcal{N}(\phi + r - \sigma_\phi^2/2, \sigma_\phi^2)$ and ϕ is the equity premium, the difference between the expected risky return and the risk-free rate on the safe asset $r = \log R$.

The problem can be represented by the recursive Bellman equation:

$$\begin{aligned}
v_t(m_t, \mathbf{p}_t) &= \max_{\{c_t, a_t, \varsigma_t\}} u(c_t) + \beta(1 - D_{t+1})\mathbb{E}_t[v_{t+1}(m_{t+1}, \mathbf{p}_{t+1})] \\
&\text{s.t.} \\
a_t &= m_t - c_t \\
0 &\leq \varsigma_t \leq 1, \\
R_{t+1} &= R + (R_{t+1} - R)\varsigma_t \\
\mathbf{p}_{t+1} &= \Gamma_{t+1}\mathbf{p}_t\psi_{t+1} \\
y_{t+1} &= \mathbf{p}_{t+1}\theta_{t+1} \\
m_{t+1} &= a_t R_{t+1} + y_{t+1}
\end{aligned}$$

where the state variables are market resources (cash-on-hand) m_t and the level of permanent income \mathbf{p}_t . The control variables that the agent can choose are current period's consumption c_t , total assets a_t , and the risky share of assets ς_t . The agents discount the future by β and face a probability of death D_{t+1} . Before the onset of next period, the agents experience permanent and transitory income shocks, distributed as $\log \psi \sim \mathcal{N}(-\sigma_\psi^2/2, \sigma_\psi^2)$ and $\log \theta \sim \mathcal{N}(-\sigma_\theta^2/2, \sigma_\theta^2)$ respectively.²

An equivalent problem that facilitates finding the solution to the above maximization problem is to separate the objective into a two-stage sequential Bellman problem, where the agent first chooses a level of saving and then chooses their risky portfolio share.

In the first stage, the agent chooses between consumption and saving given market resources, which by the market clearing condition entails only one choice.

$$\begin{aligned}
\mathbf{v}_t(m_t, \mathbf{p}_t) &= \max_{\{c_t, a_t\}} u(c_t) + \tilde{\mathbf{v}}_t(a_t, \mathbf{p}_t) \\
&\text{s.t.} \\
a_t &= m_t - c_t
\end{aligned}$$

Given a level of saving $a_t \geq 0$, the agent then chooses the proportion of assets that they are willing to invest in the risky asset. At this stage, the agent considers the risk associated with stock market participation as well as their income uncertainty next period.

$$\begin{aligned}
\tilde{\mathbf{v}}_t(a_t, \mathbf{p}_t) &= \max_{\varsigma_t} \beta(1 - D_{t+1})\mathbb{E}_t \mathbf{v}_{t+1}(m_{t+1}, \mathbf{p}_{t+1}) \\
&\text{s.t.} \\
0 &\leq \varsigma_t \leq 1 \\
R_{t+1} &= R + (R_{t+1} - R)\varsigma_t \\
\mathbf{p}_{t+1} &= \Gamma_{t+1}\mathbf{p}_t\psi_{t+1}
\end{aligned}$$

²For log-normally distributed variables, it is known that if $\log x \sim \mathcal{N}(\mu, \sigma^2)$ then $\log \mathbb{E}[x] = \mu + \frac{\sigma^2}{2}$ such that $E[\psi], E[\theta] = 1$.

$$\begin{aligned}
y_{t+1} &= \mathbf{p}_{t+1}\theta_{t+1} \\
m_{t+1} &= a_t \mathbb{R}_{t+1} + y_{t+1}
\end{aligned}$$

In summary, when determining how much of their savings to invest in the risky asset or to consume, the agent computes backwards from their expected end-of-life to maximize their expected discounted utility. The agent takes into account known risk parameters: σ_ψ , σ_θ for the income process, and $\bar{r} = \phi + r - sr^2/2$ and sr for the risky asset.

In a rational expectations model, \bar{r} and sr would be veridical, and the rate of return of the risky asset would be a lognormal distribution. In our hybrid model, the rate of return of the risky asset is instead determined by a simulation of a financial market. The agents infer μ_ϕ and σ_ϕ from a statistical process that takes into account observations of past market returns.

For the purpose of these simulations, the risky average is related to the financial market process as follows:

$$\mathbf{R}_{t+1} = \frac{P_{t+1} + d_{t+1}}{P_t} \quad (5)$$

where d_{t+1} is the dividend process and P_{t+1} is the price level of the aggregate stock³. As the simulation proceeds, agents observe a series of prices and dividends to construct a series of risky asset returns. Since our agent population includes heterogeneity in belief formation, agents with different perceptions of the market may update their beliefs of the mean market return and standard deviation differently, leading to deviations in market behavior.

B.2. Population parameters

The specific population used in this study was calibrated for a balance of realism and qualitatively interesting results. The population consisted of six subclasses differentiated by time preference factor β and risk aversion ρ .

The distribution of β was taken from [12]’s results fitting a uniform distribution of time preference to have a resulting distribution of wealth that matches that of the United States. The distribution of ρ was chosen based on that range of values that are needed to produce qualitatively interesting results in portfolio choice problems. Each uniform distribution was discretized into equiprobable points. Each combination of one of two β values and one of three ρ values comprises a population class. For each population class, there were 25 instances (or agents). All of these agents interacted with each other via the market (see below).

