

Accounting for Strategic Response in an Agent-Based Model of Financial Regulation

FRANK CHENG, University of Michigan

MICHAEL P. WELLMAN, University of Michigan

Due to complex interactions in financial markets, financial regulations can sometimes produce unexpected outcomes, and fail to achieve their macroeconomic goals. We replicate a previous agent-based simulation study which showed that the Basel banking regulations may increase financial instability, counter to their intended purpose. Our replication confirms that this is the case, following the original study's assumption that the financial firms' behaviors are fixed. We then extend the model to account for a possible strategic response, where financial firms adapt to the regulatory regime. Using empirical game-theoretic analysis, we derive equilibria with and without regulation. We find that in the new Basel-regulated equilibria, more funds stay out of default and banks lose less capital. The overall effect of regulation on financial stability becomes benign on most measures when accounting for the strategic adaptation of agents.

Additional Key Words and Phrases: Basel; leverage cycle; game theory; agent-based model; finance; financial regulation; systemic risk

1 INTRODUCTION

The financial crisis in 2008 marked a seismic shift in perceptions of how markets operate. Its reverberations are still being felt now in debt markets around the world [Shaffer, 2016]. The question of how financial crises arise time and again in the midst of sophisticated and logically motivated actors is still incompletely understood. After an initial rush in academic circles to develop new economic models to reflect forensic evidence from the crash and to recommend regulation to prevent future crashes, we have been left with a plethora of different viewpoints.

One approach that has gained influence even in crisis-unrelated research [Healy and Palepu, 2001] focuses on the fact that buyers and sellers do not have the same access to information about assets, especially in the context of debt markets. Based on this asymmetry, Bernanke et al. [1999] showed that periodic, deep financial crashes are endemic due to an overreaction to interest rate changes. The recommendation is clear: force sellers to be more transparent about their wares. Almost everyone agrees that making more information available is a good policy, and regulators commonly work towards this goal.

There are many other convincing perspectives on the crisis. The irrational exuberance [Shiller, 2015] narrative posits a deep behavioral reason for market malfunction. It is also popular to point to systemic factors like lax lending standards and oversight leading to market participants whose priorities move away from the proper valuation of assets [Commission, 2011]. Another body of work points to the seizure of credit markets at the most critical juncture of the crisis for amplifying

Authors' address: Computer Science & Engineering, University of Michigan, Ann Arbor, USA

Email: frcheng@umich.edu, wellman@umich.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. XXXX-XX/2017/1-ART1 \$15.00

DOI: <http://dx.doi.org/10.1145/3033274.3085114>

the crisis. This failure has been attributed to either too much refinancing during normal times [Roberts and Sufi, 2009] or a run on liquidity [Ivashina and Scharfstein, 2010] in response to financial panic. Agent-based simulations elaborating on traditional economic models have also been used. Bookstaber et al. [2017] introduce one where agents have fixed fire-sale behavior and the spread of the crisis can subsequently be measured through different pathways.

All of these models can explain some aspects of the financial crisis, and suggest various emphases for macroeconomic policy and regulation. Given the complexity of the financial system, however, it is difficult to draw direct conclusions from underlying causes to ideal policies. An alternative perspective is to start from the policies, and model the situations where they are beneficial or harmful. We take this perspective, focusing on international banking regulations in the **Basel framework**. Basel's salience follows from its central role in global financial policy, and its force in governing the lending policies of major financial institutions.

Basel regulations consist chiefly of a limit to **leverage**, or the ratio of gross investment to wealth. By taking on debt, this ratio can become arbitrarily large without regulation. An institution with high leverage cannot pay its obligations if its investments underperform even slightly, so it seems natural that limiting leverage may help control default risk. That leverage levels turned out to predict the financial crisis better than interest rates evidences the centrality of this variable to financial stability.¹ Regulator belief in limiting leverage is so strong that their most prominent operational regulation after the crisis was to simply implement a stricter version of Basel. But as such policies had failed to immunize the financial system, there is some uncertainty about how much the stricter version will help in the future.

The **leverage cycle** model of Geanakoplos [2003] describes the ebb and flow of leverage, shedding light on how aggregate leverage may increase to dangerous levels. Importantly, the model *endogenizes* leverage, showing how it evolves through agent decision making. Several key predictions made by this model were borne out by asset behavior during the crisis. First, price fluctuations were disproportionate to fundamental value changes. Second, the failure of a few highly leveraged investors had an outsize impact on prices. Third, the severe tightening of short term credit markets contributed to a deeper crash. All of these predictions were entailed by the model, with leverage choices at Walrasian equilibrium under perfect information.

While the model of Geanakoplos [2003] had the advantage of being fully analytical, it extended only to three periods and thus could not accommodate changing financial regulation over time. In a discrete agent-based extension to the leverage cycle model, Thurner et al. [2012] and Poledna et al. [2014] introduced a Basel-style regulator that imposes a leverage limit on financial firms. As in the latest round of Basel financial regulations, the leverage limit in the model is more stringent in times of high asset price volatility. Using this model, Poledna et al. [2014] argue that Basel regulations could, counter to their purpose, contribute to financial instability. In their model under certain settings, market participants (or **funds**) defaulted more frequently and produced more volatile asset prices when a Basel-based leverage limit was implemented. Market participants also made less profit. However, funds that defaulted were of a smaller size so that the cost per default decreased. A follow-up study by Aymanns and Farmer [2015] suggests that an inverted Basel regulation would be more effective at preventing financial crises.

Our contribution to this story is to observe that a certain fixed assumption in these models, the **aggression distribution**, might naturally be considered a strategic choice by agents. We extend the model of Poledna et al. [2014] to a game, where the agents strategically choose their aggression level in response to the regulation regime. We employ simulation-based methods to analyze this

¹A study by Geanakoplos et al. [2012] established a link between financial crashes and leverage levels that successfully predicted data from 2.2 million American households.

game, identifying aggression distributions that are in approximate Nash equilibrium. Using this approach we reverse in aggregate the finding that Basel causes more defaults. Losses due to default decrease further, and agent profits decrease less. The overall case against Basel is thus weakened by our findings.

