Discussion Paper No. 1014

COMMITMENT AND OBSERVABILITY IN GAMES*

by

Kyle Bagwell**

November 1992

Abstract

Models of commitment make two assumptions: there is a first mover, and his action is perfectly observed by the subsequent mover. The purpose of this paper is to disentangle these two assumptions, in order to see if a strategic benefit from commitment remains when the first-mover's choice is imperfectly observed. The basic finding is that the first-mover advantage is eliminated when there is even a *slight* amount of noise associated with the observation of the first-mover's selection.

---

*I thank Ken Judd, Giovanni Maggi, Roger Myerson, Wolfgang Pesendorfer, Rob Porter, Asher Wolinsky and especially Ehud Kalai and Steve Matthews for general comments and encouragement. I am also particularly grateful to Garey Ramey, who has regularly emphasized to me the general theme that small probability events can drastically affect the set of equilibrium outcomes.

**Department of Economics, Northwestern University, Evanston, Illinois 60208.

I. INTRODUCTION

One of the central insights of noncooperative game theory is the notion that a commitment to a course of action confers a strategic advantage. This idea dates back to von Stackelberg (1934), who illustrated the advantage of moving first in a duopoly context. The general significance of this insight, however, was not properly emphasized until the arrival of Schelling's (1960) celebrated book. The concept of a "first-mover advantage" is now deeply rooted in the mindset of economic theorists of all stripes. The fields of Industrial Organization, Information Economics, and International Trade, for example, all have been fundamentally influenced by the insight that commitment possesses a strategic virtue.

While the importance of commitment has surfaced in a wide range of papers, a common set of assumptions unite the various examples. Specifically, the advantage from commitment is modeled as springing from a combination of two assumptions. First, moves in the game are sequential, with some players committing to actions before other players select their respective actions, and, second, the late-moving players perfectly observe the actions selected by the first movers. These assumptions are so frequently combined that it is easy to forget that they are not equivalent. The purpose of this paper to disentangle these two assumptions, in order to determine if an advantage to moving first remains when rival players imperfectly observe the actions selected by the first movers.

To investigate this issue, I consider a noisy-leader game, in which one player moves first and then a second player observes a signal of the first-mover's actual selection before making his own move. With respect to the signal technology, I impose only a "non-moving support" assumption, whereby each of a fixed set of signals is possible for any action that the first mover might select. This allows that the signal may be very well correlated with the actual choice of the first mover, perhaps even almost perfectly so.

The basic result is striking. Assuming that the second-mover's best-response correspondence is single-valued, the set of pure-strategy Nash equilibrium outcomes for the noisy-leader game coincides exactly with the set of pure-strategy Nash equilibrium outcomes for the associated simultaneous-move game. With even the *slightest* degree of imperfection in the observability of the first-mover's selection, therefore, the strategic benefit of commitment is totally lost.

The paper proceeds as follows. I present a simple motivating example in Section 2, and characterize the result for general commitment games in Section 3. Next, in Section 4, I investigate whether an advantage to commitment might reappear in the context of mixed-strategy Nash equilibria. Concluding thoughts are offered in Section 5.

2. AN EXAMPLE

Consider a simple 2x2 setting, in which there are two players who each choose one of two actions, C and S. The payoff matrix is:

	S	C
S	(5,2)	(3,1)
C	(6,3)	(4,4)

The (S,S) outcome is the "Stackelberg outcome," since this is the unique subgame perfect equilibrium outcome for the game in which Player 1 moves first in a perfectly-observable fashion. The unique Nash equilibrium outcome of the simultaneous-move game is (C,C), and this corresponds to the "Cournot outcome."[1]

[1]The payoffs selected here bear a resemblence to those found in the traditional Cournot/Stackelberg quantity games (with S denoting the leader (follower) output for Player 1 (2)). While the results described below also apply to these traditional games, I focus on the 2x2 game because of its simplicity and because it is amenable to mixed-strategy analysis in a later section.

