



Management Science

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Collusion in Bertrand vs. Cournot Competition: A Virtual Bargaining Approach

Tigran Melkonyan, Hossam Zeitoun, Nick Chater

To cite this article:

Tigran Melkonyan, Hossam Zeitoun, Nick Chater (2018) Collusion in Bertrand vs. Cournot Competition: A Virtual Bargaining Approach. Management Science ():. <https://doi.org/10.1287/mnsc.2017.2878>

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “*Management Science*. Copyright © 2017 The Author(s). <https://doi.org/10.1287/mnsc.2017.2878>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Copyright © 2017, The Author(s)

Please scroll down for article—it is on subsequent pages



With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes. For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Collusion in Bertrand vs. Cournot Competition: A Virtual Bargaining Approach

Tigran Melkonyan,^a Hossam Zeitoun,^a Nick Chater^a

^a Warwick Business School, University of Warwick, Coventry CV4 7AL, United Kingdom

Contact: tigran.melkonyan@wbs.ac.uk,  <http://orcid.org/0000-0001-5599-6405> (TM); hossam.zeitoun@wbs.ac.uk (HZ); nick.chater@wbs.ac.uk (NC)

Received: December 16, 2016

Revised: April 3, 2017; June 1, 2017;
June 12, 2017

Accepted: June 18, 2017

Published Online in Articles in Advance:
November 22, 2017

<https://doi.org/10.1287/mnsc.2017.2878>

Copyright: © 2017 The Author(s)

Abstract. How do firms manage to collude without communicating? Why do we find more collusion in price competition than in quantity competition? Why is collusion so hard to detect? We examine strategic behavior in competitive interactions by developing and applying the concept of virtual bargaining. When decision makers virtually bargain, they mentally simulate, and choose among, agreements that they could reach if they were able to explicitly negotiate with each other. Virtual bargainers focus on agreements that offer some protection against the possibility that their counterparts may deviate and best respond to these agreements. We develop a formal account of virtual bargaining and demonstrate that it leads to collusion in Bertrand, but not in Cournot, competition. In this framework, collusion is a result of virtual bargaining as a mode of reasoning and requires neither communication nor dynamic considerations, such as rewards and punishments, between the players.

History: Accepted by Manel Baucells, decision analysis.



Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as “Management Science. Copyright © 2017 The Author(s). <https://doi.org/10.1287/mnsc.2017.2878>, used under a Creative Commons Attribution License: <https://creativecommons.org/licenses/by/4.0/>.”

Funding: N. Chater was supported by the European Research Council [Grant 295917-RATIONALITY], the Economic and Social Research Council Network for Integrated Behavioural Science [Grant ES/K002201/1], the Leverhulme Trust [Grant RP2012-V-022], and Research Councils UK [Grant EP/K039830/1].

Keywords: virtual bargaining • collusion • games-group decisions • Bertrand paradox

1. Introduction

Antitrust authorities are increasingly prioritizing the detection and punishment of price fixing. However, it is often very difficult to find concrete proof of wrongdoing, “with evidence hard to spot from the outside” (*The Economist* 2016). It appears that collusion frequently occurs with minimal or no communication among competitors. Recent reports suggest that even in online markets—where sellers are decentralized and may not know one another—there are often unspoken agreements not to undercut prices (Ungoed-Thomas and Lord 2016). Antitrust authorities typically need evidence of explicit discussion to punish colluding competitors. Therefore, *unspoken agreements* are particularly problematic. How is such collusion possible? Why is it so hard to detect?

The problem of collusion is, of course, long standing. In a famous passage in *The Wealth of Nations*, Smith (1937, p. 128) notes that, “People of the same trade seldom meet together, even for merriment and diversion, but the conversation ends in a conspiracy against the public, or in some contrivance to raise prices.” We suspect that the problem of unspoken agreements also

has a long history. Such agreements, while intuitively plausible, may seem incompatible with basic theory: collusion in oligopolistic markets without communication does not easily follow from Nash equilibrium and, hence, is difficult to analyze using standard game-theoretical tools. Finding direct evidence for the operation of unspoken collusive agreements in real markets is inevitably difficult, precisely because they are unspoken. But the existence of unspoken agreements is consistent with evidence from experimental markets, where participants’ behavior often diverges from Nash predictions. This phenomenon is especially pronounced for competition in prices, an anomaly that has been coined the “Bertrand paradox” (e.g., Engel 2007, Fatas et al. 2014, Friedman 1977, Kreps and Scheinkman 1983, Suetens and Potters 2007). Moreover, puzzling collusive behavior exists even in the absence of learning and experience effects (Dufwenberg and Gneezy 2000). The divergence between Nash predictions and observed behavior is much smaller or nonexistent when firms compete in quantities (see Potters and Suetens 2013 for an overview).

The present paper offers a novel explanation for collusive behavior in oligopolistic markets. We build on a concept of strategic behavior called *virtual bargaining* (Misyak et al. 2014), originally developed to model joint action and communication in cognitive science. The human ability to engage in virtual bargaining suggests that actors merely need to imagine a bargaining process: they can directly implement the results of a purely “virtual bargain,” without actually needing to communicate, when it is clear what the outcome of such a conversation would be. Crucially, this is possible even in one-shot interactions in the absence of learning and reputation effects. According to the virtual bargaining viewpoint, Smith’s concern that sellers’ conversations soon turn to collusion against the consumer can be broadened. Collusive behavior may occur even if the conversation can merely be imagined.

More concretely, we propose a formal definition of a virtual bargaining equilibrium that helps explain how players may coordinate on collusive outcomes. According to this theory, players identify agreements that they could potentially reach if they could enter into binding contracts (we use “agreement” and “strategy profile” interchangeably). They realize, however, that given the nature of their interaction, these agreements are not enforceable. Specifically, for each possible agreement, each player considers two scenarios for the opponent’s behavior. Under the first scenario, the opponent sticks to the agreement, whereas under the second, she¹ best responds to the agreement.

The players perceive strategic uncertainty as to which of these two options the opponent will choose. To model attitudes to strategic uncertainty, we suppose that the players are strategically pessimistic: each player uses the lowest of the payoffs under these two scenarios, called the worst payoff, to guide his behavior. This behavior can be rationalized using the maxmin model of ambiguity aversion (Gilboa and Schmeidler 1989). A player retains a strategy profile (called “feasible agreement”) if its worst payoff cannot be improved by a unilateral change of strategy. Formally, it turns out that the set of feasible agreements is equivalent to the set of Nash equilibria on the worst payoffs (as opposed to the payoffs). The set of feasible agreements contains all Nash equilibria and often includes additional strategy profiles. After identifying the set of feasible agreements via a mental simulation, the players envisage a bargaining process and choose a strategy profile from the set of possible feasible agreements. The strategy profile selected via this procedure is called the virtual bargaining equilibrium. It exists under standard regularity conditions and, in many games, the set of virtual bargaining and Nash equilibria differ.

