



**CPI** COMPETITION POLICY  
INTERNATIONAL

Volume 6 | Number 1 | Spring 2010

# Behavioral Economics as Applied to Firms: A Primer

Mark Armstrong & Steffen Huck  
University College London

# Behavioral Economics as Applied to Firms: A Primer

---

*Mark Armstrong and Steffen Huck\**

We discuss the literatures on behavioral economics, bounded rationality, and experimental economics as they apply to firm behavior in markets. Topics discussed include the impact of imitative and satisficing behavior by firms, outcomes when managers care about their position relative to peers, the benefits of employing managers whose objective diverges from profit-maximization (including managers who are overconfident or base pricing decisions on sunk costs), the impact of social preferences on the ability to collude, and the incentive for profit-maximizing firms to mimic irrational behavior.

---

\*Department of Economics, University College London. We have benefited from the non-selfish behavior of many friends and colleagues, and we are grateful for advice received from Matthew Bennett, Emanuele Giovannetti, Antonio Guarino, Joe Harrington, Bruce Lyons, Adam Rosen, David Ruck, David Sappington, Burkhard Schipper, John Vickers, Eric Van den Steen, Maurice Stucke, Mark Williams, Jidong Zhou, Liva Zorgenfrei, and the editorial team at this journal. We also want to thank the Office of Fair Trading and the Economic and Social Research Council (UK) for funding assistance.

## I. Introduction

In recent years there has been a good deal of research investigating how poor or non-standard decision making by consumers might affect market outcomes. In much of this work, the assumption is that firms are fully rational and aim to maximize their profits (and sometimes they do this by exploiting the behavioral biases of consumers). Some of this work points to situations where there is a role for policy which protects consumers from their own failings and from exploitative firms.<sup>1</sup>

In this article we focus instead on non-standard approaches to firm behavior. Consumers are kept in the background, and are present merely to generate in some fashion a demand curve for the firms' products. We present evidence—both real world and experimental—that firms (or experimental subjects playing the role of firms) sometimes depart from the profit-maximizing paradigm. For instance: firms may be content to achieve “satisfactory” rather than optimal profits; firms might rely on simple rules of thumb—such as imitating the strategies of well-performing rivals, or changing strategies only when profit falls below some acceptable threshold—rather than on explicit calculation of complex optimal strategies; firms may base pricing decisions on sunk costs as well as avoidable costs; CEOs may be overly optimistic about the profitability of mergers or other actions they undertake; managers might face incentives which induce them to care about relative rather than absolute profits; firms might punish rivals who behave “unfairly” towards them; and so on. We believe that many mainstream industrial economists and policy-makers are not yet fully aware of the substantial literature on these topics, and our aim in this survey is to bring some of the insights of behavioral economics as applied to firms to wider attention.<sup>2</sup>

There are, of course, a number of reasons why one might expect firms to be better decision makers than consumers, and this helps to explain the recent focus on consumer failings. First, there are economies of scale in making good decisions. A consumer may have to decide whether to buy a given product just a few times and it may not be worthwhile to invest much effort in making the right decision, while a firm selling to millions of customers has more at stake in getting it right. Relatedly, since firms often do the same things repeatedly, they may quickly learn how to do it right, while a consumer buying a rarely-purchased product may not have that opportunity. Second, firms compete with one another while consumers usually do not, and firms that are better at generating profits may succeed and prosper at the expense of firms which make worse decisions. In modern society, consumers rarely “exit” when they make poor market decisions.

WE PRESENT EVIDENCE—  
BOTH REAL WORLD AND  
EXPERIMENTAL—THAT FIRMS  
(OR EXPERIMENTAL SUBJECTS  
PLAYING THE ROLE OF FIRMS)  
SOMETIMES DEPART FROM THE  
PROFIT-MAXIMIZING PARADIGM.

However, there are also a number of considerations which go the other way:

- Firms typically operate within a highly complex and uncertain environment, and may need to resort to decision-making short-cuts and rules-of-thumb.
- Firms often face the added complexity stemming from strategic interaction with rivals, which consumers rarely do when buying products.<sup>3</sup> For instance, in even the simplest situations, theories of tacit collusion or reputation-building require firms to follow highly intricate strategies.
- In practice, under-performing firms may take a long time to exit, and there is often a long period of decline before an established firm actually leaves the market. More generally, the complexity of the environment may mean that crucial decisions are taken with significant delay.<sup>4</sup>
- Group decision making (as practiced more often by firms rather than by consumers) could introduce extra biases. For instance, the separation of ownership from control could leave managers free to pursue

THAT IS TO SAY, ALTHOUGH  
IT APPEARS PARADOXICAL,  
ACTUAL PROFITS MIGHT BE  
ENHANCED WHEN THE FIRM'S  
OBJECTIVE DEPARTS FROM  
PROFIT-MAXIMIZATION.

their own objectives that may differ from maximizing shareholder value. In fact, this last point can be turned on its head: In imperfectly competitive markets, shareholder returns might be enhanced by (deliberately or not) hiring managers whose objectives differ from maximization of profit. That is to say, although it appears paradoxical, actual profits might be enhanced when the firm's objective departs from profit-maximization. For instance, hiring

an aggressive or over-optimistic CEO, rewarding a CEO based on her performance relative to peers, or employing a CEO who bases pricing decisions on sunk costs, might all have strategic benefits in terms of affecting rival responses.

- Relatedly, the people who succeed in the tough career competition to manage firms might have these kinds of personality traits more frequently than the general population. Sometimes it seems that managers are motivated in part by personal animosity—or respect—towards a rival.
- Different managers have differing skills, and in particular some managers may be better at thinking strategically than others.<sup>5</sup> More generally, it seems clear that an individual manager's "style" can be important—for good or ill—for a firm's performance.<sup>6</sup>
- A manager may on occasion have a personal interest in the firm's activities (say, a sports team, newspaper, or fine wine<sup>7</sup>) beyond the profit generated, and this may cause a divergence from profit-maximization.<sup>8</sup> (Captain Ahab, in Melville's *Moby Dick*, is an extreme, if fictional, instance.)

- Rather than affecting rival responses as mentioned above, explicitly non-profit aims might act to boost consumer demand, as appears often to be the case with “fair-trade” or “green” products, for instance.
- Illegal cartels need to find ways to resolve disagreements about market shares, whether cheating has occurred, and so on, without recourse to legally binding agreements. As such, issues such as building trust and *esprit de corps* among conspirators are important.
- Finally, the potential presence of a “behavioral” type of firm in a market could induce a profit-maximizing firm to mimic irrational behavior. For instance, a firm might wish to gain a reputation to fight entry come what may, in order to deter future entry.

Milton Friedman is perhaps the most famous, and often quoted, exponent of the pervasive view that competition acts to discipline firms to maximize their profits:

---

“Let the apparent determinant of business behavior be anything at all—habitual reaction, random chance or what not. Whenever this determinant happens to lead to behavior consistent with rational and informed maximization of returns, the business will prosper and acquire resources with which to expand; whenever it does not, the business will tend to lose resources and can be kept in existence only by the addition of resources from outside. The process of ‘natural selection’ helps to validate the hypothesis [of ‘rational and informed maximization of returns’]—or, rather, given natural selection, acceptance of the hypothesis can be based largely on the judgment that it summarizes appropriately the conditions for survival.”<sup>9</sup>

---

As we discuss in our conclusions, this view has much merit when it comes to many kinds of decision errors and behavioral biases, and some decision problems exhibited by consumers (such as procrastination, for instance) are less plausible in firms. Nevertheless, at various points in this survey we will see examples of situations where firms prosper when their objective diverges from maximizing profits, when “behavioral” managers are put in place, or when firms are content to achieve merely a satisfactory profit level.

In the remainder of this article, we discuss several of these points in greater detail. Section II surveys the experimental evidence on the ability of firms to collude, and how public policy could affect this ability. Section III discusses the related issues of imitative behavior by firms and concerns for relative (not absolute) profit, both of which often make an oligopolistic market more competitive than the orthodox theory suggests. Other forms of social preferences—a

desire for vengeance if a rival “cheats” on a collusive agreement, say—are presented in section IV, which we argue may help sustain collusive agreements. Various kinds of satisficing behavior are discussed in section V, where we show how imperfectly optimizing firms may actually end up with greater profits than their profit-maximizing counterparts. Section VI collects together discussion of the impact of over-optimism by entrepreneurs and managers, the imitation of irrational behavior by profit-maximizing firms, and the possible benefits of including fixed costs in a firm’s pricing decisions. Concluding comments are contained in section VII.

We do not attempt in this article to survey the entire terrain of behavioral economics as applied to firms. In particular, our focus is on firm behavior in markets rather than on the internal organization of firms. Similarly, we do not emphasize the potential policy implications of non-standard behavior by firms, although a few remarks about this are made throughout the paper and in our conclusions.

Before continuing, however, it is worthwhile discussing the role of laboratory experiments—which supply much of the empirical evidence we present in the following discussion—as an aid to understanding firm behavior. Laboratory experiments usually involve students playing the role of firms, who make decisions in minutes and for relatively low stakes, whereas real firms employ managers who are highly remunerated, experienced, and carefully selected. Why should these experiments tell us anything about the performance of actual markets? There are a number of reasons why we cautiously believe that data from experiments are indeed useful.<sup>10</sup>

WHY SHOULD THESE EXPERIMENTS  
TELL US ANYTHING ABOUT THE  
PERFORMANCE OF ACTUAL  
MARKETS? THERE ARE A NUMBER  
OF REASONS WHY WE CAUTIOUSLY  
BELIEVE THAT DATA FROM  
EXPERIMENTS ARE INDEED USEFUL.

First, experiments have been run with business people as subjects instead of students, and the latter do not appear to perform less (or more) rationally than the former.<sup>11</sup> In addition, experimenters are careful to ensure that their subjects have a good deal of experience in playing the chosen game before they start giving weight to the generated data. Indeed, it is possible that subjects have more experience playing the highly stylized games in the laboratory than real managers have in their own naturally occurring markets. Second, as Plott<sup>12</sup> writes:

“The relevance of experimental methods rests on the proposition that laboratory markets are ‘real’ markets in the sense that principles of economics apply there as well as elsewhere. Real people pursue real profits within the context of real rules. The simplicity of laboratory markets in comparison with naturally occurring markets must not be confused with questions about their reality as markets.”

---

“The relevance of experimental methods rests on the proposition that laboratory markets are ‘real’ markets in the sense that principles of economics apply there as well as elsewhere. Real people pursue real profits within the context of real rules. The simplicity of laboratory markets in comparison with naturally occurring markets must not be confused with questions about their reality as markets.”

---

Thus, theories and models that work generally should also work in the particular cases of laboratory markets.

Third, many theories in industrial organization are very subtle, and depend on detailed assumptions about what firms observe about their rivals, and what firms can say to their rivals. Information flows can be carefully controlled in the laboratory, while the environment of naturally occurring markets is harder to pin down, and it is often impractical to test complex theories about oligopoly behavior using non-experimental data. For instance, experiments are able to examine tacit collusion, in which no communication is possible between firms, while it is hard to be sure there is no communication going on in a real market. It is also hard to measure such important parameters as marginal costs in natural markets, while these can be accurately generated in the laboratory.<sup>13</sup> Nevertheless, worries about the relevance of laboratory experiments should not be dismissed too casually: The concern that student behavior does not always closely match CEO behavior is probably more serious for oligopoly experiments than for other experiments (such as those concerning shopping or bargaining), and it is hard convincingly to replicate the detailed institutional structure of firms—such as hierarchies and group dynamics—in the laboratory.

## II. Ability to Sustain Collusion

Collusion among oligopolists occurs when firms can sustain relatively high prices by credibly threatening a price war should one firm undercut the prevailing price (or boost its output). Firms need to interact repeatedly for high prices to be sustained, since a firm must trade off high profits now (if it deviates and undercuts its rivals) with low profits in the future (after a price war is triggered). A typical “trigger” strategy to sustain high prices takes this form: Each firm sets the collusive price as long as all firms continue to do so, but if one or more firms sets a price below the collusive price, all firms subsequently set prices according to a one-shot equilibrium strategy (e.g., set price equal to marginal cost if the firms supply homogenous products). Although even this simplest setting involves complex strategies, including the need somehow to coordinate on the particular collusive price, matters become vastly more complex if: firms are asymmetric; demand varies over time; firms cannot observe each others’ price (and so cannot tell if low demand now is due to an adverse demand shock or due to one firm offering a low price); and so on.<sup>14</sup>

Fully rational firms (who are sure their rivals are also rational) cannot sustain collusion if there is a known end-point to their interaction, no matter how far in the future. (If the market ends after 100 periods, in the 100<sup>th</sup> period firms know there is no future to punish them, and so set low prices to undercut their rivals, and the one-shot equilibrium is played. In the 99<sup>th</sup> period, firms know what will happen next period, so again have no incentive to cooperate then. And the whole repeated interaction unravels, with the result that no collusion

is ever sustained.) When there is open-ended interaction, collusion is harder to sustain when there are more firms in the market. With many firms in the market, the short-run benefits of a price cut are relatively large compared to getting a small share of the on-going collusive profits. This means that collusion among many firms requires a higher discount factor (i.e., more weight placed on future

MANY EXPERIMENTS FIND THAT COLLUSION CAN STILL BE OBSERVED EVEN WHEN THERE IS A KNOWN FIXED NUMBER OF PERIODS; A FINDING THAT GOES AGAINST THE PREDICTION OF RATIONAL PLAY.

collusive profits) than does collusion among few firms. Nevertheless, theory suggests that in plausible environments, collusion is achievable with relatively large numbers of rational firms.<sup>15</sup> Collusion is also likely to be harder to sustain if firms cannot observe each others' actions.<sup>16</sup> For instance, if a firm cannot tell if low demand is due to an adverse market shock or due to an undercutting rival, then it may be

reluctant to punish the rival harshly (since punishment hurts the punisher as well), and this could make collusion hard to achieve. But if a firm can observe a rival's actual price, it can punish harshly when that price is low, and so make collusion more effective. For this reason, frequent information exchange is a vital part of many cartels.<sup>17</sup>

## A. COLLUSIVE BEHAVIOR IN THE LABORATORY

How do firms behave in repeated interactions in experiments? A typical experiment to analyze this question is conducted as follows:

- The same subjects repeatedly play a stylized oligopoly game (say, Cournot quantity competition or Bertrand price competition) for a number of periods. The number of periods is either fixed (and announced in advance) or random (with, say, a 90 percent chance of another interaction after each period's play). In the latter case, the likelihood of having another interaction plays the role of the discount factor, representing how important the future is.
- In each period, the actions available to each firm might be very restricted (e.g., with just two possibilities: "collude" or "compete" as in the Prisoner's Dilemma) or more numerous (e.g., choosing any integer output from 0 to 100).
- The payoff structure of each period's game might be described in full to each player if the actions are not too numerous (e.g., if playing a Cournot game a payoff table shows how a player's profit depends on her own chosen output along with the aggregate output of the rivals). If actions are numerous, then subjects may be given a "profit calculator" which gives the subject's profit as a function of some specified values for her own action and those of her rivals. Alternatively, the payoff structure might not be revealed at all, and after each period's play a subject sees her own realized profit, as well as (possibly) the actions and realized profits of her rivals.



