

EVALUATING REGULATORY POLICIES FOR THE US CORPORATE BOND MARKET WITH AGENT-BASED MODELS

Donald J. Berndt

Muma College of Business
University of South Florida
4202 East Fowler Avenue
Tampa, FL, USA
dberndt@usf.edu

David Boogers

FinaMetrics
5125 Palm Springs Blvd
Suite 15204
Tampa, FL, USA
david.boogers@finametrics.com

Saurav Chakraborty

College of Business
University of Louisville
110 West Brandeis Avenue
Louisville, KY, USA
saurav.chakraborty@louisville.edu

ABSTRACT

In this conference paper, we present some next steps in our work using agent-based modeling and simulation to better understand crisis dynamics in financial markets. In particular, the focus is on the US corporate bond market, which has increased dramatically in size and exhibits some signs of systemic risk. One specific risk is related to the growing role of open-end collective investment vehicles (such as mutual funds) in a market characterized by increasingly fragile liquidity conditions. In stressed markets, mutual funds can exhibit run-like behaviors as investors seek to redeem shares in a flight-to-cash. Our goal is to investigate the potential of agent-based modeling and simulation as a tool for evaluating regulatory policies aimed at reducing systemic risk. Some possible regulatory policies include minimum cash-to-asset ratios, swing pricing mechanisms, investor lock ups or gates, and alternative lending facilities. Preliminary simulation results are presented to highlight the promise of agent-based approaches for evaluating potential regulatory interventions.

Keywords: agent-based modeling, corporate bond markets, financial crises, simulation, systemic risks

1 INTRODUCTION

The financial crisis of 2008 again highlighted the evolving nature of the financial system and the potential for complex crisis dynamics to cause economic harm. In the wake of the crisis, US government and regulatory agencies took unprecedented actions in attempts to minimize the destruction of wealth, along with similar initiatives around the globe. Some new regulations aimed to curb risk taking in areas which were at the center of the crisis and limit the potential for contagion to other financial industry segments (with institutions deemed “too big to fail” being a primary concern). History shows, however, that the financial industry tends

to respond to regulation by re-allocating risks across these complex systems. As a result, crises tend to originate in new areas and rarely copy historical patterns.

Financial regulators face a daunting task when trying to develop new policies to regulate markets. Any new regulatory policies can only be evaluated using economic models before being implemented in real-world markets, where any unanticipated side effects can trigger distress and possibly a cascade of other interrelated problems. In addition, modeling crisis behaviors is very challenging since the next crisis is unlikely to look like the past. The overall goal of this research is to use agent-based modeling and simulation to better understand crisis dynamics in specific markets. Economic models are often good at describing normal aggregate behavior, but are less reliable under conditions of stress. Agent-based models can supplement more traditional economic models when evaluating regulatory policies, especially those intended to minimize destructive crisis behaviors (Bookstaber 2012).

The experiments described in this paper focus on the US corporate bond market, which has grown dramatically in size since the financial crisis of 2008. During the bond market expansion, the risks of investing in bonds increased significantly. With yields at historic lows, bonds offer little compensation for interest rate and credit risks while exhibiting heightened price sensitivity to changes in expected returns. Additionally, concerns around the deterioration of liquidity have taken center stage. The challenges are multi-faceted and include the lack of (pre-trade) price transparency, reduced investor heterogeneity (increased “herding” behavior) and the decline in dealer intermediation capacity. Given these issues, the bond market ranks high on the list of potential risks to financial stability, prompting regulators (and industry participants) to question its resilience under stress (Flood, Liechty, and Piontek 2015).

Our focus is on the potential systemic risk caused by the interaction between impaired bond market liquidity and the increased reliance on liquidity transformation provided by pooled investment vehicles, such as mutual funds (Barclays 2015). What regulatory policies might be used to control or at least alleviate these risks? Some possible policies include: 1) setting cash-to-asset ratios to provide buffers against investor redemptions, 2) providing a “lender of last resort” to bolster cash availability under conditions of stress, 3) allow “gates” for investor lock up, or 4) use “swing pricing” to discourage redemptions (by imposing discounts on the price). We intend to contribute to the conversation by taking a somewhat different agent-based simulation approach to explore the conditions under which selected regulatory policies might reduce the risk of market instability due to a run on mutual funds.

2 CORPORATE BOND MARKET AND LIQUIDITY TRANSFORMATION

The US corporate bond market has experienced remarkable growth over the past 25 years, with continued aggressive expansion following the financial crisis. Analysis of SIFMA data shows the overall US corporate bond market expanding from \$5.2 trillion in outstanding nominal (Q4 2007) to over \$8.5 trillion (Q3 2016). The expansion of the bond market coincided with decreasing risk premiums and increasing price risks. In the current low-yield environment, bond prices (which behave inversely to yields) are very sensitive to changes in expected returns, with small increases in interest rates or credit spreads triggering significant drops in bond prices (a feature known as convexity).