The rationale for the virtual bargaining equilibrium can be illustrated in a domain that we do not directly

consider here: the puzzling question of the justification of mixed-strategy Nash equilibria. Consider the game of Matching Pennies, where each of two players chooses heads, tails or a *mixed* strategy (i.e., selecting a probability of choosing each option). One of the players receives a payoff of 1 if the pennies match and a payoff of -1 if they do not match, whereas the other player receives a payoff of -1 in the case of a match and a payoff of 1 in the case of a nonmatch. Here, the equilibrium is that both players choose heads and tails with probability $\frac{1}{2}$. But if player 1 chooses this mixed strategy, player 2 can costlessly choose any strategy (including playing heads with probability 1). But if player 2 chooses any other strategy, player 1’s $\frac{1}{2}$, $\frac{1}{2}$ strategy is not a best response. So there appears to be something unstable about the justification of both players playing the mixed strategy: if one does so, the other need not, and then the next iteration of reasoning leads to a *pure* (nonprobabilistic) strategy.

This puzzle disappears when one adopts a virtual bargaining point of view.² That is, if we ask, “Which feasible agreements are there in matching pennies?”, the answer is that there is just one, the mixed-strategy Nash equilibrium profile. From a virtual bargaining perspective, this agreement is stable because if each player follows it, their worst expected payoff is zero (irrespective of any “best response” the other makes—of course, here all responses from the other are tied). Any other agreement has a strictly smaller worst payoff (e.g., if player 1 agrees to play heads, then whatever player 2 might supposedly agree, her best response—say tails—will lead me to a certain loss.)

Both players can reason that the mixed-strategy agreement is the only one that they can credibly achieve (any opponent who promises that they will play a specific move can only be attempting to bluff, or double-, or triple-... bluff). According to virtual bargaining, each player is confident that she herself will follow the “bargain” (here, the mixed strategy), whether or not the other does. Thus, in contrast to Nash reasoning, players strictly prefer mixed strategies under virtual bargaining reasoning; not just in Matching Pennies, but more generally. Thus, virtual bargaining provides a simple justification for mixed-strategy equilibria, as an alternative to highly sophisticated and controversial proposals in the literature, such as “purification” (Harsanyi 1973).

Here, we use the virtual bargaining equilibrium in a different context: to explain and predict the potential for collusive behavior in Bertrand competition (i.e., firms compete on prices) and Cournot competition (i.e., firms compete on quantities). These two models of competition are central to understanding markets and business strategies (e.g., Cabral and Villas-Boas 2005). Our predictions require neither communication nor dynamic considerations (for collusion in a dynamic environment, see, e.g., Campbell et al. 2005).

Consider Bertrand competition with differentiated goods and two players (Singh and Vives 1984, Stigler 1964). Suppose that each player envisions a collusive virtual agreement involving higher prices. Each player realizes that the opponent will have an incentive to choose a price below the collusive level. In spite of the possibility of this deviation, the collusive outcome is more attractive than Nash—both players may make high profits even if one deviates from the agreement and undercuts the other. Moreover, there exists an agreement that each player finds more attractive, in terms of her worst payoff, than all agreements resulting from her own deviations. This agreement involves a higher price than the Nash equilibrium of the game. Thus, virtual bargaining enables the players to collude and achieve higher profits than in the Nash equilibrium.

To gain further insights into this finding, consider the Traveler's Dilemma game (Basu 1994). In this game, two players simultaneously and independently choose a sum of money—e.g., an integer sum between \$1 and \$100. Both players receive the lower of the two sums; then \$2 is transferred from the player who gave the “greedier” offer to the player with the more “modest” offer (no transfer occurs if both players propose the same sum of money). This means that each player's best response is to slightly undercut the other, and an inductive argument leads to a single Nash equilibrium: that both players choose \$1 and receive this very low payoff. Thus, although both players have a great deal to gain from coordinating on a high number (e.g., by both choosing \$100), the Nash equilibrium predicts that this opportunity will be lost. Collusion in this game amounts to choosing a higher sum of money. As might intuitively be expected, in experiments, people are often able to coordinate, reasonably effectively, to obtain high payoffs in this game (Capra et al. 1999). Consider a virtual agreement that stipulates that both players choose \$100. If a player sticks to this agreement, and her opponent either follows the agreement or best responds to it, then the least that she can get is \$97. This worst payoff of the strategy profile (\$100, \$100) is thus much more attractive than a payoff of \$1 under the Nash equilibrium. It is also more attractive than other potential collusive agreements. Thus, virtual bargaining leads to collusion at the highest possible level in this game (i.e., full collusion). In Bertrand competition, however, we show that the virtual bargaining equilibrium lies strictly between the Nash equilibrium and the price vector that maximizes the total of the two players' profits. Thus, virtual bargaining leads to partial collusion.

The reasoning that enables the players to collude in the Traveler's Dilemma and Bertrand competition has similarities. In both cases, the incentive to undercut the other player exists, but it is not too detrimental to the opponent. Moreover, if one player chooses a

higher price/number, the opponent's best response is to choose a higher price/number as well (but not to the same degree), thus moving in the direction of a collusive outcome. As a result, there exists a strategy profile that improves on Nash and offers some protection against the possibility that the opponent may deviate and best respond to the agreement.

In contrast, virtual bargaining does not have a bite in Cournot competition. If one of the players contemplates a reduction in quantity, she realizes that it will be met by an offsetting increase in the opponent's quantity, if the latter chooses to best respond. Thus, when one of the players considers moving toward a more collusive outcome, she realizes that the opponent will move in the direction of higher market dominance. In contrast to Bertrand competition and the Traveler's Dilemma game, Cournot competition is not conducive to collusive behavior.

Our analysis of implicit collusion envisions collusive behavior as emerging spontaneously and implicitly. Unlike explicit regulations or laws, the implicit agreements that guide social and economic behavior are typically not written down or subject to formal sanction. People appear, nonetheless, to be averse to violating implicit agreements. This aversion has been observed in “breaching” experiments where people are instructed to violate everyday social rules (e.g., Garfinkel 1967, Milgram and Sabini 1978). Such aversion may be amplified where there is explicit consent to an agreement—e.g., through an unenforceable “handshake” (Kessler and Leider 2012). Indeed, the desire to conform to implicit agreements can motivate behavior (Krupka et al. 2017). Moreover, the tendency to follow implicit agreements tends to “spill over” between contexts (Peysakhovich and Rand 2016). Implicit agreements may do more than guide behavior merely by imposing additional costs and benefits in a conventional utility-maximizing analysis. According to March and Olsen's (2008) influential theory, such agreements have an entirely separate “logic” for guiding action (for example, according to deontological rather than consequentialist theories of ethics, see Kamm 2007).

While following implicit agreements can be seen as a foundation for virtue, it can also operate against the public interest, as in the case of collusion. Here, the implicit nature of the agreement may be particularly advantageous for colluding parties, because an explicit agreement might violate the law. In cases where it is “obvious” to both parties that following the agreement (e.g., keeping prices high) will be mutually beneficial, explicit agreement may not be required. Both parties may know what they *would* agree, if they were able secretly to communicate.