In Section 2 we describe the model used by Poledna et al. [2014], and explain the effect of Basel regulations on the leverage cycle in this model. We also describe implementation details and assumptions used in the original work that carry over into our independent implementation. In Section 3 we argue that agent aggression levels in this model should be treated as strategic variables. We propose an approach to analyze the strategic adaptations of agent aggression using Nash equilibrium. Finally, we design an experiment that distinguishes the effect of Basel on the leverage cycle for a fixed set of agents (studied in previous work) from its effect due to strategic adaptations of agents found in equilibrium. We measure the effect of Basel on financial stability by looking at default rate, agent profits, capital losses, and price volatility. The strategic adaptation of agents on its own improved every measure except price volatility, which was unchanged. In aggregate, we find that Basel decreases default rate and capital losses while also decreasing agent profits and increasing price volatility.

2 AN AGENT-BASED MODEL OF THE LEVERAGE CYCLE

Following Poledna et al. [2014], we adopt and extend the basic leverage cycle model of Geanakoplos [2010]. This model has been credited with predicting characteristics of the 2008 financial crisis years in advance. By treating leverage as an endogenous decision, it characterizes equilibrium in financial markets in terms of leverage taken by each agent in addition to the asset price and the interest rate on debt. Since Basel targets leverage as a policy variable, the leverage cycle model seems like a natural candidate for evaluating recent financial regulation.

However, the original leverage cycle model is limited in its ability to express real-world complexity. For example, it cannot be extended for an arbitrary number of periods and short asset positions are not considered. Basel regulation, since it responds to the historical volatility of asset prices, is not easily incorporated. To address such issues, Poledna et al. [2014] develop a discrete agent-based model (ABM). At a high level, this model clears the market in each period for persistent financial agents. Agents carry over their wealth to each new period, where demand undergoes a random shock and new prices are formed. Basel regulation is easily expressed in this new model since a regulator may act in each period with knowledge of historical prices. The flip side is that agents are no longer coming to an intertemporal equilibrium as in the leverage cycle model. But leverage remains an endogenous decision by each market participant.

2.1 Agent-Based Model

Poledna et al. [2014] include four types of agent, who interact in a market for a single risky asset. The asset has a *fundamental value*, V , and a market *price* $p(t)$ at time t , determined by the cumulative demand of all agents. The central actors are informed value investors called *fund managers*. Fund managers know the fundamental value and adjust their demand for the asset based on a *mispricing signal* $m = V - p_t$. The fund managers exhibit heterogeneous demand as a function of m , reflecting their differing levels of *aggression* in pursuing investment opportunities. The more aggressive a fund is, the more leverage it will take to pursue a given mispricing signal. Aggression reflects factors such as a fund's confidence that the asset price will return to its fundamental in short order, and its tolerance for risk. A fund manager that buys an asset while it is priced under its fundamental is betting that the price will move back up and make her a profit.

To buy or short-sell assets, fund managers may take loans from the **bank**, treated as an agent in the ABM of [Poledna et al. \[2014\]](#). The bank provides credit to funds but requires a set amount of **collateral** per unit debt. In effect, this collateral requirement imposes a **leverage limit**. To see this, consider how a leveraged investor keeps her books. She knows that to borrow a dollar from the bank she must commit some amount of wealth as collateral. Call this collateral requirement X . But she only has so much wealth, W , to commit in total. Thus there is a maximum amount she can borrow, W/X , which is constant given W . The ratio of this maximum amount to her current wealth is $1/X$. If the bank decides to increase X , it effectively decreases the leverage limit. In practice, we refer to the bank setting a leverage limit without relating explicitly to collateral. In general the investor is taking on debt in order to take a larger investment position, which she is free to do until she is over the leverage limit. Then she must wind down her investment position and reduce her debt until she is in compliance with the bank.

Thus, in any given time period, funds are limited in the size of the asset position they can take, either long or short, in response to asset mispricings. Note that the bank may be forced to take losses when fund managers default. These losses are recorded but the bank itself never defaults, as it is assumed that there is an unlimited bailout fund.

Mispricings are made possible by stochastic, weakly mean-reverting asset demand from aptly named **noise traders**. The noise traders represent a collection of ill-informed investors. Their collective behavior exerts a random shock to total asset demand, which in turn shifts the market-clearing asset price.

The final agent type is the **fund investor**. The fund investor's role is simply to move capital from fund managers with poor (historical) performance, toward those with good performance.

2.2 The Leverage Cycle and Basel: Intuition

It is easy to see that without any limits on leverage, the price of the asset would always return to its fundamental value, as funds would just increase their position against the direction of the shock until the mispricing disappears. It requires just one fund able to take unlimited size positions on the asset to ensure that there is never any price volatility around V . But with leverage limits, the system exhibits complete leverage cycles over time. The qualitative steps to the cycle are shown in Figure 1.

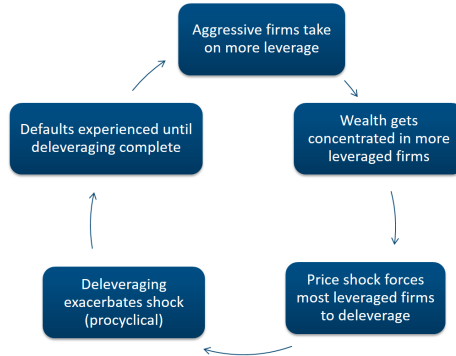


Fig. 1. Stages of the leverage cycle modeled by [Poledna et al. \[2014\]](#).

When shocks are small enough, wealth in the system becomes more and more concentrated in the most aggressive funds as they take more effective leverage in pursuing mispricing opportunities. This gradually increases the sensitivity of total wealth in the system to random price shocks until

some investors reach their leverage limit. At that point, a decrease in the asset price would push the leverage ratio over the limit,² and so the investor would be forced to sell to get back in compliance when she would actually rather be buying the asset. This action is *procyclical*, since the act of selling further depresses the price, which can lead to further deleveraging by other funds. It may even induce a mass sell-off (or buy-off for short positions) as other, formerly less leveraged funds struggle to satisfy the leverage limit due to the acceleration of the price movement against their positions. Eventually enough funds have either defaulted or deleveraged that, collectively, they have become insensitive to further price movements, and the cycle begins again. Price volatility, probability of investor default, and losses on defaulted loans are all elevated due to this leverage cycle.