- Particular market features could be introduced. For instance, players could be permitted to make (non-binding) announcements in advance about what actions they will follow, or players could decide whether to form a cartel.<sup>18</sup>

A focus of the literature has been to observe the extent of collusion in the laboratory market (say, measured as the ratio of actual profits to the theoretical one-shot Nash equilibrium profits), and how this depends on: i) the discount factor (or number of periods in the fixed-period case); ii) the number of oligopolists; iii) the kind of information revealed to participants; or iv) whether they can make pre-play announcements about their actions. Note that it is only in a sterile laboratory setting, where the information and communication flows can be precisely controlled, that one could investigate some of these questions.

This is not the place for a detailed summary of the large literature on experimental repeated oligopoly. But some highlights include the following observations.<sup>19</sup> First, the importance of the future does affect the extent of collusion. Feinberg & Husted<sup>20</sup> performed Cournot experiments with two different probabilities of continuing interaction, and observed that collusion was more prevalent with the higher probability. This is qualitatively in line with predictions of the behavior of fully rational firms in indefinitely repeated games.<sup>21</sup> However, many experiments find that collusion can still be observed even when there is a known fixed number of periods; a finding that goes against the prediction of rational play.<sup>22</sup> Of course, even with a finite horizon it is perfectly rational to cooperate for a time if one thinks that there is a chance one may be playing against a naïve opponent who is playing something like a trigger strategy (see section VI below for further discussion). But it is common to observe that collusion breaks down as the known endpoint nears.

In experimental markets, it appears that the number of oligopolists is crucial in terms of the ability to collude tacitly, i.e., where firms cannot communicate directly with each other. Huck, Normann, & Oechssler<sup>23</sup> document how there is very little tacit collusion in Cournot markets with three or more firms, which conflicts with the theoretical predictions discussed earlier. Indeed, if anything, it is more common to observe that markets with more than two firms are more competitive than the static Nash prediction. On the other hand, there is a considerable amount of collusion in duopolies.<sup>24</sup> One obvious difference between duopolies and oligopolies is that when there is a deviation from the collusive strategy, it is obvious to the two suppliers which of them deviated, presuming a deviation is known to have occurred. Moreover, it is possible to inflict punishment on that firm without hurting innocent rivals who were colluding according to plan.

IN EXPERIMENTAL MARKETS, IT APPEARS THAT THE NUMBER OF OLIGOPOLISTS IS CRUCIAL IN TERMS OF THE ABILITY TO COLLUDE TACITLY, I.E., WHERE FIRMS CANNOT COMMUNICATE DIRECTLY WITH EACH OTHER.

In Bertrand markets with homogeneous products, prices above competitive levels are routinely observed in experiments, even with more than two firms or with a one-shot market interaction.<sup>25</sup> However, these markets are very special, as in the one-shot Bertrand equilibrium firms make zero profits (when unit costs are constant), and so the only chance firms have to make positive profits is to price above cost (and they have little to lose if they fail to offer the lowest price). These markets are discussed further in section V, where we argue that satisficing behavior by firms is a plausible explanation for observed prices.

Huck, Normann, & Oechssler<sup>26</sup> and Offerman, Potters, & Sonnemans<sup>27</sup> examine the impact of the experimenter revealing information about each rival's

HOWEVER, ALL THREE STUDIES  
FIND THAT REVEALING  
INFORMATION ABOUT RIVAL  
OUTPUTS AND PROFITS  
ACTUALLY RENDERS MARKETS  
MORE COMPETITIVE.

actions and payoffs to firms. (Huck et al. investigate four-firm Cournot markets, while Offerman et al. look at three-firm Cournot markets.) In some models of tacit collusion (though not the ones studied in these laboratory experiments), when firms can directly observe each other's past actions, this enables deviations to the collusive agreement to be detected, and so helps collusion be achieved.<sup>28</sup> However, all

three studies find that revealing information about rival outputs and profits actually renders markets more competitive.<sup>29</sup> Compared to the setting in which no firm-specific information is revealed, they find significantly lower prices when firms can observe each other's actions and profits.<sup>30</sup> While standard economic theory cannot easily explain this empirical finding—that competition in the presence of information about others is often tougher than in the static Nash equilibrium, sometimes approaching the fully competitive outcome where price equals marginal cost—there is a class of models in the evolutionary and learning literature that does predict such patterns, which we discuss in section III.

A related issue concerns the impact of communication between firms, and whether the ability to engage in some form of communication before market interaction aids collusion in the laboratory. (Of course, any agreements made in the communication stage are non-binding and purely “cheap talk.”) This literature is surveyed in Potters,<sup>31</sup> who concludes that given the opportunity to communicate, firms do use this opportunity to conspire to fix prices, and this ability often has the effect of raising prices in the market (even with more than two firms). It appears that face-to-face talk has more of a collusive impact than computer-mediated communication, somewhat consistent with the camaraderie among conspirators we discuss later in section IV. Andersson & Wengstrom<sup>32</sup> report results from an experiment in which it is costly to send a message from one firm to another. Intriguingly, they find that as the cost of communication rises (which could be interpreted as the outcome of a more vigilant antitrust policy) fewer messages are sent, as is intuitive, but when messages are sent collusion is more effective. One interpretation is that firms feel more committed to a collu-

sive agreement if costs are needed to reach that agreement. The net impact of making it hard to communicate is that collusion is substantially *more* prevalent, so that certain forms of competition policy might perversely turn out to aid cartel formation and stability.<sup>33</sup>

## B. LENIENCY POLICIES

Here we continue the discussion of cartel formation in the presence of a competition authority, but in the context of leniency policies. Both U.S. and European antitrust authorities currently make use of leniency policies for whistle blowers, and in recent years the number of detected cartels has increased considerably. The idea is that providing an incentive to whistle blowers should render cartels less stable, as well as providing antitrust authorities with evidence which would otherwise be hard to obtain. While these benefits appear intuitive, theoretical work by Spagnolo<sup>34</sup> and Motta & Polo<sup>35</sup> have cast doubt on this intuition. The logic of their argument is that firms can use whistle blowing as a way to punish those members of a cartel which undercut agreed prices. Ironically, the ability to blow the whistle may then help to sustain collusion.

IRONICALLY, THE ABILITY TO BLOW THE WHISTLE MAY THEN HELP TO SUSTAIN COLLUSION.

Antitrust authorities have a number of possible policies to fight cartels, including (i) fining the participants when misconduct is proved, (ii) granting leniency to cartel members who bring evidence of the cartel to the authority, or (iii) rewarding a cartel member if it brings forward evidence of the cartel. Apesteguia, Dufwenberg, & Selten<sup>36</sup> experimentally investigated the impact of various leniency policies. They studied three-firm Bertrand markets under four experimental treatments: Standard, Leniency, Bonus, and Ideal. In Standard, Leniency, and Bonus, firms have the opportunity to form a cartel before they interact in the market. (A cartel is formed only if all three firms agree, and firms interact just once in the market to avoid issues of tacit collusion.) If a cartel is formed, the three firms can communicate in an unstructured way for 10 minutes, presumably to discuss their collusive strategy. (For instance, one possible strategy is to agree on a collusive price and to suggest that if one firm undercuts the price its rivals report the cartel.) In Standard there is no leniency clause and every firm (including the firm that blows the whistle) is fined a fraction of their turnover if a cartel has been formed and reported to the authorities. In Leniency the whistle blower gets a discount on the fine. (The discount is 100 percent if there is only one whistle blower, 50 percent if there are two simultaneous whistle blowers, and 33.3 percent if all three do so.) In Bonus the whistle blowers share the fines paid by the cartel members who kept quiet. In the fourth treatment, Ideal, the design is such that firms simply do not have the opportunity to form a cartel.

The parameters chosen in the experiments imply that with rational players the Standard and Leniency treatments can sustain collusion while Bonus and Ideal cannot. However, the data draw a different picture, and the most effective treat-

ment is not Bonus but Leniency despite its potential theoretical weaknesses. In fact, Leniency generates prices that are statistically indistinguishable from prices in Ideal where cartel formation is not possible. But theory not only fails to predict the performance of Leniency, it also gets the effect of Bonus wrong. Predicted to be the most effective remedy against collusion, Bonus, in fact, leads to prices which are as high as those seen in the Standard treatment. In addition, the Bonus treatment leads to the greatest incidence of cartel formation, while Leniency has the least. The authors conjecture that subjects in the Bonus treat-

THUS, BOTH OF THESE STUDIES  
SHOW THAT LENIENCY PROGRAMS,  
DESPITE THEIR POTENTIAL  
WEAKNESSES IN THEORY, ARE  
SUCCESSFUL IN FIGHTING CARTELS.

ment are tempted to form a cartel with an agreement to set high prices, and then also being rewarded for reporting the cartel (although this strategy is not in fact profitable when many firms form a cartel with this duplicitous intent). Apestegua, et al. conclude: “Our findings in this paper provide no reason for Gary Spratling and Mario Monti to feel disap-

pointed with the leniency clauses that have recently been incorporated into the anti-trust legislation in most member states of the OECD.”

While Apestegua, et al. study a one-shot interaction, Hinloopen & Soetevent<sup>37</sup> examine a setting where three firms repeatedly interact in the same market for 20 periods. (With a repeated interaction, firms may be more reluctant to report a cartel, since they then forgo the future benefits of collusion. In addition, even if cartel formation is impossible, there is the possibility of tacit collusion.) They report that leniency programs not only reduce the frequency with which cartels are formed but also reduce the stability of cartels that do get formed. Moreover, cartels that do get formed in the presence of the leniency programs charge relatively lower prices. Thus, both of these studies show that leniency programs, despite their potential weaknesses in theory, are successful in fighting cartels.

### III. Imitative Behavior and Concerns for Relative profit

Rather than each firm laboriously calculating its own optimal strategy, even if that were feasible, it is plausible that firms may sometimes choose to imitate the strategies of their more successful peers. As Alchian put it:

---

“[W]henver successful enterprises are observed, the elements common to these observable successes will also be associated with success and copied by others in their pursuit of profits or success. [...] What would otherwise appear to be customary ‘orthodox,’ non-rational rules of behavior turns out to be

codified imitations of observed success, e.g., ‘conventional’ markup, [...] ‘orthodox’ accounting and operating ratios, ‘proper’ advertising policy, etc.”<sup>38</sup>

---

Alchian suggests that imitation enables firms to make use of other firms’ private information and optimizing behavior, and to enjoy the benefits of conformity (as imitating firms are likely to do as well as the average of their peers). To discuss these and other points further, it is useful to distinguish between imitation by non-rivalrous firms and imitation by firms competing in the same market.

### A. IMITATION BY NON-RIVALROUS FIRMS

For now suppose that the relevant firms do not interact strategically in a market, e.g., they may be local monopolists, or they may be price-taking firms facing some given exogenous price (such as farmers deciding on an agricultural technique). It is plausible that “optimizing” and “imitating” firms might co-exist, and each firm chooses which of the two broad kinds of strategy to follow. If firms incur a cost for calculating the optimal action, then when most firms are optimizers and firms operate in a similar environment, it may benefit a firm simply to copy these optimizing firms’ actions and save itself the calculation cost. Of course, if too many firms free-ride and copy average observed behavior, there may re-emerge a benefit to investing in optimizing. Since optimizers provide a free service to imitators, in general there will be too few optimizers in equilibrium.<sup>39</sup>

In other situations, firms may be unsure about the optimal action even after significant deliberation, and instead they may have some kind of noisy signal about what is the best action. (For instance, a firm might read a trade report or attend a conference about a new technology in the market.) If firms choose their actions sequentially, and can observe the actions chosen by earlier firms, then a (rational) firm should take into account what other firms did before choosing its own action. In such cases, there is a danger that firms will become locked into taking the wrong action.<sup>40</sup> Such a framework might be used to explain inefficient technology adoption or merger waves by profit-maximizing firms. However, the predicted equilibrium depends on very sophisticated reasoning by firms. (For instance, the third firm needs to work out what the second firm’s action implies in terms of that firm’s private signal, and the required “depth of reasoning” gets progressively more onerous for firms which are further back.) In experiments, it appears that firms do not play this equilibrium, and instead put more weight on their private signal than they should.<sup>41</sup> This could be interpreted as a form of over-confidence, a topic to which we return in section VI.

Career concerns of managers might give a reason why a manager mimics the action of a peer, even if there is no extra cost to calculating the optimal action from scratch. Scharfstein & Stein<sup>42</sup> suggest a model with two managers (who do not compete in the same product market), each of whom may have private infor-

mation about the desirability of choosing action X or action Y. One manager chooses her action first, and the other then chooses after taking account of the first decision. Managers are either “smart” (i.e., they have genuine information about which strategy is better) or “dumb” (i.e., they know nothing), and managers do not know which type they are when they decide on their action. Since their private information is correlated, a pair of smart managers tends to choose the same action. Therefore, if a manager’s subsequent career depends on whether or not she is perceived to be smart, the second manager has an incentive to imitate to the first so as to be considered smart, even if that means going against her own private information. As the authors put it:

CAREER CONCERNS OF MANAGERS  
MIGHT GIVE A REASON WHY A  
MANAGER MIMICS THE ACTION OF  
A PEER, EVEN IF THERE IS NO  
EXTRA COST TO CALCULATING THE  
OPTIMAL ACTION FROM SCRATCH.

“an unprofitable decision is not as bad for reputation when others make the same mistake—they can share the blame if there are systematically unpredictable shocks.”<sup>43</sup> This managerial herding is often socially inefficient.<sup>44</sup>

However, imitation can also be socially useful. A model which several researchers have investigated supposes that each firm usually imitates its best-performing peer, but with some (perhaps small) probability it “experiments” and chooses a random strategy. A firm—if it has access to the necessary information—can observe the actions and profits of the other firms, and when it next has the opportunity to change its own action it can choose the action of the most profitable firm. In a stable and symmetric environment, this process of imitation will likely lead over time to approximately optimal behavior.<sup>45</sup>

## B. IMITATION BY COMPETING FIRMS

Additional effects come into play when firms are rivals operating in the same market, so that firms compete as well as observe.<sup>46</sup> In practice, oligopolists may have little idea of the consumer demand function, or how closely substitutable their rivals’ products are with their own. Nevertheless, they may observe their rivals’ actions and realized profits. When oligopolists imitate the most profitable actions observed in the market, it is possible that the market moves over time to a highly competitive outcome (more competitive than the one-shot Nash equilibrium).