All of the macroeconomic agents in the model shared transitory and permanent income shock parameters calibrated to a typical 40 year old male in the projections of [20]. We used conventional probability of death $D = 0.02$ and permanent growth factor $\Gamma = 1.01$. All these values were adjusted from annual to quarterly values. The entire duration of the integrated simulation is eight quarters or two years.

All agents are initialized with starting permanent income $\mathbf{p}_0 = 1$, and with $m = 1$ and $a_0 = 1 - c(1)$. In future work, these starting wealth levels can be tuned to the equilibrium values for each agent class in order to stabilize the effects of the macroeconomic changes on the

³The details of these processes are explained in later sections.

market.

C. The market simulation

The simulated market for trading the risky asset (the financial market) is written in C#. On a basic level the system simulates trading by the processing of messages added to a queue: Different parts of the simulation system subscribe to different types of messages and process them, which in turn may generate new messages to be added to the queue. The messages are placed in the queue according to their attached timestamp, and once the processing loop reaches a message it is taken off the queue and relayed to the relevant components of the system depending on the message type. An example is a market order to buy x amount of stock ABC. In this case, the buy order message is placed in the queue and relayed to the matching engine, which finds a matching passive order(s) and completes a trade. The trade, in turn, generates a new set of messages, a general message describing the trade, a message describing the change to the order book, and a message for each party of the trade, informing them of the trade.

The message types can be seen as different information streams, by which the different components of the system can subscribe to and interact with. Because of the placement in the queue according to the attached timestamp, it is possible to simulate latency between the different information streams and actors by adding a delay caused by the latency to the message timestamp.

On a higher level, we can use the system to define agents that observe and react to these information streams according to different logics and time scales based on market actors in a stock market. Market makers, for example, subscribe to the order book stream to keep track of the best bid and ask prices and adjust their own prices accordingly. Institutional trades follow the trading prices across the day as well as dividend information released at the beginning of the day. In addition to the agents described in the main document, the additional agent types are described in detail below:

C.0.1. Market Makers

When a market maker's pricing method is triggered, the agent generates a set of target passive orders for bids, B^t , and asks, A^t , each containing entries $b_n \in B^t$ and $a_n \in A^t$ for $n \in \{1, 2, \dots, l\}$ where l is the maximum number of levels for which the agent places orders on each side of the bid-ask spread, and b_1, a_1 represent the highest bid and lowest ask respectively. Each entry contains both price and quantity information, such that b_n^p, a_n^p are the respective target bid and ask prices at level n , and b_n^q, a_n^q are the respective target bid and ask quantities at level n . To populate B^t and A^t , the agent first determines where it should set its best bid price and best ask price according to:

$$\begin{aligned} b_1^p &= m - \left(\frac{s}{2} + \max\left(\phi \frac{s}{2}, \gamma^b\right) \right) \\ a_1^p &= m + \left(\frac{s}{2} + \max\left(\phi \frac{s}{2}, \gamma^a\right) \right) \end{aligned}$$

where m is the smoothed volume-weighted mid-price, s_i is the volatility-based spread, γ_i^b ,

γ_i^a are the respective position-based spread adjustments for the bid and ask sides, and ϕ is a parameter which limits how far the spread can be adjusted back toward m , where $-1 \leq \phi \leq 0$. The process for calculating the above prices can be thought of as first setting a positive spread, s_i , symmetrically about m and then applying independent adjustments to each side. For the initial spread, s , a market maker wants to set its bid-ask spread narrow enough so that its prices are competitive to get matched against marketable orders but also wide enough to avoid adverse selection where the price quickly trades through the level. A simulation (market maker) agent sets its volatility-based spread according to

$$s = \psi + \sigma \cdot \pi$$

where ψ is the minimum spread, defined as $\max\{0.01, m \cdot 0.0001\}$, π is a scalar spread-factor parameter that determines how much the agent's spread should respond to volatility, and σ is the volatility, defined as the average traded price range over a given sampling frequency.

Because queue position is a profit-center for market makers, they will have passive orders at many price levels at once. In trending markets, this strategy will create a risk of quickly accumulating a large net position in a single instrument at disadvantageous prices. Market makers respond to this risk by adjusting their spreads as a function of the position. In the simulation, the position-based spread adjustment asymmetrically widens out the spread on either the bid or ask side based on the agent's position in an instrument, λ , where $\lambda > 0$ if the agent is long and $\lambda < 0$ if the agent is short. The adjustments for the bid and ask directions are given by

$$\gamma^b = \frac{\psi \cdot \theta \cdot \max\{0, \lambda\}}{w}$$

$$\gamma^a = \left| \frac{\psi \cdot \theta \cdot \max\{0, \lambda\}}{w} \right|$$

where θ is a scalar instrument-factor parameter that determines how much the position-based adjustment should respond to the agent's position, w is the target working size for that instrument (*i.e.*, the quantity the agent starts with at each target price level).

The smoothed volume-weighted mid-price (VWMP) is calculated based on an exponential moving average (EMA) of the VWMP. Whenever a market maker witnesses an Order Book update from the Market Data channel, it removes its own active orders from the Order Book and then stores the most recent VWMP according to:

$$P_{VWMP} = \frac{P_{bid}V_{ask} + P_{ask}V_{bid}}{V_{ask} + V_{bid}}$$

Where P_{bid} and P_{ask} are the respective best bid and ask prices and V_{bid} and V_{ask} are the respective best bid and ask sizes. In real trading scenarios, the VWMP is often preferred to the simple mid-price between the best bid and ask as it reflects information about the relative sizes at the top of the order book and encodes information about which size is more likely to be taken out, allowing market makers to adjust their orders more quickly when aggressive trading begins to occur on one side of the market. The smoothed VWMPs are updated on a one-second timer that fires for each market maker. When the timer fires, the market maker adds the most

recently stored VWMP to the EMA smoothing process.