Within this model we can also incorporate Basel-style regulation, which makes leverage limits more stringent during periods of high volatility and less stringent otherwise. Under fixed fund behavior, this increases the probability that investors get into procyclical situations, since the leverage limits themselves will adapt to be procyclical. Intuitively, higher leverage limits in volatile times could allow funds to provide a voice of reason and return the price to the asset's fundamental value. So Basel, by doing the opposite, potentially makes volatility and default more likely. On the other hand, Basel may reduce the losses suffered by the bank on defaulted loans, since the doomed funds were forced to deleverage more quickly before their default.

Our study starts by reimplementing this ABM and replicating the original results. The remainder of this section describes model details; our extension to incorporate strategic response is described in Section 3. To implement the model, we used as references the description of Poledna et al. [2014] together with code provided by these authors in response to our queries. Unless otherwise stated, all of the following modeling decisions and parameters are as implemented in the prior work. We attempt to keep the model as close as possible to prior work in order to isolate the effect that our equilibrium concept has on the evaluation of Basel policies. These assumptions were justified in the original work by drawing parallels with real data on financial crashes.

2.3 Schedule

The state of the market is defined by an equilibrium asset price and the resulting holdings and debt levels of the fund managers. The agent-based simulation iterates between a price formation phase and a wealth update phase. We describe these phases at a high level below to give a sense of the scheduling of tasks.

Price Formation.

- (1) Each fund manager has a fixed demand as a function of the mispricing signal and its current wealth. These demand functions depend on the maximum leverage allowed by banks as well as the fund's idiosyncratic aggression parameter.
- (2) Noise traders demand the asset according to a stochastic process such that without fund managers the asset price would weakly mean revert around the fundamental value of the asset (an Ornstein-Uhlenbeck process). The noise trader's demand is what makes the mispricing signal nonzero in each period.
- (3) The price of the asset is set at a level that clears the market given the collection of demand functions submitted by noise traders and fund managers. This is the Walrasian equilibrium.

Wealth Update of Fund Managers.

²For example, consider the investor with \$1 who borrows \$5 to buy assets worth \$6. Her leverage ratio is currently 6. If the asset value decreases to \$5.5, the investor still owes \$5 to the bank, so her wealth is now \$0.5. Her leverage ratio increases to 11, which may be out of compliance with the bank.

- (1) The new asset price as determined in the price-formation phase induces revised wealth levels for each of the fund managers based on the market value of their holdings.
- (2) The fund investors obtain a new datapoint regarding the profitability of each fund manager by observing the wealth gained or lost in this update. They withdraw or deposit capital into funds accordingly.
- (3) Banks enforce leverage limits on funds. Funds that exceed leverage limits are subject to a **margin call**, requiring they change their position in the asset to comply with the limit. Note that the violation of leverage limits can be triggered by a change in asset price, by the movement of capital by fund investors, or some combination.

2.4 Basic Definitions

We define the key variables describing the state of a fund h at time t .

- (1) $D_h(t)$ is the amount of the asset that the fund holds at t . It may be positive, meaning the fund owns a positive amount of the asset, or negative, indicating a short position.
- (2) $M_h(t)$ is the cash position of the fund. The cash balance is changed each period based on purchases or sales of the asset. A negative $M_h(t)$ indicates a debt position.
- (3) The wealth of a fund is defined by

$$W_h(t) \equiv D_h(t)p(t) + M_h(t).$$

where recall that $p(t)$ is market price of the asset. All nondefaulted funds have $W_h(t) > 0$. This entails in particular that a short asset position is always accompanied by a positive cash position.

- (4) The **leverage** taken by each firm is defined as

$$\lambda_h(t) \equiv \begin{cases} D_h(t)p(t)/W_h(t) & \text{if } D_h(t) \geq 0 \text{ (long)} \\ M_h(t)/W_h(t) & \text{if } D_h(t) < 0 \text{ (short)} \end{cases}$$

Recalling that $W_h(t)$ within a period is constant and positive, and therefore the numerator is always positive, $\lambda_h(t) \in [0, \infty)$. It can in principle be made arbitrarily large by taking an investment position of sufficient magnitude, although due to leverage restrictions the funds are constrained in the positions they can take. Note that leverage is ill defined for $W_h(t) \leq 0$.

These definitions allow for a fund to have **short** and **long** positions on an asset. For a given wealth, a fund with a long position has a positive value of asset holdings $D_h(t)p(t)$ with cash $M_h(t)$ unrestricted. A common scenario is negative $M_h(t)$, meaning the fund has taken on debt to finance their long position. This necessarily makes $D_h(t)p(t) > W_h(t)$. Thus any fund indebted to the bank must have $\lambda_h(t) \geq 1$.

On the other hand, a fund with a short position owes future shares to its counterparty ($D_h(t) < 0$) but holds the cash it obtained from the sale of these future shares ($M_h(t) > 0$). If the price of the asset $p(t)$ increases, the wealth of a fund that is long on the asset goes up while the wealth of a fund that is short goes down. The opposite occurs if $p(t)$ decreases.

2.5 Regulatory Environments

In the somewhat misnamed **unregulated environment**, $\lambda_h(t)$ is constrained to be less than a parameter λ_{\max} for all h . The **Basel environment** adaptively selects maximum leverage $\lambda_{\max}^{\sigma(t)}$ between periods based on price volatility $\sigma(t)$. $\lambda_{\max}^{\sigma(t)}$ is calculated at the beginning of every timestep

t according to the formula

$$\lambda_{\max}^{\sigma(t)} = \min\{\lambda_{\max}, \max\{1, \frac{\sigma_b}{\sigma(t)}\}\}, \quad (1)$$

where $\sigma_b = 0.0118$ is a set benchmark level of volatility, and $\sigma(t)$ is measured as the average volatility of the log asset price in the previous 10 periods. Before 10 periods are available, all price data is used. Thus, maximum leverage requirements are adjusted downwards in periods of high volatility and are allowed to reach λ_{\max} when volatility is low. In the unregulated case, λ_{\max} is still a hard limit on leverage but it does not adapt to market conditions.

