

Financial Market Ethics

David Easley⁺ and Maureen O'Hara^{*}

+Department of Economics and Department of Information Science, Cornell University,

**Johnson College of Business, Cornell University*

June 2019

Revised February 2020

Abstract:

We use a model of psychological-games-played-on-a-network to demonstrate a role for endogenously-determined, rationally chosen ethics. Our analysis produces sharp results about the contagion of non-ethical or ethical behavior and about the possible stable configurations of each type of behavior. We find, and quantify, critical densities for clusters of each type of behavior that determine everything about contagion and stability. We then use these results to show how regulations, market structure and social opprobrium can affect whether clusters of ethical behavior can survive and how large they can be in a financial market setting.

Email addresses: dae3@cornell.edu and mo19@cornell.edu

We thank Franklin Allen, Will Cong, Phil Dybvig, Pam Moulton, Justin Murfin, Sheridan Titman, and Scott Yonker, as well as seminar participants at Cornell University, UMass (Amherst), and the China International Conference in Finance (CICF) 2019, for helpful comments.

Financial Market Ethics

Ethics is conspicuously absent in standard economic models. This failure to show up is not inadvertent; as Nobel laureate Kenneth Arrow explained, “Like many economists, I do not want to rely too heavily on substituting ethics for self-interest. I think it best on the whole that the requirement of ethical behavior be confined to those circumstances where the price system breaks down.”¹ This view of ethics as a last resort, as an overlay needed only when the system is broken, draws on a tradition in economics that markets work best when rational agents all act in their own self-interest. Yet, in actual markets, ethics (or the lack thereof) seem to play a large role not only in effecting the performance of a market, but also its viability. Moreover, the financial crisis and its aftermath suggest that non-ethical behavior can spread, with such behavior in one market inducing similar malfeasance in others. Viewing economic agents simply as autonomous members of “Homo Economicus” seems inadequate to capture these market effects.

In this paper we investigate the role of market ethics. We are interested in two basic issues. First, can rational agents care about ethical concerns and can those agents survive in equilibrium? And, second, can ethical (or non-ethical) behavior propagate in a market? Are there properties of agent networks in which non-ethical behavior spreads from node to node, and if so, how can it be stopped? Conversely, under what conditions can ethical behavior be contagious? To address these issues, we draw on results from psychological games and contagion in networks. We develop a game-theoretic model that allows rational agents to care about others’ expectations of their behavior. Agents in our model are not inherently ethical or

¹ Cited in Michael Sandel, “Market Reasoning as Moral Reasoning: Why Economists Should Re-engage with Political Philosophy,” *Journal of Economic Perspectives* 27, no. 4 (2013): 136. See also the discussion in O’Hara [2016] Chapter 5.

non-ethical but instead exhibit a type of “rational morality”. We characterize agents’ equilibrium strategies in a game in which they chose between ethical and non-ethical behavior, with a focus on when agents switch from playing ethically to playing non-ethically. We then embed this model in a network setting and examine how network density and the interaction of heterogeneous agents result in ethical or non-ethical behavior diffusing across market settings.

To motivate our analysis, consider some recent examples from financial markets. Since 2013, the European Union has fined 7 banks for manipulating euro benchmark derivatives, 6 banks for rigging interest rate derivatives, 3 banks for colluding on the Swiss franc, and early in 2019 accused 8 banks of colluding to game the Eurozone government bond market. In this last case, the contention is that “over time the alleged collusion involved roughly a third of the banks in the primary Eurobond market”.² Is this an example of contagion of unethical behavior within a network? On the other hand, consider the recent case of “manufactured defaults” in the credit default swap (CDS) market.³ Here Blackstone bought CDSs on a firm and simultaneously lent money to that firm with the pre-condition that the firm delay its next bond payment - thereby guaranteeing a hefty payoff for Blackstone when the delay triggered swap payments. While not illegal, this behavior was widely viewed as unethical, and the resultant uproar in the market forced Blackstone to back down on enforcing its claims. Is this an example of ethical behavior prevailing in equilibrium? In this paper we provide a framework for understanding when such disparate outcomes can arise and what can be done to move market equilibria from “bad” outcomes to “good” ones.

² For a discussion of all these cases see “EU accuses eight banks of collusion in sovereign bond market “ Financial Times, February 1, 2019.

³ For more details, see “The mystery trader who roiled Wall Street”, June 4, 2018. Things subsequently have gotten a bit more complicated as detailed in “Wall Street tries to clean up \$8tn market for credit derivatives”, Financial Times, March 14, 2019.

What distinguishes our analysis from more standard models is our use of psychological games. Psychological game theory, which was developed by Geanakoplos, Pearce, and Stacchetti (1989) and extended by Battigalli and Dufwenberg (2007), allows payoffs to depend directly on beliefs rather than just on actions as in traditional game theory. In our application of psychological game theory a player experiences guilt, and thus a reduction in payoff, if he disappoints others. This disappointment aversion or guilt, is modeled as a reduction in payoff based on the player's expectation of others expectation of his behavior. For example, if a diner in a restaurant believes that the waiter expects the diner to leave a tip and the diner doesn't leave a tip then the diner feels guilty. This guilt reduces the diner's payoff and it can make leaving a tip attractive even if there is no external reason to do so.

We apply psychological game theory to the choice of whether to behave ethically or non-ethically. If a player has behaved ethically in the past then his neighbors naturally expect him to behave ethically in their current interaction. The player may receive a direct (monetary) benefit from behaving non-ethically, but the guilt that he experiences from disappointing others may more than offset this monetary gain. The key point in our analysis is that this payoff reduction from behaving non-ethically is endogenous; it only occurs when others expect ethical behavior. A player who is expected to behave non-ethically experiences no guilt from non-ethical behavior as he disappoints no one with this behavior.⁴

Using this framework, we first consider a simple two-player game in which without guilt it is a dominant strategy to play non-ethically. With guilt, two pure strategy equilibria are possible (ethical and non-ethical) and they exhibit history dependence. A player's expected

⁴ Quoting Jack Sparrow from the Pirates of the Caribbean film series, "Me? I'm a dishonest man you can always trust to be dishonest. Honestly. It's the honest ones you want to watch out for, because you can never predict when they're going to do something incredibly....stupid"

behavior going forward depends on how the player acted in the past and this dependence can allow playing ethically to be an equilibrium. What matters for our purposes here is what happens to these equilibria in a setting with multiple agents, which we model as a network. Drawing on the contagion-in-networks results of Morris [2000], we ask how the network structure and the parameters in the game affect the evolution of ethical and non-ethical behavior. In particular, we are interested in analyzing the contagion of strategies across the network. We show that stable configurations of behavior involve clusters of ethical behavior and clusters of non-ethical behavior. We demonstrate that these clusters are stable if they are dense enough (where the critical density depends upon the parameters of the game and can be different for clusters of ethical or clusters of non-ethical behavior).

We then analyze the dynamics of the network and determine conditions for a complete cascade of non-ethical behavior to spread across the network. We show both that clusters of ethical behavior stop contagion of non-ethical behavior and that contagion can only be stopped by these clusters. In settings where past non-ethical behavior makes current non-ethical behavior a dominant strategy, we find that ethical behavior, if it exists, is only found in clusters whereas non-ethical behavior can be scattered or in dense clusters throughout the network. When non-ethical behavior is not so attractive, a co-ordination game emerges where two critical densities determine whether ethical or non-ethical behavior can spread across the network. In this setting, history dependence of the clusters plays a crucial role in influencing contagion. We also demonstrate how parameter values, particularly the guilt parameter, determine how difficult it is to maintain ethical behavior in the network. In an extension of the model, we consider heterogeneous agents, and allow some agents to be “hard-wired” for ethical behavior. In this setting, we find that such agents can play the role of “gatekeepers” and demonstrate how a

gatekeeper node can have a global effect on dampening the spread of non-ethical behavior in the network, even if there are few ethical nodes.

Given these results, what can be done to influence outcomes so that ethical rather than non-ethical behavior prevails? We demonstrate how regulation might be structured to accomplish this, both by changing payoffs in the game and by targeting particular nodes in the network. This role for regulation follows from two intriguing results: the importance of a tipping point in the critical density of a cluster of ethical behavior and the identity of critical nodes. First, a cluster of nodes behaving ethically is stable if its density is high enough; changing that density slightly will have no effect on the long run state of the system. But there is a critical density (which depends on payoffs) such that if the cluster's density falls below the critical level then the long run state can be dramatically different and ethical behavior can disappear entirely. Second, not all nodes in the network play the same role in maintaining ethical behavior. Nodes which provide a bridge between ethical players and non-ethical ones can be critical if they are gatekeepers. These nodes can prevent the ethical nodes from being exposed to contagious non-ethical behavior. But critical nodes also can be ones who have no partners playing non-ethically. If one of these critical nodes exogenously switches to non-ethical play it can cause the cluster of ethical behavior that it may be deeply immersed in to fall apart.

There are other factors that can also influence equilibrium outcomes. We discuss how different market structures may be more or less amenable to the proliferation of non-ethical behavior, and how this might be overcome. We also show how social opprobrium can play a role by enforcing the impact of others' expectations on individual's choices. Overall, our research suggests that market ethics need not be an oxymoron but it can be a fragile outcome; understanding why this fragility emerges is crucial for ensuring that "better" outcomes prevail.

While our focus on ethics in the context of psychological games played in a network is new, there is a variety of work that considers related issues. Glazer, Sacerdote and Scheinkman (1996) developed the idea that culture plays a role in a decision to respect rules or break them, and there is a large body of work looking at the impact of culture and its determinants (see, for example, Cremers (1993); Guiso, Sapienza, and Zingales [2006]; Grullon, Kanatas, and Weston [2008]; Liu [2016]; Lo [2016]; Mokyr [2016]). An important difference between our work and these research areas is that in this research culture and social norms are considered exogenous, while ethics in our framework arises endogenously from beliefs about others' beliefs, allowing us to characterize how such behaviors can propagate in a population. More closely related to our research is work on social norms, which relates behavior to various norms governing participants' action and often derives the stability of these norms from an evolutionary model (see, for example, Elster (1989), Ostrom (2000) and Young (2015) on social norms and their evolution). Our analysis of ethics shares motivation with this literature, but our focus on psychological games played on networks as the origin of ethical behavior leads to different insights into ethical behavior. There is also a literature examining reciprocity in games played by a network of agents (see Dufwenberger and Patel (2017) for a theoretical analysis or Leiderer et al (2009)) for an experimental analysis). This research makes use of psychological game theory, but the focus is different from ours in that we are interested in the evolution of ethical or non-ethical behavior on the network while they primarily focus on how friendship structures as represented by a network affect apparent altruism.