In summary, the market maker agent first prices its best bid and ask symmetrically according to the volatility-based spread. Then in order to avoid accumulating too much size in one instrument at a single price, it will widen out the spread in the direction that it has been filled in that instrument.

C.0.2. Institutional Investors

The institutional agent represent so-called "buy side" market participants that take large positions in the market which is executed across the trading day using a constant participation type of execution algorithm in order to minimize their market impact. Each day these agents will price the asses using the Lucas asset pricing model based on their estimate of the current value of the dividend. At the beginning of each day, a random perturbation of the current value of the dividend is revealed to each institutional agent, giving each agent an individual price to dividend ratio which is used to calculate its particular price for the asset. For institution n at day t the price of the asset is calculated as:

$$P_{n,t} = d_t \cdot r_{n,t}$$

Where d_t is the current value of the dividend at day t , and $r_{n,t}$ is the price to dividend ratio for institution n at day t . At randomly drawn intervals between 5 to 30 minutes, the institution will poll the price and calculate the discrepancy between the valuation and the current mid-price:

$$\Delta_n = (P_{n,t} - P_{open})\epsilon,$$

where ϵ is a random normally distributed variable with mean 1 and standard deviation of 0.01 representing any additional private information the institution may have. If $\Delta_n > 0$ the agent will proceed to buy the asset as it considers it undervalued whereas if the $\Delta_n < 0$ it will sell the asset. However, the institution will only take one position each day, if it has bought a position in the asset, because the asset was undervalued and by the next poll price now has risen above $P_{n,t}$, it will hold its position and not start selling.

Execution

The execution algorithm of the institutional investors aims to execute the target position by participating as a constant fraction, f_p , of the total trading. While the execution algorithm is active, it tracks the cumulative trading volume and calculates the order size needed to participate as f_p of the market:

$$V_{order} = f_p(V_{market} - V_{agent}) - V_{agent}$$

where V_{market} is the total volume traded in the period that the execution algorithm has been running, and V_{agent} is the volume executed by the agent. The algorithm initially places the volume as a passive order at the best bid or ask depending on the side and waits T_{order} , where T_{order} is drawn from a uniform distribution between 15 to 60 seconds, then it trades

any remaining volume as an active market order. Once an active order has been executed the algorithm waits T_{wait} , where T_{wait} is drawn from a uniform distribution between 1 and 5 minutes, before sending a new order.

C.0.3. Buy Sell Brokers

The broker agents form the link between the macro and micro markets by buying and selling on behalf of the macro HARK agents. Each time the market opens the brokers receive the buy and sell volume targets from the macro simulation. The buy and sell volume is executed by the buy or sell broker respectively across the day. When the agent has received its buy or sell target for the day, it calculates a "trading velocity", v_t , determining the size of the suborders it will have to execute every second in order to reach the target by the end of the day.

$$v_t = V_{remaining} / (T - T_{finish})$$

Where $V_{remaining}$ is the volume the agent needs to execute to reach its target, T is the current time, and T_{finish} is the time 5 minutes before the market closes.

The order is executed in the same way as the institutional agent, first using a passive order, then an active order for any remaining volume, and finally a pause before executing the next order. The length of such a sequence is $T_{interval} = T_{order} + T_{wait}$, and the target volume is:

$$V_{order} = v_t \cdot T_{interval}$$

The trading velocity is updated each interval in order to account for possible execution shortfalls and any remaining volume is executed at T_{finish} , in order to guarantee the broker meets its target at the end of the day. When the market closes at the end of the day the broker sends the closing price back to the macro market.

C.0.4. Technical Agents

In addition to the agents described above, the market contains a number of 'technical' agents that follow different strategies, following a common design described below.

Trading Data Representation and Memory

The price data continuously received by the agents is synthesized into so-called price "bars" that describe the trading process for a given interval, eg 5 minutes. This is a common practice among traders and coarse grains that data to a time scale that the traders are interested in (traders with long time horizons would be interested in price information on days, weeks or months). For interval i , we define the bar, b_i , as a set of values

$$b_i = \{O_i, H_i, L_i, C_i, V_i, IM_i\}$$

Where O_i , H_i , L_i , C_i , are the opening, highest, lowest and closing prices for the interval, and V_i and IM_i are the volume traded and the volume imbalance measured as the difference between the active buy and sell volume in the interval.

Each agent will have a set number of look backs, N , determining how many bars it will store defining its memory of the trading process. For agent n at time t its memory of the trading process is defined as the set of bars:

$$B_{n,t} = \{b_0, b_1, b_1, \dots, b_N\}$$

Where b_0 is the most recent bar and b_N is the bar representing the trading process N intervals ago.

Targets

Each agent also defines a *target*, which defines the conditions for entering and subsequently exiting the market. We define a target T as a set of values:

$$T_{n,t} = \{p_{entry}, t_{direction}, V, p_{stop}, p_{profit}, t_{type}\}$$

The first component of the target, p_{entry} , is the entry price defining the price required for the agent to enter the market. The second component, $t_{direction}$, defines the side of the target and is 1 for a long target and -1 for a short target, the target volume, V , similarly can be positive for a long position and negative for a short position.