Consider next the noisy-leader game. In this game, a pure strategy for Player 1 is simply an action, $a_1 \in \{C,S\}$. Let the signal received by Player 2 be denoted by ϕ, and assume for simplicity that ϕ either takes value C or S. The signal technology works as follows:

$$\text{Prob}(\phi = S \mid S) = 1 - \varepsilon = \text{Prob}(\phi = C \mid C).$$

where $\varepsilon \in (0,1)$. In other words, when Player 1 chooses a particular action, the probability that Player 2 will observe a signal specifying that same action is $1 - \varepsilon$. A pure strategy for Player 2 is then a function, $a_2(\phi)$, where $a_2 \in \{C,S\}$ for all ϕ.

The noisy-leader game admits no off-equilibrium-path information sets, and so it suffices to look for Nash equilibria. It is instructive to consider first the possibility that Player 1 selects his Stackelberg leader strategy, $a_1 = S$. Suppose, then, that this strategy is played in a Nash equilibrium of the noisy-leader game. Player 2's strategy then must be a best response against $a_1 = S$, and as a consequence Player 2 must select S no matter what signal he receives; i.e., $a_2(\phi) \equiv S$ is Player 2's best response to $a_1 = S$. But, since Player 2's best reply is to ignore the actual signal value, it follows that Player 1 can induce the higher defect payoff of 6 by deviating to C. The Stackelberg outcome therefore fails to emerge as a Nash equilibrium outcome for the noisy-leader game, and this is true no matter how precise the signal may be (i.e., no matter how small is ε).

In fact, the unique pure-strategy Nash equilibrium of the noisy-leader game occurs when Player 1 selects C and Player 2 also selects C for all signal values (i.e., $a_1 = C$ and $a_2(\phi) \equiv C$). This seems a peculiar result when ε is small. Would not Player 1 deviate to S and establish his Stackelberg advantage? The key pont is that when Player 2 expects Player 1 to follow the equilibrium strategy, Player 2 implicitly reasons that the S signal can arise only if an "incorrect" signal is generated. Given this interpretation, player 2 disregards the signal and plays C even if the S signal is received. Player 1 is then unable to affect player 2's choice, and so Player 1's first-mover advantage is completely lost.

3. A GENERAL MODEL

The results described above are easily generalized. To begin, I define a two-person simultaneous-move game G as follows. Let A_i be a finite action space for Player i, and let a_i represent an element of this set. Player i's payoff function is then $u_i(a_1,a_2)$. A pure-strategy Nash equilibrium for the simultaneous-move game G is thus any strategy pair, $(\hat{a}_1,\hat{a}_2)$, such that $u_1(\hat{a}_1,\hat{a}_2) \geq u_1(a_1,\hat{a}_2)$ and $u_2(\hat{a}_1,\hat{a}_2)$, $\geq u_2(\hat{a}_1,a_2)$ for all $a_1 \in A_1$ and $a_2 \in A_2$.

In the noisy-leader game Γ, a pure strategy for Player 1 is again an action $a_1 \in A_1$. Player 2, however, oberves a signal, ϕ, before selecting his action. If Φ is the range of possible values for ϕ, then a pure strategy for Player 2 is a function, $a_2{:}\Phi \rightarrow A_2$, and I shall write this strategy as $a_2(\phi)$. As for the signal technology, I assume that Φ is finite and define $f(\phi \mid a_1)$ to be the probability of the signal ϕ given Player 1's actual action a_1. Most importantly, a non-moving support property is assumed, whereby Φ is independent of a_1 and $f(\phi \mid a_1) > 0$ for all $\phi \in \Phi$ and $a_1 \in A_1$.

Two points should be stressed about the signal technology. First, even though the signal support is independent of Player 1's action, the signal may be very well correlated with the actual action of Player 1, and this correlation in fact may be almost perfect. Second, the non-moving support assumption implies that no off-equilibrium path information sets exist, and so the appropriate solution concept is Nash equilibrium.

The payoffs facing the players in the noisy-leader game are denoted by $v_i(a_1,a_2(\phi))$, where:

$$v_i(a_1,a_2(\phi)) = \sum f(\phi \mid a_1)u_i(a_1,a_2(\phi))$$

A pure-strategy Nash equilibrium for the noisy-leader game then consists of a strategy pair, $(\hat{a}_1,\hat{a}_2(\phi))$, such that $v_1(\hat{a}_1,\hat{a}_2(\phi)) \geq v_1(a_1,\hat{a}_2(\phi))$ and $v_2(\hat{a}_1,\hat{a}_2(\phi)) \geq v_2(\hat{a}_1,a_2(\phi))$ for all $a_1 \in A_1$ and for all functions $a_2(\phi)$.