2. Virtual Bargaining Equilibrium

Consider a game between two players, 1 and 2, who simultaneously and independently choose their strategies $\sigma_1 \in \Sigma_1$ and $\sigma_2 \in \Sigma_2$.³ Let $u_i(\sigma_i, \sigma_{-i})$ denote player i 's ($i = 1, 2$) payoff function, where player i 's opponent is denoted by $-i$ and the latter's strategy is denoted by σ_{-i} . We define the *worst payoff* of an agreement (σ_1^A, σ_2^A) for player $i = 1, 2$ as the worst-case scenario of two possibilities: (i) player i 's opponent goes through with the agreement by playing her part of the strategy profile (σ_1^A, σ_2^A) and (ii) best responds to σ_i^A . Formally, the worst payoff of an agreement (σ_1^A, σ_2^A) is given by

$$w_i(\sigma_i^A, \sigma_{-i}^A) = \min \left\{ u_i(\sigma_i^A, \sigma_{-i}^A), \sup_{\sigma_{-i} \in R_{-i}(\sigma_i^A)} u_i(\sigma_i^A, \sigma_{-i}) \right\}, \quad (1)$$

where $R_{-i}(\sigma_i^A)$ denotes the set of player $-i$'s best responses to strategy σ_i^A . Thus, player i allows for the possibility that player $-i$ will deviate from the agreement $(\sigma_i^A, \sigma_{-i}^A)$ and play a best response to strategy σ_i^A . Note that player i also believes that among all such best responses, player $-i$ will choose the best response that yields player i the highest payoff. In other words, each player believes that even if her opponent deviates from an agreement, she will not do so in a spiteful fashion. This assumption is common in many strategic environments. Examples include bargaining models, principal-agent models, and ultimatum games. In these environments, the party that receives an offer is indifferent between accepting and refusing it but chooses the former, which in turn benefits the party making the offer. The assumption of non-spitefulness ensures that the set of feasible agreements includes all Nash equilibria. However, this assumption does not influence the analytical results in the present paper, because in Cournot and Bertrand games the best response to any strategy is unique.

The players will be guided by worst payoffs if they have ambiguous beliefs about the strategies that their opponents will play and if their preference functional has Gilboa and Schmeidler's (1989) maxmin expected utility form. Suppose that for each strategy profile, each player entertains two possibilities for his opponent's behavior: "the agreement is honored" or "the opponent best responds in a nonspiteful fashion." Suppose also that each player has completely ambiguous beliefs about which of these two strategies will be played by the opponent (or, equivalently, each player considers each possible probability distribution over these two strategies). If a player with such beliefs has maxmin expected utility preferences, then his preference functional will be given by (1).⁴

Each player thinks of all strategy profiles as possible agreements that she could strike with her opponent; and for each such agreement, her choices are guided by the worst payoff. The two players narrow down the set

of all possible agreements and retain the agreements where neither player can improve her worst payoff by a unilateral deviation. We call these agreements "feasible." Formally, we will say that an agreement $(\sigma_i^F, \sigma_{-i}^F)$ is *feasible* if, for all $i \in \{1, 2\}$,

$$w_i(\sigma_i^F, \sigma_{-i}^F) \geq w_i(\tilde{\sigma}_i, \sigma_{-i}^F), \quad \text{for all } \tilde{\sigma}_i \in \Sigma_i. \quad (2)$$

The intuition is that a feasible agreement (σ_1^F, σ_2^F) is an arrangement that the two players could both credibly have reached if there is no bargain $(\tilde{\sigma}_1, \sigma_2^F)$ that would have a better "worst payoff" for player 1; and similarly, there is no bargain $(\sigma_1^F, \tilde{\sigma}_2)$ that would have a better "worst payoff" for player 2. Thus, each player uses the worst payoff function both when the player considers following an agreement and when she considers deviating from it. If an agreement is not feasible, then the two players discard it as a possible mode of behavior because at least one of the players has a more attractive alternative agreement (from the perspective of her worst payoff).

In a world where players focus on "payoffs," unilateral deviations imply that players are looking for higher individual payoffs: such deviations can potentially lead to exploitation of the other player. In a world where players focus on "worst payoffs," however, unilateral deviations have a different meaning. The question is not whether each player can deviate by best responding, thereby potentially exploiting the other player. Instead, the players consider unilateral deviations to figure out whether there is a more attractive agreement for them than the focal agreement.

Let $R_i^F(\sigma_{-i}) \equiv \arg \max_{\sigma_i \in \Sigma_i} w_i(\sigma_i, \sigma_{-i})$ denote player i 's best-response correspondence for the worst payoff function. Using this notation, the agreement $(\sigma_i^F, \sigma_{-i}^F)$ is feasible if and only if

$$\sigma_i^F \in R_i^F(\sigma_{-i}^F), \quad \text{for all } i \in \{1, 2\}. \quad (3)$$

Thus, our notion of a feasible agreement is equivalent to the Nash equilibrium, but in a world of "worst payoffs" rather than "payoffs." That is, feasible agreements ensure that no "more attractive" agreements are attainable through unilateral deviation. Therefore, it seems that for both players to consider a strategy profile as a candidate for a virtual bargain, it has to be feasible.

Note that a game may have multiple feasible agreements. We let F denote the set of feasible agreements. It follows immediately from the definition in (1) that any Nash equilibrium (NE) is a feasible agreement. However, there may be feasible agreements that violate the Nash requirements.

We use the symmetric game of Table 1 to demonstrate the concepts introduced above.⁵ Table 1(a) contains the normal form of the game, while Table 1(b) contains the worst payoffs for all pure-strategy profiles. In this game, each player $i = 1, 2$ has three pure

Table 1. Illustrative Game

Player 1	Player 2		
	H	L	N
(a) Normal form of the game			
H	(60, 60)	(33, 62)	(−10, 61)
L	(62, 33)	(36, 36)	(−6, 40)
N	(61, −10)	(40, −6)	(0, 0)
(b) Worst payoffs for pure strategies			
H	(33, 33)	(33, −6)	(−10, 0)
L	(−6, 33)	(−6, −6)	(−6, 0)
N	(0, −10)	(0, −6)	(0, 0)

strategies: H , L , and N . The set of pure-strategy feasible agreements consists only of the Nash equilibrium strategy profile (N, N) . In addition, the game has a feasible agreement in mixed strategies where each player plays strategy H with probability $\frac{6}{7}$ and strategy N with probability $\frac{1}{7}$. To demonstrate that the latter strategy profile, which is denoted by $((H, \frac{6}{7}; N, \frac{1}{7}), (H, \frac{6}{7}; N, \frac{1}{7}))$, is a feasible agreement, we verify that $\sup_{\sigma_{-i} \in R_{-i}((H, \frac{6}{7}; N, \frac{1}{7}))} u_i((H, \frac{6}{7}; N, \frac{1}{7}), \sigma_{-i}) = 34$. Since $u_i((H, \frac{6}{7}; N, \frac{1}{7}), (H, \frac{6}{7}; N, \frac{1}{7})) = 50.3$, the worst payoff for the strategy profile $((H, \frac{6}{7}; N, \frac{1}{7}), (H, \frac{6}{7}; N, \frac{1}{7}))$ is equal to 34. One can also demonstrate that any unilateral deviation from this agreement yields a worst payoff that is (weakly) smaller than 34. Thus, $((H, \frac{6}{7}; N, \frac{1}{7}), (H, \frac{6}{7}; N, \frac{1}{7}))$ is feasible. Moreover, this strategy profile is the only feasible agreement in addition to the Nash equilibrium (N, N) .