To see this, consider a Cournot market where several firms with identical constant marginal costs compete to supply a homogenous product. Suppose, whenever they have the opportunity to change their output, that firms imitate the output decision of the most profitable firm of the previous period. Then, when the market price is above cost the most profitable firm will be the one with the largest output. Hence, firms with low output will increase their output, imitating the profitable firm, which pushes the price down. (If price is below cost, the most profitable firm is the one with the smallest output, and so imitation will then drive prices up.) Thus, imitation pushes prices towards cost and the market

evolves towards the perfectly competitive outcome where price equals marginal cost. In sum, when firms myopically imitate the most profitable strategy, the industry as a whole moves to an unprofitable, highly competitive outcome. See Vega-Redondo<sup>47</sup> for more details of this model, as well as Schaffer<sup>48</sup> for a related model.<sup>49</sup>

Many markets are better modeled as firms choosing prices rather than quantities, and where firms offer differentiated products. In these markets, imitation can also induce firms to compete aggressively, although not to the extent of perfect competition as seen in the (homogeneous product) Cournot case. To illustrate, consider a duopoly with firms labeled 1 and 2, where if the two prices are  $p_1$  and  $p_2$  then firm 1's demand is

$$q_1 = 1 + p_2/2 - p_1 \quad (1)$$

(and similarly for firm 2). For simplicity, suppose costs are normalized to zero. Then profit-maximizing firms are predicted to choose the Bertrand equilibrium prices  $p_1 = p_2 = 2/3$ , and each firm sells quantity  $2/3$ . One can check that when both prices are above  $1/2$  then it is the lower-price firm which makes greater profit. (When both prices are below  $1/2$  then the higher-price firm has more profit.) Therefore, when the less profitable firm chooses the price of the more profitable firm (and firms experiment every so often as discussed in footnote 45), prices will settle over time at  $p_1 = p_2 = 1/2$ , so that the mark-up on cost falls by one-third with imitative behavior relative to profit-maximizing behavior.<sup>50</sup>

An interesting corollary of imitative behavior is that with symmetric firms the familiar distinction between price-setting and quantity-setting behavior vanishes, and the long-run performance of markets with imitative firms does not depend on whether firms choose to compete in prices or in quantities. With profit-maximizing behavior, when firms compete in quantities the outcome differs—typically it is less competitive—from when they compete in prices. (For instance, in the linear demand example of the previous paragraph, when firms compete in quantities one can show that the equilibrium involves each firm setting the output  $3/5$ , which induces each firm to set the price  $4/5$ , which is higher than when the firms compete in prices.) However, when firms choose quantities the process of imitation converges to the situation where each firm chooses output  $3/4$ , which induces each firm to set the price  $1/2$ , which is exactly the same as when firms chose prices.<sup>51</sup>

AN INTERESTING COROLLARY OF IMITATIVE BEHAVIOR IS THAT WITH SYMMETRIC FIRMS THE FAMILIAR DISTINCTION BETWEEN PRICE-SETTING AND QUANTITY-SETTING BEHAVIOR VANISHES.

As discussed earlier in the context of non-rivalrous firms, it is natural to consider situations where profit-maximizing firms compete against naïve imitators to see which type of behavior performs better. Schipper<sup>52</sup> investigates a Cournot model where “imitators” compete against “optimizers,” where the latter firms are

not really forward-looking profit-maximizers but rather myopic, and choose their output in one period as the most profitable response to the previous period's outputs. He shows that imitators make greater long-run profits than optimizers. A sterner test of the benefits of naïve imitation is when such a firm competes against firms who are forward-looking strategic players. It turns out that in a wide class of games, firms which imitate the best past performers do "essentially" as well as profit-maximizing rivals. For instance, consider a repeated Prisoner's Dilemma in which one firm plays the most profitable actions used in the previous period (starting initially by playing "collude", say). Then, unless there is very severe discounting of the future, the profit-maximizing way to play against this imitator is to collude in all periods except the last (if there is a last period), when

IT TURNS OUT THAT IN A WIDE CLASS OF GAMES, FIRMS WHICH IMITATE THE BEST PAST PERFORMERS DO "ESSENTIALLY" AS WELL AS PROFIT-MAXIMIZING RIVALS.

it is optimal to compete. The profit-maximizing firm therefore obtains identical profits to the imitator, except just for one period (if there is a last period) when it makes higher profits.<sup>53</sup>

Of course, firms can only imitate the most successful firm when they are provided with the necessary market information (namely, the profits of each firm, together with individual outputs in the Cournot case). But as soon as this information is present, a process as described by Vega-Redondo can make markets very competitive. As discussed in section II, Vega-Redondo's paper has inspired several experimental tests,<sup>54</sup> all of which found some support for the basic qualitative prediction of the imitation model. Even when firms have enough information to calculate the Cournot equilibrium, say, if they are then given the extra firm-specific information which allows them to imitate the best, this makes their behavior more competitive. However, the perfectly competitive outcome was not always achieved, indicating that not all subjects were following the "imitate the best" strategy. (An exception is, Huck et al.,<sup>55</sup> who in their treatment denoted "IMIT+" where firms did not know much about the functional forms of market demand or costs but did observe individual rival quantities and profits, found that the outcome was statistically indistinguishable from the perfectly competitive outcome.) Offerman, et al.,<sup>56</sup> suppose that a fraction of firms imitate, while the remainder are assumed to choose the most profitable output given the observed outputs of the previous period. In their data, the best fit for this mixed model is that two thirds of firms are imitators.

### C. CONCERNS FOR RELATIVE PROFITS

There is a close connection between situations in which firms imitate the most profitable action played by their rivals and situations in which firms (or their managers) have as their objective the maximization of relative profits. The reason for this close connection is as follows. Consider the differentiated product duopoly just discussed, and suppose one firm reduces its price to undercut its



rival. If the lower-price firm makes greater profit than its rival, this means that the price cut harms it less than its rival, i.e., the difference between its profit and its rival's rises. In the specific linear demand example in expression (1) above, if firm 1 wishes to maximize its relative profit, it chooses price  $p_1$  to maximize the profit difference

$$p_1(1 + p_2/2 - p_1) - p_2(1 + p_1/2 - p_2),$$

and it is therefore optimal for it to set the price  $p_1 = 1/2$ , regardless of the rival's price  $p_2$ . If both firms care purely about their relative profit, they will both set the price  $1/2$ , just as eventually occurred when firms imitate. Thus, the long-run outcome in markets in which firms imitate coincides with the (one-shot) outcome in markets in which firms are fully rational except that their objective is relative rather than absolute profit.

The importance of this result is that, while purely imitative behavior on the part of sophisticated firms perhaps seems unlikely in many markets, the assumption that firms care about their relative position in a market has strong intuitive appeal.<sup>57</sup> First, there is by now no doubt that many individuals are strongly driven by relative pay. This is evident from the behavioral and experimental economics literature, as well as from the emerging happiness literature.<sup>58</sup> Indeed, CEOs as a group may have a greater proportion of "rivalrous" people than the population as a whole, and such people may put particular weight on their relative standing. Secondly, managers often have placed on them (either explicitly or implicitly) incentives which induce them to care at least in part about relative as well as absolute performance. (One reason for this might be to insure managers against common shocks impacting the market as a whole.) For instance, Gibbons & Murphy<sup>59</sup> document empirically how a CEO's pay rise and likelihood of retention depend positively on the firm's performance and negatively on the overall industry performance. Vickers<sup>60</sup> shows that a firm in a Cournot market can improve its equilibrium profits (both in absolute and relative terms) when competing against profit-maximizing firms by inducing its managers to care about relative rather than absolute profit.<sup>61</sup> For the same reason, if potential managers differ in their intrinsic preferences for absolute or relative profit performance, Miller & Pazgal<sup>62</sup> argue that a firm may wish to hire a manager with a known behavioral bias towards relative profit so as to gain strategic advantage. By contrast, in Bertrand markets rather than Cournot markets, a firm's profits typically fall if its manager cares about relative instead of absolute profits; a manager who cares about relative profits will set a low price, which in turn will induce low prices from its rivals, and each firms' profits decline.<sup>63</sup>

THE IMPORTANCE OF THIS RESULT IS THAT, WHILE PURELY IMITATIVE BEHAVIOR ON THE PART OF SOPHISTICATED FIRMS PERHAPS SEEMS UNLIKELY IN MANY MARKETS, THE ASSUMPTION THAT FIRMS CARE ABOUT THEIR RELATIVE POSITION IN A MARKET HAS STRONG INTUITIVE APPEAL.

So far, we have argued that imitation or concern for relative profits makes the static interaction among oligopolists more competitive, but what does this behavior imply for the sustainability of collusion? Suppose that each oligopolist cares about its profits relative to its rivals. Compared to the setting with absolute profit maximization, there are two effects:

- the static Nash equilibrium—which is used to punish firms if they deviate from the collusive agreement—is more competitive, but
- the short-run benefits to deviating are larger, since by undercutting its rivals a firm not only boosts its own profits but also harms its rivals.

SO FAR, WE HAVE ARGUED THAT IMITATION OR CONCERN FOR RELATIVE PROFITS MAKES THE STATIC INTERACTION AMONG OLIGOPOLISTS MORE COMPETITIVE, BUT WHAT DOES THIS BEHAVIOR IMPLY FOR THE SUSTAINABILITY OF COLLUSION?

The first effect makes collusion easier, while the second makes it harder, and *a priori* it is unclear which dominates. However, in a Cournot setting, Lundgren<sup>64</sup> shows that when firms care purely about relative profit, collusion is impossible to sustain.

We can summarize this discussion as follows. In markets where firms offer substitute products (i.e., in which firms compete), the market is made more competitive relative to the textbook situation of firms maximizing profits if firms either (i) imitate the previous most profitable strategies or (ii) aim to maximize their relative profits. The (eventual) outcomes under scenarios (i) and (ii) coincide, and in both cases the outcome is the same whether firms compete in prices or quantities. In mixed settings, where imitators (or relative profit maximizers) compete against profit-maximizing firms, the former will often perform in absolute terms at least as well as the latter.

## IV. Vengeful Behavior and *Esprit de Corps*

The previous section discussed situations in which firms and their managers care about relative rather than absolute performance. Another kind of social preferences is present when firms care when their rivals obtain an “unfair” share of industry profits, for instance by cheating on a collusive agreement.

It is clear that many people are willing to incur costs in order to harm others who are perceived to have behaved unfairly towards them, a phenomenon which could be called vengeance or spite. Among the most famous and robust experiments in economics are those that study the “ultimatum game.”<sup>65</sup> Here, two players must share some specified prize, and one player (the “proposer”) suggests a way to share the prize between the two players. If the second player (the “responder”) agrees with the proposed shares, the prize is divided accordingly, while if the responder does not agree then neither player gets anything. If players are interested in obtaining as much of the prize as possible, the predicted outcome

(from a one-shot interaction) is that the proposer offers only a tiny share to the responder, who accepts since she prefers a small positive payoff to nothing. However, it is commonly observed that the responder will reject offers she finds too small, e.g., less than 25 percent of the prize. Because of this, proposers learn or anticipate that it is too risky to make such unfair offers, and outcomes where the proposer offers the responder 40 or 50 percent are often the norm.<sup>66</sup> These experiments have been performed (in poorer countries) when the prize is very significant in terms of monthly salary, and the results are similar.

The same effect is observed in markets. Huck, Muller, & Normann<sup>67</sup> show that a firm with a theoretical commitment advantage (a Stackelberg leader in a Cournot market) finds it hard to exploit that advantage in experimental markets. The reason is that the theoretically disadvantaged firm (the Stackelberg follower) acts more aggressively than predicted by the subgame perfect equilibrium of these market games. In fact, followers appear to punish the leader when the leader supplies a quantity above the symmetric Cournot quantity. (They punish the leader by themselves supplying a higher quantity than their most profitable response to the leader's quantity.) This behavior is in line with the vengeful behavior seen in the ultimatum game when the proposer tries to exploit his first-mover advantage.

Similar behavioral effects are observed in experiments that study strategic delegation. As already mentioned, Vickers<sup>68</sup> showed theoretically that firm owners might want to employ managers simply for strategic reasons. By writing an appropriate incentive contract for the manager, the owner can in effect commit his firm to Stackelberg-like aggressive behavior. However, if all firms do this a highly competitive outcome results: All firms are worse off than when managers care solely about profit. Huck, Muller, & Normann<sup>69</sup> tested this theory in an experiment, and do not confirm its predictions. If an owner does offer his manager the aggressive contract (and other owners do not), then managers in the weaker position are not content with simply accepting their “equilibrium fate” which would give them a lower salary than their opponent. Rather they behave more aggressively, in line with the observed behavior of the Stackelberg followers mentioned in the previous paragraph. This could mean that attempts to hire “behavioral” types of managers, as in Miller & Pazgal,<sup>70</sup> may also backfire.

A NATURAL QUESTION IS  
WHETHER VENGEANCE MIGHT  
HELP TO SUSTAIN COLLUSION

A natural question is whether vengeance might help to sustain collusion. For instance, if the collusive arrangement (a high price and equal market shares) is considered to represent the “fair” outcome, then if one firm reneges on the agreement and undercuts the price, its rivals may be offended and hence punish the deviator especially aggressively (even at extra cost to themselves). In the previous section we argued that concern for relative profits did not help sustain collusion, since it induced an extra benefit to deviating. But vengeance is subtly different from a concern for relative profit, since it only comes into effect when

someone is treated unfairly. (There is no extra benefit from harming your rival, unless he has first harmed you.) Thus, because vengeance induces more aggressive punishments for deviating, it could help sustain collusion relative to situations in which firms care only about their own profits.<sup>71</sup>

While the previous section showed how concerns for relative performance could increase competitiveness in a market (thus benefiting consumers), these examples show how vengeful behavior can reduce competitiveness (to the detriment of consumers). In settings where firms are symmetrically positioned, concerns for relative profits increase competition. On the other hand, in environments where markets are predicted to be particularly competitive due to strategic asymmetries (as in the Stackelberg model), vengeance against unfairly advantaged rivals tends to mitigate these advantages, rendering market outcomes more symmetric and less consumer-friendly.