In order to compare the equilibrium outcomes of the two games, it is useful to define the best-response correspondences in the simultaneous-move game. Let $R_1(a_2)$ denote the set of a_1 values that maximize $u_1(a_1,a_2)$ over A_1. Define $R_2(a_1)$ symmetrically for Player 2. Finally, the basic results are most easily reported under the further assumption that for any action by Player 1, a unique best-reply action exists for Player 2; in other words, I assume below that $R_2(a_1)$ is single-valued.

A couple of observations now may be made:

Observation 1: In any pure-strategy Nash equilibrium of the noisy-leader game Γ, $\hat{a}_2(\phi) \equiv R_2(\hat{a}_1)$.

In a pure-strategy Nash equilibrium, Player 2's strategy must be optimal given the equilibrium strategy of Player 1; consequently, Player 2 selects his best-reply action for all signal values, no matter how unlikely the given signal may be under the hypothesized equilibrium. Since Player 2 therefore ignores the signal's value, a deviation by Player 1 does not affect Player 2's selection, and so Player 1 must be choosing a best reply as well:

Observation 2: In any pure-strategy Nash equilibrium of the noisy-leader game Γ, $\hat{a}_1 \in R_1(R_2(\hat{a}_1))$.

With these observations in place, it is now clear that the pure-strategy Nash equilibrium outcomes for the simultaneous-move and noisy-leader games can be described simply as action pairs. Specifically, an action pair $(\hat{a}_1,\hat{a}_2)$ forms a pure-strategy Nash equilibrium outcome for G if the same action pairing constitutes a Nash equilibrium for this game. Similarly, for Γ, an action pair $(\hat{a}_1,\hat{a}_2)$ is said to form a pure-strategy Nash equilibrium outcome if the strategies $\hat{a}_1$ and $\hat{a}_2(\phi) \equiv \hat{a}_2$ form a Nash equilibrium for Γ.

The primary proposition now may be stated:

PROPOSITION: The sets of pure-strategy Nash equilibrium outcomes for the simultaneous-move game G and the noisy-leader game Γ coincide.

The proof is simple. Suppose that an action pair $(\hat{a}_1,\hat{a}_2)$ forms a Nash equilibrium outcome for Γ. Then, using Observations 1 and 2, it is necessary that $\hat{a}_1 \in R_1(\hat{a}_2)$ and $\hat{a}_2 = R_2(\hat{a}_1)$. These conditions also imply that this action pair forms a Nash equilibrium outcome for G. Next, if the action pair $(\hat{a}_1,\hat{a}_2)$ forms a Nash equilibrium outcome for G, then $\hat{a}_1 \in R_1(\hat{a}_2)$ and $\hat{a}_2 = R_2(\hat{a}_1)$ is necessary. Now define strategies for Γ so that Player 1 selects the action $\hat{a}_1$ and Player 2 selects the action $\hat{a}_2$ (for every possible signal). These strategies then form a Nash equilibrium for Γ, and so the given action pair forms a Nash equilibrium outcome for Γ.

In interpreting this proposition, it is useful to identify two kinds of first-mover advantages. First, as in Stackelberg's quantity game, a first mover can gain a strategic benefit by committing to an action that is not in his best-response correspondence. A second benefit from moving first emerges in games with multiple simultaneous-move Nash equilibria, such as the classic Battle-of-the-Sexes game, since a first mover can then select his favorite such equilibrium. The proposition above, however, indicates that both of these advantages are eliminated if the first-mover's choice is not *perfectly* observed.

4. COMMITMENT AND MIXED-STRATEGY EQUILIBRIA

Up to now, I have focused entirely on pure-strategy Nash equilibria. This raises the question as to whether an advantage to commitment might reappear in the context of the mixed-strategy Nash equilibria of the noisy-leader game. The simple 2x2 example developed above provides a tractable framework within which to explore this issue.