The example in Table 1 also illustrates that the support of feasible agreements may include strictly dominated strategies. The strategy H is strictly dominated by the strategy L for both players. However, H is played with probability $\frac{6}{7}$ in the mixed-strategy feasible agreement.

The rationale for virtual bargaining, which we explained in the introduction, implies that the players entertain the possibility that the opponent will stick to the bargain and play H . Given that virtual bargainers may not best respond and may even play strictly dominated strategies, one might refer to virtual bargaining as a form of “bounded rationality.” Within economics, bounded rationality is sometimes used as a general term for nonstandard, behavioral accounts to explain experimental findings “while retaining precision and cross-game generality” (Camerer et al. 2003, p. 194). In cognitive science, however, “bounded rationality” can generate misleading associations. Inspired by Herbert Simon, cognitive scientists typically use bounded rationality to refer to approximations to rational processes, which are “bounded” by computational limitations. Although virtual bargaining makes assumptions that deviate from Nash reasoning, it does not assume computational restrictions. Instead, virtual bargaining

assumes that players think of strategy profiles as possible agreements and associate “worst payoffs,” instead of “payoffs,” with these agreements.

What are the “rationality” properties of virtual bargaining? Virtual bargaining differs from the “rationality” assumptions underpinning the Nash equilibrium in the following ways. First, virtual bargains may involve strictly dominated strategies in terms of payoffs (but not in terms of “worst payoffs”). According to some strict conceptions of rationality (especially prevalent in standard game theory), choosing any dominated option must necessarily be irrational. However, examples of strategic interactions (e.g., the Traveler’s Dilemma, the Centipede Game, and Newcomb’s Problem) have been taken, by some, to imply that this criterion of rationality has some paradoxical consequences. Be that as it may, many theorists may view virtual bargaining as departing from strict rationality on these grounds.

Second, as in, for example, level- k reasoning, virtual bargaining involves inconsistent beliefs to some degree. On the one hand, each player considers two possibilities for each strategy profile: the other player will either stick to the bargain or best respond to it. On the other hand, each player knows for certain whether they themselves will follow the bargain or best respond. As discussed above, this “inconsistency” or “asymmetry”—is similar to that in real bargaining. This is a different type of inconsistency than in level- k reasoning, where each player assumes that their counterpart uses a lower level of reasoning than they do themselves.

When making their choices, the two players simulate a bargaining process that, given the players’ status quo positions, chooses one of the feasible agreements. The status quo position for each player in the bargaining process is given by her worst payoff from a feasible agreement that is worst for him. Formally, player i ’s fallback position is defined as $w_i^m = \min_{(\sigma_1^F, \sigma_2^F) \in F} w_i(\sigma_1^F, \sigma_2^F)$. In what follows, we will call w_i^m the minimum feasible worst payoff of player i . Note that for certain games—for example, in the Battle of the Sexes game— $\nexists (\sigma_1^F, \sigma_2^F) \in F$ such that

$$(w_1^m, w_2^m) = (w_1(\sigma_1^F, \sigma_2^F), w_2(\sigma_1^F, \sigma_2^F)).$$

That is, the minimum feasible worst payoffs of the two players may correspond to different feasible agreements. If the game has a unique feasible agreement, then both players receive their minimum feasible worst payoffs from playing the unique feasible agreement. In this case, the latter characterizes the status quo positions of both players.

The bargaining mechanism used to choose a virtual bargain is the Nash bargaining solution⁶ where the players’ status quo positions are their minimum feasible

worst payoffs. Formally, a *virtual bargaining equilibrium* (VBE) (σ_1^V, σ_2^V) is a feasible agreement that maximizes the product of differences between the players' worst payoffs under this strategy pair and the worst payoffs from the status quo subject to the constraint that both players' payoffs exceed their respective status quo payoffs:

$$(\sigma_1^V, \sigma_2^V) \in \arg \max_{(\sigma_1^F, \sigma_2^F) \in F} \prod_{i=1}^2 (w_i(\sigma_i^F, \sigma_{-i}^F) - w_i^m). \quad (4)$$

Note also that, for simplicity, we have assumed that the two players have equal bargaining powers. Our definition of the virtual bargaining equilibrium can be readily extended to the general case of arbitrary bargaining powers. Although the analysis is straightforward for the general case, for the purposes of compact notation, we maintain the assumption of equal bargaining powers throughout the paper. It follows immediately from the definition of the virtual bargaining equilibrium that under standard regularity conditions on the sets Σ_i and the payoff functions $u_i(\cdot, \cdot)$, the game has a virtual bargaining equilibrium. By the definition of the Nash bargaining solution, any virtual bargaining equilibrium will be Pareto-optimal (with respect to the players' worst payoffs) among the set of feasible agreements. Moreover, any virtual bargaining equilibrium is Pareto-undominated by any Nash equilibrium. At the same time, a game may have a strategy profile that is not feasible and whose payoff vector Pareto-dominates (with respect to the players' payoffs) all feasible agreements. The Prisoners' Dilemma is an example of such a game. Thus, our equilibrium notion is different from both Nash equilibrium and Pareto criterion (with respect to the players' payoffs).

Returning to the example of Table 1, note that the worst payoff of 34 from the feasible agreement $((H, \frac{6}{7}; N, \frac{1}{7}), (H, \frac{6}{7}; N, \frac{1}{7}))$ exceeds the minimum feasible payoff of zero. It then follows immediately from (4) that $((H, \frac{6}{7}; N, \frac{1}{7}), (H, \frac{6}{7}; N, \frac{1}{7}))$ is the unique virtual bargaining equilibrium.⁷ If both players follow through with this agreement, then they can significantly improve their payoffs compared to the Nash payoff of zero.

We now analyze Cournot and Bertrand competition in turn using virtual bargaining. The analysis provides a possible explanation why collusion arises so much more readily in the latter competitive environment.

3. Cournot Competition

Consider the following Cournot game with two firms, $i = 1, 2$, each producing output $q_i \in R_+$. We assume that the firms have constant marginal costs.⁸ Without any further loss of generality, these marginal costs are set equal to zero. The inverse demand functions for the two firms' products are given by

$$p_1 = \alpha - \beta q_1 - \gamma q_2 \quad \text{and} \quad p_2 = \alpha - \beta q_2 - \gamma q_1,$$

where $\alpha > \beta \geq \gamma > 0$. Thus, the goods produced by the two firms are substitutes and each firm i 's profit function is $\pi_i(q_i, q_{-i}) = (\alpha - \beta q_i - \gamma q_{-i})q_i$. The output choices are made by the two firms simultaneously and independently. Since $\partial^2 \pi_i / \partial q_1 \partial q_2 = -\gamma < 0$, the firms' outputs are strategic substitutes (Bulow et al. 1985, Topkis 1998).

Firm i 's best response to output q_{-i} is given by $R_i(q_{-i}) = (\alpha - \gamma q_{-i}) / (2\beta)$. The unique Nash equilibrium is given by $q_1^N = q_2^N = \alpha / (2\beta + \gamma)$. The output combination that maximizes the joint of the two firms' profits is given by $q_1^* = q_2^* = \alpha / (2\beta + 2\gamma)$. We will refer to the latter output combination as the collusive outcome.