2.6 Fund Demand

Leverage limits effectively constrain the maximum and minimum demand for each fund manager. Individually, each fund manager h with aggression parameter β_h has a demand $D_h(t)$ of the following form, at time t :

$$D_h(t) \equiv \begin{cases} (1 - \lambda_{\max}^{\sigma(t)})W_h(t-1)/p(t) & m(t) \leq (1 - \lambda_{\max})/\beta_h \\ \lambda_{\max}^{\sigma(t)}W_h(t-1)/p(t) & m(t) > \lambda_{\max}/\beta_h \\ \beta_h m(t)W_h(t-1)/p(t) & \text{otherwise} \end{cases}$$

Here the price $p(t)$ at time t is the free variable and all other quantities are fixed. The fundamental value is set at $V = 1$ and $W_h(t-1)$ denotes the total wealth accumulated by h in the last period. In the original study [Poledna et al., 2014], the aggression parameter β_h is fixed to be $h \times 5$ where $h \in \{1, \dots, 10\}$ is the index of the fund, of which there are ten in total. In our strategic analysis we allow the funds to choose the β_h 's, but for now they are fixed. Notice that in the lower and upper regions of price, demand does not depend on β_h . In particular, demand is fixed in these regions. This comes from the fact that funds have hit their leverage limit and cannot borrow further to pursue investment opportunities. An example of the demand function for two funds with different β_h is shown in Figure 2.




Fig. 2. Demand as a function of mispricing. The mispricing values beyond short saturation and long saturation on either side result in no additional demand for the asset because leverage limits have been reached.

2.7 Market Clearing Price

Every period, prices are used to coordinate asset allocation. Each fund submits their demand function truthfully and an equilibrium is found after the noise traders generate stochastic demand. Each period gets its own independent price. The only thing that transfers information between periods is the wealth that each fund ends up holding.

Given the demand functions of all funds as well as the noise trader demand $C(t)$, prices are formed via the market clearing condition

$$\sum_h D_h(t) = N - C(t), \quad (2)$$

where $C(t)$ is generated by

$$\log C_n(t) = \rho \log C_n(t-1) + \sigma_n \chi(t) + (1-\rho) \log(N).$$

The noise parameter σ_n is set to 0.035 and $\chi(t)$ is an i.i.d. standard Gaussian draw. Parameter ρ is set to 0.99, representing extremely weak mean reversion. This is the Ornstein-Uhlenbeck process, a standard model in finance for asset price dynamics.

We set $N = 10^9$. Equation (2) is a piecewise linear function of $p(t)$ which can be solved using standard methods.

2.8 Wealth Update

Using the $p(t)$ (recall it is embedded in the demand functions $D_h(t)$) obtained as solution to (2), the wealth of each fund $W_h(t)$ can now be updated. Assume for now that $D_h(t-1)$ is positive, that is, fund h takes a long position. There are three possible sources of wealth change in the model: market value of assets, interest payments, and equity flows.

The increase in asset value for each fund is simply $D_h(t-1)[p(t) - p(t-1)]$, and reflects the return on an investment made in the previous period. This is the fund's profit or loss.

In general, funds finance their investment opportunities using debt from the bank. If the amount of cash $M_h(t) = W_h(t) - D_h(t)p(t)$ is negative, then the fund has borrowed from the bank to buy $D_h(t)$ shares of the asset. In this case the fund pays a fixed interest rate $S = 0.015\%$. After paying interest, the fund's total available assets for withdrawal is now

$$M_h^*(t) = D_h(t-1)p(t) + M_h(t-1)(1+S).$$

The other source for changes in wealth is withdrawals or deposits made by fund investors. These make their decisions based on recent performance of each fund. Fund investors examine each fund's rate of return (estimated using an exponential moving average) $r_h(t)$ and withdraw/deposit an amount $F_h(t)$ as follows:

$$F_h(t) = \max\{-1, b(r_h(t) - r_b)\} \max\{0, M_h^*(t)\},$$

where the benchmark return r_b is 0.003 and b , controlling how strongly fund investors react to historical performance, is set at 0.15. The overall motion equation for wealth is then

$$W_h(t) = W_h(t-1) + D_h(t-1)[p(t) - p(t-1)] + F_h(t) + M_h(t-1)(1+S).$$

Analogously, for a short position we have

$$M_h^*(t) = D_h(t-1)p(t) + M_h(t-1) + D_h(t-1)S,$$

$$W_h(t) = W_h(t-1) + D_h(t-1)[p(t) - p(t-1)] + F_h(t) + D_h(t-1)p(t-1)S.$$

Once all fund wealths are updated, the time step is complete and we go back to price formation.

2.9 Default

Each fund starts with $W_0 = 2 \times 10^6$ in wealth. Default occurs when wealth goes below a critical value $W_e = 2 \times 10^5$. In this case, the fund's wealth, demand, and cash positions become zero for 100 timesteps, after which the fund is resuscitated at the same level of aggression β_h and given W_0 to start operations again.

3 SYSTEMIC RISK EVALUATION

The motivation for using this model is to see what effect regulation can have on the leverage cycle, and how this impacts systemic risk. We first develop an approach for evaluating the effect of regulation when fund managers are given the opportunity to adapt strategically according to changing external conditions. Then we define metrics for measuring systemic risk and perform an experiment measuring the effects of the strategic adaptations on systemic risk.

3.1 Strategic Aggression Behavior

Recall from Section 2.6 that each fund is assigned an aggression level. We call these assignments the **aggression distribution**. Poledna et al. [2014] assume that the aggression distribution is fixed. This assumption is problematic for several reasons, but fundamentally because it limits the responsiveness of fund behavior to environmental conditions, either of the market or the regulatory system.

We focus on this assumption in particular because the entire model appears to be highly sensitive to the aggression distribution. For example, if all funds have low aggression then leverage will almost never be high enough to induce defaults. As circumstantial evidence for this, in all the experiments described in Section 3.4, no fund at an aggression level of 5 ever defaulted. This may be considered an extreme setting, but it is disturbing that one can generate almost any story one wants by manipulating this fixed assumption. It is far better to derive the aggression distribution from something more fundamental, like strategic choice.

Therefore we treat aggression level as an *endogenous* variable, taking account of the dependence of each fund's payoff on the aggression of other players. That is, we cast the entire scenario as a **game** played among funds, where aggression distributions are **strategy profiles** and aggression levels are **strategies**. This allows us to search through the space of possible aggression distributions to find a **Nash equilibrium** where no fund can make more profit by changing their aggression level. This approach, notably, leaves each fund's aggression level (and thus demand function) in a given period fixed. Within a given time period we still have perfect information and complete markets, so clearing the asset market still produces a Walrasian equilibrium. The only intertemporal choice funds make is aggression level, which is chosen for all periods simultaneously. Thus our strategic analysis is limited to aggression levels, that is, we do not consider other strategic market behavior.³ This preserves the spirit of the leverage cycle story for a given equilibrium choice of aggression distribution. We see below that the aggregate effect of equilibrating prices in each period and aggression over all periods has a substantial effect on financial stability.