To this end, consider again the noisy-leader game for the 2x2 example. The interesting case is when ε is small, and so it is assumed henceforth that $\varepsilon < 1/4$. It is now possible to determine the full set of

Nash equilibria for the noisy-leader game. Let λ be the probability that Player 1 plays S; thus, λ is Player 1's (possibly mixed) strategy. Next, let $\eta(\phi)$ be the probability that Player 2 plays S after observing the signal ϕ. Recall now that there is only one Nash equilibrium of the noisy-leader game in which Player 1 does not mix; this equilibrium occurs when both players select C with probability one ($\lambda = 0 \equiv \eta(\phi)$).

The second possibility is that Player 1 mixes ($\lambda \in (0,1)$). Given $\varepsilon < 1/4$, it is direct to confirm that a Nash equilibrium does not exist in which Player 2 chooses a pure strategy ($\eta(\phi) \in \{0,1\}$ for all ϕ). It is also easily verified that a Nash equilibrium does not exist in which Player 1 mixes and Player 2 mixes for all signal values ($\eta(\phi) \in (0,1)$ for all ϕ). Thus, if additional Nash equilibria exist, they involve Player 1 mixing and Player 2 mixing for exactly one signal value.

Calculations reveal that only two such cases produce mixed-strategy Nash equilibria. Consider first the case in which Player 2 mixes when the signal is C and plays S for sure when the signal is S ($\eta(C) \in (0,1)$ and $\eta(S) = 1$). A Nash equilibrium of this sort does in fact exist, and calculations indicate that it is defined by:

$$\lambda = 1-\varepsilon \text{ and } \eta(C) = (1-4\varepsilon)/(2-4\varepsilon).$$

A novel feature of this equilibrium is that Player 1 selects C less often as the signal gets more precise. This pattern ensures that Player 2's posterior probability of the S action from Player 1 following a signal indicating C is constant and independent of ε. This in turn makes Player 2 willing to mix when the signal is C, even when the signal is very precise.

The second case occurs when Player 2 mixes when the signal is S and plays C when the signal is C ($\eta(S) \in (0,1)$ and $\eta(C) = 0$). Calculations reveal a mixed-strategy Nash equilibrium of this variety:

$$\lambda = \varepsilon \text{ and } \eta(S) = 1/(2 - 4\varepsilon).$$

Notice in this case that λ converges to zero as ε goes to zero; in other words, as the noise gets small, Player 1 plays S less and less often.

This combination of limits ensures that Player 2 maintains the critical level of uncertainty upon observing $\phi = S$, so that Player 2 is then willing to mix for all $\varepsilon > 0$.

Several points can now be made. First, the addition of noisy leadership results in the existence of mixed-strategy Nash equilibria even though the simultaneous-move game possesses no mixed-strategy Nash equilibrium. Second, one mixed-strategy Nash equilibrium outcome of the noisy-leader game converges as the noise (i.e., ε) goes to zero to the Stackelberg outcome. The other mixed-strategy Nash equilibrium outcome of the noisy-leader game, however, converges to the Cournot outcome. Third, as always, the simultaneous-move game Nash equilibrium outcome, (C,C), survives also as a pure-strategy Nash equilibrium outcome for all ε in the noisy-leader game. Thus, the Stackelberg outcome can be approximately recovered from one mixed-strategy Nash equilibrium for the noisy-leader game when observability is "close" to perfect, but there is certainly no guarantee that this mixed-strategy Nash equilibrium will in fact obtain.

5. CONCLUSION

Models of commitment or first-mover advantages have profoundly affected a number of fields in economics. These models make two assumptions: one player moves before the other players, and the first-mover's selected action is perfectly observed by the other players. The purpose of this paper is to disentangle these two assumptions and examine the strategic role of commitment when other players observe imperfectly the first-mover's choice. The key finding is that the advantage of committing to a course of action is eliminated when other players observe this commitment with even the slightest amount of imprecision.

The ideas developed above are easily extended in a variety of directions. For example, the models that I formalize here all have a single leader. While this assumption is met in a number of examples of economic interest, there also exist models in which more than one leader exists, and in which the leaders simultaneously make commitments that influence the future course of the game.[2] It is straightforward, however, to see that the basic results developed above carry over to the context of simultaneous-commitment models. Regardless of the number of leaders, when other players observe a leader's selection with even the slightest degree of imprecision, the advantage of moving before others is extinguished.