Of particular concern is a decline in liquidity when viewed in the context of the changing corporate bond investor base. Historically, corporate bonds were primarily held by institutional investors with a long investment horizon (such as life insurance companies and pension funds) seeking a steady income stream with limited principal risk. Over the past 20 years, however, the market share of “buy-and-hold” investors has continually decreased in favor of new categories of investors such as mutual funds. Using Federal Reserve data, we estimate that mutual funds accounted for roughly 1.7% of the corporate bond market at the beginning of 1981 (see Table 1). By the end of Q3 2016, the mutual fund market share had increased to around 17% (this excludes holdings of exchange traded funds which account for an additional 2.8% of the market).

Table 1: Investor ecosystem circa 1981 (Q1) and 2016 (Q3) by market share. Source: Federal Reserve Flow of Funds data, analysis by authors.

Investor Types	1981 (Q1)	2016 (Q3)
Broker/Dealers	0.43%	0.67%
Insurance Companies	39.72%	23.37%
Mutual Funds	1.68%	17.10%
Others	38.37%	23.76%
Pension Funds	15.44%	6.52%
Rest of the World	4.35%	28.58%

The emergence of mutual funds as significant players in bond markets presents some unique risks. Mutual funds contain redemption features that make them susceptible to withdrawal patterns that resemble classic “run” behaviors potentially triggering fire sales and dramatic price swings. Recent research highlights multiple mechanisms that could lead to disruptive selling by mutual funds, including a first-mover advantage associated with redemption externalities, see (Chen, Goldstein, and Jiang 2010), and the dynamic management of cash holdings to buffer against future redemption risk (Morris, Shim, and Shin 2017). When faced with major redemptions, a fund may need to sell some holdings in order to pay out redemptions. Forced selling of bonds in an illiquid market causes further price drops which in turn leads to more investor redemptions. In a downward market, bond funds can therefore introduce a powerful negative feedback loop which can—through various contagion mechanisms—spill over to other asset markets and impact the real economy (ECB 2016).

3 AN AGENT-BASED MODEL OF THE CORPORATE BOND MARKET

The bond market agent-based model implements a somewhat stylized investor ecology, with participants trading a limited universe of bonds (here 5 representative bonds across common maturities) through dealers that provide transaction immediacy on a principal basis using a request-for-quote (RFQ) protocol. In selecting the set of agents, we aim to model representative corporate bond investor heterogeneity. While there are multiple ways to segment the investor base, our selection of buy-side agents was guided by overall investment mandate and the nature of their liabilities.

1. **Investment mandate:** The investment mandate or policy is perhaps the most important criterion since it encompasses the goals of the organization. Key factors here include the investment horizon and any performance benchmarks used for decision making. Is the investment vehicle passively or actively managed? Can both long and short positions be held or is the emphasis on long only positions? Basically, the mandate constrains the decision-making options and as such is central to implementing an agent class.
2. **Nature of liabilities:** The key factor is leverage with respect to the nature of any liabilities. For example, the hedge fund agent class relies heavily on leverage, while the insurance company does not borrow at all. Are the liabilities more or less runnable? Clearly, an open-ended vehicle such as a mutual fund is subject to investor redemptions at any time, creating inflows and outflows based on performance (or anything else that affects the whims of skittish investors).

As a result, the market model includes three buy-side agent classes representative of a mutual fund, an insurance company and a hedge fund. Out of the three types of buy-side agents, two represent “real money” investors (using no leverage), while the third maintains leveraged positions. Real money investors cannot maintain short positions (must be “long only”); these include an insurance company (typical value investor)

and a mutual fund (using passive index tracking). The leveraged investor represents an unconstrained participant (such as a hedge fund) that can maintain both long and short positions (see Figure 1). The sell-side includes a group of leveraged broker-dealers (three in the minimal model). The dealers are the price makers, maintaining an internal bid/ask spread and responding to trading requests with price quotes. As price takers, the buy-side agents shop for the best price by getting quotes from multiple dealers. While the minimal model includes just six agents, more robust experiments will be conducted with many instances of each buy-side agent class. For more details on this model, see (Berndt, Boogers, Chakraborty, and McCart 2017).