In addition to the entry price, the target contains a ‘take profit’ price, p_{profit} , and a ‘stop loss’ price, p_{stop} , defining the price at which the agent will exit its position, either to capture the profits of its investment (‘take profit’) or to avoid further losses (‘stop loss’).

The target can be a momentum type entry, in which case the agent will buy when the asset is trading above the entry price for a long target position, or sell if the asset is trading below the target price and the target position is short. Inversely, for a target with a reversion type entry, the agent will sell when the asset is trading above the target price and the target position is short or buy when the asset is trading below the target price and the target position is long.

The ‘take profit’ and ‘stop loss’ prices are set symmetrically around the target price.

$$p_{profit} = p_{entry} + vol \cdot c_{stoploss}$$

$$p_{stop} = p_{entry} - vol \cdot c_{stoploss}$$

for a momentum type target, and

$$p_{profit} = p_{entry} - vol \cdot c_{stoploss}$$

$$p_{stop} = p_{entry} + vol \cdot c_{stoploss}$$

for a reversion type target. vol_t is a measure of the price volatility at time t , defined as:

$$vol_t = \frac{1}{N} \sum_0^N (H_i - L_i)$$

Using the base framework we define 6 agent classes that use different strategies to define

their target price. In addition to how they define the target price, the agents are heterogeneous with respect to the interval by which they set their sample bars, how many bars they keep and their stop loss factor defining when to exit their positions.

Target Size

For agent i , the target size at time t , is given by:

$$V = C/Vol_{i,t}.$$

$Vol_{i,t}$ is current annualised daily volatility, C is a cash amount assigned to the trader (100000) and F_t is drawn from a folded normal distribution:

$$F_t = (0.005 + (0.025 - 0.005)|\mathcal{N}|)$$

Where \mathcal{N} is a normal distribution with zero mean and unit variance. F_t represents the fraction of the assigned amount of cash the trader use for the trade and F_t is capped at 0.5 to avoid potentially very large values.

Zero Information Traders

This type of agent is essentially a liquidity trader that seeks to buy or sell the asset for reasons exogenous to our model, generating a fluctuating flow of orders into our market. At random intervals, they will enter the market by manually triggering an entry into the market, with a take profit/stop loss that is 20 times the current volatility, meaning that they will not trigger a take profit or stop loss. They do not simulate traders following a specific strategy, but rather simulate the fluctuating order flow of a large number of traders following a multitude of strategies with wide range of investment horizons. The agent will randomly choose to buy or sell with 50% probability. In order to avoid having these agents hold a long term position in the market, they have a mean-reverting behavior. If an agent is long and wants to sell, but the size of the sell order is not large enough to enter a short position, the size of the sell order is set so that the new position is $(1 - f_r) \cdot I_t$, where f_r is a reversion factor and I_t is the current inventory. In the current simulation the reversion factor is set to 0.5. This means if the agents build a position that's outside the size range of the order sizes they are more likely to start to reduce their inventory rather than increase it further.

Aggressor Trend

This agent is a momentum type agent, using the imbalance between buy and sell trades and the direction of the price change to determine when the price is trending and follows the trend. The target direction and entry price are determined from the collected trading volume imbalances for the N bars stored by the agent. If the trade volume imbalance of the last bar is significantly above(below) the mean, as determined by a z test, and the difference between the opening and closing price of the current bar is positive(negative) it will set a target for a long(short) position at the closing(opening) price of the current bar:

$$t_{direction, p_{entry}} = \begin{cases} 1, C_0 & \text{if } Z_i < \frac{IM_0 - \langle IM \rangle}{\sigma_{IM}} \text{ and } C_0 - O_0 > 0 \\ -1, O_0 & \text{if } -Z_i > \frac{IM_0 - \langle IM \rangle}{\sigma_{IM}} \text{ and } C_0 - O_0 < 0 \end{cases}$$

Where $\langle IM \rangle$ and σ_{IM} are the average and standard deviation of trade volume imbalance. This means that if the agent observes a change in the buying volume and an increase in price it will buy if the price continues to increase, reversely if it observes an increase in the selling volume and a decrease in price it will (short)sell. Each agent, n , has an individual z-score threshold, Z_i , determining when it will consider the volume imbalance significant.

Breakout Trend

The breakout trend agent is a momentum type agent that monitors when the price moves outside a historic range. Every time the agent closes a bar it will calculate a long and short breakout value as the last bar opening price plus and minus the exponential moving average (EMA) of the price ranges (high-low) for the stored bars multiplied by a factor. If the low of the current bar is less than the open of the last bar and the current bar high is above or equal to the long breakout value, it will target a long position at the current bar high. Inversely if the current bar high is greater or equal to the last bar open and the current bar low is less or equal to the short breakout value, it will target a short position at the current bar low:

$$t_{direction, p_{entry}} = \begin{cases} 1, H_0 & \text{if } L_0 \leq O_1 \text{ and } H_0 > O_0 + \langle H - L \rangle_{exp} \cdot C_{BT,n} \\ -1, L_0 & \text{if } H_0 \geq O_1 \text{ and } L_0 < O_0 - \langle H - L \rangle_{exp} \cdot C_{BT,n} \end{cases}$$

Where $\langle H - L \rangle_{exp}$ is the exponential average of the difference between the highest and lows prices of each stored bar. The agent sets a target if only one of these conditions is met. Each agent has an individual factor, which determines its breakout prices.

Pullback Reversion

The pullback reversion agent is a reversion type agent that tries to trade on price reversion or pull back. If the agent identifies a price reversion and then observes a deviation from this reversion it will take a position in the direction of the reversion. The algorithm follows a two step process.