3.2 Experimental Setup

In our approach we evaluate the introduction of Basel into the model on two dimensions. First, given a particular aggression distribution, we measure the effect of regulatory environment on systemic risk measures. This isolates the effect of Basel on the leverage cycle story given fixed aggression levels. Second, given a particular regulatory environment, we look at the effect of

³Since funds are not atomic, there could be room for exercising market power. In our setting we view the market as competitive enough at that level to ignore such options.

changing the distribution of fund aggression levels. This isolates the effect of Basel on the strategic choice of aggression levels while leaving its effect on the leverage cycle constant.

Nash Equilibrium. First, some definitions and implementation details. A *pure profile* is an H -dimensional vector β containing a single aggression level, or **strategy**, for each of H funds. For our study we chose a number of funds $H \in \{10, 21\}$ and a set of 7 possible aggression levels/strategies $\Gamma = \{5, 10, 15, 20, 30, 40, 50\}$. Thus a particular fund's **mixed strategy** is a 7-dimensional vector of positive real numbers summing to 1, corresponding to the probabilities that the fund will pick each of the seven strategies for use over the entire simulation. A mixed strategy D induces a probability distribution over pure profiles. Drawing from Γ using the multinomial distribution specified by D assigns a β_h to each of the H funds.

There are many strategy profiles and we cannot evaluate them all. Our focus is on profiles that are strategically stable. A mixed strategy D^* is stable if no fund has an incentive to deviate from drawing its strategy from D^* . In other words, D^* is a **mixed strategy symmetric Nash equilibrium** over fund aggression levels. The problem becomes finding D^* in a computationally feasible way, since the number of profiles increases exponentially with the number of players and strategies. To do this, we first need to define how the game is played and how payoffs are generated.

Payoffs and Game Definition. Each run of the model provides a ready made payoff: the fee earned by each fund. To find this fee, we first run the model for $T = 50,000$ periods using $H \in \{10, 21\}$ funds and a given β vector representing the aggression level for each fund. At the end of the simulation, we annualize into 1000 years containing 50 periods each as before. In each year i starting at periods t , we calculate the average size of fund h , $W_h(i) = \sum_{t=t}^{t+50} W_h(t)/50$ as well as the profit $\zeta_h(i) = W_h(t+50) - W_h(t) - \sum_{t=t}^{t+50} F_h(t)$, which is simply the change in wealth net of investor withdrawals and deposits over the course of the year. The fee for year i is then defined as $0.2\zeta_h(i) + 0.02W_h(i)$ and the total fees over the entire simulation for fund h are

$$\sum_{i=1}^{1000} 0.2\zeta_h(i) + 0.02W_h(i).$$

This is the specification used by Poledna et al. [2014], based on standard practice in the financial industry. It captures the common 2-and-20 fee structure, which pays hedge funds 2% of assets under management and 20% of value returned to investors annually. This annualization (1000 years that are each 50 periods long) is done according to the original paper [Poledna et al., 2014], where the model was calibrated to real yearly data. Note that the payoff for a given mixed strategy D depends on whether or not Basel is active as well as λ_{\max} and is an average over many 50,000-period simulations where pure profiles β are generated from D .

Environmental Settings. In summary, we can view the model input as choice of regulation R , a particular β vector of strategies representing a pure profile, and parameters λ_{\max} . The model itself can be viewed as a black box function, mapping to the model output which is a vector of rewards, or fees, for each fund. In reality, remember that we are relying on separate Walrasian equilibria in each of the 50,000 periods to set prices.

Finding the Nash equilibrium distribution $D^*(R, \lambda_{\max})$ of such a system requires searching through input β and evaluating fees until an equilibrium is found. We adopt a heuristic approach, inexhaustively exploring a space of strategy profiles and estimating their payoffs through agent-based simulation. This method goes by the name of **empirical game-theoretic analysis** [Wellman, 2016], as it applies standard game-theoretic concepts to a game model derived by empirical means. We use a suite of tools developed by Cassell and Wellman [2013] to automate the simulation and

data management. We were able to obtain approximate $D^*(R, \lambda_{\max})$ for $R \in \{\text{Basel}, \text{Unregulated}\}$ and $\lambda_{\max} \in \{8, 20\}$. These levels of λ_{\max} were chosen to reflect a wide range of conditions.

Note that we chose to show results for the same set of strategies Γ as that used by Poledna et al. [2014]. This was mainly for comparability, but we did test lower minimum (down to 1), higher maximum (up to 70), and a finer resolution in between for aggression levels. We found that the results were not qualitatively different. For example, a beta of 1 was not adopted in the mixed strategy equilibrium when added to the reported set of strategies. In all tested cases, Basel encouraged less average aggression in equilibrium.

Empirical Game-Theoretic Analysis. Our empirical game approach imposes two major approximations. First, the set of available strategies Γ is restricted to a modest number of enumerated choices. As noted above, we limit attention to $|\Gamma| = 7$ choices of aggression level. Since the number of strategy profiles is exponential in the number of strategies, we have no choice but to impose such a limit. Second, for the $H = 21$ setting we used a technique called *deviation preserving reduction* [Wiedenbeck and Wellman, 2012], which employs aggregation to model the 21-player game in terms of a reduced six-player game.

3.3 Systemic Risk Metrics

Our metrics for systemic risk differ from those provided by Poledna et al. [2014] in one key respect. Rather than focusing on the risk of the most aggressive fund manager, we instead evaluate risk aggregated over all fund managers. The aggregate measure reflects a more direct evaluation of the entire economy. Moreover, as we allow distribution of aggression to vary based on endogenous choices, we need a way to evaluate the riskiness across situations with different maximum aggression levels. This difference is only relevant for measuring probability of default and capital losses.⁴

Probability of Default. To calculate probability of default of a single fund for a single run of the model, we record the years during which a default occurred. Then we divide the number of defaults by the number of years during which default was possible, K_h . $K_h \leq 1000$ because defaulted funds go out of operation for 2 years, or 100 time steps, as specified in Section 2.3. Thus, there are entire years during which default is impossible. Note that there are also years during which default is only partially possible as the fund is still in a defaulted state at the beginning of the year. We choose to exclude these years from K_h as well. The overall probability of default is then

$$\sum_h \# \text{defaults}_h / (K_h H).$$

Here, recall that H is the total number of funds.