The impact of vengeful behavior is less likely to be present when agents are very asymmetric or if some are perceived to be entitled to a bigger share of the surplus.<sup>72</sup> Thus, the potential for spiteful behavior may affect conduct between rival firms, but not so much between a firm and its consumers, for instance, or between a large supermarket and a small supplier.

WHILE VENGEFUL BEHAVIOR MAY SOMETIMES BE A WAY TO SUSTAIN COLLUSION, AN ALTERNATIVE METHOD IS TO FOSTER A SENSE OF LOYALTY, OF *ESPRIT DE CORPS*, AMONGST THE CONSPIRATORS.

Consider modifying the ultimatum game so that there are many responders. (More precisely, the single proposer offers a share, and the various responders simultaneously decide whether to accept. If some responders accept, one is picked at random and given her offered share; if none accepts then no one gets anything.) Then orthodox theory predicts—and experiments confirm—that the single proposer will be able to offer the responders very little and still find a willing responder. In a sense, competition forces the responders to act as if they were purely self-interested, and so “vengeance” is a rent which competition dissipates.<sup>73</sup>

If some responders accept, one is picked at random and given her offered share; if none accepts then no one gets anything.) Then orthodox theory predicts—and experiments confirm—that the single proposer will be able to offer the responders very little and still find a willing responder. In a sense, competition forces the responders to act as if they were purely self-interested, and so “vengeance” is a rent which competition dissipates.<sup>73</sup>

While vengeful behavior may sometimes be a way to sustain collusion, an alternative method is to foster a sense of loyalty, of *esprit de corps*, amongst the conspirators. Although information sharing is an important ingredient for collusive schemes to work, a useful by-product of regular meetings is that loyalty and friendship may be inculcated, which may make it socially costly to cheat on agreements. Moreover, since illegal cartels cannot enforce agreements with legally binding contracts, trust plays a central role in their operation.<sup>74</sup> (Recall from section II that face-to-face communication seemed to foster collusion in the laboratory more effectively than computer-mediated communication.) An important role for antitrust, via leniency programs and the like, is indeed to foster distrust among conspirators.

A clear example of an attempt to run a cartel (or cartel-like operation) by means of encouraging a sense of camaraderie is the U.S. steel industry during the years 1907-1911, in which cooperation was fostered through a series of social events and meetings which have become known as the “Gary Dinners” after the chairman of U.S. Steel, Judge Elbert Gary. It is worth quoting one of Gary’s speeches at length:<sup>75</sup>

---

“[W]e have something better to guide and control us in our business methods than a contract which depends upon written or verbal promises with a penalty attached. We as men, as gentlemen, as friends, as neighbors, having been in close communication and contact during the last few years, have reached a point where we entertain for one another respect and affectionate regard. We have reached a position so high in our lines of activity that we are bound to protect one another; and when a man reaches a position where his honor is at stake, where even more than life itself is concerned, where he cannot act or fail to act except with a distinct and clear understanding that his honor is involved, then he has reached a position that is more binding on him than any written or verbal contract.”

---

In essence, if conspirators can find a way to increase the social or psychological cost of cheating, collusion will be more readily achieved.

Further historic evidence of the impact of social preferences on anticompetitive behavior comes from the UK shipping cartels in the period 1879 to 1929. Podolny & Scott Morton<sup>76</sup> document how established cartels behaved significantly more aggressively toward entrants who had low social status (or who were foreign), relative to the situation where an entrant had a knighthood, say. (The study controls for the correlation between social status and “deep pockets,” which may make predation less profitable.) The authors argue that the likely reason is not so much pure snobbishness on the part of the cartel (who typically comprised high status individuals), but rather that social status was used as a signal of the entrant’s likelihood of cooperating within the cartel, and the expected “transactions costs” of having the entrant as a member of the cartel, if admitted.

## V. Satisfactory, Not Maximum, Profits

Rather than denying the importance of (absolute) profits, one might instead question whether firms really *maximize* profits. For reasons of complexity, ignorance, or the “easy life,” firms might instead engage in satisficing behavior to secure a target level (or “aspiration level”) of profit. In its starkest form, the util-

ity of a firm then has just two values: good enough, and not good enough. This idea goes back at least to Rothschild, Gordon, and Simon.<sup>77</sup> Just as a consumer might not change her bank, say, until the perceived level of service quality she receives falls below some threshold, a firm might not revise its strategy while it continues to obtain satisfactory profits. Only if the realized profit falls below the target level will a firm resort to experimenting with an alternative strategy. Such behavior could give rise to a degree of price rigidity for instance, even if underlying cost or demand conditions vary.<sup>78</sup> Experimentation could be purely random in extreme cases, or firms may be able to target their attention towards new strategies which will likely boost profits. (For instance, an idea in the early literature on satisficing was that firms often had a degree of “managerial slack,” and when times were hard managers could focus their energies on cutting costs or expanding their markets.) What level of profits a firm considers to be “satisfactory” is likely to depend on its historical returns as well as the performance of its peers and the economy as a whole.

Cyert & March<sup>79</sup> discuss some implications of satisficing behavior and test these ideas empirically analyzing data from (a few) manufacturers of farm implements. They conclude that a firm is more likely to resort to aggressive strategies (seeking cost reductions or sales expansions) if it is operating close to its

RATHER THAN DENYING THE IMPORTANCE OF (ABSOLUTE) PROFITS, ONE MIGHT INSTEAD QUESTION WHETHER FIRMS REALLY MAXIMIZE PROFITS.

breakeven point or if its costs are higher than the industry average; behavior which accords with the idea that firms are most likely to change strategy after they realize low profits.

Huck, Konrad, Muller, & Normann<sup>80</sup> document related effects in an experiment concerning mergers in Cournot markets. As shown by

Salant, Switzer, & Reynolds,<sup>81</sup> bilateral mergers in linear Cournot markets with more than two firms are predicted to reduce the joint profits of merging firms when firms are profit-maximizers. Since this prediction appears counterintuitive at first sight, the result has been dubbed the “merger paradox.” However, the intuition for the result is simple. In Cournot markets firms’ outputs exert negative externalities on their rivals. Hence, after a merger the two insiders have an incentive to internalize some of these negative externalities; that is, they have an incentive to reduce their combined output. (Obviously, they would keep their pre-merger profits if they simply kept their output constant, but this is not the most profitable response to the un-merged firms’ outputs.) Further, as Cournot interactions are games with strategic substitutes, the un-merged firms’ optimal responses to this contraction in output are to increase their own outputs (in response to which the merged firm will reduce its output even further, and so on). This leads into a new equilibrium where the market price is higher, the outsiders are better off, and the two merging insiders are worse off. The reduced market share of the merging firms (from  $2/n$  to  $1/(n - 1)$  if there are  $n$  firms initially) outweighs the benefits from increased market concentration.<sup>82</sup>

Huck, et al. test this prediction in laboratory markets with initially three or four firms. Mergers are implemented exogenously after 25 periods in which firms learn to play the Cournot equilibrium. Firms are randomly selected to merge and then play for a further 25 periods. The merged firm is managed by one of the two subjects playing the role of the two merging firms. Profits are shared equally between the manager-owner and the passive owner. After the merger, profit-maximizing firms are predicted to behave as if one of the firms had vanished, and in particular, a merged firm is predicted to choose the same output as the non-merged firms. However, a different behavior emerges in the laboratory; the merged firm systematically supplies more than the outsiders. (The outsiders are seen to respond approximately optimally to the behavior of the merged firm.) This implies that the losses from the merger are smaller than predicted. In fact, when there are four firms initially present the merged firm does not lose at all. Rather, it experiences a significant increase in profits for a few periods and then, when the outsiders are fully adjusted, its per-period profits fall slightly compared to their pre-merger levels.

Huck, et al. discuss a number of reasons for this effect, and by conducting suitable control experiments they suggest something akin to satisficing behavior explains observed behavior.<sup>83</sup> After two firms merge, the firms perceive the danger that their joint profits fall, and pursue aggressive strategies to avoid this (as the firms in Cyert & March did). The result is that the merged firm's output is significantly higher than its rivals' outputs, even though post-merger a merged firm and an unmerged firm are in a symmetric position. (However, the merged firm's output is lower than their combined outputs in the pre-merger phase, and so their strategies are not entirely inert.)

Dixon and Oechssler discuss an interesting application of satisficing behavior.<sup>84</sup> They consider a number of oligopolies (rather than just one), and suppose that firms follow the following rule of thumb: "If my profits are no lower than the average profits observed across all markets, I keep my strategy unchanged; if my profits are lower than average profits across all markets, I experiment and choose a random strategy." Thus, firms use the average payoff of all firms as their aspiration level. Such a policy requires information about average profitability, but no information about individual firm actions or profits or about consumer demand. It turns out that when firms behave in this manner (with a small amount of noise added) then behavior in each market eventually becomes collusive. Consequently, practices that look innocent from the orthodox point of view—myopic adjustments towards better strategies—may lead to undesirable outcomes. There can be collusive effect without any collusive intent.<sup>85</sup>

To understand this surprising result, consider for simplicity a pair of identical duopolies (i.e., there are four firms in all), which interact over time in a Prisoner's Dilemma manner using the two possible actions of "compete" or "collude." Suppose that all firms are initially colluding, in which case all firms are satisfied with their profits. Next, imagine one firm "trembles" and changes its

action to “compete.” A duopoly in which both firms collude has higher profits than a duopoly where one firm colludes and the other competes, and so firms in the colluding duopoly remain satisfied with their profits and do not experiment. The only firm dissatisfied with its profit is the colluding firm playing against the competing firm, which will experiment and end up playing “compete.” We then have two firms in a market playing “compete,” both of whom will be dissatisfied with their profits, and will experiment and end up both playing “collude.” Thus all firms colluding is the stable steady state of this process (unlike all firms competing, where one “tremble” will eventually induce all firms to collude). In contrast to Adam Smith’s invisible hand, which guides profit-maximizing firms

IN CONTRAST TO ADAM SMITH’S  
INVISIBLE HAND, WHICH GUIDES  
PROFIT-MAXIMIZING FIRMS  
TOWARDS OUTCOMES WHICH  
BENEFIT CONSUMERS, WITH THIS  
FORM OF SATISFICING BEHAVIOR, A  
SECOND INVISIBLE HAND GUIDES  
FIRMS TOWARD OUTCOMES  
WHICH EXPLOIT CONSUMERS.

towards outcomes which benefit consumers, with this form of satisficing behavior, a second invisible hand guides firms toward outcomes which exploit consumers.

A concept related to satisficing behavior is “approximately” optimal behavior. Firms may not find it worthwhile to calculate their optimal strategy precisely, but instead to cease their search for a good strategy when they get to within some tolerance of the optimal strategy. In technical terms, the outcome of approximately optimal behavior is termed an “ $\epsilon$ -equi-

librium,” where  $\epsilon > 0$  is the (perhaps small) tolerance which firms have for short-fall in optimal profits. An important insight is that even a small tolerance by firms for sub-optimal strategies—that is, a small departure from fully optimizing behavior—can result in significant departures from the outcomes corresponding to optimizing behavior. (In the neighborhood of a firm’s optimal strategy its profits are roughly flat, and so a moderate change in its strategy might have little impact on its own profit, and yet could have a significant impact on its rivals’ profits and strategies.)

To illustrate, consider two symmetric firms supplying a homogeneous product engaged in quantity competition. To be concrete, suppose that the consumer demand curve is  $q(p) = 1 - p$  and that production is costless. Then the symmetric collusive outcome in this market involves each firm supplying quantity  $1/4$  inducing the monopoly price  $p = 1/2$ . How small can a firm’s tolerance for sub-optimal behavior be for this monopoly outcome to be an  $\epsilon$ -equilibrium? The most profitable response to the rival’s monopoly output of  $1/4$  is to supply output  $3/8$  rather than  $1/4$ . However, its gain in profit from pursuing the optimal strategy rather than the satisficing strategy of supplying output  $1/4$  is only  $1/64$ . Thus, the collusive outcome (where each firm supplies quantity  $1/4$ ) is an  $\epsilon$ -equilibrium provided that  $\epsilon > 1/64$ . Note that  $1/64$  is about 6 percent of a firm’s share of the monopoly profit, and so if firms are prepared to optimize to within 6 percent of their exact optimal profits, the monopoly outcome can be sustained without



any collusive intent.<sup>86</sup> Of course, though, a problem with the notion of  $\varepsilon$ -equilibrium is its lack of predictability: firms could also accidentally play more competitively than the Cournot duopoly equilibrium.

Baye & Morgan<sup>87</sup> present an interesting account of how almost-optimal behavior can lead to significant departures from the predictions of fully-optimal models. They analyze a static homogenous product Bertrand market, so that profit-maximizing firms are predicted to set prices at marginal cost. They obtain an  $\varepsilon$ -equilibrium involving mixed strategies for choosing prices that yields relatively high profits even for small  $\varepsilon$ . (With duopoly, when  $\varepsilon$  is just 1 percent of monopoly profit, their chosen  $\varepsilon$ -equilibrium yields expected profit which is more than 25 percent of the monopoly profit.)

Baye & Morgan also discuss an alternative model of bounded rationality, which is that firms play a so-called quantal-response equilibrium.<sup>88</sup> This model supposes that a firm is more likely to choose a price which yields higher profit, but is not sure to choose the most profitable price. (This model nests the fully rational model and a model with purely random behavior as special cases.) While the details of the equilibria in the two models of bounded rationality differ, the broad conclusions—that prices are above cost, profits are positive, and profits fall with more rivals—coincide. Baye & Morgan run laboratory experiments to generate data, and estimate which model of firm behavior best fits the data. Especially for duopoly, they find that both the  $\varepsilon$ -equilibrium and the quantal-response equilibrium model fit the data better than either of the extreme fully-rational or fully-random models.<sup>89</sup>

## VI. Other Topics

### A. OVER-OPTIMISM

As is well-documented in the psychology literature, over-optimism (or over-confidence) about one's own ability or about the probability of favorable outcomes is apparently common in the population.<sup>90</sup> Adam Smith<sup>91</sup> wrote: "the chance of gain is by every man more or less over-valued, and the chance of loss is by most men under-valued." There are good reasons for thinking over-optimism is still more common among entrepreneurs than the population as whole. There is a "winner's curse" aspect to launching a new business: It is likely that others have already thought about launching a similar product (be it a new restaurant in a particular locale, or something more ambitious), and the entrepreneur who actually decides to start the new business is likely to be more optimistic than others. As a result, even if beliefs about the likely return from the investment are unbiased on average, the entrepreneur will typically be overly optimistic. Unless entrepreneurs rationally take full account of the fact that others have decided not to enter this market, we expect to see: (i) high failure rates for new businesses, and (ii) credit rationing, or loans being offered only with collateral.<sup>92</sup>

Internal promotion procedures may also have a tendency to favor the over-optimistic, so that CEOs as well as entrepreneurs may be disproportionately over-optimistic. Consider a situation where management promotions are driven through rank-order tournaments, as in Lazear & Rosen.<sup>93</sup> To be promoted to a higher level, managers tend to require both skill and luck. In Lazear & Rosen's original model luck was exogenous noise, but the model can be extended so that managers can affect the riskiness of projects they undertake. The consequence is that the best performing manager in a tournament is likely to have high skill and to have chosen a risky set of projects. Over-optimism in a manager might mean that that manager downplays the true riskiness of projects. In this setting, managers who rise to the top of a firm are likely to be the highly-skilled optimists who were lucky.<sup>94</sup>

Using a more abstract selection mechanism, Heifetz, Shannon, & Spiegel<sup>95</sup> also argue that optimists will systematically outperform realists in competitive environments, and so will predominate in the pool of successful agents. The key intuition is that optimism can serve as a commitment device if observed by rivals (much as making a manager maximize relative profits does in Vickers).<sup>96</sup> Translated into a market context, Heifetz et al.'s analysis suggests that particular forms of managerial over-optimism may sometimes act to soften competition.

