Who's on First? Commitment and Observability with Move-order Uncertainty

Kyle Bagwell^{*}

Stanford University

November 27, 2023

Abstract

This paper explores a two-player game in which actions are imperfectly observed and players are uncertain about move order. We study two examples, a "commitment" example and a "battle-of-the-sexes" example, where in each example each player has two possible actions and observes the realization of a binary signal before moving. If a given player is the actual second mover, then the distribution of the signal observed by that player is affected by the preceding move of the other player. We study first-mover advantages within a single game by examining how, for each Nash equilibrium, player 1's payoff changes as she becomes more likely to move first. For pure-strategy Nash equilibria, no first-mover advantage or disadvantage is present. By contrast, for mixedstrategy Nash equilibria, payoffs vary with the probability that player 1 moves first. We find that whether a first-mover advantage exists varies across mixed-strategy equilibria in the commitment example and across parameter regions in the BoS example. We also provide a new perspective on Newcomb's paradox.

^{*}I am grateful to Steve Callander, Chris Sanchirico and Takuo Sugaya for helpful discussions. I thank Ava Bochner and Luiz Carpizo for excellent RA work.

1 Introduction

A fundamental insight of game theory is that an observable commitment to an action can confer a strategic advantage. A common way of illustrating this advantage is by comparing equilibria across different two-person games. The standard comparison is between the pure-strategy Nash equilibria (PSNE) of a simultaneous-move game and the pure-strategy subgame perfect equilibria (SPE) of a counterpart sequential-move game in which player 1 moves first.

Under this comparison, a first-mover advantage arises in two main ways. First, assuming player 2's best-response correspondence is single-valued, player 1 can ensure that the outcome associated with her preferred PSNE of the simultaneous-move game obtains by committing to the action that is assigned to player 1 in that equilibrium. Thus, for example, the classic 2x2 battle-of-the-sexes (BoS) game has two PSNE, and the unique SPE outcome of the associated sequential-move game is the outcome that obtains in player 1's preferred PSNE of the simultaneous-move game. Second, in some settings, player 1 obtains strictly higher payoffs in the SPE of the sequential-move game than she enjoys in any of the Nash equilibria of the simultaneous-move game. This possibility can be illustrated with a simple 2x2 commitment ("Stackelberg") setting such that in the unique SPE of the sequential-move game player 1 commits to an action that is strictly dominated in the simultaneous-move game.

The role of observability is emphasized by Bagwell (1995). He considers a "noisy-leader game" with two players, in which player 1 moves first and player 2 moves after observing an imperfect signal of player 1's action. Assuming that player 2's best-response correspondence is single-valued and that the signal support is "nonmoving," he finds that the set of PSNE outcomes for the noisy-leader game coincides with the set of PSNE outcomes for the associated simultaneous-move game. This result holds even if the signal is arbitrarily precise. For a 2x2 commitment example, he also characterizes the two mixed-strategy Nash equilibria that arise in the associated noisy-leader game. As the noise goes to zero, one of the equilibrium outcomes converges to the PSNE ("Cournot") outcome of the simultaneous-move game while the other converges to the unique SPE ("Stackelberg") outcome of the sequential-move game. van Damme and Hurkens (1997) show more generally that a Nash equilibrium outcome for the noisy-leader game always exists that converges to the Stackelberg outcome as the noise vanishes. They also devise a new equilibrium selection theory that selects this equilibrium.¹

In this paper, we modify the noisy-leader game in a fundamental way. The assumption that actions are imperfectly observed is maintained, but we now assume that players are uncertain about the order of moves. When a player is called upon to move, she is thus

¹Their theory combines elements from Harsanyi and Selten (1988) and Harsanyi (1995).

uncertain whether the other player has already moved or will move subsequently. Intuitively, a player may be tempted to select an action that could confer a strategic advantage if she believes that the other player is likely to move subsequently; however, if a player believes it likely that the other player has already moved, then she may focus on selecting an action that is a best reply to the conjectured behavior of the other player.