First, it will identify a pullback:

$$D = \begin{cases} 1 & \text{if } C_1 < \langle C \rangle_{exp} \text{ and } C_0 > \langle C \rangle_{exp} \\ -1 & \text{if } C_1 > \langle C \rangle_{exp} \text{ and } C_0 < \langle C \rangle_{exp} \end{cases}$$

If the agent closes a bar where the closing is higher (lower) than the EMA of the closing prices and the closing price of the previous bar was less (more) than EMA of the closing prices, *i.e.*, the current bar is moving in the opposite direction than the average compared to the last bar, the agent considers the price reverting up (down).

If the pullback direction, D is positive, a target direction and entry price are set

$$t_{direction} = 1, p_{entry} = H_0 \quad \text{if } C_1 > \langle C \rangle_{exp} \text{ and } C_0 < \langle C \rangle_{exp}$$

For a negative pullback direction, the target direction and entry price are set as:

$$t_{direction} = -1, p_{entry} = L_0 \quad \text{if } C_1 < \langle C \rangle_{exp} \text{ and } C_0 > \langle C \rangle_{exp}$$

When the agent has determined an upwards (downwards) pullback, it will set a target to enter the market if it observes a deviation from the pullback. If a bar closes below (above) the EMA of the stored bars, while the previous bar closed above (below), a long (short) position target is set at the high (low) of the current bar. Once the target is set it will set, D is set to 0.

RSI Reversion

The RSI reversion agent is a reversion type agent that uses the relative strength index (RSI) to determine if the asset is over bought or over sold, and is likely to revert. The agent calculates the RSI using the open to close price returns of the stored bars as:

$$RSI = 100 - \frac{100}{1 + \frac{\langle R_{up} \rangle}{\langle R_{down} \rangle}},$$

where $\langle R_{up} \rangle$ and $\langle R_{down} \rangle$ are the exponential averages of the observed positive and negative price changes determined as $O_i - C_i$. Each agent has a two RSI thresholds; RSI_+ , determining the price it considers the asset over bought, and RSI_- below which it considers the asset over sold. For each agent RSI_+ and RSI_- are determined as:

$$RSI_+ = 100 - P_n, \quad RSI_- = P_n$$

$$t_{direction}, p_{entry} = \begin{cases} -1, C_0 + R(RSI_+)_{up} & \text{if } RSI > 50 \\ 1, C_0 + R(RSI_-)_{down} & \text{if } RSI < 50 \end{cases}$$

Where $R(RSI_+)_{up}$ and $R(RSI_-)_{down}$ are the positive and negative price changes required for the RSI to reach RSI_+ and RSI_- respectively:

$$R(RSI_+)_{up} = \frac{1}{a} \left(\frac{100(1-a)\langle R_{down} \rangle}{100 - RSI_+} - (1-a)\langle R_{down} \rangle - (1-a)\langle R_{up} \rangle \right)$$

$$R(RSI_-)_{down} = \frac{-1}{a} \left(\frac{(100 - RSI_-)(1-a)\langle R_{up} \rangle}{RSI_-} - (1-a)\langle R_{down} \rangle \right)$$

Each RSI reversion type agent differs by the RSI threshold values used to determine if the asset is over-bought or over-sold.

Scalper Reversion

The Scalper Reversion agent is an agent that trades on the assumption that if the price starts trading significantly away from the mean trend, the price is going to revert back to the mean trend. Each time the agent closes a bar it calculates the difference between the closing price and the mean of the closing price of all bars stored. For each bar, it will keep this difference. If the current difference (between the current bar close and the mean close price) is significant, as determined by a z-test, from the average difference of the stored bars it will set a long or short target depending on whether the current difference is negative or positive. In case of a negative difference, it will set a long reversion target price at the current close price and if the difference is positive it will set a short reversion target at the current close price.

C.1. Sensitivity of the Market Simulation

One key element of the integration of the HARK macro model and the financial market ABM, AMMPS execution is to determine the relative size of the markets, or put differently to determine when an imbalance in buy and sell orders, aggregated at the macro-level will have a sizable impact on the price movement of the risky asset. We explored this in two ways: 1) we ran a short, single-day price impact, simulation study (which is reported here) and 2) we ran the full simulation using a range of parameter values, where the values were sufficient to 'break' the market portion of the simulation platform (which is reported in the main body of the manuscript).

For the single-day price impact study, we used the same AMMPS market configuration as was used for the full study and let the simulation run for 30 days, then allowed the *Buy Broker* and *Sell Broker* to trade, using a range of different order sizes, for a single day and measured the percent change in the price of the risky asset. The results are given in Figure C.1.

We generated 5 years worth of daily data ($5 * 250$, or 1,250) for each combination of *Buy Broker* and *Sell Broker* settings (determined by a base number of shares: 500, 1,200, 5,000 and 12,000 and a sell to buy ratio, where the base number of shares was associated with the number 1—e.g., a 3 to 1 ratio with a base number of shares of 500 would mean that the *Sell Broker* sold 1,500 shares and the *Buy Broker* bought 500 shares).

The rate of return is the percent change of the starting and ending price for the final trading period. We adjusted each of the percent changes so that on average a 1 to 1 scenario would have the same mean and standard deviation as the *S&P 500* index over the last five years. We did not match any higher moments, and the tail behavior of the simulated data versus the *S&P 500* index are clearly driving differences in these simulated distributions versus the *S&P 500* index. We anticipate that more advanced agent-based market simulations for the risky asset will allow us to have a distribution of outputs that is more similar to the *S&P 500* dynamics.