The worst payoff of an agreement (q_1, q_2) for firm $i = 1, 2$ is given by

$$\begin{aligned} w_i(q_i, q_{-i}) &= \min \left\{ \pi_i(q_i, q_{-i}), \sup_{q'_{-i} \in R_{-i}(q_i)} \pi_i(q_i, q'_{-i}) \right\} \\ &= \min \left\{ \pi_i(q_i, q_{-i}), \pi_i(q_i, R_{-i}(q_i)) \right\} \\ &= \begin{cases} (\alpha - \beta q_i - \gamma q_{-i})q_i, & \text{if } q_i > \frac{\alpha - 2\beta q_{-i}}{\gamma}; \\ \frac{((2\beta - \gamma)\alpha - (2\beta^2 - \gamma^2)q_i)q_i}{2\beta}, & \text{if } q_i \leq \frac{\alpha - 2\beta q_{-i}}{\gamma}. \end{cases} \end{aligned} \quad (5)$$


Thus, when the weighted average of two players' quantities is large, player i 's worst payoff for (q_1, q_2) is equal to profit $\pi_i(q_i, q_{-i})$ for the scenario where the opponent follows her part of the agreement. Conversely, when the weighted average is small, player i 's worst payoff is equal to profit for the scenario where the opponent best responds. This property of the worst payoff function is intuitive. An agreement is relatively attractive when the market is not saturated. In this case, a best-response by an opponent will have a detrimental effect on a player relative to the payoff from the agreement. Conversely, when the market is saturated, the agreement itself is relatively unattractive and an opponent's best response to the agreement has a positive impact on a player.

We demonstrate in the appendix that the best-response function of player i for payoff function $w_i(q_i, q_{-i})$ is given by

$$R_i^F(q_{-i}) = \begin{cases} \frac{(2\beta - \gamma)\alpha}{2(2\beta^2 - \gamma^2)}, & \text{if } q_{-i} \leq \frac{\alpha}{2\beta} - \frac{\alpha(2\beta - \gamma)\gamma}{4\beta(2\beta^2 - \gamma^2)}; \\ \frac{\alpha - 2\beta q_{-i}}{\gamma}, & \text{if } \frac{\alpha}{2\beta} - \frac{\alpha(2\beta - \gamma)\gamma}{4\beta(2\beta^2 - \gamma^2)} < q_{-i} < \frac{\alpha}{2\beta + \gamma}; \\ \frac{\alpha - \gamma q_{-i}}{2\beta}, & \text{if } q_{-i} \geq \frac{\alpha}{2\beta + \gamma}. \end{cases} \quad (6)$$

Thus, the best-response function $R_i^F(q_{-i})$ is constant for relatively small values of q_{-i} and decreasing afterward.

Figure 1. (Color online) Best-Response Functions $R_1^F(q_2)$ and $R_2^F(q_1)$, NE, and VBE: Cournot Competition



The functions $R_1^F(q_2)$ and $R_2^F(q_1)$ are given by the kinked lines that pass through points $ADNEG$ and $BENDH$, respectively, in Figure 1. Point C in Figure 1 represents the collusive output combination (q_1^*, q_2^*) , while point N represents the Nash equilibrium (q_1^N, q_2^N) .

From the definition of a feasible agreement in (3), an agreement (q_1^F, q_2^F) is feasible if and only if $q_i^F \in R_i^F(q_{-i}^F)$ for $i = 1, 2$. It follows immediately from this fact and (6) that the set of feasible agreements for the Cournot game is given by

$$F = \left\{ (q_1, q_2): q_2 = \frac{\alpha - 2\beta q_1}{\gamma}, \right. \\ \left. \frac{\alpha}{2\beta} - \frac{\alpha(2\beta - \gamma)\gamma}{4\beta(2\beta^2 - \gamma^2)} \leq q_1 \leq \frac{\alpha}{2\beta + \gamma} \right\} \\ \cup \left\{ (q_1, q_2): q_1 = \frac{\alpha - 2\beta q_2}{\gamma}, \right. \\ \left. \frac{\alpha}{2\beta} - \frac{\alpha(2\beta - \gamma)\gamma}{4\beta(2\beta^2 - \gamma^2)} \leq q_2 \leq \frac{\alpha}{2\beta + \gamma} \right\}.$$

The set F is given by the union of the line segments EN and ND in Figure 1, which correspond to the first and second elements of the above union.

Using (4), we obtain that the unique virtual bargaining equilibrium coincides with the Nash equilibrium (point N in Figure 1).⁹ We summarize our findings in the following:

Proposition 1. *The unique VB equilibrium of the Cournot game coincides with the NE.*

The intuition why the VB collapses to Nash is as follows. Suppose that the players consider reducing their outputs below the Nash level. Because of strategic substitutability between the players' choices, each

player realizes that if she decreases her output and the opponent foresees the reduction, the latter player will compensate for the reduction by increasing her own output; and the more a player decreases her own output, the more the opponent will compensate by increasing her output. Hence, the worst payoff associated with smaller outputs will be determined by the scenario where the opponent increases her output by best responding. This, however, is a relatively unattractive scenario since the decrease in output is met by a noncollusive output increase by the opponent. Formally, an agreement with outputs that are lower than the Nash level for both players is not feasible. As a result, the VB equilibrium coincides with the NE.

4. Bertrand Competition

Virtual bargaining and Nash behavior are different when firms compete in prices. Consider the following Bertrand game with linear demand functions. There are two firms, $i = 1, 2$, in the market each producing output $q_i \in R_+$ and selling its output at price $p_i \in R_+$. As in Cournot competition, the firms have constant marginal costs that are normalized to zero. The direct demand functions for the two firms' products are given by

$$q_1 = a - bp_1 + cp_2 \quad \text{and} \quad q_2 = a - bp_2 + cp_1,$$

where $b \geq c > 0$. Thus, the goods sold by the two firms are substitutes. Each firm i 's profit function can be written as $\pi_i(p_i, p_{-i}) = (a - bp_i + cp_{-i})p_i$. The prices are chosen by the two firms simultaneously and independently. Since $\partial^2 \pi_i / \partial p_1 \partial p_2 = c > 0$, the prices of the two firms are strategic complements.

Firm i 's best response to price p_{-i} is given by $R_i(p_{-i}) = (a + cp_{-i}) / (2b)$. The unique Nash equilibrium is given by $p^N \equiv p_1^N = p_2^N = a / (2b - c)$. The Nash equilibrium output is given by $q^N \equiv q_1^N = q_2^N = ab / (2b - c)$.

The price combination that maximizes the total of the two firms' profits is given by $p^* \equiv p_1^* = p_2^* = a / (2b - 2c)$. It is straightforward to verify that the joint profit-maximizing price p^C exceeds the Nash equilibrium price p^N . The corresponding output levels have the reverse relationship: $q^* = a/2 < q^N$. As in the previous section, we refer to this outcome as collusive.