Another large, recent body of work examines contagion in financial networks (see Benoit et al (2017) and Glasserman and Young (2016) for recent surveys). Much of this literature analyzes the role of the network between financial firms in creating systemic risk. Most papers

consider the effect of more or less densely connected, but exogenously-given networks (see for example, Allen and Gale (2000), Eisenberg and Noe (2001), Freixas et al (2000), Allen and Babus (2008), Gai et al (2011), Acemoglu et al (2013), and Elliot et al (2014)), while a few (Blume et al (2013), Erol and Vohra (2014) and Babus (2016)) consider endogenously-determined networks. Our analysis of the spread of non-ethical behavior is also about contagion on a financial network, however the underlying interactions between agents are quite different in our setting. Our agents are playing a psychological game on the exogenous network, and we are interested in contagion of behavior, rather than contagion of defaults leading to systemic risk. Our approach can apply more broadly than just in financial markets; a direction we plan to pursue in future work.

An aspect of our work not addressed in any of these related research areas is our focus on an important policy question: How can market design and regulation facilitate ethical behavior? This question is of obvious importance, and there is a growing body of work empirically investigating financial misconduct. Research here focuses on behaviors such as option back-dating (Lie [2005]; Bernile, Genovova, and Jarrell [2009]), earnings management (Kedia, Koh, and Rajgopal [2015]), SEC violations (Karpoff, Koester, Lee and Martin (2017); Parsons, Sulaeman, and Titman [2018]), and financial advisor fraud (Dimmock, Gorken, and Graham [2015]; Liu [2016]; Egan, Matvos, and Seru [2019]). Our theoretical work complements and provides an explanation for some of the important findings here. For example, Parsons et al (2018) find empirically that financial misconduct occurs in clusters (by cities) and that there are ebbs and flows in the propensity to commit misconduct over time. Our results offer an explanation for why such clusters emerge and how network dynamics determine their movement over time. Similarly, a variety of papers find that financial misconduct appears to be

contagious.⁵ For example, Dimmock et al [2015] find evidence of contagion effects, showing that fraud by financial advisors is more likely if new co-workers have a history of fraud. Liu [2016] finds a similar effect when insiders move from companies with a low corruption index to a high corruption index. Egan, Matvos, and Seru [2019] show that one-third of advisors with misconduct are repeat offenders and that firms that hire such advisors are those with higher rates of misconduct. Such effects are consistent both with the psychological games framework developed here and with the predictions of our model with respect to the propagation of non-ethical behavior in networks.

I. Games, Networks and Contagion

In this section, we set out the psychological game in which players chose to play ethically or non-ethically. We characterize the set of equilibria in the game with a particular focus on how the parameters of the game affect the potential for ethical play. We then embed these psychological-game-playing agents in a network and we determine the equilibria in the network. Our focus here is on two issues: contagion, or more specifically, determining when non-ethical behavior can spread across the network; and stability, what conditions are needed for at least pockets of ethical behavior to persist.

A useful first step is to discuss what we mean by ethical and non-ethical behavior. In our setting, ethical behavior arises when agents care about other's expectations of their behavior. These ethical considerations can result in agents choosing actions that lead to lower monetary rewards so as to avoid disappointing other's expectations. Such considerations do not arise from

⁵ There is an extensive literature in accounting showing how financial reporting misbehavior appears to spread across firms (see, for example, Kedia, et. al. [2015]).

non-ethical play.⁶ The specific manner in which this is implemented is discussed below. Note that we use the terminology non-ethical instead of unethical. It seems likely that many non-ethical behaviors are also unethical, but we use the broader terminology to capture a potentially broader range of behaviors.

I.A. The Game

Our analysis of the diffusion of ethical or non-ethical behavior through a network is built on a basic model of the interaction between agents in two-player games. These agents are embedded in an underlying network in which they each have a choice in each period between two types of behavior: ethical (E) and non-ethical (N). If there is an edge between two agents they play the game with each other and they each receive payoffs from their joint play.⁷ We assume that each agent must choose one strategy to use in all of its interactions.

Although agents will play this game many times, we assume that they do not view their interaction as part of a repeated game. Instead, they play in each period a one-period game. In Appendix B we analyze the infinitely repeated game and show that for sufficiently low, but non-zero discount factors, the play that we describe in the text is consistent with a subgame perfect psychological Nash equilibrium of the infinite horizon game.⁸ In this equilibrium, agents do not build and exploit reputations. Instead, they act as if they were playing the one-period game repeatedly with each game played in isolation.

⁶ More formally, ethical behavior is play in the game such that if a player is expected to choose this play but the player doesn't do so then the player feels guilty and this guilt can overcome the potential individual gain from alternative play. Alternatively, deviating from other's expectations of non-ethical play does not generate guilt.

⁷ We view the nodes in the network as agents who play the game with each other if they are connected by an edge. One can also view the nodes as groups of agents or locations representing groups with an agent selected from the group at each moment to play the game. For example, each node may represent a unit mass of agents and at each moment, an agent selected uniformly at random from each group plays the game and receives the payoff. In this case, even if agents are aware that the game is infinitely repeated they have no incentive to care about reputations or to play other than naively.

⁸ The critical discount factor is the ratio of the payoff of correctly anticipated joint non-ethical play (b) to the payoff to joint ethical play (a).

The game that each pair of linked agents play is described by the following payoff matrix.

		Player 2	
		<i>E</i>	<i>N</i>
Player 1	<i>E</i>	a, a	$0, c - g \hat{\beta}$
	<i>N</i>	$c - g \hat{\alpha}, 0$	$b - g \hat{\alpha}, b - g \hat{\beta}$

If both players chose ethical behavior (E), then they each receive a payoff of $a > 0$. If one player behaves non-ethically (N) and the other behaves ethically, then the ethical player receives a payoff of 0 (a normalization). The non-ethical player has a greater material payoff $c > a$ than he would obtain by responding with ethical play, but he may experience guilt which reduces his payoff. This payoff reduction is the guilt parameter $g > 0$ times the player’s expectation of the other player’s expectation of the probability that they will play ethically.⁹

Let α be player 1’s probability of playing ethically and β be player 2’s probability of playing ethically. Then player 1’s expectation of player 2’s beliefs about 1’s play is $\hat{\alpha} = E_1[E_2[\alpha]]$ where subscripts on the expectation operators indicate to which player’s expectation we are referring. Similarly, 2’s expectation of 1’s beliefs about 2’s play is $\hat{\beta} = E_2[E_1[\beta]]$. The payoff matrix encapsulates the idea that payoffs in the game directly depend on expectations about behavior, hence the term “psychological game”.

⁹ A player who does not experience guilt has $g=0$ and so expectations do not affect his payoffs. Note that if $g=0$ (for both players), $c > a$ and $b > 0$, then non-ethical behavior is a strictly dominant strategy and the only Nash equilibrium in the game is (N, N). In this uninteresting case, every agent in the network plays non-ethically regardless of the network structure. In our setting in which non-ethical play is dominant in terms of material payoffs, ethical play is possible only because of guilt. It’s guilt that induces our rational agents to pass up the opportunity to exploit the other player by playing non-ethically.

If both players chose non-ethical behavior (N) then they receive payoffs of $b - g \hat{\alpha}$ and $b - g \hat{\beta}$, where we assume that $b > g$ so that N is the best response to N for all beliefs. We also assume $a > b$ reflecting the idea that cooperating on ethical behavior leads to greater material payoffs than the non-cooperative outcome in which both players play non-ethically. However, the actual payoff to any player who behaves non-ethically is reduced by the guilt parameter $g > 0$ times the player's expectation of the other player's expectation of the probability that they would play ethically. These parameter assumptions are summarized below.

Base Model Assumption: The parameters of the game satisfy the restriction: $c > a > b > g > 0$.

Note that with these parameters the welfare implications of play in the game are clear. The outcome (E,E) yields the highest overall benefit. However, the pursuit of individual optimization, particularly if the guilt factor is low or non-existent, can lead to the (N, N) outcome in which overall welfare is lower. Our particular interest is in which of these outcomes can occur in equilibrium. In a later section, we consider an alternative parameter structure in which the returns to non-ethical play are less dominant, and examine how this affects equilibrium outcomes.

A **psychological Nash equilibrium** consists of a pair of strategies, one for each player, beliefs for each player about the other player's strategy and beliefs for each player about the other player's beliefs $(\hat{\alpha}, \hat{\beta})$ such that each player best responds to the strategy of the other player and each player's beliefs are correct. We derive all equilibria of the game but we will focus on pure strategy psychological Nash equilibria, hereinafter referred to as Nash equilibria.

Suppose that each player expects the other player to pay ethically and each player's beliefs about the other player's beliefs are correct ($\hat{\alpha} = \hat{\beta} = 1$). For these beliefs to be part of a Nash equilibrium E must be a best response to E given these beliefs. This requires $g \geq c - a$. Note that, if $g < c - a$ then N is strictly dominant and the unique Nash equilibrium is (N,N). Note also that asymmetric play, (E,N) or (N,E), is not possible in a Nash equilibrium as N is the unique best response to N for all beliefs.

However, for $g \geq c - a$ there are multiple Nash equilibria. To see this, first note that if an agent knows that he is expected to play non-ethically, then he experiences no guilt from non-ethical play, which makes non-ethical play a strictly dominant strategy. So, for any value of g, (N, N) is an equilibrium. Calculation shows that in the region in which $g > c - a$ there is also a symmetric mixed strategy equilibrium in which $\alpha = \hat{\alpha} = \beta = \hat{\beta} = \frac{b}{(a - c + b + g)}$. So if $g > c - a$ there are three Nash equilibria: (E,E), (N, N) and the mixed strategy equilibrium above. Figure 1 illustrates the structure of the equilibrium correspondence.

I.B Network

We view financial markets as networks. Some, such as interbank networks and networks of dealers, are explicitly networks. Others, such as exchanges and trading rooms, are less explicitly networks. But they can all be described by a collection of nodes (the banks, dealers or traders) connected by edges representing who-interacts-with-who.

The nodes (players in the game described above) are indexed by $i \in \mathcal{J} = \{1, \dots, I\}$ where $I < \infty$. An edge exists between two nodes, i and j, if nodes i and j interact--- i.e. play the game specified above. Figure 2 depicts a piece of a network indicating that i and j play the game with each other, both play with node k, and both play with other distinct nodes. Each node has some

number of incident edges with this number ranging from 0 if the node has no neighbors to d if the node is connected to every other node in the network. The number of edges that a node has is called its degree, d . In Figure 2 the degree of nodes i and j is 4 while the degree of node k is 2.