To capture this setting in a coherent way, we analyze the following game. The game begins with a move by Nature, which chooses whether player 1 or player 2 moves first. Both players understand that Nature selects player 1 to go first with probability $\eta \in (0, 1)$ and that neither player directly observes Nature's selection. A player observes a realization of a binary signal before making a binary action choice. We thus refer to low and high values for the signal and action, respectively. If a player is selected to move first, then the binary signal realization observed by the player is induced by an unconditional distribution such that the low and high signal values occur with equal probability. If the player is selected to move second, however, then Nature uses a conditional distribution: a high (low) signal is generated with probability $\sigma \in (1/2, 1)$ when the other player has selected a high (low) action. Thus, if a player turns out to be the first mover, then the player can strategically manipulate the signal realization probabilities for the other player. A player can never be certain, however, that she moves first. We refer to this game as the "who's-on-first game."²

Our analysis of the who's-on-first game focuses on strategic environments with two actions per player and two possible signal realizations. The resulting strategic-form game is then a 4x4 game, since each player decides which of two actions to take after observing each of two possible signal realizations. We further focus on two particular underlying strategic settings corresponding to 2x2 commitment and BoS settings, respectively. These examples enable analysis of the two first-mover advantages mentioned above while also providing a tractable framework in which to develop a detailed analysis of the full set of (pure- and mixed-strategy) Nash equilibria.

The who's-on-first game provides a novel framework for analyzing certain strategic applications. For example, in the economic context, a firm contemplating a strong move into a given product group or regional area may wonder if higher-than-expected factor prices reflect random considerations or the initiation of a prior move by a rival firm. Similar examples can be described in R and D races or in political and marketing contexts. More generally, the relevance of the game is best assessed within the context of specific applications.

As a game structure, the who's-on-first game offers at least two key benefits. First, it enables the study of the strategic advantage of commitment within a single game. In particular, for any (pure or mixed) Nash equilibrium of the who's-on-first game, we can

²The name is a lighthearted reference to a famous comedy routine by Abbott and Costello.

analyze how the outcome changes as η changes. We can thus explore - in a continuous way - when or even whether a player gains a strategic advantage as the player becomes *more likely* to be the first mover. A second benefit is that when appropriate limits are taken, the who's-on-first game can be used as a unified framework in which to capture other standard games: the (perfect-information) sequential-move game corresponds to a limit case where η and σ equal one, the noisy-leader game corresponds to a limit case where η equals one, and the simultaneous-move game corresponds to a limit case where η and σ equal one half.

The who's-on-first game has strong implications for the set of PSNE. As we show for the two examples, the PSNE outcomes of the who's-on-first game coincide with the PSNE outcomes of the simultaneous-move game. The PSNE outcomes are thus independent of η . For PSNE, therefore, the strategic value of commitment disappears, even if the signal is arbitrarily precise (i.e., even if σ is arbitrarily close to 1). This generalizes the finding in Bagwell (1995) for noisy-leader games, wherein $\eta = 1$. The key intuition in both games is that players ignore the signal when under the Nash equilibrium hypothesis they already believe themselves to know the other player's action. For the two examples studied here, the PSNE are also strict Nash equilibria.

Pure- and mixed-strategy Nash equilibria can have different implications for the strategic value of commitment. Intuitively, players may respond to the signal realization in a mixed-strategy Nash equilibrium, since the equilibrium hypothesis is then insufficient for the inference of the other player's action. The different implications of pure- and mixedstrategy Nash equilibria are emphasized in the theoretical and experimental literature that studies the noisy-leader game.³ We characterize the full set of Nash equilibria here.

In our first example, the underlying strategic environment corresponds to a 2x2 commitment setting. When we embed this setting in the who's-on-first game form, we find as expected that a strict PSNE exists for all parameter values that yields the same outcome as the PSNE of the associated simultaneous-move game. We also characterize a nonlinear and negative relationship between η and σ such that for all η and σ that are below this boundary the who's-on-first game can be solved by two rounds of eliminated of strictly dominated actions in the strategic form. This process leads to the PSNE outcome of the simultaneous-move game without invoking the Nash equilibrium concept for the who's-onfirst game. When η and σ are above this boundary, the who's-on-first game also admits two mixed-strategy Nash equilibria. The outcomes of these equilibria approach the outcomes of the respective mixed-strategy Nash equilibria of the corresponding noisy-leader game as η goes to one. In line with Bagwell's (1995) analysis of a 2x2 commitment example, as η and σ go to one, the outcome of one mixed strategy Nash equilibrium converges to PSNE outcome

³References for this literature are provided in the penultimate paragraph of this section.

of the simultaneous-move game and the outcome of the other converges to the SPE outcome of the sequential-move game.