The worst payoff of an agreement (p_1, p_2) for firm $i = 1, 2$ is given by

$$w_i(p_1, p_2) = \min \left\{ \pi_i(p_i, p_{-i}), \sup_{p'_{-i} \in R_{-i}(p_i)} \pi_i(p_i, p'_{-i}) \right\} \\ = \min \{ \pi_i(p_1, p_2), \pi_i(p_i, R_{-i}(p_i)) \} \\ = \begin{cases} (a - bp_i + cp_{-i})p_i, & \text{if } p_i > \frac{2bp_{-i} - a}{c}; \\ \left(a - bp_i + \frac{c(a + cp_i)}{2b} \right) p_i, & \text{if } p_i \leq \frac{2bp_{-i} - a}{c}. \end{cases} \quad (7)$$

Thus, when a player's price is relatively large compared to the opponent's price, the player's worst payoff for (p_1, p_2) is equal to profit $\pi_i(p_1, p_2)$ for the scenario where the opponent follows her part of the agreement. In this case, the opponent's best response is a price that is large relative to the agreement's prescription. This, in turn, results in a profit that is larger than the profit $\pi_i(p_1, p_2)$ under the agreement. Conversely, when a player's price is relatively small compared to the opponent's price, the player's worst payoff is equal to her profit for the case where the opponent best responds.

We demonstrate in the appendix that the best-response function of player i for the worst payoff function $w_i(p_1, p_2)$ is given by

$$R_i^F(p_{-i}) = \begin{cases} \frac{a + cp_{-i}}{2b}, & \text{if } p_{-i} \leq \frac{a}{(2b-c)}; \\ \frac{2bp_{-i} - a}{c}, & \text{if } \frac{a}{(2b-c)} < p_{-i} < \frac{a(4b^2 - c^2 + 2bc)}{4b(2b^2 - c^2)}; \\ \frac{a(2b+c)}{2(2b^2 - c^2)}, & \text{if } p_{-i} \geq \frac{a(4b^2 - c^2 + 2bc)}{4b(2b^2 - c^2)}. \end{cases} \quad (8)$$

Thus, the best-response function $R_i^F(p_{-i})$ is constant for relatively small values of p_{-i} and for relatively large values. The functions $R_1^F(p_2)$ and $R_2^F(p_1)$ are given by the kinked lines that pass through points $ANDVG$ and $BNEVH$, respectively, in Figure 2. Point C in Figure 2 represents the combination of prices (p_1^*, p_2^*) that maximizes the total of the two firms' profits. The Nash equilibrium (p_1^N, p_2^N) is given by point N in Figure 2.

From the definition of a feasible agreement, an agreement (p_1^F, p_2^F) is feasible if and only if $p_i^F \in R_i^F(p_{-i}^F)$ for $i = 1, 2$. It follows immediately from this characterization of feasible agreements and (7) that the only

feasible agreements in the case of Bertrand competition are the Nash equilibrium (p_1^N, p_2^N) (point N in Figure 2) and the price combination (p_1^V, p_2^V) (point V in Figure 2) where the latter is given by

$$p^V \equiv p_1^V = p_2^V = \frac{a(2b+c)}{2(2b^2-c^2)}. \quad (9)$$

Using (4), we obtain that the unique virtual bargaining equilibrium is given by (p_1^V, p_2^V) . We summarize our findings in the following:

Proposition 2. (a) *The unique VB equilibrium price of the Bertrand game is strictly greater than the NE price and strictly smaller than the price vector that maximizes the joint profits of the two firms:*

$$p^N < p^V < p^*.$$

(b) *The joint profits under the VB equilibrium in the Bertrand game are strictly greater than the joint profits of the two firms under the NE but strictly smaller than the joint profits under the collusive outcome.*


This result is in stark contrast with the finding for Cournot competition where virtual bargaining leads to the noncollusive Nash equilibrium outcome. Virtual bargaining allows the two players to partially collude by raising the price above the Nash level. To see the intuition behind this finding, suppose that the players consider moving toward a collusive scenario and increasing the prices above the Nash level. Because of strategic complementarity, each player realizes that an increase in price will be met by a higher price of her opponent. Hence, the worst payoff associated with higher prices will be determined by the scenario where the players follow the agreement to increase their prices. Since this is a relatively attractive option, the players will favor a collusive price increase.

Both the VB and the collusive prices decrease with increases in b and decreases in c , and the latter price exhibits greater reductions. Thus, the VB price is closer to the collusive price (the ratio of the two prices $p^V/p^* = 1 - bc/(2b^2 - c^2)$ is large) when the slope of the direct demand curve is large with respect to own price (large b) and small with respect to the price of the competitor (small c).

Similarly, both the NE and VB prices decrease with increases in b and decreases in c and the latter price exhibits greater reductions. Thus, the NE price is closer to the VB price (the ratio of the two prices $p^N/p^V = 1 - c^2/(4b^2 - c^2)$ is large) when the slope of the direct demand curve is large with respect to own price (large b) and small with respect to the price of the competitor (small c).

The changes in the ratio of the profits under different outcomes with respect to parameters b and c mimic the changes in prices. The ratio of the profits under the VB

Figure 2. (Color online) Best-Response Functions $R_1^F(p_2)$ and $R_2^F(p_1)$, NE, and VBE: Bertrand Competition



equilibrium to the profits under the collusive outcome, $\pi_i(p^V, p^V)/(\pi_i(p^*, p^*)) = 1 - b^2 c^2 / (2b^2 - c^2)^2$, increases with the slope b of the direct demand with respect to own price and decreases with the slope c of the direct demand with respect to the competitor's price. Analogously, the ratio of the profits under the NE to the profit under the VB equilibrium, $\pi_i(p^N, p^N)/(\pi_i(p^V, p^V)) = 1 - c^4(3b - c)/((2b - c)^3(2b + c)(b + c))$, increases with b and decreases with c .

5. Comparison with Alternative Accounts

It is illuminating to contrast the concept of virtual bargaining with alternative non-Nash accounts of strategic behavior, such as the cognitive hierarchy and level- k models (e.g., Camerer et al. 2004, Nagel 1995, Stahl and Wilson 1994) and quantal response equilibrium (QRE) (e.g., Goeree et al. 2002, 2003; McKelvey and Palfrey 1995). The most crucial difference is that virtual bargaining offers a cognitive account of actors being guided by “unspoken agreements” (e.g., Carlin et al. 2007, Jones 1921). Thus, virtual bargaining makes a distinctive contribution in explaining the underlying cognitive mechanism and ensuing behavior.

A second important difference is that the standard versions of level- k reasoning and QRE rely on players choosing from among discrete alternatives. Virtual bargaining, in contrast, enables the players to make continuous choices. In the real world, we can think of quantities, prices, and effort levels as quasi-continuous measures. Equally importantly, if we model such measures as discrete, then the choice of discretization can generate artifacts in some theories. For example, a level- k reasoner might draw back from complete collusion by k iterations—but whether this corresponds to k cents or k dollars fewer than the completely collusive solution depends on the precision with which amounts are defined.

In addition to these two crucial differences, it is unclear what level- k theory and QRE would predict for the competitive settings considered in our paper. We are not aware of any published papers that apply these two theories to Cournot and Bertrand competition. In the case of level- k reasoning, the predictions would be highly sensitive to the starting point (i.e., to the assumption made about level-0). If players face a fine-grained grid of choices (in terms of quantities and prices), the level- k prediction will be very close to level-0, because most players do not make more than three iterative steps of reasoning (e.g., Arad and Rubinstein 2012).