The payoff that a player/node receives in the network depends on its play, the play of each of its neighbors and on what those neighbors expect it to do. A node plays the same strategy, E or N, in its interaction with each of its neighbors so its payoff from either strategy can be written as a function of the number of its neighbors playing either strategy and the expectations those nodes have about its play. In an equilibrium, expectations about the play of others must be correct, but during the adjustment process we suppose that each player assumes that a node will play now as it did in the previous period and that each player knows this fact.¹⁰ Suppose that fraction p of a node's d neighbors play E. Then the node's payoff to ethical play is pd and its payoff to non-ethical play is $pd + (1-p)b - dg\hat{\alpha}$. So the node will choose ethical play if

$$pd \geq pd + (1-p)b - dg\hat{\alpha} \quad (1)$$

The ranking of these payoffs does not depend on the number of neighbors, d .¹¹ So the inequality above reduces to

$$p \geq p + (1-p)b - g\hat{\alpha} \quad (2)$$

Thus, there is a critical value of the fraction of neighbors playing ethically such that the node will play ethically if and only if at least this fraction of its neighbors play ethically. This critical value is¹²

¹⁰ Of course, we are also implicitly assuming that every player knows how each of its neighbors played last period.

¹¹ Because d and the number of agents in the network do not affect this ranking of payoffs they will play no role in the evolution of behavior or in the possible stable configurations of behavior.

¹² Note that p^* need not be between 0 and 1. If $p^* < 0$ the node plays ethically for all p and if $p^* > 1$ then the node plays non-ethically for all p .

$$p^* = \frac{b - g\hat{\alpha}}{a + b - c} \quad (3)$$

Our model induces a threshold rule, as in Morris [2000]. A node's play is determined by the fraction of its neighbors playing ethically; the number of neighbors does not matter and payoffs affect play only through their effect on p^* . However, ours is not a simple threshold as it depends on expectations of play and we have seen that multiple correct expectations are possible in the underlying game. In the standard use of the threshold rule (see Morris [2000] for the development of the theory or Easley and Kleinberg [2010] for an exposition of it), nodes form naive expectations of others' play in that they assume that the fraction of their neighbors who currently play a strategy will be the fraction who played it last period. This expectation is correct in a steady state, but may be wrong while play is evolving in the network. We use the same approach to determine the expectations parameter that enters guilt, i.e. each node expects all other nodes to play as they did last period.¹³

If a node played non-ethically last period, then $\hat{\alpha} = 0$, making $p^* > 1$ and the node plays non-ethically this period regardless of the play of its neighbors. If the node played ethically last period, then $\hat{\alpha} = 1$ and the node will play ethically this period if and only if enough of its neighbors are expected to play ethically this period, i.e. if $p \geq p^{**} = \frac{b - g}{a + b - c}$, where p^{**} is p^* evaluated at $\hat{\alpha} = 1$. So, for example, in Figure 2 if nodes i and k played ethically last period, and node j played non-ethically, then node k will play ethically this period if and only if $p^{**} \leq 1/2$.

¹³ If play converges then in the limit expectations of play, the determinants of guilt, will be correct. One could instead require expectations to be correct at all moments, but then there would be no dynamics and analysis of the stability of various configurations of play in response to shocks to behavior.

I.C. Dynamics

The analytical question we begin with is: starting from a network in which everyone is ethical or everyone is non-ethical, can a small number of nodes exogenously flipping to the other type of behavior cause that new behavior to spread through the network possibly flipping everyone to the new behavior?¹⁴ Obviously, this cannot happen if every node begins with non-ethical play, as given past non-ethical play, playing non-ethically is dominant. So, a network of non-ethical play is stable. For ethical play to be possible, we need guilt to be strong enough, i.e. $g \geq c - a$, and that is the case we focus on. In this case, if every node initially plays ethically, and a small number of nodes flip to non-ethical play, the eventual outcome depends on the parameters of the game and the structure of the network. Most importantly, whether these flipped nodes cause other nodes to flip depends on what fraction of their neighbors are playing non-ethically.

The dynamic process induced by these flipped nodes evolves as follows:

1. Initially each node is labeled with the same play: all E or all N.
2. Next the labels for some set of nodes are exogenously switched from E to N if all were initially labeled E, or from N to E if all were initially labeled N.
3. Then each node expecting its neighbors to play as labeled and knowing that it's neighbors expect it to play as labeled selects a best response.
4. The nodes are then labeled with these best responses and step (3) is repeated.
5. The process is declared to have stopped if in two successive labelings no labels are modified.

¹⁴ In actual markets, this flip or mutation could occur if ethical (non-ethical) employees are replaced by ones who turn out to be the other type (for evidence of this see Liu (2016)).

To state the results for the outcome of this dynamic process we first need a few definitions. A **cluster of density p** is a set of nodes such that each node in the set has at least fraction p of its neighbors in the set. A dense cluster is a set of nodes for which most of their interaction with other nodes occurs within the set rather than with nodes outside of the set. Of course, the entire network is a cluster of density one.

Consider a network in which every node is initially playing E. We say that a set of nodes, S , changing their behavior to N causes a **complete cascade** if when the nodes in S change to N then eventually every node in the networks plays N. Morris [2000] shows that, for the structure he analyzes, whether complete cascades occur or not is determined entirely by whether sufficiently dense clusters exist or not. His result, modified slightly, applies to our setting.

Theorem 1: Suppose the parameters of the game satisfy the basic model assumption. Consider a network in which every node initially plays E. Suppose that a set S of nodes switches to behavior N.

- (i) If the remaining network, consisting of nodes in $\mathcal{J} - S$ and the edges between them, contains a cluster of density at least p^{**} then a complete cascade does not occur.
- (ii) If a complete cascade does not occur, then the remaining network must contain a cluster of density at least p^{**} .

All proofs are provided in Appendix A. This theorem is a modification of the argument in Easley and Kleinberg [2010] to our setting. Essentially this result says that clusters stop contagion and contagion can only be stopped by clusters.

Figure 3 illustrates the evolution of non-ethical behavior. Suppose that $p^{**} = 0.6$ and that initially every node in the network plays ethically.¹⁵ Now suppose that in period 1 as depicted in the upper-left graph one node switches to non-ethical behavior---which we represent by coloring this node red. Then in period 2 the upper and lower nodes connected to this first mover have one-half of their neighbors playing non-ethically, and thus less than fraction p^{**} playing ethically, so they switch to non-ethical play in period 2, and are colored red in the upper-right graph. The node which has an edge into the cluster on the left has only one of its five neighbors who are expected to play non-ethically so it plays ethically in period 2. However, in period 3 this node now has three of its five neighbors playing non-ethically so it too switches to non-ethical play, and is colored red in the lower-left graph. Then in period 4 the remaining node not in the cluster on the left switches to non-ethical play, and is colored red. However, the node in the cluster on the left with an edge to a non-ethical node never switches to non-ethical play as only one of its three neighbors plays non-ethically. So none of the nodes in the cluster switch to non-ethical behavior and the spread of non-ethical behavior stops.

Figure 4 illustrates how clusters stop cascades. Non-ethical behavior spreading through the network in Figure 3 from the right stops when it reaches the clusters of Es on the left if $p^{**} < 2/3$. In this case the node in the cluster with a neighbor playing non-ethically has more than p^{**} of its neighbors playing ethically and so it won't switch to non-ethical behavior. Because it won't switch, non-ethical behavior cannot invade the cluster and the cascade stops. Alternatively, if $p^{**} > 2/3$ then the node with a non-ethical neighbor will flip to non-ethical behavior. In the next period, its two ethically playing neighbors will flip to non-ethical play as

¹⁵ For this particular network, the critical p^{**} can be anywhere in the open interval $(1/2, 2/3)$. We focus here on a specific point for clarity of exposition.

they have less than p^{**} of their neighbors playing ethically. Then in the following period, the two other nodes flip to non-ethical behavior and the cascade is complete.

The intuition for part (i) of Theorem 1 is straightforward. Suppose that the contagion of N could invade a cluster of nodes all playing E of density at least p^{**} . Then there must be some first node in this cluster who flips to N, but at the time of the flip at least fraction p^{**} of its neighbors are playing E. Therefore, this node will not flip and thus no node in the cluster flips. The intuition for part (ii) is similar. If a complete cascade does not occur then some nodes remain playing E. Unless those nodes are in a cluster of nodes, all playing E, of density at least p^{**} , then at least one of them wants to switch to N. But then the cascade has not stopped.

The second analytical question we are interested in is: does the process stop and if it does stop what limit structure of ethical and non-ethical play is possible. We imagine a network in which exogenous shocks are rare and first ask if the dynamic process stops. In our framework any node that exogenously switches to N continues to play N and any node that is induced to play N because of the play of its neighbors won't switch back to E. There are a finite number of nodes and as any node that plays N cannot change back to E, the process stops in finite time either with all nodes playing N or some playing N and the remaining nodes playing E. This guarantees that following an exogenous shock with some nodes switching their behavior the process of behavioral adjustment converges in finite time.

We say that a configuration of play, $\{a_i\}_{i \in \mathcal{J}}$ where each $a_i \in \{E, N\}$, is **stable** if: (i) each player is playing a best response in the network game, and, (ii) all player's beliefs (both first and

second order) about the play of their neighbors are correct.¹⁶ Clearly, every node playing N is a stable configuration of play. For a configuration of play in which some nodes play E to be stable each of those nodes must have fraction at least p^{**} of their neighbors playing E as otherwise they would switch to playing N. So the subset of nodes playing E must be a cluster with density at least p^{**} as otherwise the dynamic cannot have stopped. No other configuration of play is stable as in any other case some node wants to switch its play.

Corollary 1: The dynamic process stops in finite time and once it is stopped the configuration of play is stable.

Corollary 2: The stable configurations of play are:

1. All nodes play N.
2. Any configuration of clusters of nodes each with density of at least p^{**} playing E and all other nodes playing N.

Note that if $g < c - a$ then $p^{**} > 1$, so no cluster can have density p^{**} and the only stable configuration of nodes is all playing N. If $g > c - a$ then there may be multiple stable configurations of play in which some nodes play E, but in each of those stable configurations each node playing E is in a cluster of density at least p^{**} .

Our explanation for clusters of ethical behavior is that no other configuration of ethical behavior is stable. Scattered nodes playing ethically face too much payoff pressure to persist, so if ethical behavior is observed at all it will be in clusters. This provides an explanation for the geographic cluster of ethical and non-ethical behavior found in Parsons, Sulaeman and Titman

¹⁶ The play in a stable configuration and associated beliefs also constitute a Nash equilibrium for the game played by the nodes with payoffs determined by equation (1) and the constraint that each player uses the same strategy in every interaction.

[2018]. They find that financial misconduct occurs in geographical clusters and these clusters cannot be explained by geographic variation in firm characteristics or enforcement. Their explanation is geographic variation in social norms. If edges exist between all firms in the same geographic location and there are fewer edges connecting firms across locations, then our network can be viewed as a geographical network in which clusters arise in cities. With this interpretation of the network, our explanation for clusters of behavior provides an endogenously determined explanation for the geographic variation these authors find in social norms.