Capital Losses to Banks. Capital losses are the amount of loss banks suffer when a fund manager defaults. It is defined as the wealth $W_h(t)$ of the defaulting fund at default. This is an important metric for systemic risk because these losses are often paid for through bank bailouts, which detract from social welfare.

Capital losses were measured on an annual basis, in the same way as for probability of default. Yearly capital losses are totaled, then averaged across the number of funds and number of years during which default was possible.

⁴We also calculated the original study's metrics and verified that our implementation matched the results reported by Poledna et al. [2014].

Price Volatility. Price volatility is detrimental to the extent it contributes to financial uncertainty. One of the functions of financial markets is efficient price discovery, and one might hope that in a single asset market where some investors have perfect information about the fundamental value of the asset, this function would be performed well. Price volatility is a good proxy for how well price discovery is performed. Indeed, if we allow $\sum_h D_h(t)$ to be large compared to $C(t)$ for all t , we would expect price volatility to be near zero as the effect of the noise traders is minimized. Price volatility is the average of the annual variances of $p(t)$. This metric does not depend on K_h like the other two, since prices can be formed regardless of the default status of any fund.

3.4 Systemic Risk Results

First we note that in every setting we evaluated, the strategy profile that maximized total fees was for every fund to use $\beta_h = 5$. This is the least aggressive strategy available, and is the most limiting in the range of investment choices available to the fund. The only benefit to playing this strategy is that it mitigates the leverage cycle and may prevent crashes from happening. This is circumstantial evidence that defaults are very harmful not only to banks and investors, but also to the funds themselves. It also provides a basis for believing that strategy profiles with low aggression may be better for systemic risk.

Regulation	numFunds	λ_{\max}	5	10	15	20	30	40	50	β_{weighted}
Unregulated	21	20	35%	4%	0%	0%	0%	0%	60%	32.3
Basel	21	20	40%	12%	6%	0%	0%	0%	42%	25.3
Unregulated	21	8	32%	15%	7%	2%	0%	0%	43%	26.3
Basel	21	8	46%	11%	9%	5%	10%	15%	4%	16.7
Unregulated	10	20	46%	5%	0%	0%	0%	0%	49%	27.4
Basel	10	20	51%	14%	8%	0%	0%	0%	26%	18.4
Unregulated	10	8	40%	21%	8%	0%	0%	0%	31%	20.8
Basel	10	8	53%	18%	7%	1%	11%	10%	0%	13.0

Fig. 3. Mixed-strategy symmetric equilibria for different regulatory regimes and leverage restrictions.

Mixed Strategy Nash Equilibria. We found an approximate mixed strategy Nash equilibrium for each of 8 environmental settings, shown in Figure 3. We also report β_{weighted} , which is the average value of β_h drawn from each distribution over strategies. Notice that decreasing λ_{\max} and imposing Basel both decrease the β_{weighted} of the Nash equilibrium. When Basel is imposed, it is no longer as lucrative for funds to be aggressive since they will reach their leverage limits immediately. Decreasing λ_{\max} has a similar effect. This strategic response has ramifications on our systemic risk measures.

Interestingly, all equilibria are far from the social optimum. This is because it takes coordination to prevent defaults and maximize profits. Unilaterally playing a less aggressive strategy has positive externalities, namely a lower default rate and crash rate for *all* agents, that are not internalized by fund payoffs.

Systemic Risk under Equilibria. When Basel is imposed on a previously unregulated market, two things happen in equilibrium. First, the mechanics of fund behavior change as their leverage limits become more stringent in volatile periods. Second, funds adapt their aggression to fit the regulation, shifting to less aggressive (on average) Nash equilibria over mixed strategies as in Figure 3. To evaluate these two effects on financial stability, we first select an mixed strategy equilibrium. Next, we take 200 draws from the distribution. Each draw is a valid strategy profile. We calculate the

metrics from Section 3.3 on each of these 200 strategy profiles both with and without Basel. The averages are presented in Figure 5 and Figure 4.

↑	Profile	Volatility	Setting	
		$\lambda_{\max} = 20$	Basel	Unregulated
		Basel	0.0222	0.0182
		Unregulated	0.0221	0.018

↑	Profile	Volatility	Setting	
		$\lambda_{\max} = 8$	Basel	Unregulated
		Basel	0.0338	0.0205
		Unregulated	0.0341	0.0208

↑	Profile	Default Probability	Setting	
		$\lambda_{\max} = 20$	Basel	Unregulated
		Basel	3.67%	2.99%
		Unregulated	5.02%	3.73%

↓	Profile	Default Probability	Setting	
		$\lambda_{\max} = 8$	Basel	Unregulated
		Basel	1.43%	2.20%
		Unregulated	3.11%	3.55%

↓	Profile	Capital Loss	Setting	
		$\lambda_{\max} = 20$	Basel	Unregulated
		Basel	2.66E+06	3.42E+06
		Unregulated	3.21E+06	3.88E+06

↓	Profile	Capital Loss	Setting	
		$\lambda_{\max} = 8$	Basel	Unregulated
		Basel	5.45E+05	2.13E+06
		Unregulated	9.15E+05	2.76E+06

↓	Profile	Fees	Setting	
		$\lambda_{\max} = 20$	Basel	Unregulated
		Basel	1.80E+05	1.90E+05
		Unregulated	1.78E+05	1.87E+05

↓	Profile	Fees	Setting	
		$\lambda_{\max} = 8$	Basel	Unregulated
		Basel	1.43E+05	2.03E+05
		Unregulated	1.37E+05	1.93E+05

Fig. 4. Three systemic risk metrics evaluated at two levels of λ_{\max} for 21 funds. Analogous table for 10 funds in Figure 5. Within a table, each row fixes an equilibrium mixed strategy and applies a different regulatory setting, each column fixes a regulatory setting and switches to a different mixed strategy equilibrium. Down arrows signify that the metric decreased when Basel was applied in the study by Poledna et al. [2014] while up arrows signify that the metric increased. Blue lettering means changing that property results in a metric change that agrees with Poledna et al. [2014] while red means changing that property results in a metric change that goes against their finding. Black means that there was no significant change in the variable. Example: For $\lambda_{\max} = 20$'s default probability table, up arrow means that Poledna et al. [2014] found Basel increased default probability. Fixing profiles and changing settings agrees with this finding, indicating successful replication, while fixing setting and changing profiles goes against this finding.