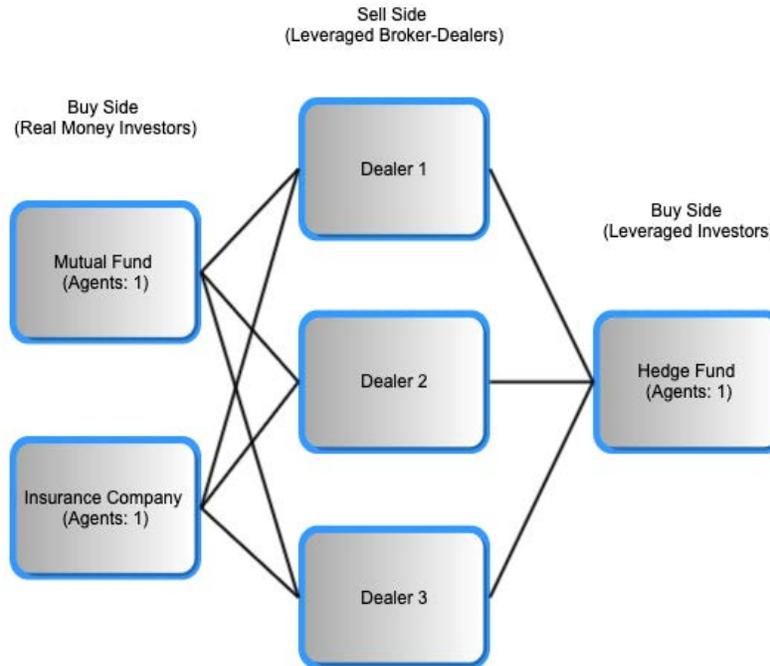


Figure 1: The minimal model includes six agents, three buy-side agents (mutual fund, insurance company and hedge fund) and three sell-side broker-dealers.

3.1 Mutual Fund Class

A mutual fund acts as a real money investor that aims to replicate the performance of a defined benchmark which includes the full universe of bonds available in the initial model. Any dividends or capital gains distributions are assumed to be re-invested. As noted, the fund is long only and does not leverage positions. The fund also maintains a dynamic cash balance as a buffer against investor redemptions (limiting forced sales) and to minimize transaction costs by parking cash until sizable orders can be made. The cash buffer is continually re-evaluated based on an agent-specific cash-to-asset ratio, which also serves as one of the potential regulatory mechanisms.

The cash management behaviors of the mutual fund class are of particular interest, especially since possible regulatory approaches include mandating cash-to-asset ratios or investor lock ups. Algorithm 1 outlines the cash management approach (using pseudo-code). A high/low watermark is maintained for asset sales and purchases based on the cash-to-asset ratio, triggering cash adjustments. Most of the existing literature on fund manager's portfolio liquidation practices assumes the use of either a "vertical" (all bonds are sold in proportion to their weights in the portfolio) or "horizontal" (targeting maximum sales of bonds deemed

Algorithm 1 Basic cash management algorithm in the mutual fund class.

```

1: if get_ca_ratio(bonds) > high_buy_car then                                ▷ Need to invest cash. How much?
2:   Code for asset purchases (not shown).
3: ...
4: else if get_ca_ratio(bonds) < low_sell_car then                            ▷ Need to raise cash. How much?
5:   while cash_to_raise > min_cash_raise do
6:     cusip ← bond_to_sell(bonds)      ▷ Pick bond based on biggest departure from the benchmark.
7:     best_bid ← 0.0                    ▷ Clear the value for best price (bid).
8:     best_none ← None
9:     for i in range(len(sell_side)) do                                       ▷ Request bids from broker-dealers.
10:      bid ← sell_side[i].rfq_bid(cusip)
11:      if bid > best_bid then          ▷ Check that the dealer returns a non-zero RFQ response.
12:        best_bid ← bid                ▷ Keep the best bid and dealer!
13:        best_dealer ← i
14:      end if
15:    end for
16:    sell_side[best_dealer].buy_bid(cusip)    ▷ Trade with dealer who quoted the best price.
17:    proceeds = nominal_amt * best_bid      ▷ Cash proceeds of the sale = bond face value * price.
18:    cash_to_raise ← cash_to_raise – proceeds    ▷ Keep track of the cash raised.
19:  end while
20: end if

```

liquid) approach, without consensus on the prevalence of either approach. Given the investment mandate of our mutual fund class (passive management against a broad market index with composition as outlined in Table 3, Index Weight), redemptions are handled using an approach that minimizes the fund’s tracking error. Bonds to be sold are selected based on the departures (or overages) as compared to the benchmark index. Each bond sale adds to the total cash raised.

3.2 Insurance Company Class

The insurance company agent implements a long-term value investor with a liability driven investment strategy. This agent manages an investment portfolio across equity and fixed income markets and shifts the allocation between these markets depending on market-wide factors such as volatility. The insurance company is a “long only” investor with additional constraints limiting the concentration of risk in any specific bond. In the initial model, leveraged positions cannot be established (hence another real money investor) and we further assume there are no external inflows or outflows in the form of premiums or claims.