For applied theorists, the key message of the paper is that the many predictions derived from models with commitment may require reconsideration. Apparently, these predictions are valid only for settings in which the committed action is in fact *perfectly* observed by subsequent players. This requirement is quite stringent, and it would seem to be violated in a number of real-world settings to which popular commitment models are thought to apply.

A role for commitment may reappear, however, in the context of incomplete-information games. For instance, if other players are incompletely informed of the first-mover's type, then the equilibrium behavior of the first-mover in the associated noisy-leader game is likely to differ considerably from that which would occur were moves made simultaneously (Matthews and Mirman, 1983). Further, in repeated games of incomplete information, a long-lived player may repeatedly play the Stackelberg strategy, in order to develop the reputation for being a "commitment type," who always selects this strategy (Fudenberg and Levine, 1992).

Finally, the findings developed here also highlight a new discontinuity for the sequential equilibrium (Kreps and Wilson, 1982) correspondence. For example, in the 2x2 game, when Player 1 moves first and his selection is perfectly observed, the unique sequential equilibrium outcome is the Stackelberg outcome, (S,S). With a slight amount of noise, however, the Cournot outcome, (C,C), emerges as a sequential equilibrium outcome. Sequential equilibrium thus fails in the limit to be upper hemicontinuous with respect to the level of observability (i.e., ε) in the game.[3]

---

[2]See, for example, Fershtman and Judd (1987), Kreps and Scheinkman (1983) and Spencer and Brander (1983).

[3]As Fudenberg and Tirole (1991, p. 342) and Myerson (1991, p. 189) discuss in other contexts, upper hemicontinuity of sequential equilibrium can be

The possibility of discontinuities in the sequential equilibrium correspondence has arisen previously in other games. Most notably, in the traditional signaling game, a slight amount of uncertainty as to the sender's payoffs can induce separating equilibria in which the sender's behavior differs radically from that which would occur were the sender's preferences commonly known.[4] The new feature of the discontinuity emphasized in the present paper is that it occurs in the transition from an imperfect- to a perfect-information game whereas the signaling-game discontinuity emerges in the transition from an incomplete- to a complete-information game.

---

restored if the definition is extended to allow nature to choose some branches with zero probability and if nature is allowed to tremble.

[4]Kreps, Milgrom, Roberts and Wilson (1982) and Fudenberg, Kreps and Levine (1988) explore further implications of this general theme.

REFERENCES

Fershtman, C. and K. Judd (1987), "Equilibrium Incentives in Oligopoly," *American Economic Review*, 77:927-940.

Fudenberg, D. and D. Levine (1992), "Maintaining a Reputation When Strategies are Imperfectly Observed," *Review of Economic Studies*, 59:561-580.

Fudenberg, D. and J. Tirole (1991), *Game Theory*. Cambridge, Massachusetts: The MIT Press.

Fudenberg, D. , D. Kreps, and D. Levine (1988), "On the Robustness of Equilibrium Refinements," *Journal of Economic Theory*, 44:354-380.

Kreps, D., P. Milgrom, J. Roberts and R. Wilson (1982),"Rational Cooperation in Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27:245-252.

Kreps, D. and J. Scheinkman (1983), "Quantity Pre-commitment and Bertrand Compeition Yield Cournot Outcomes," *Bell Journal of Economics*, 14:326-337.

Kreps, D. and R. Wilson (1982), "Sequential Equilibrium," *Econometrica*, 50:863-894.

Matthews, S. and L. Mirman (1984), "Equilibrium Limit Pricing: The Effects of Private Information and Stochastic Demand," *Econometrica*, 51:981-995.

Myerson, R. (1991), *Game Theory*. Cambridge, Massachusetts: Harvard University Press.

Schelling, T. (1960), *The Strategy of Conflict*. New York: Oxford University Press.

Spencer, B. and J. Brander (1983), "International R&D Rivalry and Industrial Strategy," *Review of Economic Studies*, 50:707:722.

von Stackelberg, H. (1934), *Marktform und Gleichgewicht*. Vienna: Julius Springer.