The network we analyze is one in which there is a tendency toward non-ethical play in the sense that nodes play E only if they are in a sufficiently dense cluster of Es, but if a node flips from E to N then it will never flip back. So ethical behavior is only observed in dense clusters rather being scattered throughout the network. However, scattered nodes can play N. Those nodes playing N may infect others unless the contagion is stopped by a sufficiently dense cluster of nodes all playing E. How dense a cluster of ethical behavior needs to be in order to be resistant to the spread of non-ethical behavior depends on the payoffs in the underlying game. If guilt becomes more important (g increases) or if the material benefit from deviating to non-ethical behavior ($c-a$) becomes smaller, then p^{**} decreases and ethical behavior is easier to maintain. The implications of these results for the types of public policy that can be effective in maintaining ethical behavior are discussed in Section II.

I.D. What if N is not a dominant strategy?

In the game described in section I.A, with $c > a$, if players do not experience guilt, then non-ethical behavior is a dominant strategy. This makes non-ethical behavior a stable strategy in the network regardless of the network structure. We now consider the evolution of play when, perhaps because of regulation or legal or social standards, non-ethical behavior is not so

powerful. In this section, we analyze the network with the assumption $a > c$ while maintaining the other parameter assumptions as listed below.

Alternative Model Assumption: The parameters of the game satisfy the restriction:

$$a > c > b > g > 0 .$$

Given this assumption on parameters, there are two pure strategy Nash equilibria in the underlying game: (E, E) and (N, N). So we now have a coordination game in which both players playing ethically yields greater payoffs than if they both play non-ethically. The game is complicated, however, by the presence of endogenous guilt. There are now two critical values for cluster density.

First, consider a node that previously played non-ethically. For this node $\hat{\alpha} = 0$, so the node will play ethically this period if and only if

$$p > \frac{b}{a+b-c} = p_N^{**} .$$

For a node that previously played ethically, we have $\hat{\alpha} = 1$. Therefore, this node will play ethically this period if and only if

$$p > \frac{b-g}{a+b-c} = p_E^{**} .$$

Note that, if $g > 0$, then $1 > p_N^{**} > p_E^{**} > 0$. These results are summarized in Figure 4. If

$p \in [p_E^{**}, p_N^{**}]$, then whether a node plays ethically or non-ethically depends on how it played previously.

Clusters are again critical for understanding the spread of behavior in the network. However, whether nodes begin with play of E or N now matters as there are two critical densities, each of which is dependent on past play. Consider a network in which all nodes

initially play ethically. Any cluster with density p_E^{**} of nodes all playing E is sufficient to stop the spread of N if initial play of N begins outside of this cluster. To see this, consider the first node in the cluster which has a neighbor outside of the cluster who plays N at some date. This node in the cluster played E at that date so it will play E again as long as at least fraction p_E^{**} of its neighbors played E at that date. However, as it is in a cluster of density at least p_E^{**} it does have enough neighbors playing E. Therefore, the node will continue to play ethically and no node in this cluster will flip to play of N.

Alternatively, consider a network in which all nodes initially play N. Any cluster with density $1 - p_N^{**}$ of nodes all playing N is sufficient to stop the spread of E if initial play of E begins outside of this cluster. To see this consider the first node in the cluster which has a neighbor outside of the cluster who plays E at some date. This node in the cluster played N at that date so it will switch to E only if more than fraction p_N^{**} of its neighbors played E at that date. However, as it is in a cluster of density at least $1 - p_N^{**}$, it cannot have enough neighbors playing E. Therefore, the node will continue to play non-ethically and no node in this cluster will flip to play of E.¹⁷

Theorem 2: Consider a network in which every node initially plays E (or all play N). Suppose that a set S of nodes switches to behavior N (or to E) and that the process of nodes adjusting their behavior converges.

¹⁷ It is also possible for nodes to continually flip back-and-forth between E and N and thus for the process not to converge. To see this, consider a network of two agents initially playing E. Suppose that one of them exogenously flips to N. Then in the next period, the E-agent will change to N and the N-agent will change to E. This repeats every period. It's caused by agents responding naively to past play just as in fictitious-play dynamics which need not converge to a Nash equilibrium in simple two-player games. We are interested in stable configurations of behavior so we focus on what happens if the process converges.

- (i) If the remaining network, consisting of nodes in $\mathcal{J} - S$ and the edges between them, contains a cluster of density at least p_E^{**} (or $1 - p_N^{**}$) a complete cascade does not occur.
- (ii) If a complete cascade does not occur then the remaining network must contain a cluster of density at least p_E^{**} (or $1 - p_N^{**}$).

Thus, both clusters of E and clusters of N can be stable and the density necessary to make each of them stable can be small as both p_E^{**} and $1 - p_N^{**}$ can be small. For a cluster of density $d \in [p_E^{**}, p_N^{**}]$ there is history dependence in the sense that the behavior of this cluster in response to the spread of a behavior depends on the past behavior of the cluster. If the nodes in it previously played E they won't flip to N in response to the spread of N beginning outside of the cluster, but if they previously played N they won't flip to E in response to the spread of E beginning outside of the cluster.

A related form of history dependence in an overlapping-generations economy is demonstrated in Acemoglu and Jackson [2015]. In their model, overlapping-generations create a network in which each agent is connected to the one who played before and the one who plays after. These agents play a cooperation game with two equilibria: one desirable and one not. History dependence occurs because each agent's payoff directly depends on the play of the agent who moves before and the one who moves after. As in our analysis, history dependence also arises through its effect on expectations of play. The main differences, other than interpretation, are that we consider a general network and generate history dependence through a psychological game in which guilt depends on expectations about play, which in turn depend on past play.

The difference between p_E^{**} and p_N^{**} is determined primarily by the guilt parameter g . If $g=0$, then we have a standard coordination game with a single critical value for cluster density. As g grows, p_E^{**} declines, eventually approaching zero. For large g , it is difficult for the play of non-ethical behavior to spread, while for small g it spreads more easily.

I.E Heterogeneous Agents

All of our agents play ethically if and only if it is in their best interest to do so as reflected in their payoffs. In our base game, without guilt, non-ethical play is a dominant strategy, so we see ethical play only if guilt is strong enough and ethical players are in sufficiently dense clusters. It is also interesting to ask how play evolves if some agents are hardwired to play ethically. These agents play ethically regardless of the play of others. One could model this behavior by attributing different payoffs to these agents, but we find the story more compelling if they are thought of as agents who simply will not play non-ethically. In this subsection, we demonstrate how they influence the play of other agents in the network.

Clusters remain critical for understanding the evolution of play in the network. However, in computing the density of a cluster of nodes all playing E the fraction of ethical playing neighbors for nodes that are hardwired to play E plays no role. Suppose that all nodes in $\mathcal{E} \subset \mathcal{J}$ are hardwired to play ethically. Then the relevant density d_E^C of a cluster, C , of nodes all playing ethically is the minimum over nodes $i \in C/\mathcal{E}$ of the fraction of node i 's neighbors in C . So densities of clusters of nodes playing ethically weakly increase due to the presence of ethical nodes. This can have a significant effect on the spread of non-ethical behavior even if there are very few ethical nodes. To see this, consider again the network in Figure 4 and suppose that the node in the cluster on the left who has a neighbor in the non-ethical cluster is hardwired to play ethically. Then the relevant density for this cluster is one. Non-ethical behavior cannot invade.

Corollary 3: Consider a network in which nodes in $\mathcal{E} \subset \mathcal{J}$ are hardwired to play ethically.

Suppose that every node initially plays E (or all nodes $i \in \mathcal{C}/\mathcal{E}$ play N). Suppose that a set

$S \subset \{1, \dots, I\} / E$ of nodes switches to behavior N (or to E).

- (iii) If the remaining network, consisting of nodes in $\mathcal{J} - S$ and the edges between them, contains a cluster C of density d_E^C at least p_E^{**} (or a cluster of density $1 - p_N^{**}$) then a complete cascade does not occur.
- (iv) If a complete cascade does not occur then the remaining network must contain a cluster C of density d_E^C at least p_E^{**} (or a cluster of density $1 - p_N^{**}$).

The implication of clusters of ethical behavior is most striking if there are gatekeepers who are hardwired to play ethically. A node i in a network is said to be a **gatekeeper** if there are two other nodes j and k in the network such every path between j and k passes through node i . An ethical gatekeeper can have a global effect on the spread on non-ethical behavior. Suppose that node i is an ethical gatekeeper for two nodes j and k . Non-ethical behavior starting at a single node j cannot spread to node k , as non-ethical behavior spreads along paths and any path from j to k has to pass through i . At worst, j 's non-ethical behavior can spread to all nodes l such that there is a path between j and l that does not pass through node i . It cannot spread to any node l such that there is a path between l and k that does not pass through node i , as if it did then i could not be a gatekeeper for the j, k pair.

Consider a network in which all nodes initially play ethically. Suppose that node i is a gatekeeper for nodes j and k and that node i always plays ethically. Suppose that node j exogenously switches to non-ethical behavior. This may cause other nodes eventually to play non-ethically. However, it cannot cause any node with a path to node k that does not pass

through node i to switch. Consider the graph with node i deleted and let $C_i(k)$ be the set of nodes in the graph with node i deleted such that there is a path between the node in the set and node k . For non-ethical behavior to enter this set there would have to be a path between some node in the set and node j , but then there would be a path between node j and node k that does not pass through node i . So non-ethical behavior cannot spread from node j into the set of nodes $C_i(k)$.

II. Building and sustaining ethical markets

Our model demonstrates that ethical outcomes are, at best, fragile. This fragility arises because an individual agent's behavior can induce changes in other agents' behavior, leading to a contagion of non-ethical behavior. In this section we consider how to influence the outcomes away from non-ethical equilibria and towards ethical ones. Our model suggests that a wide range of factors such as social opprobrium, regulation, and network structure can all play a role in affecting market ethics. It's useful to note that, although we analyze agents playing two-player games with their neighbors in a network, the number of agents and the number of neighbors any agent interacts with play no role in determining the evolution of behavior or the stable configurations of behavior. If we view an agent's size relative to the entire market as depending on the agent's degree and the number of agents in the network then the size of agents is not one of the factors that plays a role in affecting market ethics. Instead, what matters from this network point of view are the payoffs in the game and the structure of the network.

II.A. Payoff structures and the guilt factor

As noted in the introduction, agents in our model (at least those not hardwired to play ethically) exhibit a type of "rational morality" in that they play ethically or non-ethically depending upon their payoff expectations and network structure. That payoffs play a critical role

is not unexpected. In the initial setting considered in Section I, if agents are expected to play non-ethically then the payoff structure makes non-ethical behavior a dominant strategy and so cascades of non-ethical behavior spread easily through the network. But even when non-ethical behavior is not a dominant strategy, the best we can hope for are parameter and network settings in which multiple equilibria can prevail.