Trading activity for the insurance company results from changes in portfolio allocation between equity and fixed income markets. Macro allocation decisions are driven by a number of variables, including equity market volatility as well as the current level and slope of the yield curve. Time series of equity volatility measures such as the VIX are loaded into the model, so different time periods can be used to anchor the simulations in realistic business cycles.

3.3 Hedge Fund Class

The hedge fund agent acts as a short-term tactical trader that follows a relative value trading strategy. As such the hedge fund maintains both long and short positions and makes active use of leverage. In the real

world, fixed income relative value hedge funds have historically been among the most leveraged market participants.

The hedge fund agent is not subject to external inflows (basically a closed end fund) or redemptions (assume investor lock up); its trading capacity is constrained only by the availability of secured financing (leverage) from broker-dealer agents. We assume the hedge fund finances all positions on margin through prime-brokerage style arrangements with some of the dealers. Broker-dealers limit leverage using security-specific haircuts that can be dynamically adjusted depending on market conditions.

All margining is assumed to occur on an overnight basis. At the start of each trading day (a tick in the agent-based simulation), margin requirements are calculated based on current market prices and security-specific haircuts, as set by the broker-dealer agents. The difference between margin requirements and current wealth determines the trading capacity. If the new margin requirements exceed current wealth, the hedge fund is forced to liquidate positions (deleverage) to meet margin calls. Any excess wealth is free to be invested.

3.4 Broker-Dealers and the RFQ Protocol

Dealers respond to a “request for quote” (RFQ) from buy-side agents and trade with clients on a principal basis (there is no inter-dealer market in the initial model). Asset owners must trade with the dealer offering the lowest price. Dealers can maintain both long and short positions. Dealer behavior is limited through regulatory constraints, market discipline and internal risk management directives (the latter expressed through limits on position size and concentration).

The implementation of dealer pricing heuristics follows the model outlined by Treynor, in which the buy-side community provides the “outside spread” (essentially the price at which dealers can offload to the buy-side if forced to cut risk positions) (Treynor 1987). Under normal conditions, dealers make markets by maintaining a narrower “internal spread” (dealers stand ready to trade at prices within the outside spread, subject to their constraints). To start, dealers are instantiated with zero positions (no dealer has a long or short position) and with a specialization based on maturity. Dealer specialization coefficients model each dealer’s reach across the buy-side community and provide the dealer’s specific outside spread (the wider the dealer’s reach for a given bond, the lower his outside spread). See Table 2, listing dealer specializations based on bond maturities. Dealer specialization leads to variations in the prices quoted by different dealers from the outset. In addition to specialization, dealer pricing also factors in current holdings which therefore provides another source of quote variation. In response to a “request for quote” from the buy-side, dealers respond either with a “no quote” (in case the resulting trade would lead to a violation of the dealer’s constraints) or a “full quote” (i.e. the price at which the dealer stands ready to trade for the full amount requested).

Table 2: Dealer specialization across bond maturities.

Dealers/Bonds:	MM101	MM102	MM103	MM104	MM105
Dealer 1	90%	90%	75%	50%	50%
Dealer 2	50%	75%	90%	75%	50%
Dealer 3	50%	50%	75%	90%	90%

3.5 Market Universe and T-Zero Starting Conditions

The market universe consists of five tradable bonds (see Table 3 for details). The bonds are identical with respect to structure, form and major covenants including issuer, redemption (bullet redemption at maturity without optionality clauses) and rate provisions (fixed coupon). The bonds differ along only the following three dimensions.

1. Outstanding nominal amount ranges from 500M to 2B.
2. Maturities cover major points on the yield curve (1, 2, 5, 10 and 25 years).
3. Coupon rates range from 1.75% to 4.00%.

Table 3: Characteristics of tradable bonds.

Bonds	MM101	MM102	MM103	MM104	MM105
Nominal	500M	500M	1B	2B	1B
Maturity	1 Year	2 Year	5 Year	10 Year	25 Year
Coupon	1.75%	2.50%	2.25%	2.40%	4.00%
Yield	1.50%	1.75%	2.50%	2.60%	4.21%
Price	100.247	101.468	98.832	98.249	96.772
Index Weight	10%	10%	20%	40%	20%

At any point in time, all asset owners perceive the same fundamental value for a specific bond. That is, all asset owners use the same valuation model and observe the same input prices (maintaining homogeneous beliefs). The value for the above five bonds is fully reflected in five data points (a par yield curve with five rates) and a simple calculation of a bond's price given its par yield.

The starting conditions for the agent-based simulation include the following items.