For instance, consider a Bertrand oligopoly with differentiated products where there is some uncertainty about demand. The more optimistic a manager is—say, about the scale of demand, or the extent of product differentiation among suppliers—the higher the price she will charge. If a manager's optimism is observable or if competitors at least have some informative signal about her degree of optimism, they will rationally anticipate these higher prices and optimally adjust their own prices upwards as well. Hence, one optimist is enough to increase the prices charged by all firms. And crucially the optimist will earn more money for her firm than the realist would have done (although her rivals will earn still more).<sup>97</sup>

IN WHAT IS PERHAPS THE BEST-KNOWN PAPER ABOUT MANAGERIAL IRRATIONALITY, ROLL SUGGESTS THAT EXCESSIVE MERGER ACTIVITY MAY BE GENERATED BY MANAGERIAL OVER-CONFIDENCE.

In what is perhaps the best-known paper about managerial irrationality, Roll<sup>98</sup> suggests that excessive merger activity may be generated by managerial over-confidence. In his theory, a CEO is too confident in the accuracy of his information about the potential profitability of a takeover, and so will be too quick to launch a takeover bid. Predictions of this theory are that the combined gain to bidder and target will be close to zero, and that the bidder's value will fall on announcement of the bid. Roll suggests that the available empirical evidence does not reject these predictions. Over-confidence by CEOs in their ability to choose investment projects and merger opportunities has been detected by Malmendier & Tate.<sup>99</sup> CEOs will differ in the extent of their overconfidence, and Malmendier & Tate identify as over-confident those CEOs who hold on to their stock options

until they expire. They interpret this as evidence of the CEO's over-estimation of their firm's future performance. They show that this group of CEOs is 55 per cent more likely to undertake a merger.<sup>100</sup> Managerial over-confidence could explain why companies that undertake mergers seem to under-perform.

Another form of potential over-confidence concerns a manager's faith that he will not be caught if he pursues an illegal activity such as price-fixing.<sup>101</sup> If many cartelists downplay the risk of detection, then policy measures to increase the penalties for illegal cartel activity could be less effective than would be predicted by an approach based on rational profit-maximizing behavior.

ANOTHER FORM OF POTENTIAL OVER-CONFIDENCE CONCERNS A MANAGER'S FAITH THAT HE WILL NOT BE CAUGHT IF HE PURSUES AN ILLEGAL ACTIVITY SUCH AS PRICE-FIXING.

There is also a rich literature in finance about whether overconfident securities traders can survive in the long term or whether, a la Friedman, they are driven out by more rational traders. If an overconfident trader underestimates the riskiness of an asset, she will buy more of the asset than rational traders. If the risky asset also has higher expected return, then the over-confident traders can become wealthier than more rational investors (although their expected utility is lower).<sup>102</sup> If an over-confident trader believes that her estimate of the expected return of a risky asset is more precise than it really is, such a trader will trade more aggressively than an unbiased trader, and since she trades based on useful information, her expected wealth is higher than unbiased traders (although her expected utility is lower). As a result, over-confident traders can persist in, and even dominate, the market.<sup>103</sup> In such cases, Darwinian selection actually selects the biased traders. Kyle & Wang<sup>104</sup> present a model with just two traders, one of whom is over-confident and known to be so by the rational trader. They show, in a very similar manner to our previous discussion about Heifetz, et al., that the over-confident trader gains strategic advantage by being known to trade aggressively, which induces the rational trader to scale back her own trades. The result is that the overconfident trader performs better than the rational trader.

## B. ACCOUNTING ANOMALIES

It is a fundamental tenet of profit-maximizing behavior that fixed and sunk costs, while they are important for entry and exit decisions, should not play a role in the determination of prices to customers. For instance, competition authorities do not always put much weight on claimed synergies between merging firms which act to reduce fixed costs, since such synergies are not expected to feed through into lower prices. However, this policy seems to be widely flouted by managers. Long ago, Hall & Hitch<sup>105</sup> interviewed 38 business executives about their methods for setting prices. Instead of equating marginal revenue to marginal cost, the authors concluded that: "... there is a strong tendency among business men to fix prices directly at a level which they regard as their 'full cost'." Al-

Najjar, Baliga, & Besanko<sup>106</sup> describe a number of more recent surveys that report similar attitudes among managers, the majority of whom claim to take fixed and sunk costs into account when setting prices. (They also report a managerial accounting textbook which argues against basing prices on marginal costs.)

Experimental work confirms that supposedly irrelevant sunk costs can have an impact on how prices are actually chosen. Offerman & Potters<sup>107</sup> conduct an experiment to investigate pricing in a Bertrand-type duopoly market with product differentiation. In one treatment, there are no sunk costs and pricing behavior is observed to converge to the Bertrand equilibrium. In a second treatment, participants must pay a sunk entry fee to join the market. (This is designed to model auctioning a license for the right to enter a market.) In this second treatment, the average mark-up of prices over marginal cost is substantially higher than in the first treatment. Interestingly, when they perform the same pair of experiments but with monopoly instead of oligopoly, there is no impact of sunk costs on the chosen monopoly price (which is observed to be close to the profit-maximizing level). This is somewhat reminiscent of our discussion of communication costs in section II, where we saw that

EXPERIMENTAL WORK CONFIRMS  
THAT SUPPOSEDLY IRRELEVANT  
SUNK COSTS CAN HAVE  
AN IMPACT ON HOW PRICES  
ARE ACTUALLY CHOSEN.

firms were more inclined to stick to a non-binding collusive agreement if they incurred costs to reach that agreement.

This observation is consistent with Al-Najjar, et al., who present a theoretical model to show how the use of full-cost pricing policies might persist in the long-run in oligopoly markets. The main result is by now a familiar one in this survey: By introducing a behavioral bias in managerial decision-making, a firm can gain strategic advantage. In more detail, Al-Najjar et al. suppose that firms compete in a Bertrand market with product differentiation. In such a market, if a manager somehow commits to set a high price, its rivals will also set a high price, and all firms will make higher profits.<sup>108</sup> (The effect is akin to hiring an over-optimistic manager or, as previously discussed, placing a suitable incentive scheme on the manager.) But a manager who bases prices in part on fixed and sunk costs effectively commits to set a high price, and so hiring a manager who practices this “naïve” pricing policy (or instilling a corporate culture where this form of pricing is used) boosts the firm’s profits.<sup>109</sup> The effect is akin to the strategic tax policy analyzed by Eaton & Grossman,<sup>110</sup> where a country has an incentive to tax the output of a home firm in order to relax competition with a foreign rival. If, for whatever reason, firms take fixed costs into account when setting their prices, this has potentially important implications for merger policy when the merger is expected to generate only synergies in fixed costs. Such mergers would in fact lead to lower prices than would be predicted in a profit-maximizing world.

Other alleged accounting anomalies might be explained by similar myopic learning or evolutionary pressures, or reasons of strategic delegation. For

instance, a firm might be organized into separate profit centers, each of which is given the task of maximizing its own profit (despite the competitive or contracting externalities which might exist between these profit centers). Thus, a manufacturer might choose to supply its products through an independent retailer rather than sell directly to consumers via an integrated retailer. To see one reason why this might be so, despite the apparent dangers of double marginalization, suppose that two manufacturers are competing to sell to consumers. If they sell directly to consumers, then in the absence of collusion we expect to see the Bertrand equilibrium emerge. However, if each firm delegates its retailing operation to a separate division (and each retailing division sees the terms at which the rival retailer sources its product), then by setting a wholesale price above its production cost a firm can induce its retailing division to price high, which softens competition and boosts profits.<sup>111</sup>

### C. UNCERTAINTY ABOUT THE RATIONALITY OF RIVALS

Even if a firm is fully rational, its behavior will be affected when it believes that its rivals may not be rational. Because of this, a rational firm may have an incentive to mimic non-rational behavior so as to induce its rival to think it may be a non-rational type of firm. Consequently, even the potential for behavioral biases can have a substantial effect on market outcomes.

CONSEQUENTLY, EVEN THE POTENTIAL FOR BEHAVIORAL BIASES CAN HAVE A SUBSTANTIAL EFFECT ON MARKET OUTCOMES.

As we have discussed already, if two rational firms (who know for sure that each other is rational) play a Prisoner's Dilemma game for a known finite number of rounds, with the two actions "collude" and "compete," there is no way to sustain collusion even at the start. But if one firm is unsure about the rationality of its rival, it may be unwise to "compete" in each round. For instance, suppose the rival is for some reason believed to be using the strategy of "tit-for-tat," i.e., this rival starts off colluding, and then, in each subsequent round, imitates the previous action of the rival. (Note that the "tit-for-tat" strategy cannot be a rational strategy for a profit-maximizing firm, since in the final round it is a dominant strategy to "compete" regardless of the rival's penultimate action.) If one firm thinks it is playing against a "tit-for-tat" firm, it will typically be optimal to collude for initial periods so as to induce the rival to continue with collusion. Only toward the endpoint will the rational firm start to compete. A tempting strategy is for one or both players to mimic the behavior of "tit-for-tat," so as to induce the rival to believe she is more likely to be this irrational type. This will induce collusive behavior for some rounds.<sup>112</sup> Thus, the (perhaps small) possibility that one or both firms are irrational can induce play that is better for both firms (and worse for consumers).<sup>113</sup>

Another way in which a rational firm might wish to mimic an irrational firm is in the context of predatory pricing, as analyzed in the model of Kreps & Wilson.<sup>114</sup> If a monopoly incumbent faces a sequence of potential entrants to its

market, it may wish to establish a reputation for fighting entry whenever it occurs, even when fighting is actually more costly than accommodating entry. The predictions of the Kreps-Wilson model of predatory pricing are broadly (i) there is little entry in the early period, (ii) when there is entry in early periods, the incumbent “fights” even if it is a rational firm, and (iii) towards the endpoint there is more entry and less willingness to fight on the part of the rational incumbent. This model, which involves highly sophisticated reasoning on the part of firms, is tested in an experiment by Jung et al.<sup>115</sup> The experiment found widespread predatory pricing—defined to be either no entry in the early stages, or a rational incumbent fighting entry if it does occur in the early stages—although the more detailed predictions of the Kreps-Wilson model did not fit the data. (For instance, experimental subjects did not enter more frequently if there had been no previous entry as compared to the situation where there had been fought entry, whereas the Kreps-Wilson model predicts there should be more entry in the former case.)

Rather than go to the trouble of mimicking the behavior of behavioral type of firm, which in any case may only deter entry for some of the time, a more straightforward method to deter entry could be simply to hire an “aggressive” manager, who is known by potential entrants intrinsically to wish to fight entry whenever it occurs. (For instance, an aggressive manager might be someone who is over-optimistic about the ease by which entry can be successfully fought, and so is willing to fight more often than an unbiased manager would.) The impact of this policy, assuming it is credible to keep this aggressive manager in place after entry occurs, is akin to hiring a “conservative” central banker who is less

THUS, IT CAN BE RATIONAL  
TO BEHAVE IRRATIONALLY.

likely to be tempted to cause inflation when unemployment rises.<sup>116</sup> As Schelling<sup>117</sup> wrote: “the conspicuous delegation of authority to a military commander of known motivation, exemplifies a common means of making credible a response pattern that the original source of decision might have been thought to shrink from or to find profitless.” Thus, it can be rational to behave irrationally.<sup>118</sup>

## VII. Discussion

This paper has discussed a number of reasons why firms might not pursue maximum profits:

1. In some complex, uncertain environments, the optimization problem is simply too hard, and firms must resort instead to satisficing and the use of rules of thumb. Decision shortcuts included imitating the actions of well-performing peers, satisficing, or being content to achieve profits to within “ $\epsilon$ ” of the maximum. These rules of thumb appear attractive since in many situations that require strategic sophistication, such as herding, reputation-building, or collusion among

many firms, the complex strategies needed to maximize profits are rarely observed in even in the most simplified laboratory settings and even if participants are highly experienced.

2. Alternatively, optimization might still occur, but with alternative aims or under mistaken beliefs. Thus, a manager might maximize her profits relative to those of her peers, or a manager could be over-optimistic about the profitability of some action. The reason why managers have aims different from maximizing profits could be due to selection effects (e.g., only “competitive” or over-optimistic people rise to become CEOs, or because firms which aim to maximize their relative standing actually obtain greater absolute profits than their profit-maximizing rivals), or because profit-maximizing principals choose to give their managers distorted incentives to gain strategic advantage.
3. Social preferences (other than caring about relative profit as above) may play a role, and a firm might punish a rival if that rival obtains an “unfair” share of profits. Alternatively, face-to-face communication between firms, or having a similar social background among firms, may generate solidarity among conspirators that makes it socially hard to cheat on collusive agreements.

IN SOME SITUATIONS, MARKETS  
ARE MORE COMPETITIVE  
WHEN FIRMS DO NOT AIM TO  
MAXIMIZE THEIR PROFITS.

In some situations, markets are more competitive when firms do not aim to maximize their profits. For example, if firms in a market myopically imitate the action of the most profitable rival, then the market may paradoxically move towards a highly competitive outcome. Alternatively, we saw at least in laboratory settings that firms were often unable to achieve tacit collusion, despite this being an equilibrium option for profit-maximizing firms. Some of the more complex strategies that foster collusion in theory are perhaps too subtle to matter empirically.