For each equilibrium, we then assess the first-mover advantage by characterizing the change in player 1's payoff as η rises.⁴ For PSNE, strategies are independent of η and thus payoffs do not change with the likelihood that player 1 is the first mover. By contrast, when η and σ are sufficiently large so that mixed-strategy Nash equilibria exist, we find that the payoffs for those equilibria depend on η . For the mixed-strategy Nash equilibrium outcome that converges to the SPE outcome of the sequential-move game, an increase in η leads to strict increases in player 1's payoff and player 2's payoff. This equilibrium thus exhibits a continuous and within-game notion of a first-mover advantage. For the mixed-strategy Nash equilibrium that converges to the PSNE outcome of the simultaneous-move game, however, the results are different: as η rises, player 1's payoff strictly falls while player 2's payoff is convex in η and minimized at $\eta = 1/2$. For this equilibrium, therefore, we find a first-mover disadvantage.

We also use this example to provide a new perspective on Newcomb's paradox.⁵ This paradox concerns a thought experiment in which two boxes are presented to an agent who must decide whether to open only box 2 or to open both box 1 and box 2. In the numerical formulation considered here, Box 1 is certain to contain 1 dollar, and Box 2 may contain 10 dollars or be empty. A superior being of some kind either places 10 dollars in box 2 or not. If the being thinks that the agent will open both boxes, or if the being thinks that the agent is randomizing her choice, then the being places no money in box 2. If the being thinks that the agent will open box 2 only, then the being places 10 dollars in box 2. The agent seeks to maximize monetary return. Traditional formulations of the paradox assume that the being moves first and has almost perfect foresight about the agent's choice, and then ask what the agent should do. The paradox is that dominance arguments suggest opening both boxes, but then a being with excellent foresight would likely leave box 2 empty.

Our approach is to engage with this paradox in the context of a game with uncertain move order and imperfect observability. We thus consider a who's-on-first game in which the agent plays the role of player 1 and the being plays the role of player 2. When called upon to move, the agent is therefore unsure if the being has already decided whether to place 10 dollars in box 2 or if it will subsequently decide whether to do so. The being is

⁴For the commitment example, we focus on the first-mover advantage for player 1 since player 1 is then the player who gains from moving first in the associated sequential-move game. We note, though, that our results also indicate how player 2's equilibrium payoffs vary as she becomes more likely to move first (i.e., as η falls). For the BoS example discussed below, the players' payoffs take a symmetric (mirror-image) form, and it is without loss of generality to assess the first-mover advantage from player 1's perspective.

⁵See, for example, Gardner (1973, 2001), Nozick (1969), Weber (2016) and Wolpert and Benford (2013). See Weber (2016) for a valuable and detailed review of this literature.

likewise uncertain about the move order and bases its inferences on its signal observation and knowledge of the Nash equilibrium. Our approach differs from the conventional approach in three key respects: we do not assume that the being necessarily moves first, we assume that the being, if moving second, imperfectly observes the agent's action as opposed to the agent's (perhaps mixed) strategy, and we assume that the being is a player in the game and thus makes Nash conjectures as to agent behavior.

Our findings imply that, if the likelihood that the agent moves first is sufficiently low and/or the signal is sufficiently imprecise (i.e., if η and σ are below the aforementioned boundary), then dominance arguments alone imply that the agent should open both boxes. But in the contrary case where the likelihood that the agent moves first is sufficiently high and/or the signal is sufficiently precise, then there exist multiple Nash equilibria. In the PSNE, the agent again opens both boxes; but in the mixed-strategy Nash equilibria, the agent may or may not open both boxes. As the likelihood that the agent moves first goes to one and the precision of the signal moves toward perfection, one mixed-strategy Nash equilibrium outcome converges to the PSNE outcome in which both boxes are opened whereas the other mixed-strategy Nash equilibrium outcome converges to the SPE outcome of the sequential-move game in which the agent commits to opening only box 2.

We thus offer a unified, game-theoretic perspective on Newcomb's paradox, under which dominance logic and the associated prediction that the agent opens both boxes can be linked to parameter values under which the being is a sufficiently likely to be the first mover and/or observes with sufficient imprecision. Under the opposite parameter environment, dominance arguments no longer suffice, and the being's inferences are guided by the Nash equilibrium hypothesis. Multiple Nash equilibria are then possible, offering potential justification for both actions for the agent. We note, though, that the PSNE in which the agent opens both boxes is the only strict Nash equilibrium.