How might regulators influence the market towards outcomes with more players adopting ethical behavior? As our model shows, which outcome prevails is determined by the critical density p^{**} and reducing p^{**} makes ethical behavior easier to maintain. It is easy to demonstrate that if $g > c - a$ then

$$\frac{\partial p^{**}}{\partial c} > 0 \text{ and } \frac{\partial p^{**}}{\partial b} > 0 .$$

So reducing the payoffs to playing non-ethically (the b and c) will result in larger stable clusters of ethical behavior. Regulation can do so directly by imposing costs to playing non-ethically (via increased supervision, and fines, jail time, and the like for more egregious cases) but it can only be successful if it is done on an *ex ante* not *ex post* basis. Punishing people after the fact is standard in financial markets, but it need not deter future bad behavior or move the market to an ethical outcome. Similarly, sporadic enforcement with a large fine in the unlikely event that you get caught may do little to change expectations of how your neighbors will play in the future. Deferred prosecution agreements, probationary periods, or even announcements of target areas of interest to regulators are possible strategies to influence these *ex ante* payoffs.

In a psychological games setting, the guilt parameter also plays an important role. It is easy to show that $\frac{\partial p^{**}}{\partial g} < 0$. So, an increase in g , the guilt parameter, will result in larger

stable clusters of ethical behavior. Thus, inducing greater guilt (or aversion to disappointing others) would result in an ethical equilibrium being more likely to prevail.

How might this be accomplished? Here social opprobrium, rather than regulation, may prove more effective. Such social opprobrium can be encouraged more formally, as captured by the social credit system in Nanjing, China which categorizes some residents as “discredited persons” if they commit too many unsocial actions.¹⁸ Another example is the quaint practice our local library had of posting prominently a list patrons with overdue library fines. But it may be even more effective when it arises more informally. As the college admissions scandal demonstrates, paying a proctor \$15,000 dollars to correct an offspring’s SAT test is not the crime of the century, but the resulting social outrage has wreaked havoc on the social and professional lives of those involved. Similarly, the Occupy Wall Street movement cast a pall over the financial services industry, inducing at least some to shy away from careers as “banksters” and spawning a variety of “culture” initiatives for extant bankers.¹⁹ Because expectations of what you expect others to do affects the equilibrium, shifting beliefs to expect others to behave ethically can have a substantial effect on the incidence of ethical behavior in the network.

II.B. Utilizing the network structure to target enforcement

Network structure, and particularly the role of clusters, is fundamental to understanding the spread and incidence of non-ethical behavior. While the decisions of individual agents’ matter, how much they matter differs depending upon where in the network a node resides.

¹⁸ For a discussion of the social credit system see “Society: Nanjing targets jaywalking in credit list,” China Daily, July 9, 2019. Whether this social credit system simply enforces greater civic behavior or veers into trying to push more politically correct conformity is unclear. For a discussion of these negative concerns see “Chinese Social Credit System Prevents 2.5 Million “Discredited Entities from Buying Plane Tickets”, www.infowars.com, August 19, 2019.

¹⁹ This negative social view of financial services persists today as the 2019 Edelman Global Trust Index shows yet again that financial services placed last among all industries in terms of trust.

Much like the “influencers” in social media, some nodes have an out-sized importance in affecting the spread or curtailment of non-ethical behavior.²⁰ Viewed from this perspective, enforcement actions can be much more effective if they are modeled on the tactics used to stop the spread of a disease or false information rather than as simply punishing miscreants.

Drawing on research on disease dynamics, stopping or preventing a cascade of non-ethical behavior in a market requires supporting clusters of ethical behavior. A cluster of ethical nodes can disappear either due to an invasion of non-ethical behavior from outside the cluster or because some of the nodes inside the cluster exogenously switch to non-ethical behavior. To keep this from happening, regulators need to maintain “critical” nodes as ethical, but determining this criticality can be complex. For example, critical nodes need not be the ones connected to nodes playing non-ethically (i.e. ones on the boundary of a cluster of Es), but can be nodes deep inside of the cluster of Es. To see why, consider a cluster C of nodes playing ethically with density greater than p^{**} . Non-ethical behavior cannot invade this cluster from outside of it. However, switch some node $n \in C$ to non-ethical behavior. Now consider the cluster $C / \{n\}$. If this cluster has density less than p^{**} , then non-ethical behavior can invade.

Hence, what matters for criticality is the effect of a node’s switching on the density of the cluster of Es to which they belong. If switching to N does not reduce density below p^{**} , then only the switched node changes its behavior. If it does reduce this density below p^{**} , then this single switch can trigger a cascade. To see how this can occur consider Figure 6. Suppose $1/2 < p^{**} < 2/3$. The figure shows that the cluster C on the left cannot be invaded by non-ethical behavior originating outside of C . Even if node i which is connected to the outside of C

²⁰ This point that some nodes are better than others at causing cascades has been extensively explored in the networks science literature. See, for example, Kempe, Kleinberg and Tardos (2003, 2005),

switches to behavior N, then the remaining cluster still has density $2/3$. Therefore, there is no spread of N to the rest of the cluster; it's still a sufficiently dense cluster of nodes playing ethically. If, instead, node k switches to behavior N, then the density of the remaining cluster decreases to $1/3$. This is less than p^{**} so N invades. In fact, eventually all nodes switch to N. This example shows that even though node i is a gatekeeper, and can play an important role in preventing the spread of non-ethical behavior from outside of C into the cluster, it is not the most important node to protect when considering mutation to non-ethical behavior from within cluster C.

Understanding these dynamics may shed light on some of the rampant misbehavior in financial markets noted earlier. In particular, in charging eight banks of colluding to rig the 7 trillion Euro government bond market, EU anti-trust authorities alleged that “the price rigging was not bank-wide but involved specific traders at various lenders who in some cases moved between institutions during the period.”²¹ Over time, the alleged collusion spread to a third of the banks operating in that market. The contagion of non-ethical behavior, abetted by the movement of non-ethical players to previously ethical banks, is what our model would predict.

Certainly, targeting these critical nodes for early enforcement actions would have been helpful, but this may not be easily accomplished. To see why, consider a group of nodes all playing N. Regulators could attempt to switch a node from playing N to E or they could simply remove the N node from the market. Switching some of them (or just one of them) to playing E will not generate a cascade of Es in our base model – i.e. it will not affect the other N's in the cluster. In our more complex model, it also will not generate a cascade of Es if the density of the remaining group of Ns is greater than $1 - p_N^{**}$. Alternatively, removing the node is not a better

²¹ Cited in “EU Accuses eight banks of collusion in sovereign bond market”, Financial Times, February 1, 2019.

solution as it will have even less of an effect on the remaining density. Unless the enforcement action drops the density of the remaining group playing N to less than $1 - p_N^{**}$, there is no multiplier effect of N's switching to E's, and so no general removal of the problem. For regulators, the implication is clear: either make a significant intervention or don't bother.

II.C. Market Structure Effects

Finally, we turn to the influence of market structure on market ethics. In our psychological games setting, ethical behavior can be an equilibrium outcome if people care about the expectations of others regarding their behavior. But in anonymous market settings, where knowing your specific counterparty is precluded, such expectations may be hard to form, or more likely undefined. This suggests that anonymity in markets can foster more non-ethical outcomes and a greater propensity to spread across the network.

To investigate this idea in the context of our model, we view an anonymous market as equivalent to $g=0$ in the game. In our base model, the only possible outcome is for everyone to play N (regardless of the network). In the alternative model in which N is not a dominant strategy, setting $g=0$ results in $p_E^{**} = p_N^{**} = \frac{b}{a+b-c}$. This higher cutoff makes sustaining ethical behavior more difficult. Clusters of ethical behavior have to be denser than otherwise to survive, and if there aren't any clusters that are sufficiently dense, then eventually everyone switches to N.

Given this, we might expect more ethical issues to arise in exchange settings where anonymity prevails. Conversely, direct markets, such as dealer-to-customer markets for bonds or swaps, may support ethical outcomes because it is easier to form expectations about the behavior

of counter-parties in this setting.²² The Blackstone CDS episode noted earlier is a case in point. Blackstone's actions in precipitating a manufactured default were not what market participants expected, and because of the non-anonymous nature of the market were visible to all. The resulting uproar in the market, along with the CFTC's noting that such antics might be viewed as manipulation, induced Blackstone to forego profiting from its trade. Subsequent actions by ISDA (The International Swap and Derivative Association), the industry group setting standards in this area, have sought to clarify what behavior is expected from market participants so as to avoid a recurrence of such problems in the future.²³ In at least this setting, expecting others to behave ethically can have a positive effect on the market.

III. Conclusion

We demonstrate a role for endogenously determined, rationally chosen ethics in a financial network setting. Absent the guilt that non-ethical behavior induces, the agents we consider would rationally choose to behave non-ethically and have lower payoffs than if they behaved ethically. However, we show that guilt combined with a sufficiently dense set of neighbors who play ethically makes socially desirable ethical behavior both optimal and stable. This sufficient density, the critical level of density, determines everything: clusters of ethical behavior with at least this density stop the spread of non-ethical behavior and only these sufficiently dense clusters can stop it.

Ethical behavior is good for society in the sense that it results in higher physical (monetary) payoffs than non-ethical behavior, so it's natural to ask how society can make ethical behavior more stable. Our psychological-games-played-on-a-network model allows us to provide

²² We note, however, that transparency in markets can be a two-edged sword in that it can limit contagion but also foster collusion.

²³ See "Wall Street tries to clean up \$8tn market for credit derivatives," Financial Times, March 14, 2019.

answers to this question. We find that reducing the payoffs to non-ethical behavior or increasing the guilt that players experience from unanticipated non-ethical play matters if and only if it moves cluster densities across the critical level; otherwise policies that change these payoffs have no effect on the spread of non-ethical behavior. We also find that ethical behavior is more likely to exist if society identifies critical nodes, ones whose behavior is critical for the density of ethical clusters or who play the role of gatekeepers, and maintains them as ethical players. These results are (intentionally) reminiscent of the advice that network analysis generates for the control of contagious disease. The details are different (payoffs versus probabilities of infection), but the lesson is the same: interventions need to be big in order to be effective, small interventions are in the long run a waste of effort, and it's important to identify and protect nodes that are critical for the transmission of whatever harmful attribute is spreading through the network.

Our analysis, of course, has several limitations, with perhaps the most important one being the fixed, exogenously given network of homogenous nodes. If some nodes are more likely to play ethically than others, and the network results from nodes' choices of whom to interact with, then clusters of ethical behavior may arise also because of the endogeneity of the network. More ominous is that non-ethical agents may attract other non-ethical agents to join the network, turning hitherto ethical markets into unethical venues. We hope to explore these implications of endogenous network formation on market ethics in future research.