In other situations, when firms cannot maximize profits, their realized profits are actually increased. For these firms, following Voltaire’s dictum, the best is the enemy of the good. This can be clearly seen in the case of static Bertrand competition with homogenous products. Here, the only equilibrium involves firms setting prices equal to marginal costs, leaving them with no profits. But if satisficing firms are content to choose actions which are only approximately optimal, they may all be able to enjoy substantial profits. Likewise, when satisficing firms alter their actions only when they under-perform relative to average performance, the result may be as if firms were successfully colluding.

If a firm adopts a non-standard objective, it may gain strategic advantage in the market since competitors’ behavior will change in a desirable way. As such, it could be rational to behave irrationally. For instance, in Cournot markets a firm which aims to maximize its relative profits may do better in equilibrium

than a profit-maximizing rival. Alternatively, a firm which chooses to base its price on “full costs” rather than marginal costs may do better than if it followed textbook profit-maximizing precepts. Or a firm led by an aggressive manager may deter entry more often than a profit-maximizing manager. These non-standard objectives could be put in place by far-sighted profit-maximizing shareholders (as emphasized in the strategic delegation literature) or they could arise myopically due to evolutionary selection of better performing managers and/or firms.<sup>119</sup> Regardless of the reason why firms have non-profit objectives, the presence of non-profit objectives is likely to affect competitive interaction. For instance, a manager operating under an incentive scheme which rewards relative performance is likely to behave more competitively than orthodox analysis would indicate, and this is a factor which a competition authority should take into account.

FOR INSTANCE, A MANAGER OPERATING UNDER AN INCENTIVE SCHEME WHICH REWARDS RELATIVE PERFORMANCE IS LIKELY TO BEHAVE MORE COMPETITIVELY THAN ORTHODOX ANALYSIS WOULD INDICATE, AND THIS IS A FACTOR WHICH A COMPETITION AUTHORITY SHOULD TAKE INTO ACCOUNT.

(The effect is the reverse of the situation when there is cross-ownership in an oligopoly, which competition authorities already recognize leads to a blunting of competition.)

Thus we see there are several situations in which Friedman’s Darwinian critique of non-profit maximizing behavior appears to fail. Nevertheless, there are a number of situations in which market competition and market experience do seem to diminish those behavioral biases which do not confer evolutionary advantage. For instance, competitive versions of the ultimatum game appear in experiments to conform to more orthodox models of selfish behavior. In addition, as shown in the work of List,<sup>120</sup> market experience can dampen the bias known as the endowment effect, where the valuation of a good increases when it is owned. Such a bias cannot easily improve own performance in markets with many traders (rather it hinders agents from making otherwise beneficial trades). Likewise, we are not aware of many situations in which procrastination, impulsive behavior, and other self-control problems play a major role in firm behavior (e.g., price fixing is not a “crime of passion”), although these behaviors are prominent in discussions of behavioral economics as applied to consumers.

We have seen a number of reasons why firms may not maximize their profits, and this potentially has implications for empirical studies of markets, including the use of merger simulation for competition policy. Empirical market studies typically assume profit-maximizing behavior on the part of firms to produce their estimates of, say, marginal costs. In such cases, the analysis may lead to biased estimates of the welfare impact of a merger if, in fact, the merging firms were not profit-maximizers. To illustrate, consider the merger situation in Huck et al.<sup>121</sup> which we discussed in section V. If one took the data from this experiment pre- and post-merger and, in line with the structural approach, assumed these data were generated by profit-maximizing firms, one would conclude that the merger



must have induced substantial cost synergies. Only with reduced costs would it be possible for the merged firm to have higher outputs than the outsiders. This would affect estimates of the merger's welfare consequences. While the true consequences are unambiguously negative (after all, in the experiment the merged firm does operate with the same costs), sufficiently big synergies could offset the loss in consumer welfare. Hence, if there are systematic deviations from orthodox profit-maximizing assumptions, structural approaches that assume profit-maximization might detect increases in welfare when, in fact, welfare is reduced. Likewise, a merger simulation exercise may be inaccurate if the assumption is that firms maximize profits.

More fundamentally, much of competition policy is founded on an assumption of profit-maximizing behavior by firms. Some go so far as to say that the “entire antitrust enterprise is dedicated to the proposition that business firms behave rationally.”<sup>122</sup> Courts and regulators in some jurisdictions may not consider seriously conduct (such as predatory pricing, for instance) which does not appear to make “business sense” according to their judgment. Leslie<sup>123</sup> reports that “if a plaintiff’s complaint describes a conspiracy that the judge concludes is irrational, then the court rules that the conspiracy must not have happened as a matter of law, regardless of the evidence presented by the plaintiff to support its claim.” In the light of the theories and evidence reported in this article, we suggest that a dogmatic attitude towards the pervasiveness of business rationality may lead to instances where harmful behavior goes unpunished. Behavioral economics may sometimes pose at least as many questions as it answers, and complicate antitrust debate. But it also sheds light on important market phenomena, and if competition policy is to reflect market realities, behavioral economics cannot be ignored.

WE SUGGEST THAT A DOGMATIC ATTITUDE TOWARDS THE PERVASIVENESS OF BUSINESS RATIONALITY MAY LEAD TO INSTANCES WHERE HARMFUL BEHAVIOR GOES UNPUNISHED.

Finally, while this article has surveyed behavioral economics as it applies to firms, in future it would be interesting to investigate how it applies also to policy-makers.<sup>124</sup> Competition authorities, like firms, operate within a complex and strategic environment, and may need to resort to rules-of-thumb and satisficing behavior. Instances of this could include an authority’s use of *per se* rules, or a reliance on relatively rigid market definitions and market share thresholds. It may also be advantageous to induce competition authorities to have an objective or institutional focus which differs from social welfare, in order to alter the response from the firms subject to regulation.<sup>125</sup> Imitative strategies may sometimes be employed, and safety in numbers may be enjoyed, as policy-makers look around the world for current “best practice.” (Indeed, the recent emphasis on behavioral economics may be an instance of this.) Friedman’s point about competitive pressures may have less force in terms of constraining good decision-making by public officials, and the result may be that behavioral biases are more prevalent among policy-makers than in the firms they oversee. ▼

- 1 See Armstrong for a review of some of this literature and its implications for consumer protection policy, Mark Armstrong, *Interactions Between Competition and Consumer Policy*, COMPETITION POL'Y INT'L 4, 97-148 (2008). See Spiegler for an account of economic models which assume profit-maximizing behavior by firms but bounded rationality on the part of consumers, RAN SPIEGLER, BOUNDED RATIONALITY AND INDUSTRIAL ORGANIZATION (in press). As observed by Ellison, fifty years ago the focus was instead very much on non-optimizing behavior by firms rather than consumers, Glenn Ellison, *Bounded Rationality in Industrial Organization*, ADVANCES IN ECONOMICS AND ECONOMETRICS: THEORY AND APPLICATIONS, 145, (R. Blundell, W. Newey, & T. Persson, eds.) (2006).
- 2 The recent financial crisis may well have stemmed in part from a variety of behavioral biases of some of the banks' executives. For instance, some managers may have been over-optimistic about the risks they were taking in their lending strategies, and there may have been a herd mentality among some managers, who imitated apparently successful lending strategies and who may have felt there was "safety in numbers."
- 3 Hayek argued that a central merit of competitive markets, specifically price-taking behavior by consumers and firms, is that agents' strategies are then relatively simple. (Agents need to know only their endowments, preferences, and the market prices to optimize.) Frederick Hayek in *The Use of Knowledge in Society*, AMER. ECON. R. 35, 519-530 (1945), at §VI. Gale & Sabourian argue that in oligopolistic markets, where optimal strategies may be extremely complex, if agents incur "complexity costs" when they pursue complex strategies then the outcome might be more competitive than the standard theory suggests. Douglas Gale & Hamid Sabourian, *Complexity and Competition*, ECONOMETRICA 73, 739-769 (2005).
- 4 The take-up of technological improvements in agriculture provides useful evidence. Ellison & Fudenberg quote a historian of the English agricultural revolution as writing "land tilled in very ancient ways lay next to fields in which crop rotations were followed." They report that the agricultural practices known as the "new husbandry" diffused through England and France at the rate of just one mile per year. Glenn Ellison & Drew Fudenberg, *Rules of Thumb for Social Learning*, J. POL. ECON. 101, 612-643 (1993).
- 5 A recent pair of papers by Goldfarb & Yang and Goldfarb & Xiao illustrate this. For instance, it appears that older and better-educated managers tended to enter markets with fewer competitors after the U.S. telecommunications market was deregulated in 1996. Avi Goldfarb & Botao Yang, *Are All Managers Created Equal?*, J. MARKETING RES. 46, 612-622 (2009); Avi Goldfarb & Mo Xiao, *Who Thinks About the Competition? Managerial Ability and Strategic Entry in US Local Telephone Markets*, working paper University of Toronto (2010).
- 6 See Bertrand & Schoar for an empirical analysis of the importance of "manager fixed effects." For instance, they find that managers with an MBA tend to follow strategies that are more aggressive. They find that managers differ in their attitude to mergers, dividend policy, and cost-cutting policy. Marianne Bertrand & Antoinette Schoar, *Managing with Style: The Effect of Managers on Firm Policies*, Q. J. ECON. 118, 1169-1208 (2003).
- 7 Apparently, a proportion of wine producers in California do not care purely about the profit they generate, and instead enjoy producing high-quality, high-price wine. (Profit-maximizing wineries tend to offer lower quality wine.) See Fiona Scott Morton & Joel Podolny, *Love or Money? The Effects of Owner Motivation in the California Wine Industry*, J. INDUS. ECON. 50, 431-456 (2002).
- 8 An interesting example of this was seen in the Genzyme-Novazyme merger-to-monopoly which was approved by the FTC in 2004. This was a merger of two firms both engaged in R&D for treating a rare disease, where the prime danger from the merger was whether the discovery of a successful treatment would be delayed relative to the duopoly outcome. One factor in the decision was that the proposed CEO of the merged entity had two children with the disease, who may therefore not have wished to delay discovery. See the statement by the then FTC chairman Timothy Muris, available at [www.ftc.gov/os/2004/01/murisgenzymestmt.pdf](http://www.ftc.gov/os/2004/01/murisgenzymestmt.pdf).

- 9 MILTON FRIEDMAN, *ESSAYS IN POSITIVE ECONOMICS*, 22 (1953).
- 10 For further details, see Charles Holt, *Industrial Organization: A Survey of Laboratory Research*, THE HANDBOOK OF EXPERIMENTAL ECONOMICS, §3 (John Kagel & Alvin Roth eds.) (1995). For a collection of papers discussing the use of experiments for competition policy, see EXPERIMENTS AND COMPETITION POLICY, (Jeroen Hinloopen & Hans-Theo Normann eds.) (2009).
- 11 For instance, see Douglas DeJong, Robert Forsythe, & Wilfred Uecker, *A Note on the Use of Businessmen as Subjects in Sealed Offer Markets*, J. ECON. BEHAVIOR AND ORG. 9, 87-100 (1988). More recently, for an illustration of how competition agency practitioners behaved on average in a similar way to students in a particular regulatory decision problem, see Bruce Lyons, Gordon Menzies, & Daniel Zizzo, *Professional Interpretation of the Standard of Proof: An Experimental Test on Merger Regulation*, Centre for Competition Policy Working Paper 10-2 (2010).
- 12 Charles Plott, *An Updated Review of Industrial Organization Applications of Experimental Methods*, THE HANDBOOK OF INDUSTRIAL ORGANIZATION, Vol. II, at 1165, (Richard Schmalensee & Robert Willig eds.) (1989).
- 13 In much empirical work on naturally occurring markets, marginal costs are inferred from observed data such as prices by *assuming* firms maximize their profit. At the end of this paper we discuss the dangers of this method when firms might potentially not be optimizers. The fact that marginal costs are rarely directly observable makes it hard to perform tests of the profit-maximization hypothesis. Steven Levitt in *Bagels and Donuts for Sale: A Case Study in Profit Maximization*, mimeo (2008), is a rare example of such a test. He obtained data from a relatively "simple" firm which supplies bagels and donuts to businesses, and where marginal costs are known. He finds that the firm is extremely good at predicting demand for given prices, but apparently prices too low given the estimated demand elasticity. (Interestingly, the decision-maker for this firm is a well-trained economist, who has published in the *Journal of Political Economy*.)
- 14 For an account of the theory of collusion, assuming rational firms, see Marc Ivaldi, Bruno Jullien, Patrick Rey, Paul Seabright, & Jean Tirole, *The Economics of Tacit Collusion*, Report for DG Competition, European Commission, (2003).
- 15 For instance, in the case of Bertrand price competition and homogenous products, the fully collusive outcome can be sustained in an infinitely repeated interaction with  $n$  symmetric suppliers if the discount factor  $\delta$  satisfies  $\delta > 1 - 1/n$ . With reasonable choices for the discount factor and say, monthly price adjustment, collusion should be possible in oligopolies consisting of a hundred firms.
- 16 George Stigler, *A Theory of Oligopoly*, J. POL. ECON. 72, 44-61 (1964).
- 17 For instance, see Margaret Levenstein & Valerie Suslow, *Cartel Bargaining and Monitoring: The Role of Information Sharing*, THE PROS AND CONS OF INFORMATION SHARING (Mats Bergman ed.), (2006).
- 18 Some experiments randomly match subjects in each period, so that firms play against different rivals in each period. However, real markets do not frequently operate like this, and so we mainly focus on experiments where subjects interact repeatedly in the same groups.
- 19 For a brief survey, see Steffen Huck, Hans-Theo Normann, & Jorg Oechssler, *Two are Few and Four are Many: Number Effects in Experimental Oligopolies*, J. ECON. BEHAVIOR AND ORG. 53, 435-446 §§ 2 and 3 (2004).
- 20 Robert Feinberg & Thomas Husted, *An Experimental Test of Discount Rate Effects on Collusive Behavior in Duopoly Markets*, J. INDUS. ECON. 41, 153-160 (1993).