- Bond index composition with weights based on nominal amount (Table 3, Index Weights).
- Initial par yield curve and bond prices (Table 3, Yield and Price).
- Starting bond holdings are allocated across one or more agents from the two buy-side agent classes: 1) mutual funds and 2) insurance companies. The mutual funds hold varying market shares from 15% to 35% of the outstanding nominal in the market, with the remaining nominal held by the insurance agent. Mutual funds are given different market shares depending on the simulation experiments. A mutual fund is invested across all bonds in the index based on the index weights. All dealers start with square (zero) inventory positions, apart from any specialization established at the start.

In addition, initial endowments for the buy-side agents are set as follows.

- Mutual fund: In addition to its bond holdings, the fund has an opening cash position reflecting a 5% cash-to-asset ratio.
- Insurance company: Initial portfolio allocation includes a 60/40 split between fixed income and equity markets (assumed to be invested in a broad market index like the S&P 500).

4 EVALUATING REGULATORY POLICIES

As noted earlier, the overall goal of this research is to supplement current regulatory policy analysis tools with agent-based models and simulation. In this section, several policy alternatives are outlined as examples, including the following.

1. Setting forth simple cash-to-asset ratio (CAR) regulations for different actors in the market is among the easiest policy options to implement.
2. Another way to increase the availability of cash under conditions of stress is to provide a “lender of last resort” possibly under government control.

3. The use of “gates” for temporarily locking up investors is a very direct way of controlling (or avoiding) redemptions.
4. A fourth option is to allow open-end funds to impose “swing pricing” whenever an investor withdraws funds under difficult market conditions. These investors would be able to redeem some portion of their investments at reduced valuations.

4.1 Regulating the Cash-to-Asset Ratio

Financial regulators often impose capital and liquidity requirements to bolster the ability of firms to weather uncertain times. For example, the banking industry is subject to a wide array of solvency and liquidity rules that constantly evolve. To begin exploring the use of cash-to-asset ratios in the US corporate bond market, the mutual fund agent is configured to take a cash-to-asset ratio (CAR) parameter as part of the class constructor.

To evaluate the impact of cash-to-asset ratio policies, simulations with changing ratios were run, basically increasing the size of the cash cushion to handle investor-driven redemptions. The simulations traded 5 representative bonds (across common maturities) for 252 trading days (or simulation ticks), with a 100 basis point interest rate shock at tick 50. The first experiments were run with a mutual fund market share of 35%.

Figures 2 and 3 present price trends with different cash-to-asset ratios (6%, 8%, 10% and 12%). In Figure 2 (left) with a 6% CAR, the initial shock-induced price change can be seen as a straight-line drop, with further declines as a result of the redemption-driven feedback loop. This is exactly the type of phenomenon that can be usefully investigated using agent-based modeling and simulation. After the feedback loop reaches an inflection point, the price drop actually accelerates until the market ceases to function. Figure 2 (right) shows a similar pattern, but the initial feedback loop lengthens due to the higher cash-to-asset ratio at 8%.

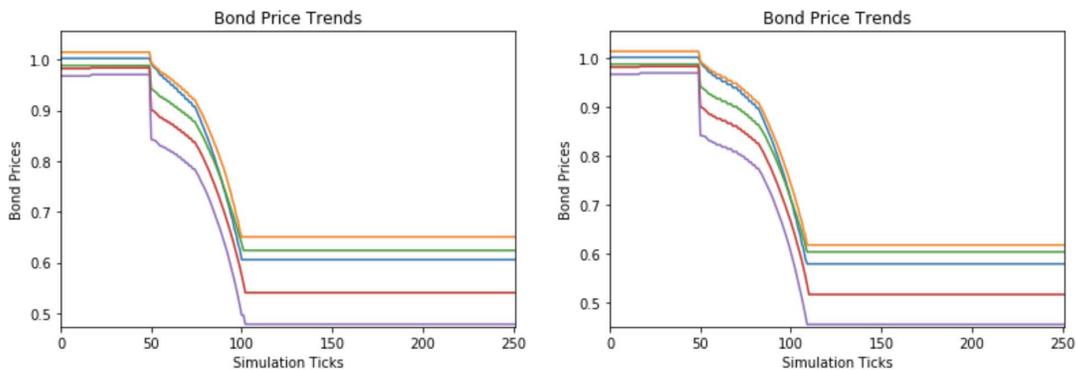


Figure 2: Bond price trends with a 6% CAR (left) and an 8% CAR (right) at a 35% mutual fund market share.

Figure 3 (left) presents an even more drawn out fire sale, with a CAR of 10%. Finally, a CAR of 12% changes the situation as seen in Figure 3 (right). The shock-induced price drop is followed by an extended redemption-driven feedback loop and a subsequent market recovery. Note that in the right panel the y-axis is on a different scale to showcase the shape of the feedback loop. So, the results of our agent-based model and simulations confirm the effectiveness of cash-to-asset ratios as tried and true regulatory policies, along with quantitative information on setting adequate CAR levels.