FIGURES:

$$\alpha = \Pr(E)$$

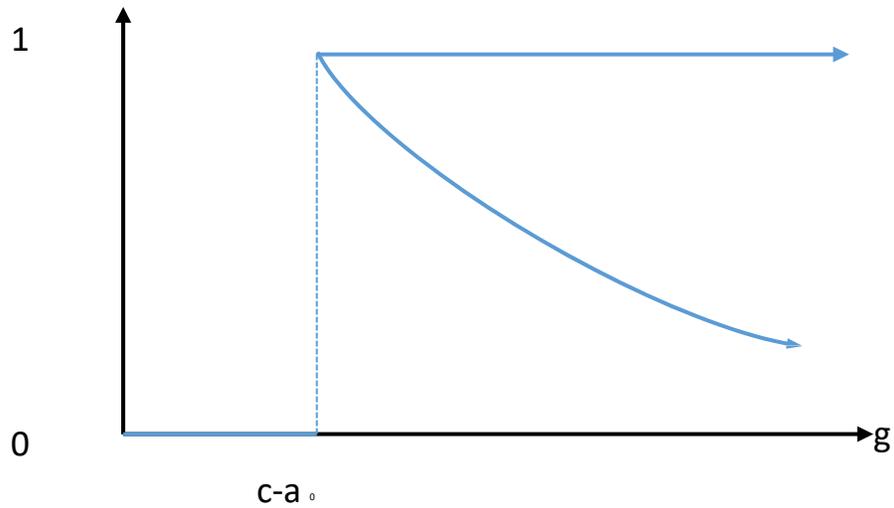


Figure 1: Graph (in blue) of symmetric Nash equilibria as a function of the guilt parameter g . There is one equilibrium (N, N) for g less than $c-a$, and there are three equilibria for g greater than $c-a$: (N, N), (E, E) and the mixed strategy equilibrium illustrated in the graph.

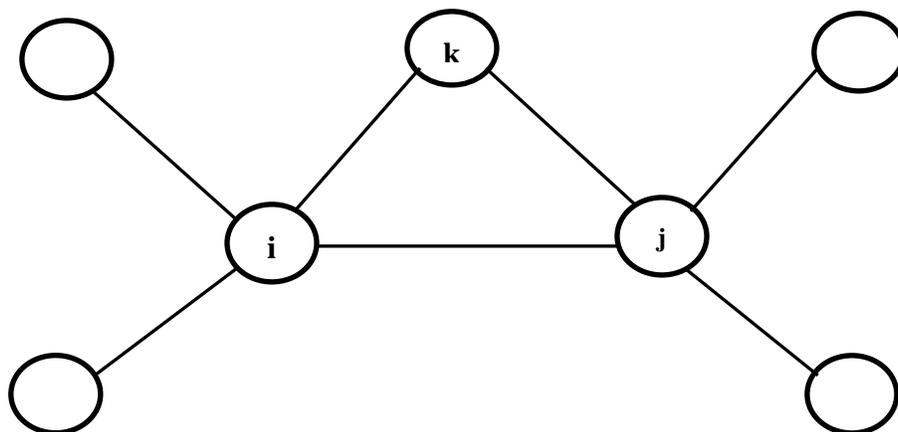


Figure 2: A piece of the network containing nodes i and j and the edge between them. Nodes i and j play the game with each other and both play it with node k . In addition, each of i and j play the game with two other nodes.

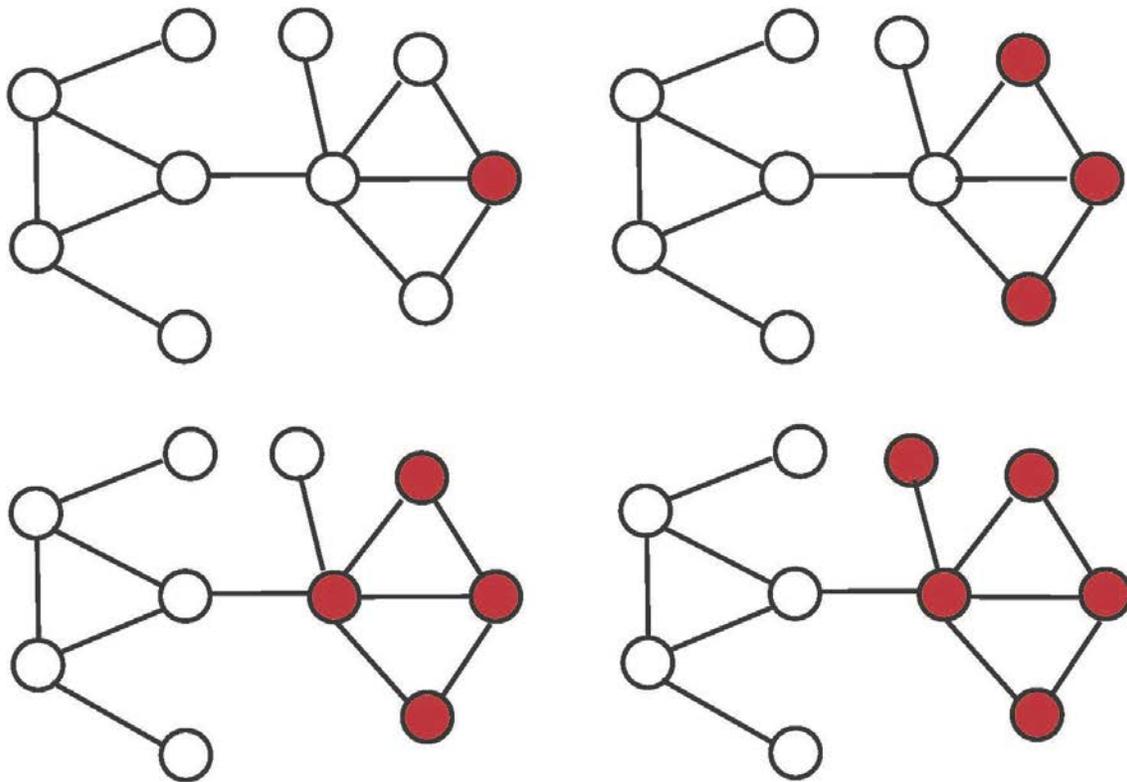


Figure 3: These graphs show the spread of non-ethical behavior represented as red filled nodes in a network in which each node initially plays ethically. In the first figure (upper-left) one node exogenously switches. Then in period 2 (the upper-right figure) the upper and lower nodes connected to this first mover have one-half of their neighbors playing non-ethically so they switch to non-ethical play in period 2. The node which has an edge into the cluster on the left has only one of its five neighbors who are expected to play non-ethically so it stays with ethical play in period 2. However, in period 3 (the lower-left figure) this node now has three of its five neighbors playing non-ethically so it too switches to non-ethical play. Then in period 4 (the lower-right figure) the remaining node not in the cluster switches to non-ethical play. However, the node in the cluster on the left with an edge to a non-ethical node never switches to non-ethical play as only one of its three neighbors plays non-ethically. So none of the nodes in the cluster switch to non-ethical behavior and the spread of non-ethical behavior stops.

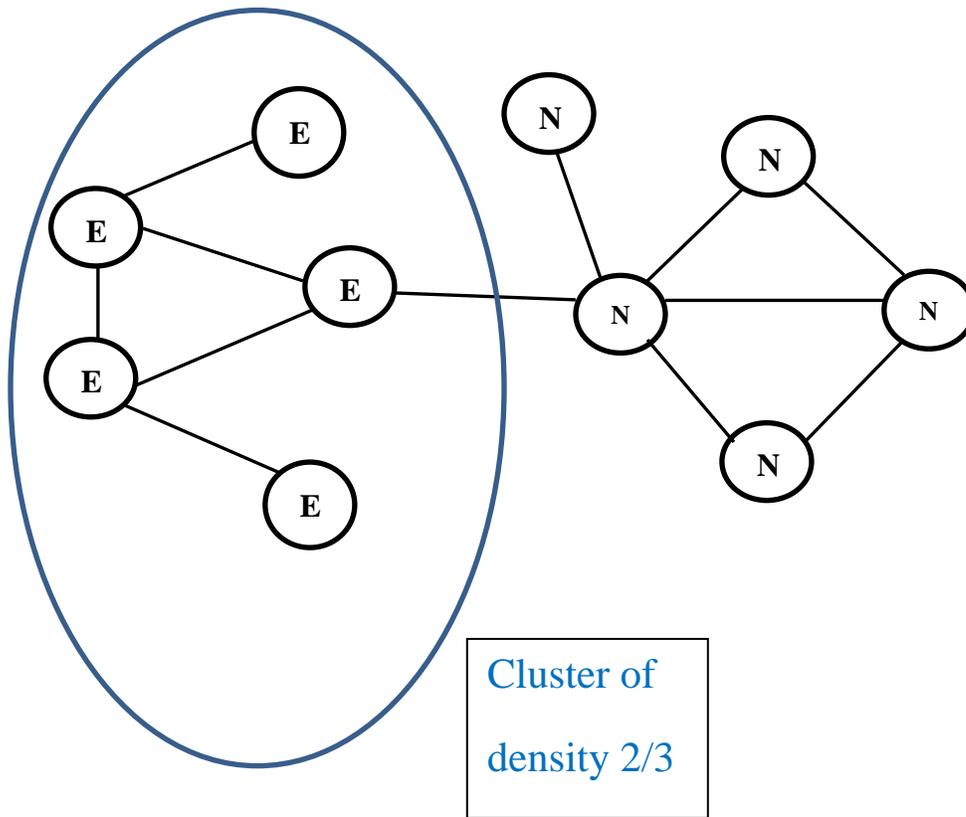


Figure 4: Spread of N on right stopped by cluster on left of density $2/3$ where $2/3 > p^{**}$

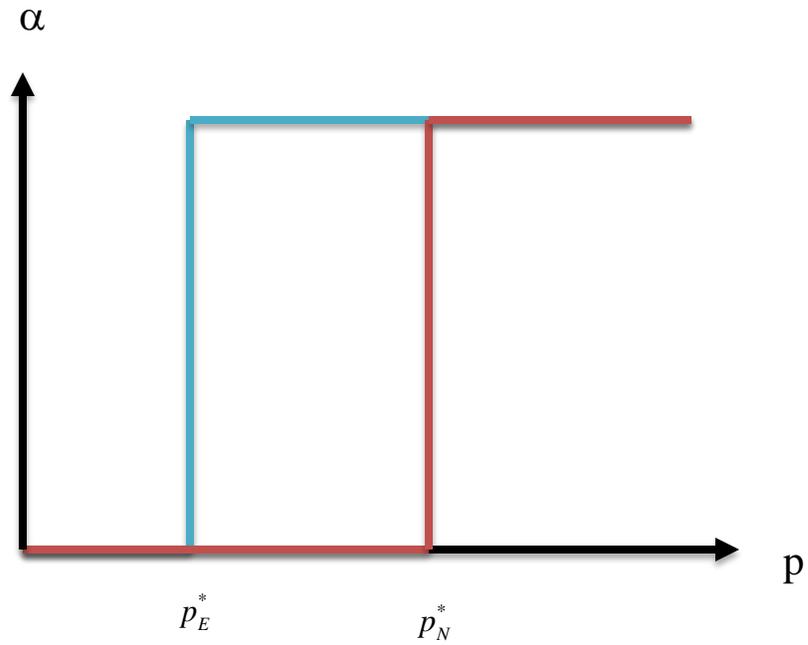


Figure 5: This figure describes optimal play as a function of the fraction of neighbors expected to play ethically. The blue curve provides this relationship for a node that previously played ethically; the red curve describes the relationship for a node that previously played non-ethically.