- 21 Feinberg argues that social welfare may sometimes be higher when firms' managers are relatively myopic or short-termist, perhaps because of high managerial turnover or the kinds of incentive schemes they are offered, since collusion thereby becomes harder to sustain. For the same reason, it seems plausible that shareholders who wish to achieve collusion would not wish to put in place a manager who was myopic or had hyperbolic time preferences. Robert Feinberg, *In Defense of Corporate Myopia*, *MANAGERIAL AND DECISION ECON.* 16, 205-210 (1995).
- 22 It might seem that this apparent falsification of rational play is not necessarily important in practice, as most real markets do not have a known "endpoint." Nevertheless, the unravelling argument applies even if the endpoint is uncertain, but it is known for sure that the interaction will have ceased by some date (e.g., if subjects in the laboratory do not believe that the experiment could possibly go on for more than a day, or if it is common knowledge that the world will have ended in 10 billion years).
- 23 *Supra* note 19.
- 24 As an aside, it is interesting in this regard to recall Robert Bork's assertion that any merger that left at least three rivals should be presumptively lawful. However, tacit collusion is not the only problem for concentrated markets; even the one-shot interaction could be insufficiently competitive when there are few firms. See ROBERT BORK, *THE ANTITRUST PARADOX: A POLICY AT WAR WITH ITSELF* (1978).
- 25 For instance, see Martin Dufwenberg & Uri Gneezy, *Price Competition and Market Concentration: An Experimental Study*, *INT'L J. INDUS. ORG.* 18, 7-22 (2000); Michael Baye & John Morgan, *Price Dispersion in the Lab and on the Internet: Theory and Evidence*, *RAND* 35, 449-466 (2004); and Jeroen Hinloopen & Adriaan Soetevent, *Laboratory Evidence on the Effectiveness of Corporate Leniency Programs*, *RAND* 39, 607-616 (2008).
- 26 Steffen Huck, Hans-Theo Normann, & Jorg Oechssler, *Learning in Cournot Oligopoly: An Experiment*, *ECON. J.* 109, C80-C95 (1999) and Steffen Huck, Hans-Theo Normann, & Jorg Oechssler, *Does Information About Competitors' Actions Increase or Decrease Competition in Experimental Oligopoly Markets?*, *INT'L J. INDUS. ORG.* 18, 39-57 (2000).
- 27 Theo Offerman, Jan Potters, & Joep Sonnemans, *Imitation and Belief Learning in an Oligopoly Experiment*, *REV. ECON. STUDIES* 69, 973-997 (2002).
- 28 An interesting real-world experiment is reported in Svend Albaek, Peter Mollgaard, & Per Overgaard, *Government-Assisted Oligopoly Coordination? A Concrete Case*, *J. INDUS. ECON.* 45, 429-443 (1997). The Danish antitrust authority changed policy so that transaction prices in the concrete market were published, and subsequently the average prices rose significantly. The authors argue that the most plausible reason is that this enabled the firms to coordinate their prices at a high level.
- 29 Offerman et al., *supra* note 27, also investigated the case where rival outputs but not rival profits were reported, and found that the market was less competitive than when no firm-specific information was reported. It would be interesting in future work to see what happens when rival firms' profits (or even just the average profits achieved by firms) but not their actions were revealed. It is possible that firms with higher profits than average might keep their strategy unchanged, but firms who do poorly revise their strategies, in accordance with a model of satisficing behavior discussed in section V. Bigoni allows firms to choose the kinds of information they see (e.g., aggregate output of rivals, individual rival outputs, individual rival profits) when operating under a time constraint, see Maria Bigoni, *Information and Learning in Oligopoly: An Experiment*, mimeo, (2008).
- 30 Huck et al., (2000), *supra* note 26, find that when firms set prices rather than quantities, the revelation of firm-specific data reduces prices only slightly.
- 31 Jan Potters, *Transparency About Past, Present and Future Conduct*, *EXPERIMENTS AND COMPETITION POLICY* (Jeroen Hinloopen & Hans-Theo Normann eds.) (2009).

- 32 Ola Andersson & Erik Wengstrom, *Do Antitrust Laws Facilitate Collusion? Experimental Evidence of Costly Communication in Duopolies*, SCANDINAVIAN J. ECON. 109, 321-339 (2007).
- 33 Explicit collusion in naturally occurring markets appears to be feasible with large numbers of participants. For instance, in their study of 41 cartels in Europe, Levenstein & Suslow, *supra* note 17, find that 18 involved more than five firms.
- 34 Giancarlo Spagnolo, *Divide et Impera: Optimal Leniency Programs*, CEPR Discussion Paper 4840 (2004).
- 35 Massimo Motta & Michele Polo, *Leniency Programs and Cartel Prosecution*, INT'L J. INDUS. ORG. 21, 347-379 (2003).
- 36 Jose Apesteguia, Martin Dufwenberg, & Reinhard Selten, *Blowing the Whistle*, ECON. THEORY 31, 143-166 (2007).
- 37 Hinloopen & Soetevent, *supra* note 25.
- 38 Armen Alchian, *Uncertainty, Evolution, and Economic Theory*, J. POL. ECON. 57, 211-22, at 218, (1950).
- 39 For a model along these lines, see John Conlisk, *Costly Optimizers Versus Cheap Imitators*, J. ECON. BEHAVIOR AND ORG. 1, 275-293 (1980). The situation is somewhat related to Grossman & Stiglitz's analysis of the incentives for investors to become better informed about the return of an uncertain asset: If all investors choose to be better informed, the asset's price reflects the information, and there is no need for any individual investor to become informed if there is a cost to doing so, Sanford Grossman & Joseph Stiglitz, *Information and Competitive Price Systems*, AMER. ECON. REV. 66, 246-253 (1976). It is also related to Burdett & Judd's analysis of a consumer's incentive to search for a low price: If all consumers choose to search then the market is highly competitive, and there is no need for any individual consumer to search if there is a cost to doing so, Kenneth Burdett & Kenneth Judd, *Equilibrium Price Dispersion*, ECONOMETRICA 51, 955-969 (1983).
- 40 Banerjee & Bikhchandani, et al., consider a situation in which similar agents need to decide between two options (say, whether to adopt technology X or technology Y), and each agent has a private signal as to which option is better and can also observe the previous choices (but not the payoffs) made by other agents. (The order in which agents have to make their choice is pre-determined in these models.) Even if agents are completely rational, it is possible that they become locked into the wrong choice. For instance, if technology X is in fact superior, but by chance the first few agents have private signals which induce them to choose Y, then subsequent agents will infer that the superior action is likely to be Y despite their own private signals to the contrary. (If instead, agents could observe the private signals of the earlier adopters, this inefficient herding could not occur.) Abhijit Banerjee, *A Simple Model of Herd Behavior*, Q. J. ECON 107, 797-817 (1992), and Sushil Bikhchandani, David Hirshleifer, & Ivo Welch, *A Theory of Fads, Fashion, Custom, and Cultural Change as Information Cascades*, J. POL. ECON. 100, 992-1026 (1992).
- 41 See Georg Weizsacker, *Do We Follow Others When We Should? A Simple Test of Rational Expectations*, AMER. ECON. REV. (forthcoming), for an analysis of data from several experiments on the Bikhchandani et al. model, *Id.*
- 42 David Scharfstein & Jeremy Stein, *Herd Behavior and Investment*, AMER. ECON. REV. 80, 465-479 (1990).
- 43 *Id.* at 466.
- 44 The authors (*Id.*) quote Keynes as writing: "Worldly wisdom teaches that it is better for reputation to fail conventionally than to succeed unconventionally." Tse & Tucker empirically investigate the timing of

earnings warnings, and find that a manager is more likely to issue an earnings warning if a peer has done so in the previous days. They conclude that the data are better explained by managers attempting to maintain their reputations than by the impact of a common shock. Senyo Tse & Jennifer Wu Tucker, *Within-Industry Timing of Earnings Warnings: Do Managers Herd?* Rev. of ACCT. STUDIES (forthcoming).

- 45 It is important that firms experiment occasionally, otherwise the process will grind to a halt after one period when all firms imitate the most profitable firm in the first period (which is unlikely to have chosen the optimal price immediately), and then all prices are unchanged thereafter. An alternative framework is presented in Ellison & Fudenberg, *supra* note 4, §II where firms choose between one of two technologies, and the relative payoff from using one technology rather than the other is uncertain. If in any period firms choose the technology which performed best in the previous period, then the chosen technology will flip over time depending on which one happened to work best one period earlier. Thus, instead of converging to the consistent use of the superior technology, the outcome is merely that the better technology is used more frequently. The authors go on to investigate less naïve rules of thumb—where a firm conditions its choice on how many firms use that technology—which have superior efficiency properties. (The reason why market shares matter for firms is that they reveal information about the relative performance of the two technologies for more than just the single previous period.)
- 46 David Ridley analyzes a model in which a second firm sometimes decides to enter a market only if its rival has first entered, in order to save on the costs of acquiring its own market information. He provides some anecdotes about how competitors of McDonald's often locate near a new McDonald's franchise, and he quotes a manager of a coffee shop chain as saying: "The reason we want to open across the street from every Starbucks is they do a great job at finding good locations." David Ridley, *Herding Versus Hotelling: Market Entry with Costly Information*, J. ECON. AND MGMT. STRATEGY 17, 607-631 (2008).
- 47 Fernando Vega-Redondo, *The Evolution of Walrasian Behavior*, ECONOMETRICA 65, 375-384 (1997).
- 48 Mark Schaffer, *Are Profit-Maximizers the Best Survivors?*, J. ECON. BEHAVIOR AND ORG. 12, 29-45 (1989).
- 49 Karl Schlag proposes an alternative model of imitation which yields a contrasting long-run prediction. In Schlag's model, there are many distinct oligopolies, and a firm in one market imitates the best-performing firm in another oligopoly, not the best-performing rival in the same market (as in Vega-Redondo's model). The result of this alternative specification is that firms move towards the Cournot-Nash equilibrium rather than the perfectly competitive outcome. Karl Schlag, *Why Imitate, and If So, How? A Boundedly Rational Approach to Multi-Armed Bandits*, J. ECON. THEORY 78, 130-156 (1998), and Karl Schlag, *Which One Should I Imitate?*, J. MATH. ECON. 31, 493-522 (1999). See Jose Apesteguia, Steffen Huck, & Jorg Oechssler, *Imitation: Theory and Experimental Evidence*, J. ECON. THEORY 136, 217-235 (2007) for further discussion of the two approaches. In addition, Vega-Redondo's result that imitation leads to competitive outcomes does not necessarily hold when firms have longer memory. If, say, firms choose the most profitable action over the previous two periods, the Cournot equilibrium is then a stable outcome (although the perfectly competitive outcome remains stable as well). See Carlos Alos-Ferrer, *Cournot Versus Walras in Dynamic Oligopolies with Memory*, INT'L J. INDUS. ORG. 22, 193-217 (2004).
- 50 See Paul Rhode & Mark Stegeman, *Non-Nash Equilibria of Darwinian Dynamics with Application to Duopoly*, INT'L J. INDUS. ORG. 19, 415-453 (2001) for more details.
- 51 See Yasuhito Tanaka, *Stochastically Stable States in an Oligopoly with Differentiated Products: Equivalence of Price and Quantity Strategies*, J. MATH. ECON. 34, 235-253 (2000) and Rhode & Stegeman, *Id.* for a general argument. If firms are not symmetrically placed, then imitation may not tend towards competitive outcomes, nor is imitation equivalent to the maximization of relative profits. See, Jose Apesteguia, Steffen Huck, Jorg Oechssler, & Simon Weidenholzer, *Imitation and the Evolution of Walrasian Behavior: Theoretically Fragile but Behaviorally Robust*, J. ECON. THEORY (forthcoming).

- 52 Burkhard Schipper, *Imitators and Optimizers in Cournot Oligopoly*, J. ECON. DYNAMICS AND CONTROL 33, 1981-1990 (2009).
- 53 See Peter Duersch, Jorg Oechssler, & Burkhard Schipper, *Unbeatable Imitation*, mimeo, (2009) for further details about how imitators do almost as well against even the smartest opponents in a wide class of games. The main kind of game where imitators do poorly is something like "rock-paper-scissors," where a smart player can systematically trick the naïve imitator into always playing the wrong action. (These kinds of games do not seem common in market situations, however.)
- 54 Huck, Normann, & Oechssler (*supra* note 26); Offerman, Potters, Sonnemans (*supra* note 27); and Apesteguia, Huck, and Oechssler (*supra* note 49).
- 55 Huck et al., *supra* note 26.
- 56 Offerman et al., *supra* note 27 at 989.
- 57 Goel & Thakor propose a model in which merger waves can be caused by concerns for relative compensation by CEOs. If CEO compensation is based in part on firm size, then if one merger occurs which boosts that CEO's pay, other CEO feel envious and set about finding their own take-over targets. Anand Goel & Anjan Thakor, *Do Envious CEOs Cause Merger Waves?*, REV. FIN. STUDIES 23, 487-517 (2010).
- 58 For instance, see Andrew Clark, Paul Frijters, & Michael Shields, *Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles*, J. ECON. LIT. 46, 95-144 (2008).
- 59 Robert Gibbons & Kevin Murphy, *Relative Performance Evaluation for Chief Executive Officers*, INDUS. AND LABOR REL. REV. 43, 30-51 (1990).
- 60 John Vickers, *Delegation and the Theory of the Firm*, ECON. J. 95, 138-147, at 143, (1985).
- 61 The impact of making the firm wish to maximize relative profits is that the firm then behaves as the Stackelberg leader, even though both firms in fact choose quantities simultaneously. It is important that rivals observe the incentive scheme so that they know the firm's objective and can react to it accordingly.
- 62 Nolan Miller & Amit Pazgal, *Relative Performance as a Strategic Commitment Mechanism*, MANAGERIAL AND DECISION ECON. 23, 51-68 (2002).
- 63 Vicente Salas Fumas, *Relative Performance Evaluation of Management*, INT'L J. OF INDUS. ORG. 10, 473-489 (1992).
- 64 Carl Lundgren, *Using Relative Profit Incentives to Prevent Collusion*, REV. INDUS. ORG. 11, 533-550 (1996).
- 65 See Güth et al.'s path-breaking experiment and the many papers which followed, Werner Güth, Rolf Schmittberger, & Bernd Schwarze, *An Experimental Analysis of Ultimatum Bargaining*, J. ECON. BEHAVIOR AND ORG. 3, 367-388 (1982). See Ernst Fehr & Klaus Schmidt, *A Theory of Fairness, Competition, and Cooperation*, Q. J. ECON. 114, 817-868, §IIIA (1999) for a review of this literature.
- 66 The proposer might offer a significant proportion of the prize to the responder because he is purely self-interested and is afraid of lower offers being rejected by spiteful responders, or because he cares directly for fair allocations and is willing to sacrifice some of the prize to achieve a more equitable outcome. A variant of the ultimatum game—the dictator game—can discriminate between these two hypotheses. The dictator game does not allow the responder to reject the offer. Experimental comparisons of the two games reveal that offers are considerably less generous (and often zero) in the dictator game than in the ultimatum game, suggesting that generosity on the part of many proposers is

purely strategic. For more details, see Robert Forsythe, Joel Horowitz, N. E. Slavin, & Martin Sefton, *Fairness in Simple Bargaining Experiments*, GAMES AND ECON. BEHAVIOR 6, 347-369 (1994).