In the case of QRE, we can view QRE as a “generalization of the Nash equilibrium, and it converges to a Nash outcome as [the ‘error’ parameter] μ goes to zero (perfect rationality)” (Goeree et al. 2003, p. 101). QRE's deviation from Nash equilibria crucially depends on the parameter μ , and the predictions for Cournot and Bertrand competition would be highly sensitive to the

assumption about this parameter. Beyond determining μ , we would also need conceptual arguments to clarify whether μ is a universal parameter, or whether it should differ for Cournot versus Bertrand settings.

A notion related to virtual bargaining equilibrium is that of maximin equilibrium introduced by Ismail (2014a, b). He proposes a model of equilibrium behavior where the worst-case payoff from a strategy profile is defined as the minimum payoff the player receives (i) from that profile or (ii) when the opponent better-responds. According to his procedure, the game is transformed into another game in which the original payoffs are replaced with the worst-case payoffs. Accordingly, an agreement is called a maximin equilibrium if its worst-case payoffs are Pareto-optimal (Ismail 2014b), or it is a Nash equilibrium of the transformed game (Ismail 2014a). Note that virtual bargaining equilibrium chooses an agreement that must be a specific Pareto-optimal Nash equilibrium in the transformed game corresponding to the feasible payoffs, whereas maximin equilibrium chooses all Pareto-optimal profiles or all the Nash equilibria in the transformed game corresponding to the procedure in Ismail (2014a, b).

In sum, we can conclude that virtual bargaining makes a distinctive contribution in explaining and predicting collusion in Cournot and Bertrand settings. And these two settings are, of course, foundational for understanding competitive behavior. This discussion should not imply that we want to discard level- k reasoning and QRE as models of strategic behavior. Both approaches have generated a great many useful insights. Moreover, these models have significant explanatory power in many strategic situations. We leave to future research the interesting research question of comparing, both theoretically and empirically, level- k reasoning and QRE with reasoning based on virtual bargaining.

6. Conclusion

This paper offers a novel account of collusive behavior that distinguishes between outcomes in Bertrand and Cournot competition. Virtual bargaining contains elements of both noncooperative and cooperative reasoning. It sustains collusive outcomes without third-party enforcement. Collusion can be a result of virtual bargaining as a mode of reasoning, thereby requiring neither communication nor dynamic considerations, such as rewards and punishments, between the players. The model provides insights into why collusion in price competition can be sustained without communication and why it is hard to detect.

Acknowledgments

The authors thank department editor Manel Baucells, the associate editor, and the anonymous referees for many valuable comments that led to significant improvements in the paper.

Appendix

Derivation of (6)

We argued in Section 3 that an agreement is feasible if and only if it is a best response for the worst payoff function. To determine this best-response correspondence, consider the properties of $w_i(q_i, q_{-i})$ as a function of q_i . Using (5), we obtain the following. When $q_{-i} \geq \alpha/(2\beta + \gamma)$, $w_i(q_i, q_{-i})$, as a function of q_i , increases on the interval $[0, (\alpha - 2\beta q_{-i})/\gamma]$, has a kink at $(\alpha - 2\beta q_{-i})/\gamma$, continues to increase on the interval $[(\alpha - 2\beta q_{-i})/\gamma, (\alpha - \gamma q_{-i})/(2\beta)]$, achieves its unique maximum at $(\alpha - \gamma q_{-i})/(2\beta)$, and decreases for the values of q_i exceeding $(\alpha - \gamma q_{-i})/(2\beta)$. When $\alpha/(2\beta) - (\alpha(2\beta - \gamma)\gamma)/(4\beta(2\beta^2 - \gamma^2)) < q_{-i} < \alpha/(2\beta + \gamma)$, $w_i(q_i, q_{-i})$ increases on the interval $[0, (\alpha - 2\beta q_{-i})/\gamma]$, has a kink and a unique maximum at $(\alpha - 2\beta q_{-i})/\gamma$, and decreases for the values of q_i exceeding $(\alpha - 2\beta q_{-i})/\gamma$. When $q_{-i} \leq \alpha/(2\beta) - (\alpha(2\beta - \gamma)\gamma)/(4\beta(2\beta^2 - \gamma^2))$, $w_i(q_i, q_{-i})$ increases on the interval $[0, (2\beta - \gamma)\alpha/(2(2\beta^2 - \gamma^2))]$, achieves its unique maximum at $(2\beta - \gamma)\alpha/(2(2\beta^2 - \gamma^2))$, decreases on the interval $((2\beta - \gamma)\alpha/(2(2\beta^2 - \gamma^2)), (\alpha - 2\beta q_{-i})/\gamma]$, has a kink at $(\alpha - 2\beta q_{-i})/\gamma$, and continues to decrease for the values of q_i exceeding $(\alpha - 2\beta q_{-i})/\gamma$.

Thus, $w_i(q_i, q_{-i})$ is single peaked as a function of q_i for all three possible ranges of q_{-i} spelled out in the preceding paragraph. The unique peak in each of these cases corresponds to the best-response function in expression (6).

Derivation of (8)

Consider the properties of $w_i(p_i, p_{-i})$ as a function of p_i . When $p_{-i} < a/(2b - c)$, $w_i(p_i, p_{-i})$, as a function of p_i , increases on the interval $[0, (2bp_{-i} - a)/c]$, has a kink at $(2bp_{-i} - a)/c$, increases on the interval $[(2bp_{-i} - a)/c, (a + cp_{-i})/(2b)]$, achieves its unique maximum at $(a + cp_{-i})/(2b)$, and decreases for the values of p_i exceeding $(a + cp_{-i})/(2b)$. When $a/(2b - c) < p_{-i} < a(4b^2 - c^2 + 2bc)/(4b(2b^2 - c^2))$, $w_i(p_i, p_{-i})$ increases on the interval $[0, (2bp_{-i} - a)/c]$, has a kink and a unique maximum at $(2bp_{-i} - a)/c$, and decreases for the values of p_i exceeding $(2bp_{-i} - a)/c$. When $p_{-i} > a(4b^2 - c^2 + 2bc)/(4b(2b^2 - c^2))$, $w_i(p_i, p_{-i})$ increases on the interval $[0, a(2b + c)/(2(2b^2 - c^2))]$, achieves its unique maximum at $a(2b + c)/(2(2b^2 - c^2))$, decreases on the interval $(a(2b + c)/(2(2b^2 - c^2)), (2bp_{-i} - a)/c]$, has a kink at $(2bp_{-i} - a)/c$, and continues to decrease for the values of q_i exceeding $(2bp_{-i} - a)/c$.

Note that $w_i(p_i, p_{-i})$ is single peaked as a function of p_i in all three cases characterized in the preceding paragraph. The unique peak in each of these cases corresponds to the best-response function in expression (8).

Endnotes

¹ We refer to the opponent as “she.”

² In the matching pennies game, and in constant-sum games more generally, by playing a mixed strategy, a player minimizes the opponent’s ability to distinguish and exploit systematic patterns of behavior.

³ We define the notion of a virtual bargaining equilibrium for normal form games with arbitrary number of players elsewhere.

⁴ We are indebted to the department editor for suggesting this interpretation of the worst payoffs.