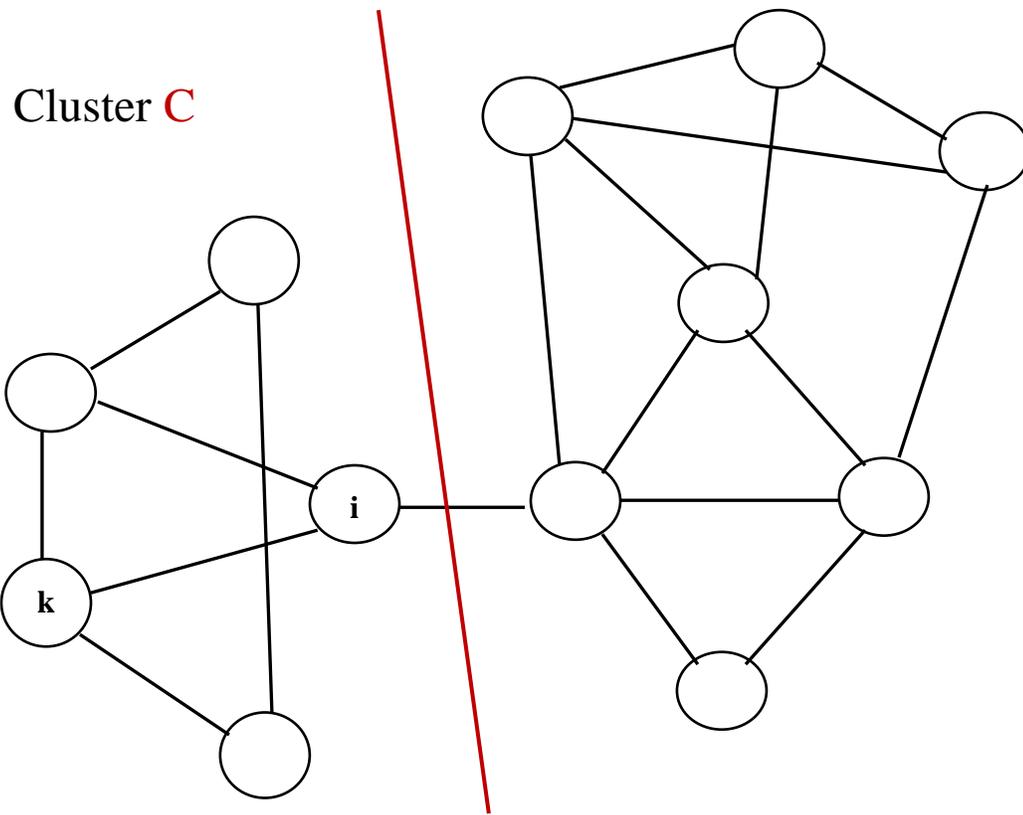


Figure 6: This Figure illustrates the idea that although node i is a gatekeeper for cluster C it is not the most critical node in preventing non-ethical behavior from spreading as a result of a mutation within C . If all nodes in C play E and all others play N , then C is a cluster of density $2/3$ playing E . If node i switches to N the density of $C - \{i\}$ is still $2/3$. If node k switches to N the density of $C - \{k\}$ falls to $1/3$.

Appendix A: Proofs

Proof of Theorem 1:

- (i) A complete cascade requires all of the nodes in the cluster of density p^{**} to switch from E to N. For a complete cascade to occur there must be some first node in the cluster to switch to N. At the time t at which this node switches to N it and all of the other nodes in the cluster had played E at time $t-1$. Each of the nodes in the cluster thus has at least fraction p^{**} of its neighbors who are expected to play E at time t . However, with at least fraction p^{**} of its neighbors expected to play E this first node plays E and does not switch to N. Thus, no node in the cluster switches to N and a complete cascade does not occur.
- (ii) If a complete cascade does not occur then there is some remaining set of nodes, R , playing E. We will show that R is a cluster of nodes playing E with density at least p^{**} . Consider any node $n \in R$. This node chooses to play E so it must have at least fraction p^{**} of its neighbors who are expected to play E, that is, who played E in the previous step and because the cascade has stopped who will play E again. The only nodes left playing E are the nodes in the set R . So the fraction of n 's neighbors who are in R must be at least p^{**} . This holds for all $n \in R$ so R is a cluster of nodes playing E with density at least p^{**} .

Proof of Corollary 2:

All nodes playing N is clearly stable as a node will play E if and only if at least p^{**} fraction of neighbors are expected to play E. So if every node expects every other node to play N then every node actually plays N. Next, consider any cluster of nodes playing E with density at least p^{**} . As every node in the cluster expects every other node in the cluster to play E each of these nodes actually plays E. Finally, no other configuration of play is stable as in any other configuration of play there are some nodes playing E and those nodes do not form a cluster with density at least p^{**} . So some node in this cluster would want to switch to N.

Proof of Theorem 2:

- (i) Consider a network in which all nodes initially play E and a set S of these nodes switch to N. A complete cascade of Ns requires all of the nodes in the cluster of

density p_E^{**} to switch from E to N. For a complete cascade to occur there must be some first node in the cluster to switch to N. At the time t at which this node switches to N it and all of the other nodes in the cluster had played E at time $t-1$. Each of the nodes in the cluster thus has at least fraction p_E^{**} of its neighbors who are expected to play E at time t . However, with at least fraction p_E^{**} of its neighbors expected to play E this first node plays E and does not switch to N. Thus, no node in the cluster switches to N and a complete cascade does not occur.

Alternatively, consider a network in which all nodes initially play N and a set S of these nodes switch to E. A complete cascade of Es requires all of the nodes in the cluster of density $1 - p_N^{**}$ to switch from N to E. For a complete cascade to occur there must be some first node in the cluster to switch to E. At the time t at which this node switches to E it and all of the other nodes in the cluster had played N at time $t-1$. Each of the nodes in the cluster thus has at least fraction $1 - p_N^{**}$ of its neighbors who are expected to play N at time t . So fraction less than p_N^{**} of its neighbors plays E.

However, with fraction less than p_N^{**} of its neighbors expected to play E this first node plays N and does not switch to E. Thus, no node in the cluster switches to E and a complete cascade does not occur.

- (ii) Consider a network in which all nodes initially play E and a set S of these nodes switch to N. If a complete cascade does not occur then there is some remaining set of nodes, R , playing E. We will show that R is a cluster of nodes playing E with density at least p_E^{**} . Consider any node $n \in R$. This node chooses to play E so it must have at least fraction p_E^{**} of its neighbors who are expected to play E, that is, who played E in the previous step and because the cascade has stopped who will play E again. The only nodes left playing E are the nodes in the set R . So the fraction of n 's neighbors who are in R must be at least p_E^{**} . This holds for all $n \in R$ so R is a cluster of nodes playing E with density at least p_E^{**} .

Alternatively, consider a network in which all nodes initially play N and a set S of these nodes switch to E. If a complete cascade does not occur then there is some

remaining set of nodes, R , playing N . We will show that R is a cluster of nodes playing N with density at least $1 - p_N^{**}$. Consider any node $n \in R$. This node chooses to play N so it must have fraction less than p_N^{**} of its neighbors who are expected to play E and thus at least fraction $1 - p_N^{**}$ of its neighbors who are expected to play N , that is, who played N in the previous step and because the cascade has stopped who will play N again. The only nodes left playing N are the nodes in the set R . So the fraction of n 's neighbors who are in R must be at least $1 - p_N^{**}$. This holds for all $n \in R$ so R is a cluster of nodes playing N with density at least $1 - p_N^{**}$.

Proof of Corollary 3:

Follows immediately from Theorem substituting d_E^C for the density in Theorem 2.

Appendix B: Infinite Horizon Game

In this appendix we demonstrate that for sufficiently low, but non-zero discount factors, the play of the two-player game we discuss in Section I.A is also equilibrium play of the infinitely repeated game with non-myopic players. In doing so we use strategies that are consistent with the dynamics examined in Section I.C, so those dynamics are consistent with an equilibrium outcome of the low-discount factor, infinite horizon game.

There are two players $i \in \{1, 2\}$ who in each period, $t = 1, 2, \dots$, chose actions $a_i \in \{E, N\}$. Player's payoffs depend on both actions and on expectations about each other's beliefs; we let (α_1, α_2) be the player's expectations of the other player's belief that the player will play action E and we refer to the objects simply as beliefs. Payoffs in the stage game are given by the payoff matrix on page 10 and we assume that these parameters satisfy the base model assumption. We represent the payoffs as $\pi_i((a_i, a_j), \alpha_i)$. Players discount payoffs at rate $0 \leq \delta < 1$.

A partial history of play through period $t-1$ is $h^{t-1} = ((a_1^1, a_2^1), \dots, (a_1^{t-1}, a_2^{t-1}))$ where the superscripts refer to play in each period. Strategies and beliefs both may depend on the history of play and both need to be specified to describe the actual play of the game. For player i beliefs are given by a sequence of functions $\{\alpha_i^t\}_{t=1}^\infty$ where, abusing notation slightly, $\alpha_i^t(h^{t-1})$ gives player i 's belief at period t as a function of the partial history of play through period $t-1$. Similarly, strategies are $\{a_i^t\}_{t=1}^\infty$ where, again abusing notation slightly, $a_i^t(h^{t-1})$ specifies player i 's action at period t as a function of the partial history of play through period $t-1$. Strategies and belief functions induce an infinite path $(\{a_1^t, a_2^t\}_{t=1}^\infty, \{\alpha_1^t, \alpha_2^t\}_{t=1}^\infty)$ of actions and beliefs that will actually be played and believed. Paths induce payoffs for player 1, $\sum_{t=1}^\infty \delta^{t-1} \pi_1(\{a_1^t, a_2^t\}_{t=1}^\infty, \{\alpha_1^t\}_{t=1}^\infty)$, and similarly for player 2.

Definition 1: The pair of actions and beliefs $((a_1^*, a_2^*), (\alpha_1^*, \alpha_2^*))$ is a pure strategy psychological Nash equilibrium of the stage game if:

- (i) $\pi_1((a_1^*, a_2^*), \alpha_1^*) \geq \pi_1((a_1, a_2^*), \alpha_1^*)$ for all $a_1 \in \{E, N\}$, and similarly for player 2, and,
- (ii) $\alpha_1^* = 1$ if $a_1^* = E$ and $\alpha_1^* = 0$ if $a_1^* = N$, and similarly for player 2.

Definition 2: The pair of strategy and belief functions, $(\{a_1^t, a_2^t\}_{t=1}^\infty, \{\alpha_1^t, \alpha_2^t\}_{t=1}^\infty)$, is a subgame perfect psychological Nash equilibrium of the infinite horizon game if they induce a Nash equilibrium (the play induced by the strategy functions are best responses) in the infinitely repeated game following any partial history (including those not on the outcome path) and the beliefs induced by the belief functions are correct on the outcome path induced by the strategy and belief functions.