It appeared that imbalances starting as small low as 500 shares for the base level can produce noticeably large movements in the daily price of the risk asset, particularly in the presence of imbalances. We have done initial explorations of determining the optimal mapping between HARK dollars and AMMPS dollars, with limited success. We anticipate making this type of exploration a point of future study.

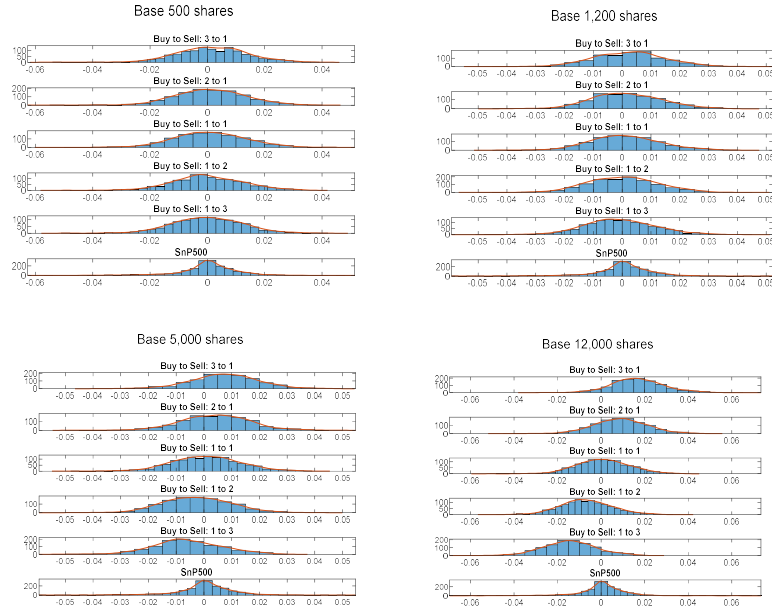


Figure 2: Distributional summary of a simulation study to determine the impact of imbalances between the relative and absolute size of *Buy Broker* and *Sell Broker* orders over a single AMMPS trading period (after a 30-day burn-in period). The Base Shares indicate the number of shares executed over the trading period by both brokers in the 1 to 1 state. In the 3 to 1 scenario, the *Sell Broker* sells 3 times the base number of shares and the *Buy Broker* buys the base number of shares and vice versa for a 1 to 3 scenario.

D. Details of Simulation Design

In this section, we give the details of the algorithm that governs how the software for the HARK model and the financial market AMMPS model interact.

The basic concept is that at the beginning of each trading day, a random subset of the HARK agents pay attention to how much of their wealth is allocated to the risky asset. They then calculate a mean and variance for the risky asset (estimating \bar{r}_t and sr_t), using a weighted combination of the historic mean and variance and the price of the risky asset generated so far in the simulation. (Each simulation assumes 8 quarters, each with 60 trading days.) For each of these attending HARK agents, using their utility (risk aversion) and time value of money, we calculate their optimal holdings in the risky asset, as described above, using the closing price from the previous trading day in the simulation to calculate their current wealth. The

HARK agents desired adjustments (buy/sell) to their holdings are aggregated and passed to a *Buy Broker* and *Sell Broker* respectively. These desired levels of holding of the risky asset by an attending HARK Agent is based on their earnings, consumption, and investment, which are updated on dates spread throughout each quarter.

The simulation platform has three primary scripts or blocks of software:

1. *SHARKFin*. This is python code that controls the overall simulation. It sets up global parameters for each simulation and interacts with the HARK and AMMPS simulation code.
2. The HARK Model, which solves agent-based macroeconomic calculations. It takes inputs from the *SHARKFin* and returns relevant output based on the price of the risky asset. In particular, it determines the aggregate change (buy and sell) in the desired holding of the risky asset by the HARK agents who are paying attention to the financial market.
3. AMMPS, which takes inputs from SHARKFin (aggregate buy, sell order, and dividend rates) and simulates a trading day based on those inputs and the AMMPS agents' state from the previous day. Because the AMMPS agents' states are used, each simulated day becomes path dependent and the state of the market the previous days affect the simulated trading the next. When the market closes for the day, the closing price is returned to SHARKFin and the AMMPS simulation waits for the next input from SHARKFin.

Detailed Algorithm

A. Start - Initialize HARK model

- Create HARK Agents: $HA_i, i = 1, \cdot, n_H$.
 - There are n_H agents.
- Set initial economic conditions, as discussed in Section B.2: $HEC_{i0}, i = 1, \cdot, n_H$.
 - This is done in the HARK program, these parameters are set and stored inside of the HARK program.
 - The amount of wealth that each agent holds in the risky asset is updated at the end of each trading day based on the price of the risky asset and the dividend. The updating is done independently of whether an agent pays attention to the market and wants to change the amount of wealth that they have of the risky asset or whether they are ignoring the market that day.
 - Each agent is assigned a macroeconomic update day d_i . This is the day of the quarter (out of 60) on which the agent undergoes their earning (subject to exogenous shocks), consuming, and investment cycle as described in Section B.1.
- Set risk preferences: $HRP_i, i = 1, \cdot, n_H$.
 - These are parameter values, e.g. ρ or the CRRA parameter, set in the HARK program and they do not change throughout the simulation.