Each table shows two-dimensional variation across regulatory environments. Along rows, we keep the mixed strategy equilibrium fixed and change the environment. Along columns, we keep the regulatory setting fixed and change the equilibrium. For every metric and setting we tested, shifting between Basel and unregulated environments under a fixed equilibrium resulted in the same effect on the metric as that reported by Poledna et al. [2014]. This should not be surprising since with a fixed aggression distribution, we are essentially repeating their experiment with slightly different aggression settings. This is confirmation that we have done a successful replication.

Something interesting happens when we examine these tables along columns. This is equivalent to fixing the regulatory environment and switching between mixed strategy equilibria. We have already seen that the average β_h for Basel equilibria is substantially lower. It turns out that less aggressive funds increase financial stability by decreasing the probability of default and bank capital losses. For example, at $\lambda_{\max} = 20$, even though Basel makes default more likely for a fixed aggression distribution, the shift towards the Basel equilibrium counteracts and reverses this effect for both $H = 10$ and $H = 21$.

↑ Profile	Volatility	Setting	
	$\lambda_{\max} = 20$	Basel	Unregulated
	Basel	0.0232	0.0188
	Unregulated	0.0229	0.0188

↑ Profile	Volatility	Setting	
	$\lambda_{\max} = 8$	Basel	Unregulated
	Basel	0.0317	0.0205
	Unregulated	0.0313	0.0219

↑ Profile	Default Probability	Setting	
	$\lambda_{\max} = 20$	Basel	Unregulated
	Basel	2.24%	2.11%
	Unregulated	4.17%	3.65%

↓ Profile	Default Probability	Setting	
	$\lambda_{\max} = 8$	Basel	Unregulated
	Basel	0.52%	1.35%
	Unregulated	1.23%	2.46%

↓ Profile	Capital Loss	Setting	
	$\lambda_{\max} = 20$	Basel	Unregulated
	Basel	1.43E+06	2.24E+06
	Unregulated	1.91E+06	2.79E+06

↓ Profile	Capital Loss	Setting	
	$\lambda_{\max} = 8$	Basel	Unregulated
	Basel	3.09E+05	1.41E+06
	Unregulated	3.52E+05	1.56E+06

--- Profile	Fees	Setting	
	$\lambda_{\max} = 20$	Basel	Unregulated
	Basel	4.25E+05	4.34E+05
	Unregulated	4.28E+05	4.26E+05

↓ Profile	Fees	Setting	
	$\lambda_{\max} = 8$	Basel	Unregulated
	Basel	3.48E+05	4.63E+05
	Unregulated	3.28E+05	4.34E+05

Fig. 5. 10 funds. See Figure 4 for key.

In aggregate (that is, moving from the Basel/Basel cell to the Unregulated/Unregulated one), the story has changed. Basel decreases probability of default in every measured case. Capital losses, which were already decreasing in the original study, decrease even further. Price volatility still increases with Basel, since the shift in equilibrium did not have an effect. Fees are lower in the Basel setting, but by less than was predicted by [Poledna et al. \[2014\]](#). Shifting between equilibria increases fees in most cases to cancel out some of the fee loss found in the original study.

3.5 Discussion

So what explains Basel's performance in decreasing systemic risk? We break this question into two parts. First, how does Basel allow less aggressive funds to thrive compared to aggressive ones? And second, why do less aggressive funds allow for better financial stability?

Consider this model with a single fund manager. Suppose she is currently operating in the unregulated environment with $\lambda_{\max} = 20$ and she hears that Basel will soon be implemented. This can only lower her leverage limit. We posit that even without any strategic interactions with other funds, she will be inclined to lower her aggression levels. A fund gets more fees by generating higher returns, and higher returns are gotten by taking high leverage when there are large mispricing signals. When leverage limits are high, an aggressive fund will just snap up all the mispriced assets at and profit off of all of them. But when leverage limits become more stringent under Basel, too much aggression will quickly lead to paralysis due to hitting the leverage limit. In contrast, less aggression will preserve some investment capacity for the truly large mispricing signals, generating a better return.

Now consider two funds who each know that Basel is coming. If they both decide to be aggressive, they can mitigate more price movements, reduce price volatility, and operate as if Basel was never implemented. However, if fund *A* decides to shirk its price enforcing duties by being less aggressive, Basel will restrict maximum leverage since price volatility will rise. This doesn't affect fund *A* much

since it doesn't take much leverage anyway. In fact, in all of our simulations, the funds playing strategy $\beta_h \in \{5, 10\}$ never defaulted a single time. But the leverage limit does affect the more aggressive fund B , and will cause it to default eventually. When this happens, A will have the entire market to itself. When this strategic advantage is coupled with the fact that less aggressive funds stand to make more profit regardless of what other funds do, it is easy to see why Basel induces funds to be less aggressive.

Now why does the shift to a less aggressive set of funds result in better default rates, bank losses, and fees while not having any effect on price volatility? Note that here we are only giving intuition for the effect of the shift in aggression distribution, not the aggregate effect of Basel.

Less aggressive funds simply do not take much leverage before the mispricing signal disappears randomly. Thus, they never get to the point where a price shock can cause them to lose all their wealth. Again we quote the result that not a single low-aggression fund defaulted in our experiments. Thus, default rates decrease when funds are less aggressive.

Funds that are aggressive under equilibrium that end up defaulting also cause less capital loss to banks. Recall the feedback cycle that occurs with Basel. First, aggressive firms accumulate all the wealth in good times since Basel hasn't kicked in yet. Then, a price shock happens that forces the aggressive funds to wind down their positions to comply with collateral requirements. This procyclical unwinding deepens the initial shock, increasing volatility and forcing Basel to make leverage limits more stringent. When there are many less aggressive funds, none of this volatility can be absorbed and the Basel requirements will get extremely stringent. Thus, the funds that end up defaulting will have been trying to reach a lower leverage limit. This reduces their size when they default.

Fees are slightly helped by less aggression for the same reason funds decided to be less aggressive in the first place. And price volatility is definitely hurt in normal times by less aggressive funds, but the decrease in default rate reflects a lower incidence of financial crashes, which helps price volatility. It seems these two effects cancel out.