We then turn to our second example, where the underlying strategic environment corresponds to a 2x2 BoS setting. When we embed this setting in the who's-on-first game, we find as expected that two strict PSNE exist for all parameter values, where the respective PSNE yield the same outcomes as the two PSNE of the associated simultaneous-move game. Under a parameter restriction on payoff values, we also show that, generically, a single mixed-strategy Nash equilibrium exists for all η and σ . The form of the Nash equilibrium, however, varies across parameter regions for η and σ . As expected, the mixed-strategy Nash equilibrium when $\eta = 1$ is imposed corresponds to the mixed-strategy Nash equilibrium of the noisy-leader game. Interestingly, as σ goes to one, the outcome of this equilibrium converges to the preferred PSNE outcome of *player two*.

For the BoS setting, the simultaneous-move game admits three Nash equilibria. The

outcomes of the two PSNE are directly captured by the PSNE of the corresponding who'son-first game. We show as well that, as η and σ both go to 1/2, the associated mixed-strategy Nash equilibrium outcome of the who's-on-first game converges to the mixed-strategy Nash equilibrium outcome of the simultaneous-move game. We thus capture all Nash equilibrium outcomes of the simultaneous-move game as Nash equilibrium outcomes of the who's-onfirst game when $\eta = 1/2$ and σ goes to 1/2.⁶ Intuitively, when these parameter values are imposed, each player is equally likely to move first and neither player can learn anything from the signal. As such, the who's-on-first game with this parameter configuration provides an extensive-form representation for a simultaneous-move game that does not rely on the asymmetric assumption that one player moves first, as in standard treatments.

For each equilibrium, we again assess the first-mover advantage by characterizing the change in player 1's payoff as η rises. As before, strategies in the PSNE are independent of η and thus payoffs do not change with the likelihood that player 1 is the first mover. On the other hand, in the mixed-strategy Nash equilibrium, behavior and payoffs are sensitive to model parameters. For any $\eta \in (0, 1)$, if the precision of the (informative) signal is modest, then we find that the mixed-strategy Nash equilibrium is such that a small increase in η generates a strictly lower (higher) equilibrium payoff for player 1 (2), indicating a first-mover disadvantage for player 1. For η near 1/2, this result also obtains when the signal exhibits high precision. Finally, when the signal precision is sufficiently high and player 1 is sufficiently likely to move first, the mixed-strategy Nash equilibrium is such that a small increase in η results in strictly higher equilibrium payoffs for both player 1 and 2, indicating a first-mover advantage for player 1. We find, however, for the high-precision scenario that player 1's mixed-strategy equilibrium payoff when she is very likely to move first is lower than when she moves first with probability 1/2 or is very unlikely to move first.⁷

The first-mover advantage explored in this paper is defined in terms of the change in a player's equilibrium payoff as the player becomes more likely to move first. It is thus a *marginal* concept. Alternative definitions that compare equilibrium payoff levels are also of interest. Notably, we might ask whether a player enjoys a higher level of equilibrium payoff when the player is almost certain to move first in comparison to other benchmarks such as when the player is almost certain to move second or moves first with probability one half.

⁶This results holds as well for the commitment setting; however, in that example, the result is immediate since the corresponding simultaneous-move game has only one Nash equilibrium outcome (namely, the PSNE outcome that also exists for all parameter values in the who's-on-first game).

⁷This result arises since, when σ is high, there exists a critical $\hat{\eta}(\sigma) > 1/2$ such that player 1's equilibrium payoff is strictly higher as η approaches $\hat{\eta}(\sigma)$ from below than from above. (There exists a continuum of mixed-strategy Nash equilibrium at the measure-zero scenario where $\eta = \hat{\eta}(\sigma)$.) Finally, as we show, given the symmetry of the BoS game, exactly related statements apply for player 2. For each equilibrium, we can assess the first-mover advantage for player 2 by characterizing the change in her payoff as η falls.

Previous work addresses these questions by comparing equilibrium outcomes for different games. For a given Nash equilibrium, we can make these comparisons within the unified who's-on-first game by examining payoffs for different values of η . At the end of Section 3, we return to this distinction in the context of the BoS example and observe that the different definitions sometimes give different answers as to the existence of a first-mover advantage.