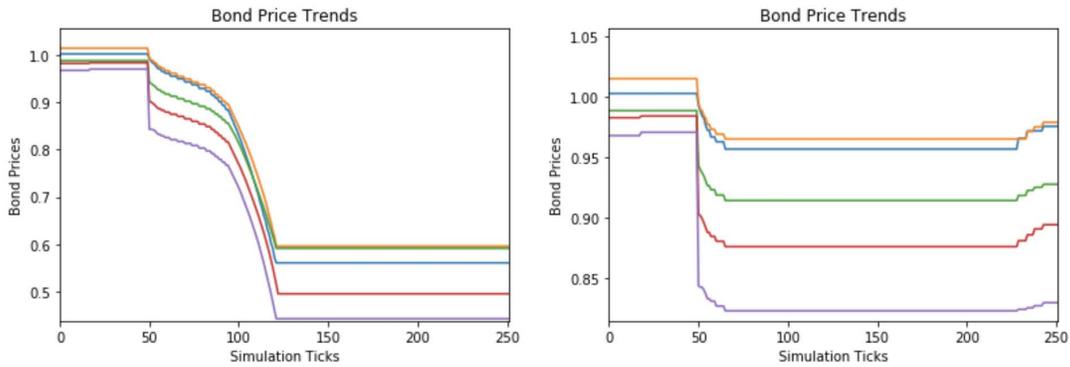


Figure 3: Bond price trends with a 10% CAR (left) and a 12% CAR (right) at a 35% mutual fund market share.

Figure 4 depicts a similar experiment at a mutual fund market share of 30%. The goal here is find more precise thresholds for the cash-to-asset ratio. Figure 4 (left) shows the same shock-induced straight line drop followed by a fire sale and then an inflection point toward a market crash with a 7% CAR. The right panel shows similar drops with a market recovery using just a slightly larger cash buffer provided by a 8% CAR. Note that the overall cash-to-asset ratios required are lower since the mutual fund market share was reduced.

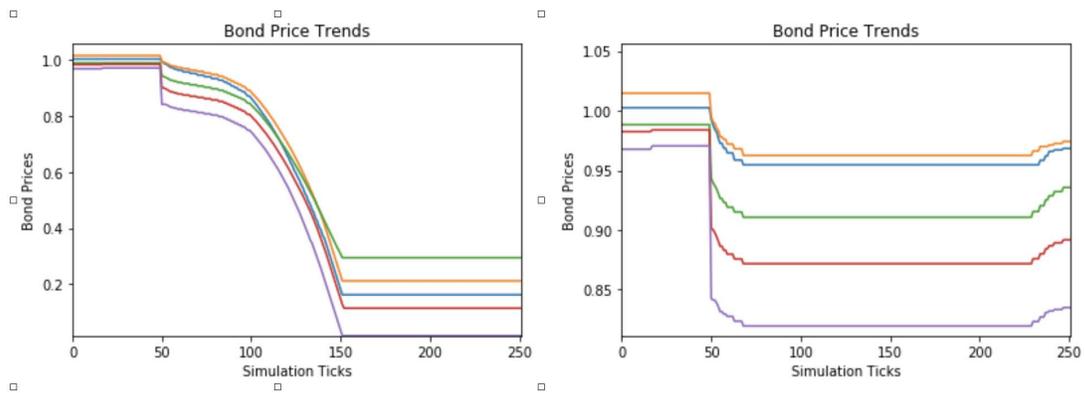


Figure 4: Bond price trends with a 7% CAR (left) and an 8% CAR (right) at a 30% mutual fund market share.

4.2 Providing a Lender of Last Resort

We have also run some preliminary simulations for regulatory policies based on a “lender of last resort.” Basically, this policy introduces a lending facility (possibly government run or at least sanctioned) that would provide cash during conditions of market stress. This could be based on direct loans or by providing Treasury notes (T-notes) in exchange for corporate bonds as collateral. These highly liquid T-notes could be sold for cash in the secondary market, thereby avoiding any corporate bond price effects. Our initial simulations show a very similar pattern to the cash-to-asset ratio experiments. This make intuitive sense since the mechanism is similar, simply infusing cash that can be used to meet investor redemption demands without forcing corporate bond sales, thereby avoiding potential fire sales. One of the advantages of this

approach is that reduced cash-to-asset ratios could be enforced during normal market conditions, with cash infusions used only to stave off crisis conditions.