4 CONCLUSION

When evaluating the effect of financial regulation on markets, we should always be aware that agents in this arena are highly sophisticated and resourceful. It is a good bet that any possible path to profit will eventually be explored. So when economic models make fixed assumptions on behavior, we should always be on guard for the day these assumptions start break down. To become more robust, models can endogenize important behavioral assumptions so that agent response to the environment is as realistic as possible.

We argue for endogenizing the aggression level of funds in an agent based leverage cycle model used to evaluate Basel regulations. We then propose a strategic analysis that endogenizes an important fixed assumption, the aggression distribution. After making appropriate approximations to make finding equilibria feasible, we perform a series of experiments to measure the aggregate effect of Basel on financial stability under our new equilibrium concept. Our findings suggest that the pessimism surrounding Basel's exacerbation of leverage cycles may be overstated. We find that irrecoverable losses on the part of banks and defaulted funds are reduced substantially under Basel. Although fund profits fall, our new equilibrium concept lessens the fall compared to the one without endogenous aggression. Volatility remains a problem, as funds still cannot prevent day to day fluctuations in price effectively.

There is still much work to be done, both in evaluating Basel and in building richer economic models. For the former, the obvious question is how to improve on Basel so that we get its benefits (less aggressive agents on average) without its costs (agents constrained from taking volatility

out of prices). There has been some work on this, but not under endogenous aggression levels [Aymanns and Farmer, 2015, Poledna et al., 2015, Poledna and Thurner, 2016]. There are also many other directions where computer science techniques can aid in understanding Basel. For example, using network science and graph theory to model financial networks could help us understand how contagion works and how regulation affects it.

The way we solve for equilibria in agent based models can also be studied further. For example, how close does clearing the market in each period independently come to preserving Pareto efficiency in a multi-period setting? How does imperfect information affect regarding noise traders play into this? In addition, in our analysis we fixed aggression to be constant throughout the entire simulation. This was mainly to not disturb the market clearing process within each time step. But perhaps there is a way to combine the market clearing mechanism and the strategic aggression adjustments so that they interact with each other. This would make for a less interpretable model but may be more realistic for certain markets.

ACKNOWLEDGMENTS

We are grateful to Sebastian Poledna and Stefan Thurner for correspondence that assisted us in replicating their results. Thanks to anonymous reviewers for constructive suggestions about the presentation. This work was supported in part by the National Science Foundation and the Office of Financial Research under NSF grant IIS-1440360.

REFERENCES

- Christoph Aymanns and J. Doyne Farmer. 2015. The dynamics of the leverage cycle. *Journal of Economic Dynamics and Control* 50 (2015), 155–179.
- Ben S. Bernanke, Mark Gertler, and Simon Gilchrist. 1999. The financial accelerator in a quantitative business cycle framework. In *Handbook of Macroeconomics*, J. B. Taylor and M. Woodford (Eds.). Vol. 1. Elsevier, 1341–1393.
- Richard Bookstaber, Mark Paddrik, and Brian Tivnan. 2017. An agent-based model for financial vulnerability. *Journal of Economic Interaction and Coordination* (2017).
- Ben-Alexander Cassell and Michael P. Wellman. 2013. EGTAOnline: An experiment manager for simulation-based game studies. In *Multi-Agent Based Simulation XIII. Lecture Notes in Artificial Intelligence*, Vol. 7838. 85–100.
- Financial Crisis Inquiry Commission. 2011. *The financial crisis inquiry report: Final report of the national commission on the causes of the financial and economic crisis in the United States*. PublicAffairs.
- John Geanakoplos. 2003. Liquidity, default, and crashes: Endogenous contracts in general equilibrium. In *Advances in Economics and Econometrics: Theory and Applications, Eighth World Congress*, Vol. 2. Cambridge University Press, 170–205.
- John Geanakoplos. 2010. The leverage cycle. In *NBER Macroeconomics Annual*. Vol. 24. University of Chicago Press, 1–65.
- John Geanakoplos, Robert Axtell, J. Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan M. Palmer, and Chun-Yi Yang. 2012. Getting at systemic risk via an agent-based model of the housing market. *American Economic Review* 102, 3 (2012), 53–58.
- Paul M. Healy and Krishna G. Palepu. 2001. Information asymmetry, corporate disclosure, and the capital markets: A review of the empirical disclosure literature. *Journal of Accounting and Economics* 31, 1 (2001), 405–440.
- Victoria Ivashina and David Scharfstein. 2010. Bank lending during the financial crisis of 2008. *Journal of Financial Economics* 97, 3 (2010), 319–338.
- Sebastian Poledna, José Luis Molina-Borboa, Serafín Martínez-Jaramillo, Marco van der Leij, and Stefan Thurner. 2015. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *Journal of Financial Stability* 20 (2015), 70–81.
- Sebastian Poledna and Stefan Thurner. 2016. Elimination of systemic risk in financial networks by means of a systemic risk transaction tax. *Quantitative Finance* 16, 10 (2016), 1599–1613.
- Sebastian Poledna, Stefan Thurner, J. Doyne Farmer, and John Geanakoplos. 2014. Leverage-induced systemic risk under Basle II and other credit risk policies. *Journal of Banking and Finance* 42 (2014), 199–212.
- Michael R. Roberts and Amir Sufi. 2009. Renegotiation of financial contracts: Evidence from private credit agreements. *Journal of Financial Economics* 93, 2 (2009), 159–184.
- Leslie Shaffer. 2016. Debt is holding back the global economic recovery, say central bankers Dudley, Rajan and Zeti. *CNBC* (Aug. 2016).
- Robert J. Shiller. 2015. *Irrational Exuberance*. Princeton University Press.

- Stefan Thurner, J. Doyne Farmer, and John Geanakoplos. 2012. Leverage causes fat tails and clustered volatility. *Quantitative Finance* 12, 5 (2012), 695–707.
- Michael P. Wellman. 2016. Putting the agent in agent-based modeling. *Autonomous Agents and Multi-Agent Systems* 30 (2016), 1175–1189.
- Bryce Wiedenbeck and Michael P. Wellman. 2012. Scaling simulation-based game analysis through deviation-preserving reduction. In *11th International Conference on Autonomous Agents and Multiagent Systems*. 931–938.