The paper is related to the literature on commitment and observability, including Bagwell (1995), Bhaskar (2009), Guth, et al. (1998), Maggi (1999), Morgan and Vardy (2007, 2013), Oechssler and Schlag (2000), van Damme and Hurkens (1997) and Vardy (2004). Move-order uncertainty also arises in an example that Kreps and Ramey (1987) consider, although their goal is to explore the relationship between sequential rationality and structurally consistent beliefs when information sets cross. Nishihara (1997) also considers a model with move-order uncertainty. For a model with N players and considering only pure strategies, he shows that cooperation attains in equilibrium in the Prisoners' Dilemma setting for an information structure with a moving support such that each player knows nothing about the order of moves, except that a player knows if in a preceding move a player selected the "defect" action. Gallice and Monzon (2019) explore related themes in a model with position-order uncertainty wherein each player directly observes a sample of her predecessors' actions.⁸ Move-order uncertainty is also related to a literature that considers the possibility of revision strategies. Related papers include Bagwell (2018), Henkel (2002), Kamada and Kandori (2020) and Weber (2016), among others. Finally, the paper is related to a literature that explores endogenous move-order determination when commitments are perfectly observed. This literature includes van Damme and Hurkens (1993), Hamilton and Slutsky (1990, 1993) and Rosenthal (1991), among others.

The paper is organized as follows. We begin with the commitment example. After characterizing equilibrium behavior in the associated benchmark games, we study the associated who's-on-first game and the associated possibility of a first-mover advantage. We next apply our findings to Newcomb's paradox. Finally, we turn to the BoS example, where we again study the associated benchmark games, who's-on-first game and possibility of a first-mover advantage. A final section offers concluding thoughts.

2 The Commitment Example

In this section, we focus on an underlying 2x2 strategic environment featuring an advantage to commitment. After characterizing equilibrium behavior for the associated benchmark

⁸See also Doval and Ely (2022) for general methods to summarize the distribution of equilibrium outcomes in incomplete information games when the information structure and extensive form may vary.

scenarios of a simultaneous-move game, a sequential-move game and the noisy-leader game, we explore the who's-on-first game. We then examine the first-mover advantage in terms of a comparative-statics exercise with respect to the parameter η . We conclude the section with a new perspective on Newcomb's paradox.

2.1 Benchmark games

The commitment example refers to a setting with two players, where player 1 has two actions denoted l and h while player 2 has two actions denoted λ and H. Let $u_i(x, y)$ denote the payoff to player i, i = 1, 2, when player 1 selects action x and player 2 selects action y, where $x \in \{l, h\}$ and $y \in \{\lambda, H\}$. The payoffs satisfy the following rankings:

$$u_1(h,\lambda) > u_1(l,\lambda) > u_1(h,H) > u_1(l,H)$$

 $u_2(l,\lambda) > u_2(l,H), u_2(h,H) > u_2(h,\lambda).$

As discussed below, the key feature of the payoff structure is that player 1 can enjoy a firstmover advantage in the sequential-move game by committing to an action (namely, l) that is strictly dominated in the simultaneous-move game. Notice that player 2 seeks to match λ to l and H to h, respectively.

To simplify the discussion and ensure that the commitment example provides a convenient framework for interpreting Newcomb's paradox, as discussed below, we further assume for player 1 that $u_1(h, \lambda) = 11$, $u_1(l, \lambda) = 10$, $u_1(h, H) = 1$ and $u_1(l, H) = 0$, and we further assume for player 2 that $u_2(l, \lambda) = u_2(h, H) = B$ and $u_2(l, H) = u_2(h, \lambda) = 0$, where B > 0.9These payoffs are illustrated in Table 1.

$$\begin{array}{c|c} \lambda & H \\ l & 10, B & 0, 0 \\ h & 11, 0 & 1, B \end{array}$$

 Table 1: Commitment Example

Consider first the simultaneous-move game. The unique PSNE of the simultaneous-move game is (h, H) with player 1 receiving a payoff of 1 and player 2 receiving a payoff of B. Notice further that the action l is strictly dominated by the action h for player 1. The simultaneous-move game thus does not admit a Nash equilibrium in which player 1 selects

⁹Recalling the discussion of Newcomb's paradox from the Introduction, we may anticipate that the agent will play the role of player 1 and the being will play the role of player 2, where the agent will either open just box 2 (l) or open both boxes (h) while the being will either put 10 dollars in box 2 (λ) or no money in box 2 (H).

the action l with positive probability. As a consequence, it is easy to see that the PSNE (h, H) is the unique Nash equilibrium of the simultaneous-move game.