4.3 Closing Gates (or Investor Lockup)

The use of “gates” for temporarily locking up investors is a very direct way of controlling (or avoiding) redemptions. This would certainly affect the cash position of mutual funds and therefore, have an impact much like raising the cash-to-asset ratio. Currently, a simplified approach is being implemented, allowing mutual funds to skip any sale that imposes an unacceptable cost (see Algorithm 2). This leads to some partial sales based on the differing prices returned by the broker-dealers. Basically, the different bonds may experience varying levels of price drops, making some acceptable for sale and others not. In the basic approach, mutual funds can forego any sales for as long as market turbulence persists. Given this unrestricted ability to “close the gates,” the results are similar to those in the CAR experiments, depending on the gate “cutoff” imposed. This cutoff controls when a redemption is skipped, because a quoted price is too low in comparison with recent trades.

Algorithm 2 Modifications to mutual fund cash management for gates or investor lockups.

```

1: for  $i$  in range(len(sell_side)) do                                ▷ Request bids from broker-dealers.
2: ...
3: end for
4: if best_bid > recent_trade_prices * gate_cutoff then           ▷ Close the gate if price is too low.
5:   sell_side[best_dealer].buy_bid(cusip)                          ▷ Make the bond sale if price is over the gate.
6:   cost = nominal_amt * best_bid
7:   cash_to_raise ← cash_to_raise – cost                            ▷ Keep track of the cash raised.
8: end if

```

Algorithm 2 outlines the updates necessary to implement gates in pseudo-code, with any single bond sale made conditional on receiving a “reasonable” price (assuming the original code in Algorithm 1). A gate cutoff, expressed as a percentage, is used to determine whether to proceed with a specific trade. This approach leads to partial cash raises as certain bonds are skipped due to unacceptably large price drops. This bond-by-bond evaluation makes sense since market volatility is not likely to be evenly distributed, at least in the early stages of a fire sale. Therefore, avoiding the first few trades at fire sale prices may be helpful in stabilizing the market. Of course, there are many possible approaches to investor lockup. For instance, should a group of trades be evaluated before any one is executed? This more expansive view does offer some possible advantages, but for now such comparisons will be left to future work. A straightforward trade-by-trade approach offers a good starting point.

Several simulation runs are used to evaluate the use of gates for investor lockups. Table 4 presents the results of simulation runs with differing gate cutoffs. The gate cutoff is expressed as a price percentage, with a 0.01 meaning that up to a 1% loss is acceptable. Anything beyond a 1% loss are grounds for closing the gates and skipping trades. Other descriptive statistics include the start and end of gate closures (expressed in simulation ticks), the number of skipped trades and the aggregate nominal amount. The same 100-bps rate shock described above (at tick 50) is used here to cause sudden price drops and stress the system. The subsequent price trends are easy to see as the gate cutoff is progressively tightened from 20% (0.2) to 1% (0.01). With losses of up to 20% considered tolerable, gate closures are few in number (7) and start later in the simulation. There is a dramatic change at gate cutoffs in the 1-2% range, with much larger numbers of skipped trades (at 1% the gates are closed hundreds of times).

Figure 5 shows bond price trends for the entire simulation, plotted with a fixed y-axis minimum (0.0) to make price comparisons easier. The price trends for gate cutoffs of 5% (left), 2% (middle) and 1% (right)

Table 4: Simulation results for differing levels of gate cutoffs.

Gate Cutoff	Starting Tick	Ending Tick	Skipped Trades	Nominal Amount
0.2	148	149	7	\$15,890,000
0.1	143	144	7	\$25,539,000
0.05	135	136	8	\$41,794,000
0.02	118	121	17	\$181,328,000
0.01	102	252	752	\$10,549,433,000

show successively smaller market selloffs. The 5% cutoff (rightmost panel) shows the initial shock-induced price drops at tick 50 (straight downward drops), followed by a deep feedback-driven fire sale and large decreases in prices. In fact, at gate cutoffs of 5% or larger, the price trends look much like the complete market crashes shown in Section 4.1. However, gate cutoffs in the 1-2% range start limiting the losses with prices remaining above 0.6. These more restrictive policies result in a prolonged, but more shallow fire sale followed by price stabilization. Mutual funds are able to avoid most losses, preserving market prices at the cost of locking up investors! This is an interesting trade-off for regulatory policymakers to consider.