- 67 Steffen Huck, Wieland Muller, & Hans-Theo Normann, *Stackelberg Beats Cournot: On Collusion and Efficiency in Experimental Markets*, ECON. J. 111, 749-765 (2001).
- 68 *Supra* note 60.
- 69 Steffen Huck, Wieland Muller, & Hans-Theo Normann, *Strategic Delegation in Experimental Markets*, INT'L J. INDUS. ORG. 22, 561-574 (2004).
- 70 *Supra* note 62.
- 71 See Doruk Iris & Luis Santos-Pinto, *Tacit Collusion Under Fairness and Reciprocity*, mimeo, (2009), for a model along these lines. However, we are aware of no experimental test of this hypothesis.
- 72 For instance, if the ultimatum game is modified so that the proposer has first to run around a running track, or if the proposer is chosen by his performance in a general knowledge quiz, then the responder may feel that the proposer "owns" the initial stake, and so be more willing to accept a lower share of the prize.
- 73 See §§ IIIB and IIIC of Fehr & Schmidt (*supra* note 65) for further details and references.
- 74 See Levenstein & Suslow, *supra* note 17, and, especially, Christopher Leslie, *Trust, Distrust and Antitrust*, TEXAS L. REV. 82, 515-680 (2004), for further discussion of these points and illustrations of the role of trust and distrust in cartel stability.
- 75 William Page, *The Gary Dinners and the Meaning of Concerted Action*, SMU L. REV. 62, 597-619, at 608, (2009).
- 76 Joel Podolny & Fiona Scott Morton, *Social Status, Entry and Predation: The Case of British Shipping Cartels 1879-1929*, J. INDUS. ECON. 47, 41-67 (1999).
- 77 For instance, Gordon writes:
- The fear of bankruptcy and the even more widespread fear of temporary financial embarrassment are probably more powerful drives than the desire for the absolute maximum in profits. [...] Given the fog of uncertainty within which [the businessman] must operate, the limited number of variables his mind can juggle at one time, and his desire to play safe, it would not be at all surprising if he adopted a set of yardsticks that promised reasonably satisfactory profits, R. A. Gordon, *Short-Period Price Determination in Theory and Practice*, AMER. ECON. REV. 3, 265-280 at 271 (1948)
- See also K.W. Rothschild, *Price Theory and Oligopoly*, ECON. J. 57, 299-320 (1947) and Herbert Simon, *A Behavioral Model of Rational Choice*, Q. J. ECON. 69, 99-118 (1955).
- 78 Simon makes a joke about those economists who believe that departures from the predictions of rational behavior do not matter if the predictions of the standard models are good enough: "economists who are zealous in insisting that economic agents maximize turn around and become satisficers when the evaluation of their own theories is concerned." Herbert Simon, *Rational Decision Making in Business Organizations*, AMER. ECON. REV. 69, 493-513, at 495, (1979).
- 79 Richard Cyert & James March, *Organizational Factors in the Theory of Oligopoly*, Q. J. ECON. 70, 44-64 (1956).



- 80 Steffen Huck, Kai Konrad, Wieland Muller, & Hans-Theo Normann, *The Merger Paradox and Why Aspiration Levels Let it Fail in the Laboratory*, *ECON. J.* 117, 1073-1095 (2007).
- 81 Stephen Salant, Sheldon Switzer, & Robert Reynolds, *Losses from Horizontal Merger: The Effects of an Exogenous Change in Market Structure on Cournot-Nash Equilibrium*, *Q. J. ECON.* 98, 185-199 (1983).
- 82 The prediction is typically reversed when firms have differentiated products and compete in prices. See Raymond Deneckere & Carl Davidson, *Incentives to Form Coalitions with Bertrand Competition*, *RAND* 16, 474-486 (1985).
- 83 Levin presents a related theoretical analysis of mergers in Cournot markets, in which he allows that a merged firm changes its behavior from a Cournot-Nash player to a Stackelberg player (among other possible behavioral changes), in which case a merger becomes profitable. Daniel Levin, *Horizontal Mergers: The 50-Percent Benchmark*, *AMER. ECON. REV.* 80, 1238-1245 (1990).
- 84 Huw Dixon, *Keeping Up with the Jones's: Competition and the Evolution of Collusion*, *J. ECON. BEHAVIOR AND ORG.* 43, 223-238 (2000), and Jorg Oechssler, *Cooperation as a Result of Learning with Aspiration Levels*, *J. ECON. BEHAVIOR AND ORG.* 49, 405-409 (2002).
- 85 An alternative adjustment mechanism which does not require any information about rival actions or profits, or indeed about the demand and cost functions, but which may nevertheless lead again to collusive outcomes, is discussed by Huck, Normann, & Oechssler. They consider the rule of thumb: "If your last increase in output [or price] increased your profit, increase your output [or price] again; if your last increase in output [or price] decreased your profit, now decrease your output [or price]." They show that if firms are constrained to change their strategy at a fixed rate over time, then the market moves towards the collusive outcome. Steffen Huck, Hans-Theo Normann, & Jorg Oechssler, *Through Trial & Error to Collusion*, *INT'L ECON. REV.* 45, 205-224 (2004).
- 86 See Ellison (*supra* note 1 at 170-171). Note that if there are  $n$  firms instead of just two, the monopoly outcome can be implemented as an  $\varepsilon$ -equilibrium provided that  $\varepsilon > (n - 1)^2 / (16n^2)$ , and so with more firms it does become harder to sustain the monopoly outcome with satisficing behavior. Relatedly, approximately optimal behavior by firms can sustain tacit collusion when they interact only finitely often, at least for the early periods of the interaction. (In simple models, fully rational firms cannot sustain any collusion when they meet a known, finite number of times.) For further details, see Roy Radnor, *Collusive Behavior in Non-Cooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives*, *J. ECON. THEORY* 22, 136-154 (1980).
- 87 Baye & Norman, *supra* note 25.
- 88 For further details of this form of approximately optimal behaviour, see Richard McKelvey & Thomas Palfrey, *Quantal Response Equilibria for Normal Form Games*, *GAMES AND ECON. BEHAVIOR* 10, 6-38 (1995).
- 89 Renou & Schlag propose an alternative model for the Bertrand market, which is that firms are unsure about the rationality of their opponents, and a firm aims to minimize the maximum "regret" it experiences when competing with rivals. (A firm feels regret if it sets a price far below the minimum price of its rivals, since it could have made more profit with a higher price, and it feels regret if it sets a price above the minimum of its rivals' prices.) Using this model, they predict that firms set random prices, and expected profits are positive but decreasing with the number of competitors. They argue that the data reported in Baye & Morgan, *supra* note 25, conform well with their predictions. See, Ludovic Renou & Karl Schlag, *Minimax Regret and Strategic Uncertainty*, *J. ECON. THEORY* 145, 264-286 (2010).
- 90 For example, Svenson finds that 93 percent of respondents report that they are above the median in terms of driving ability. Ola Svenson, *Are We All Less risky and More Skilful Than our Fellow Drivers?*, *ACTA PSYCHOLOGICA* 47, 143-148 (1981).

- 91 For a persuasive account of how Adam Smith anticipated many of the central ideas of behavioral economics, see Nava Ashraf, Colin Camerer & George Loewenstein, *Adam Smith, Behavioral Economist*, J. ECON. PERSP. 19, 131-145 (2005).
- 92 See De Meza & Southey for empirical evidence and a model which predicts that entrepreneurs are more optimistic than the general population, David De Meza & Clive Southey, *The Borrower's Curse: Optimism, Finance and Entrepreneurship*, ECON. J. 106, 375-386 (1996). Camerer & Lovallo document in an experimental study that excessive entry can be driven by over-confidence in own ability, Colin Camerer & Dan Lovallo, *Overconfidence and Excess Entry: An Experimental Approach*, AMER. ECON. REV. 89, 306-318 (1999). See Tor for an extended discussion of behavioral economics and entry decisions, and the implications for competition policy, Avshalom Tor, *The Fable of Entry: Bounded Rationality, Market Discipline, and Legal Policy*, MICH. L. REV. 101, 482-568 (2002).
- 93 Edward Lazear & Sherwin Rosen, *Rank-Order Tournaments as Optimal Labor Contracts*, J. POL. ECON. 89, 841-864 (1981).
- 94 For a model along these lines, see Anand Goel & Anjan Thakor, *Rationality, Overconfidence and Leadership*, mimeo, (2000).
- 95 Aviad Heifetz, Chris Shannon, & Yossi Spiegel, *The Dynamic Evolution of Preferences*, ECON. THEORY 32, 251-286 (2007).
- 96 Vickers, *supra* note 60.
- 97 See Florian Englmaier, *A Strategic Rationale of Having Overoptimistic Managers*, mimeo, (2007) for such a model. In a Cournot market, hiring a manager who is overoptimistic (say, about the scale of market demand) again confers strategic advantage to the firm, but if all firms do this they will all be worse off relative to the situation with unbiased managers.
- 98 Richard Roll, *The Hubris Hypothesis of Corporate Takeovers*, J. BUS. 59, 197-216 (1986).
- 99 Ulrike Malmendier & Geoffrey Tate, *CEO Overconfidence and Corporate Investment*, J. FIN. 60, 2661-2700 (2005) and Ulrike Malmendier & Geoffrey Tate, *Who Makes Acquisitions? CEO Overconfidence and the Market's Reaction*, J. FIN. ECON. 89, 20-43 (2008).
- 100 This correlation between holding onto stock options and corporate behavior could also be due to the CEO's inside information. However, Malmendier & Tate argue that this is unlikely to be the explanation since the CEOs who hold onto their stock options do not gain money from doing so, *Id.*
- 101 For further discuss of this point, see Maurice Stucke, *Am I a Price-Fixer? A Behavioral Economics Analysis of Cartels*, mimeo, University of Tennessee College of Law (2010).
- 102 For more details, see Bradford De Long, Andrei Shleifer, Laurence Summers, & Robert Waldmann, *The Survival of Noise Traders in Financial Markets*, J. POL. ECON. 64, 703-738 (1991).
- 103 See Luo Hirshleifer & Guo Ying Luo, *On the Survival of Overconfident Traders in a Competitive Securities Market*, J. FIN. MARKETS 4, 73-84 (2001).
- 104 Albert Kyle & Albert Wang, *Speculation Duopoly with Agreement to Disagree: Can Overconfidence Survive the Market Test?*, J. FIN. 52, 2073-2090 (1997).
- 105 Robert Hall & Charles Hitch, *Price Theory and Business Behavior*, OXFORD ECON. PAPERS 2, 12-45 (1939).
- 106 Nabil Al-Najjar, Sandeep Baliga, & David Besanko, *Market Forces Meet Behavioral Biases: Cost Misallocation and Irrational Pricing*, RAND 39, 214-237 (2008).

- 107 Theo Offerman & Jan Potters, *Does Auctioning of Entry Licences Induce Collusion? An Experimental Study*, R. ECON. STUDIES 73, 769-792 (2006).
- 108 A numerical example may help to fix ideas. Suppose there is a duopoly with differentiated products where the two firms choose prices and where a firm's demand is given by expression (1) as in section III. But instead of profit-maximizing behavior, suppose that firm  $i = 1, 2$  behaves as if its marginal cost is  $c_i$  rather than the true cost (which is zero). Suppose that far-sighted profit-maximizing share-holders of firm  $i$  (or some form of evolutionary pressures) choose a corporate culture such that their manager behaves as if cost was  $c_i$ , and both groups of share-holders do this before price competition takes place. Then one can show that the equilibrium chosen "costs" are  $c_1 = c_2 = 1/11$ , and costs are artificially boosted so as to raise equilibrium prices. In this example, prices rise by approximately 10 percent relative to the situation in which managers base their prices on the true marginal costs.
- 109 Al-Najjar et al., *supra* note 106, suppose that the oligopolists reach equilibrium prices via a myopic adjustment process, and this means that rivals do not need to observe the "bias" of the manager. In addition, they assume that each firm observes only its own profits and not the profits of its rivals, and are more likely to choose cost methodologies that are performed well for it in the past.
- 110 Jonathan Eaton & Gene Grossman, *Optimal Trade and Industrial Policy Under Oligopoly*, Q. J. ECON. 101, 383-406 (1986).
- 111 See Giacomo Bonanno & John Vickers, *Vertical Separation*, J. INDUS. ECON. 36, 257-265 (1988) for this model.
- 112 See David Kreps, Paul Milgrom, John Roberts, & Robert Wilson, *Rational Cooperation in the Finitely Repeated Prisoners' Dilemma*, J. ECON. THEORY 27, 245-252 (1982).
- 113 However, if the potential irrationality took a different form then collusion might not be sustained. For instance, if the irrational player has a strategy of always colluding, then a rational player's best response to this is always to defect (just as if she were playing against a rational agent).
- 114 See David Kreps & Robert Wilson, *Reputation and Imperfect Information*, J. ECON THEORY 27, 253-279 (1982).
- 115 Yun Joo Jung, John Kagel, & Dan Levin, *On the Existence of Predatory pricing: An Experimental Study of Reputation and Entry Deterrence in the Chain-Store Game*, RAND 25, 72-93 (1994).
- 116 Kenneth Rogoff, *The Optimal Degree of Commitment to an Intermediate Monetary Target*, Q. J. ECON. 100, 1169-1189 (1985).
- 117 THOMAS SCHELLING, *THE STRATEGY OF CONFLICT*, 142 (1960).
- 118 For more analysis of this and related points, see Christopher Leslie, *Rationality Analysis in Antitrust*, U. PA. L. REV. 158, 262-353 (2010).
- 119 However, as discussed in section III, the effect may be limited if rivals react punitively when one firm puts in place an aggressive manager or an incentive scheme which induces aggressive behavior by its manager.
- 120 John List, *Does Market Experience Eliminate Market Anomalies?* Q. J. ECON. 118, 41-71 (2003) and John List, *Neoclassical Theory Versus Prospect Theory: Evidence From the Marketplace*, ECONOMETRICA 72, 615-625 (2004).
- 121 Huck et al., *supra* note 80.

- 122 HERBERT HOVENKAMP, THE ANTITRUST ENTERPRISE: PRINCIPLE AND EXECUTION, 134 (2005).
- 123 Leslie, *supra* note 118 at 269.
- 124 See also Leslie, *supra* note 118 at § III.F.
- 125 For instance, Armstrong & Vickers show how it can be optimal for a competition authority to permit only those mergers which do not harm consumers, even if society places equal weight on profit and consumer surplus. The reason is that a "consumer standard" affects the merger opportunities considered by firms, in a way which enhances total welfare. Mark Armstrong & John Vickers, *A Model of Delegated Project Choice*, *ECONOMETRICA* 78, 213-244 (2010).