It's well known that repeated play of a pure strategy Nash equilibrium of the stage game induces a subgame perfect Nash equilibrium of the infinite-horizon game. That result applies here too, if $((a_1^*, a_2^*), (\alpha_1^*, \alpha_2^*))$ is a pure strategy Nash equilibrium of the stage game then the strategy and belief functions specified by:

$$(\{a_1^t(h^{t-1}) = a_1^*, a_2^t(h^{t-1}) = a_2^*\}_{t=1}^\infty, \{\alpha_1^t(h^{t-1}) = \alpha_1^*, \alpha_2^t(h^{t-1}) = \alpha_2^*\}_{t=1}^\infty), \text{ for all } h^{t-1}$$

is a subgame perfect psychological Nash equilibrium of the infinite-horizon game. Although these beliefs and strategies describe an equilibrium, they are not satisfactory for our purposes as in the dynamic analysis we consider strategies in which each player expects the other player to play the action played last period and each player naively best responds given these beliefs. Fortunately, for sufficiently low discount factors these belief functions and strategy functions are a subgame perfect Nash equilibrium of the infinite-horizon game.

In the theorem below we consider the case in which multiple equilibria are possible, i.e. $g > c - a$. If, instead, $g < c - a$ then N is a strictly dominant strategy in the stage game. So strategies that always play N and beliefs that always anticipate N are obviously equilibria in the infinite horizon game.

Theorem A.1: Suppose that $g > c - a$, $\delta < \frac{b}{a}$ and that the initial play used to set expectations is either (E, E) or (N, N). The belief and strategy functions:

$$\beta_i^t(h^{t-1}) = \begin{cases} 1 & \text{if } a_i^{t-1} = E \\ 0 & \text{if } a_i^{t-1} = N \end{cases} \text{ for all } h^{t-1}, \text{ and}$$

$$a_i^t(h^{t-1}) = E \text{ if}$$

$$\pi_i((E, a_j^{t-1}), \beta_i^t(h^{t-1})) \geq \pi_i((a_i, a_j^{t-1}), \beta_i^t(h^{t-1})) \text{ for all } a_i \in \{E, N\}$$

$$a_i^t(h^{t-1}) = N \text{ otherwise}$$

are a subgame perfect psychological Nash equilibrium of the infinite horizon game.

Proof: We need to show that no deviations from these strategies are profitable, given the beliefs, and that given the strategies the beliefs are correct on the equilibrium path of play. First, we consider deviations. For any partial history h^{t-1} these belief functions and strategy functions depend only on play in period $t-1$. There are four possible values for that play: (E, E), (N, N), (E, N) and (N, E). We consider each of the four types of subgames these pairs of actions induce.

1. Suppose that at time t the partial history h^{t-1} ends with play (E, E). The proposed belief and strategy functions imply that each player is expected to play E. As $g > c - a$ both players play E. Consider a deviation by player 1 (wlog). The proposed strategy yields player 1 a payoff of a in period t . The only greater payoff in the stage is c which is possible only if player 1 expects player 2 to expect player 1 to play N and player 2 actually plays E. Note, however, that 2 expects 1 to play N only if 1 played N in the previous stage. In this case, the strategy function has 2 playing N, not E. So 1's maximum stage payoff following any partial history is a . Thus, no deviation by 1 is profitable.
2. Suppose that at time t the partial history h^{t-1} ends with play (N, N). The proposed belief and strategy functions imply that each player is expected to play N and that each player does play N. This yields a payoff of b for each player. Consider a deviation by player 1 (wlog) to a play of E at period t . As 2 plays N this gives player 1 a payoff of 0 in period t . Player 1's deviation can at best generate payoffs of a in each future period. Thus the deviation is not profitable if $\frac{b}{1-\delta} > 0 + \frac{\delta a}{1-\delta}$ which follows from our assumption that $\delta < \frac{b}{a}$.
3. Suppose that at time t the partial history h^{t-1} ends with play (N, E). The proposed belief functions imply that player 1 expects that 2 expects 1 to play N, player 2 expects that 1 expects 2 to play E. The strategy functions imply that each player plays N. This yields a payoff of b for player 1 in period t and payoff of $b-g$ for player 2 in period t , and payoffs of b for each player in all future periods. Consider a deviation by player 1 to a play of E at period t . Since player 2's strategy function has 2 play N at t , player 1 expects a payoff of 0 at period t . The partial history h^t now ends with play (E, N). The strategy functions give play at period $t+1$ of (N, N) and payoff to player 1 at $t+1$ of $b-g$. Following a play of (N, N) the strategy functions specify that all future play is (N, N) yielding a payoff of b in all future periods for each player. Thus player 1's deviation is costly in periods t and $t+1$ without yielding any future benefit. The argument for a deviation by player 2 is similar.

Finally, note that according to cases (1) and (2) above beliefs are correct along any equilibrium path beginning from initial play of either (E, E) or (N, N).

References

- Acemoglu, D., and M. O. Jackson, 2015, History, expectations, and leadership in the evolution of social norms, *Review of Economic Studies* 82, 1–34.
- Acemoglu, D., A. Oxdaglar, and A. Tahbaz-Salehi, 2015, Systemic Risk and Stability in Financial Networks, *American Economic Review*, 105, 2, 564-608.
- Allen, F. and A. Babus, 2008, Networks in Finance, Tech. Report 08-07, Wharton Financial Institutions Center.
- Allen, F. and D. Gale, 2000, Financial Contagion, *Journal of Political Economy*, 108, 1, 1-33.
- Babus, A., 2016, The Formation of Financial Networks, *RAND Journal of Economics* 47, 239–272.
- Battigalli, P. and M. Dufwenberg, 2007, Guilt in games, *American Economic Review*, 97(2), 170-176.
- Benoit, S., J. Colliard, C. Hurlin, and C. Perignon, 2017, Where the Risks Lie: A Survey on Systemic Risk,”*Review of Finance*, 21, 109–152.
- Bernile, Gennaro, and Gregg A. Jarrell, 2009, The impact of the options backdating scandal on shareholders, *Journal of Accounting and Economics* 47, 2–26.
- Blume, L., D. Easley, J. Kleinberg, R. Kleinberg, and E. Tardos, 2013, Network Formation in the Presence of Contagious Risk, *ACM Transactions on Economics and Computation*, 1,2, 1-20.
- Cremer, J., 1993, Corporate Culture and Shared Knowledge, *Industrial and Social Change*, 2, 351-386.
- Dimmock, S. G., W. C. Gerken, and N. P. Graham, 2018, Is fraud contagious? Co-worker influence on misconduct by financial advisors, *Journal of Finance* 73, 1417– 1450.
- Dufwenberger, M. and A. Patel, 2017, Reciprocity Networks and the Participation Problem, *Games and Economic Behavior*, 101, 260-272.
- Easley, D. and J. Kleinberg, 2010, *Networks, Crowds, and Markets*. (Cambridge Books: UK).
- Eisenberg, L. and T. H. Noe, 2001, Systemic Risk in Financial Systems, *Management Science* 47, 2, 236–49.
- Elliott, M., B. Golub, and M. Jackson, 2014, Financial Networks and Contagion, *American Economic Review*, 104, 10, 3115-53.

- Egan, M., G. Matvos, and A. Seru, 2019, The Market for Financial Advisor Misconduct, *Journal of Political Economy*, *forthcoming*.
- Elster, J., 1989, Social Norms and Economic Theory, *Journal of Economic Perspectives*, 3, 4, 99-117.
- Erol, S. and R. Vohra, 2014, Network Formation and Systemic Risk, working paper, University of Pennsylvania.
- Freixas, X., B. Parigi, and J.-C. Rochet, 2000, Systemic Risk, Interbank Relations, and Liquidity Provision by the Central Bank, *Journal of Monetary Economics*, 32, 3, 611-638.
- Gai, P., A. Haldane, and S. Kapadia, 2011, Complexity, Concentration and Contagion, *Journal of Monetary Economics*, 58,5, 453–70.
- Geanakoplos, J., Pearce, D., & Stacchetti, E. (1989). Psychological games and sequential rationality. *Games & Economic Behavior*, 1(1), 60-79.
- Glasserman, P. and H. P. Young, 2016, Contagion in Financial Networks, *Journal of Economic Literature*, 54, 3, 779-831.
- Guiso, L. P. Sapienza, and L. Zingales, 2006, Does culture affect economic outcomes? *Journal of Economic Perspectives*, (20), 23-48.
- Karpoff, J. M., A. Koester, D. S. Lee, and G. S. Martin, 2017, Proxies and databases in financial misconduct research, *Accounting Review* 92, 129–163.
- Kedia, S., K. Koh, and S. Rajgopal, 2015, Evidence on contagion in earnings management, *Accounting Review* 90, 2337–2373.
- Kempke, D., J. Kleinberg and E. Tardos, 2003, Maximizing the Spread of Influence through a Social Network, *KDD, Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining 2003*, 137-146.
- 2005, Influential Nodes in a Diffusion Model for Social Networks, *International Colloquium on Automata, Languages, and Programming*, 1127-1138.
- Leider, S., M. Möbius, T. Rosenblat, and Q-A Do, 2009, Directed Altruism and Enforced Reciprocity in Social Networks, *The Quarterly Journal of Economics*, 124, 4, 1815–1851.
- Lie, E., 2005, On the timing of CEO stock option awards, *Management Science*, 51(5), 802-812.
- Liu, X., 2016, Corruption culture and corporate misconduct, *Journal of Financial Economics* 122, 307–327

- Lo, A. W., 2016, The Gordon Gekko Effect: The Role of Culture in the Financial Industry, *Economic Policy Review*, Aug, pp. 17-42.
- Mokyr, J., 2016, *A Culture of Growth*, (Princeton University Press; Princeton NJ).
- Morris, S. 2000, Contagion, *Review of Economic Studies*, 67, 57-78.
- O'Hara, M., 2016, *Something for Nothing: Arbitrage and Ethics on Wall Street*, (W.W. Norton: NY)
- Ostrom, E., 2000, Collective Action and the Evolution of Social Norms, *Journal of Economic Perspectives*, 14, 3, 137-158.
- Parsons, C. A., Sulaeman, J. and Titman, S., (2018), *The Geography of Financial Misconduct*, *Journal of Finance*, **73**, issue 5, p. 2087-2137.
- Sandel, M., 2013, Market Reasoning as Moral Reasoning: Why Economists Should Re-engage with Political Philosophy, *Journal of Economic Perspectives* 27, no. 4: 136
- Sugden, R., 1989, Spontaneous order, *Journal of Economic Perspectives* 3, 85-97.
- Young, H. P., 2015, The Evolution of Social Norms, *Annual Review of Economics*, 7,1, 359-387.