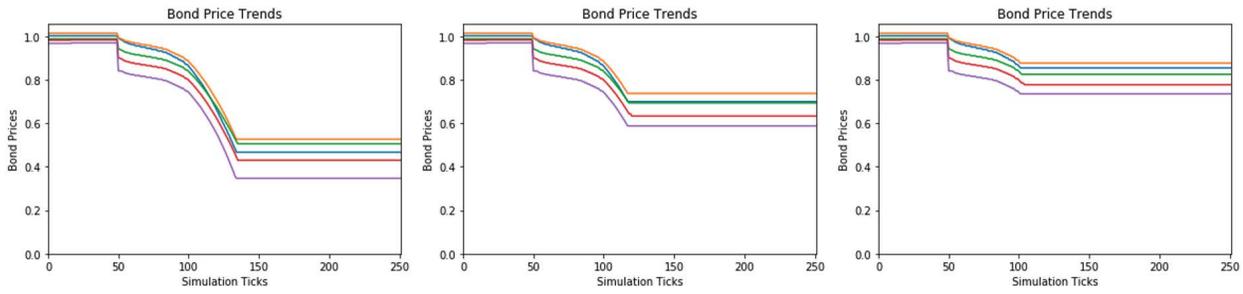


Figure 5: Bond price trends with gate cutoffs of 0.05% (left), 0.02% (middle) and 0.01% (right) with successively smaller fire sales and price drops.

4.4 Swing Pricing (or Investor Haircuts)

Swing pricing mechanisms offer the most sophisticated regulatory mechanism of those considered here. Essentially, the mutual fund would make an offer to meet a redemption request, using a discounted price and returning only a portion of investor funds. This imposes a “haircut” based on the level of stress in the market. This feature is still being refined based on the synthesis of relevant financial literature. Some preliminary findings should be available soon.

5 CONCLUSIONS AND FUTURE WORK

This paper reports on our on-going work related to agent-based modeling and simulation for financial markets. One of the overall goals of this research is to create tools for evaluating regulatory policies that complement traditional economic models. The preliminary experiments presented here use our US corporate bond market model. The US corporate bond market has increased dramatically in size, with an evolving micro-structure that leads to more risk. In particular, mutual funds that provide needed liquidity transformation especially for smaller investors, introduce the risk of redemption-driven feedback loops. Regulatory policies such as cash-to-asset ratios, alternate lending facilities, investor lock ups or gates, and swing pricing

ing mechanisms are all possible ways to reduce these risks. Our initial simulations highlight the promise of agent-based modeling and simulation approaches for evaluating such regulatory policies.

ACKNOWLEDGEMENTS

We gratefully acknowledge the past support of the National Science Foundation (NSF Award 1445403) and the Office of Financial Research (OFR).

REFERENCES

- Barclays 2015. “Mutual Funds and Credit Liquidity, Nobody Move”. Barclays Credit Research.
- Berndt, D. J., D. Boogers, S. Chakraborty, and J. McCart. 2017. “Using Agent-Based Modeling to Assess Liquidity Mismatch in Open-End Bond Funds”. *Systems* vol. 5 (4). Available online: <https://doi.org/10.3390/systems5040054>.
- Bookstaber, R. 2012. “Using Agent-Based Models for Analyzing Threats to Financial Stability”. Technical Report 0003, Office of Financial Research.
- Chen, Q., I. Goldstein, and W. Jiang. 2010. “Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows”. *Journal of Financial Economics* vol. 97, pp. 239–262.
- ECB 2016. “Shadow Banking in the Euro Area: Risks and Vulnerabilities in the Investment Fund Sector”. ECB Occasional Paper 174.
- Flood, M. D., J. C. Liechty, and T. Piontek. 2015. “Systemwide Commonalities in Market Liquidity”. Technical Report 15-11, Office of Financial Research.
- Morris, S., I. Shim, and H. S. Shin. 2017. “Redemption Risk and Cash Hoarding by Asset Managers”. Technical Report 608, Bank for International Settlements.
- Treynor, J. L. 1987. “The Economics of the Dealer Function”. *Financial Analysts Journal* vol. 62 (3), pp. 27–34.

AUTHOR BIOGRAPHIES

DONALD J. BERNDT is on the faculty of Muma College of Business at the University of South Florida (USF). He received his Ph.D. in Information Systems from the Stern School of Business at New York University and also holds a M.S. in Computer Science from the State University of New York at Stony Brook. Dr. Berndt’s research interests include the intersection of artificial intelligence and database systems, data and text mining, agent-based modeling and simulation, as well as a long-standing interest in parallel programming. His email address is dberndt@usf.edu.

DAVID BOOGERS is a principal at FinaMetrics where he advises financial institutions on the optimization of decision processes and analytic infrastructures. As a financial industry consultant, David has worked with broker-dealers, institutional money managers, hedge funds and non-financial corporations across all major financial markets and asset classes. David holds graduate degrees from KU Leuven (Belgium) and Facultés Universitaires Saint-Louis (Belgium). His email address is db@finametrics.com.

SAURAV CHAKRABORTY is on the faculty of the College of Business at the University of Louisville. He received his Ph.D. in Information Systems from the Muma College of Business at the University of South Florida (USF). Dr. Chakraborty’s research interests focus on the application of agent-based modeling and simulation to complex financial systems, including the interbank loan market. His email address is saurav.chakrabortya@louisville.edu.