⁵ It is straightforward to verify that, similarly to Bertrand competition, the game of Table 1 is supermodular (Milgrom and Shannon 1994, Topkis 1998).

⁶ One could allow for alternative bargaining mechanisms to arrive at the virtual bargain.

⁷ This holds for any distribution of bargaining powers among the players.

⁸ Our assumptions of constant marginal costs and linear demands are made to save on notation and to facilitate illustration of the virtual bargaining equilibrium. Our main findings hold for a general Cournot model with standard assumptions on the demand and cost functions. The same comment applies to the model of Bertrand competition.

⁹ Note that if we relaxed the assumption of equal bargaining power between the players, then the VB equilibrium could be an agreement in the set F that was different from the Nash equilibrium.

References

- Arad A, Rubinstein A (2012) The 11–20 money request game: A level-k reasoning study. *Amer. Econom. Rev.* 102(7):3561–3573.
- Basu K (1994) The traveler’s dilemma: Paradoxes of rationality in game theory. *Amer. Econom. Rev.* 84(2):391–395.
- Bulow JI, Geanakoplos JD, Klemperer PD (1985) Multimarket oligopoly: Strategic substitutes and complements. *J. Political Econom.* 93(3):488–511.
- Cabral LMB, Villas-Boas M (2005) Bertrand supertraps. *Management Sci.* 51(4):599–613.
- Camerer C, Ho T, Chong K (2003) Models of thinking, learning, and teaching in games. *Amer. Econom. Rev.* 93(2):192–195.
- Camerer CF, Ho TH, Chong JK (2004) A cognitive hierarchy model of games. *Quart. J. Econom.* 119(3):861–898.
- Campbell C, Ray G, Muhanna WA (2005) Search and collusion in electronic markets. *Management Sci.* 51(3):497–507.
- Capra CM, Goeree JK, Gomez R, Holt CA (1999) Anomalous behavior in a traveler’s dilemma? *Amer. Econom. Rev.* 89(3):678–690.
- Carlin BI, Lobo MS, Viswanathan S (2007) Episodic liquidity crises: Cooperative and predatory trading. *J. Finance* 62(5):2235–2274.
- Dufwenberg M, Gneezy U (2000) Price competition and market concentration: An experimental study. *Internat. J. Indust. Organ.* 18(1):7–22.
- Economist*, The (2016) No truck with cartels. *The Economist* (June 2), <http://www.economist.com/news/business/21699957-expensive-times-companies-accused-collusion-no-truck-cartels>.
- Engel C (2007) How much collusion? A meta-analysis of oligopoly experiments. *J. Competition Law Econom.* 3(4):491–549.
- Fatas E, Haruvy E, Morales AJ (2014) A psychological reexamination of the Bertrand paradox. *Southern Econom. J.* 80(4):948–967.
- Friedman J (1977) *Oligopoly and the Theory of Games* (North-Holland, Amsterdam).
- Garfinkel H (1967) *Studies in Ethnomethodology* (Prentice Hall, Englewood Cliffs, NJ).
- Gilboa I, Schmeidler D (1989) Maxmin expected utility with non-unique prior. *J. Math. Econom.* 18(2):141–153.
- Goeree JK, Holt CA, Palfrey TR (2002) Quantal response equilibrium and overbidding in private-value auctions. *J. Econom. Theory* 104(1):247–272.
- Goeree JK, Holt CA, Palfrey TR (2003) Risk averse behavior in generalized matching pennies games. *Games Econom. Behav.* 45(1):97–113.
- Harsanyi JC (1973) Games with randomly disturbed payoffs: A new rationale for mixed-strategy equilibrium points. *Internat. J. Game Theory* 2(1):1–23.
- Ismail M (2014a) Maximin equilibrium. Unpublished manuscript.
- Ismail M (2014b) Maximin equilibrium. GSBE Research Memoranda; No. 037. Maastricht: GSBE.
- Jones E (1921) *The Trust Problem in the United States* (Macmillan, New York).
- Kamm FM (2007) *Intricate Ethics: Rights, Responsibilities, and Permissible Harms* (Oxford University Press, Oxford, UK).
- Kessler JB, Leider S (2012) Norms and contracting. *Management Sci.* 58(1):62–77.

- Kreps DM, Scheinkman JA (1983) Quantity precommitment and Bertrand competition yield Cournot outcomes. *Bell J. Econom.* 14(2):326–337.
- Krupka EL, Leider S, Jiang M (2017) A meeting of the minds: Informal agreements and social norms. *Management Sci.* 63(6): 1708–1729.
- March JG, Olsen JP (2008) The logic of appropriateness. Goodin RE, Moran M, Rein M, eds. *The Oxford Handbook of Public Policy* (Oxford University Press, Oxford, UK), 689–708.
- McKelvey RD, Palfrey TR (1995) Quantal response equilibria for normal-form games. *Games Econom. Behav.* 10(1):6–38.
- Milgram S, Sabini J (1978) On maintaining urban norms: A field experiment in the subway. *Adv. Environment. Psych.* 1:31–40.
- Milgrom P, Shannon C (1994) Monotone comparative statics. *Econometrica* 62(1):157–180.
- Misyak JB, Melkonyan T, Zeitoun H, Chater N (2014) Unwritten rules: Virtual bargaining underpins social interaction, culture, and society. *Trends Cognitive Sci.* 18(10):512–519.
- Nagel R (1995) Unraveling in guessing games: An experimental study. *Amer. Econom. Rev.* 85(5):1313–1326.
- Peysakhovich A, Rand DG (2016) Habits of virtue: Creating norms of cooperation and defection in the laboratory. *Management Sci.* 62(3):631–647.
- Potters J, Suetens S (2013) Oligopoly experiments in the current millennium. *J. Econom. Surveys* 27(3):439–460.
- Singh N, Vives X (1984) Price and quantity competition in a differentiated duopoly. *RAND J. Econom.* 15(4):546–554.
- Smith A (1937) *An Inquiry Into the Nature and Causes of the Wealth of Nations* (The Modern Library, New York).
- Stahl DO, Wilson PW (1994) Experimental evidence on players models of other players. *J. Econom. Behav. Organ.* 25(3):309–327.
- Stigler GJ (1964) A theory of oligopoly. *J. Political Econom.* 72(1): 44–61.
- Suetens S, Potters J (2007) Bertrand colludes more than Cournot. *Experiment. Econom.* 10(1):71–77.
- Topkis DM (1998) *Supermodularity and Complementarity* (Princeton University Press, Princeton, NJ).
- Ungoed-Thomas J, Lord A (2016) It's a bargain on Amazon—Or is it? *The Sunday Times* (November 13), <http://www.thetimes.co.uk/edition/news/its-a-bargain-on-amazon-or-is-it-57w52rg90>.

CORRECTION

In this article, “Collusion in Bertrand vs. Cournot Competition: A Virtual Bargaining Approach” by Tigran Melkonyan, Hossam Zeitoun, and Nick Chater (first published in *Articles in Advance*, November 22, 2017, *Management Science* 64(12):5599–5609, DOI: 10.1287/mnsc.2017.2878), a paragraph has been added to Section 5, “Comparison with Alternative Accounts,” describing a notion related to virtual bargaining equilibrium, maximin equilibrium.