- Start with an initial stock price for the risky asset: $RAP_0 = 100$, the Risky Asset Price at time 0.
 - This value of the Risky Asset Price is maintained in *SHARKFin*.
 - As a convention set the price to 100.
- Use this stock price to determine the initial number of shares of the risky asset held by each agent:
 - $RAS_{i0} = (HEC_{i0}.RAW)DPH/RAP_0$.
 - $HEC_{i0}.RAW$ is the Risky Asset Wealth held by the i^{th} HARK agent at time 0 and RAS_{i0} is the Risky Asset Shares held by the i^{th} HARK agent at time 0. This value is strategically computed as a function of the HARK agent's assets. It is stored and updated by each HARK agent object. It is related to the total wealth held by an agent and the share of wealth that the agent puts in the risky asset, ς_t .
 - DPH is a scaling factor, *dollars per HARK money unit*, which is set as a parameter that is not changed throughout the simulation.

B. Update prices and dividend for each trading day in each Quarter

- At the beginning of each trading day, $t + 1$ (during a quarter), select a random group of HARK agents (with replacement—meaning that the same agent could trade multiple times during a quarter—even back to back days).
 - HA_{i^*} , where i^* indicates that the HARK agent was selected and HA_i , where i indicates that the HARK agent was not selected.
 - Note that time, t , starts at the beginning of the simulation and does not reset each time a new quarter starts.
- For the selected group of HARK agents, HA_{i^*} , calculate the desired level of wealth that the agent wants to have allocated to the risky asset based on the previous day's trading price, RAP_{t-1} , then determine the total number of shares that each agent wants to trade, buy or sell.
 - *First*, determine the expected return and standard deviation for the risky asset (\bar{r}_t, sr_t) to be used for decision in period $t + 1$.
 - To do this the agents use a weighted average of rates of return (ROR_0, \dots, ROR_{t-1}) and a historic (e.g. S&P 500) average return and standard deviation (\bar{r}_0, sr_0) to calculate the 'current' expected return and standard deviation. The rate of return includes the returns based on price movements as well as an exogenous dividend process.
 - The formula for the mean and standard deviation used by the HARK agent is given by

$$\bar{r}_T = \sum_{t=1}^T w_t ROR_t + w_o \bar{r}_0 \quad (6)$$

and

$$sr_T = \sqrt{\sum_{t=1}^T w_t (ROR_t - \bar{r}_T)^2 + w_o (sr_0)^2}, \quad (7)$$

where

- * $w_t = (1 - S_T) \frac{\exp\{a * t\}}{D_T}$, for $T \geq t > 0$ and $w_0 = S_T$ for all t .
- * Further $D_T = \sum_{t=1}^T \exp\{a * t\}$, and $S_T = \exp\{b * T\}$.
- * The parameter $a = -\frac{\ln\{p_1\}}{\delta t_1}$, where $1 > p_1 > 0$ is the proportion (or percentage of decay) for the weights over time δt_1 —with the weight with the largest time (corresponding to the most recent price information) being larger. So $w_t = p_1 w_{t+\delta t_1}$. This captures the decay of the weights or the influence of the most recent percent returns over time. For example, if $p_1 = 0.1$ and $\delta t_1 = 120$, then the percent return of the risky asset two quarters ago has 10% of the impact of the percent return of the risky asset today.
- * The parameter $b = \frac{\ln\{p_2\}}{\delta t_2}$, where $1 > p_2 > 0$ is the proportion of the weight that is assigned to the historical mean and standard deviation, or it is the value of w_0 , when the amount of data that has been seen is equal to δt_2 . So $w_0 = p_2$, when $t = \delta t_2$. For example, if $p_2 = 0.1$ and $t = \delta t_2 = 120$, then $\sum_{l=1}^1 20w_l = 0.9$ and $w_0 = 0.1$.
- * See Figure 3 for numerical examples.
- *Second*, using this (\bar{r}_t, s_t) , calculate the total number of shares of the risky asset that the selected HARK agents want to buy or sell.
 - * Have HARK calculate the amount of wealth that each of the selected agents wants to hold in the risky asset (this is an update or new number), or determine $HEC_{i*t+1}.RAW$.
 - * Using the closing price from the previous period, RAP_t , determine the new number of shares that each selected agent wants to hold, or

$$RAS_{i*t+1} = \frac{HEC_{i*t+1}.RAW \cdot DPH}{RAP_t}. \quad (8)$$

- * Aggregate the total number of shares of the risky asset that the selected agents want to buy and the total (aggregate) number of shares of the risky asset that the selected agents want to buy, $BS_{t+1} = \sum (RAS_{i*t+1} - RAS_{i*t})^+$, where $(\cdot)^+ = \max(\cdot, 0)$, and the total (aggregate) number of shares of the risky asset that the selected agents want to sell, $SS_{t+1} = \sum (RAS_{i*t+1} - RAS_{i*t})^-$, where $(\cdot)^- = \min(\cdot, 0)$.
- * For HARK agents that were not selected leave the number of shares in the risky asset unchanged, or $RAS_{it+1} = RAS_{it}$.

- * For any HARK agents whose macro-update day d_i is the current day of the quarter, update their economic conditions based on the equations in Section B.1:
 - a) The agent solves their consumption choice problem given their current expectations of the market. Note that if the agent was not selected this day, their expectations of the market may be ‘old’ – last updated the previous time that they were selected.
 - b) The agent undergoes permanent and transitory income shocks.
 - c) The agents’ market resources m_t are updated as a function of their income, as well as returns on their saved assets. The return on saved assets in the macro-update step is the *dividend rate*, as capital gains are handled separately.
 - d) The agent chooses their optimal allocation to the risky asset and level of consumption according to their previously computed solution.
 - e) The agent recomputes their target number of shares RAS_{i^*t+1} and adjusts the buy and sell aggregate numbers BS_{t+1} and SS_{t+1} accordingly.
- *Third*, sample the next dividend from the exogenous dividend process. For a *dividend growth rate* G and *dividend standard deviation* σ_η :

$$\eta \sim \text{Lognormal}(0, \sigma_\eta) \quad (9)$$

$$d_{t+1} = d_t G \eta S \quad (10)$$

- *Fourth*, using (BS_{t+1}, SS_{t+1}) and d_{t+1} as the input to the Buy Broker and Sell broker run a day of trading and determine the end of day price of the risky asset RAP_{t+1} and the rate of return ROR_{t+1}
- Use the final price to calculate the amount of wealth each (all) HARK agent has allocated to the risky asset, $HEC_{it+1}.RAW = RAS_{it+1} * RAP_{t+1}$. Additionally, award each agent wealth according to the dividend $d_{t+1} RAS_{it+1}$.
- Repeat for each trading day until the end of the quarter, then repeat for each quarter until the end of the simulation.

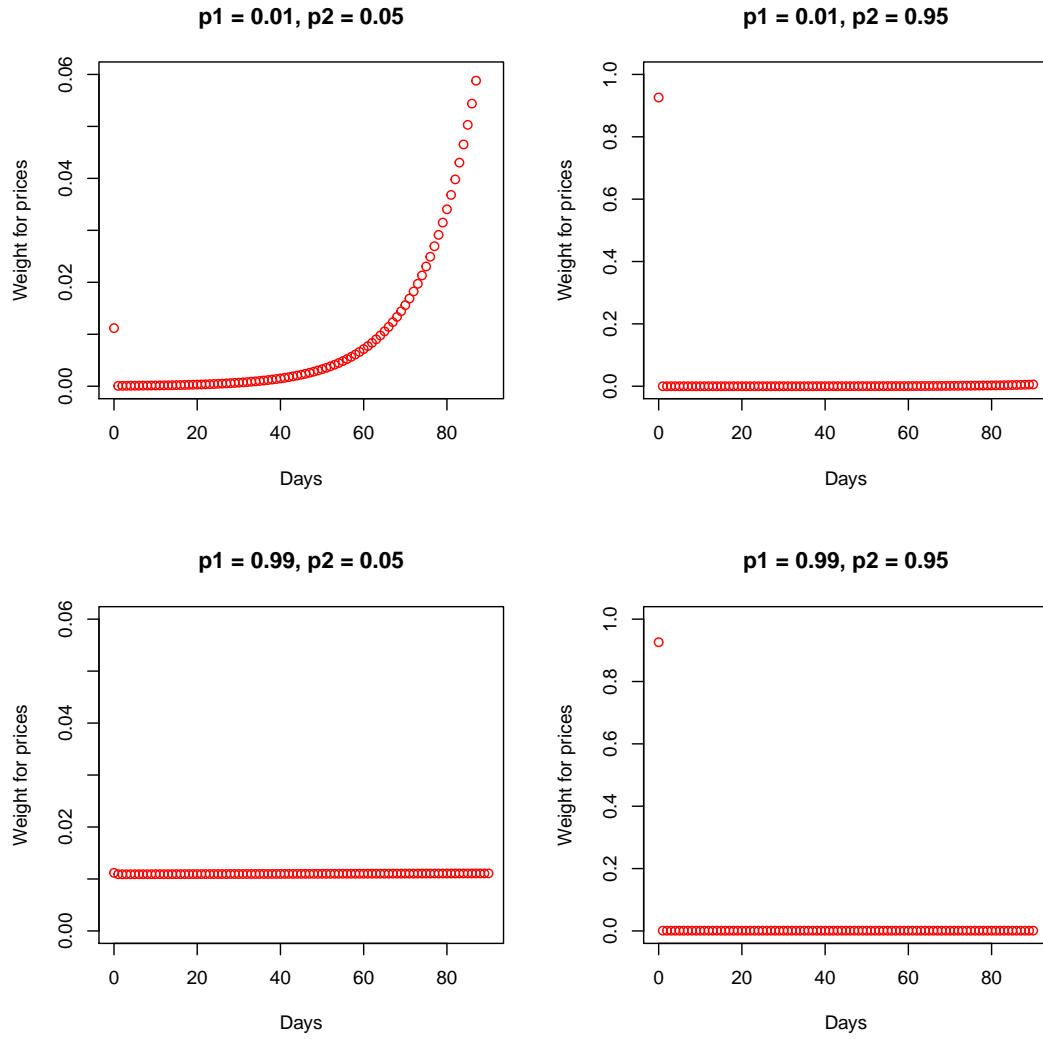


Figure 3: Example of the relative importance of the historical mean and variance and recent price process as a function of p_1 and p_2 . This example shows how 90 days of price data from the market simulation would be weighted. The weight for day 0 represents the importance of the historical mean and variance and the weights for the remaining days show how the return of the risky asset, from day 1 to day 90 are incorporated. When p_1 and p_2 are close to 0 the HARK agents discount past price movements and are focuses on the relative recent past. When p_2 is close to 1 the agents tend to ignore the recent past and place most of their weight on the historic mean and variance. The explicit formula for how these weights are used is given in (6) and (7).