Now consider the sequential-move game in which player 1 moves first, player 2 moves second, and player 2 perfectly observes player 1's choice. In the sequential-move game, the unique SPE outcome entails a choice of l by player 1 and a choice of λ by player 2. The resulting payoffs are (10, B). Thus, in the unique SPE, player 1 achieves a strategic advantage by committing to an action in the sequential-move game that is strictly dominated in the simultaneous-move game.

Finally, consider the *noisy-leader game*. In this game, a pure strategy for player 1 entails a choice of l or h. Once player 1 selects an action, a signal s is generated where s either takes value a or b with a < b. We can thus think of a as the low signal and b as the high signal. The signal s is generated according to the following probability distribution:

$$Prob(s = a|l) = \sigma = Prob(s = b|h)$$

where $\sigma \in (1/2, 1)$. Note that a higher value for σ corresponds to a more precise signal.

After observing the signal realization, player 2 selects action λ or H. Since the signal can take two values, player 2 has four pure strategies: $\lambda\lambda, \lambda H, H\lambda$ and HH, where the first (second) symbol in any pair indicates the action that is taken following the observation of the signal value a (b).

We turn now to the strategic form for the noisy-leader game. Using $\sigma > 1/2$, we find for player 2 that the strategy $H\lambda$ is strictly dominated by the strategy λH . We can thus focus on the 2x3 strategic form as depicted in Table 2.

	$\lambda\lambda$	λH	HH
l	10, B	$10\sigma, \sigma B$	0,0
h	11, 0	$11(1-\sigma) + \sigma, \sigma B$	1, B

Table 2: Noisy-leader 2x3 strategic-form game

For this strategic form, we see that for any given action by player 1, player 2 has a strict ranking over her actions. Similarly, given a selection by player 2 of either $\lambda\lambda$ or HH, player 1 has a strict ranking over her actions. Finally, if player 2 selects λH , then player 1 strictly prefers l to h if $1 > \sigma > 11/20$ and strictly prefers h to l if $1/2 < \sigma < 11/20$. If player 2 selects λH and $\sigma = 11/20$, then player 1 is indifferent between l and h.

Suppose first that $1/2 < \sigma < 11/20$. In this scenario, the signal has low precision. Since player 1 then strictly prefers h to l when player 2 selects λH , it is straightforward to confirm that h strictly dominates l for player 1. If we eliminate l from the strategic form so that the resulting strategic form is 1x3, then it is clear that HH strictly dominates $\lambda\lambda$ and λH for player 2. Thus, after two rounds of elimination of strictly dominated strategies, we arrive at the implication that the players must select the pair (h, HH) with the corresponding payoff pair (1, B). Thus, for the scenario of low signal precision, simple dominance arguments ensure the same outcome and payoffs as found in the PSNE of the simultaneous-move game.

We consider now the set of PSNE. Looking at the 2x3 strategic form in Table 2, we may confirm that the unique (and strict) PSNE of the noisy-leader game is (h, HH) leading to a payoff pair of (1, B). The unique PSNE outcome is thus that player 1 selects the action hwhile player 2 selects the action H. This aligns with Bagwell's (1995) finding that the set of PSNE outcomes for the noisy-leader game coincides with the set of PSNE outcomes for the associated simultaneous-move game. This result holds for all levels of signal precision.

Next, we identify the set of mixed-strategy Nash equilibria. Given the dominance logic described above, a mixed-strategy Nash equilibrium does not exist for the first scenario of low signal precision. We thus now suppose that $11/20 \leq \sigma < 1$, which we refer to as the scenario of high signal precision.

Considering all possible supports for player 2's mixed strategies, we can confirm the existence of exactly two mixed-strategy Nash equilibria under high signal precision. The first mixed-strategy Nash equilibrium arises when player 2 plays the strategy $\lambda\lambda$ with zero probability. Let p_l (p_h) denote the probability that player 1 selects l (h), and let $q_{\lambda H}$ (q_{HH}) be the probability that player 2 selects λH (HH). The equilibrium takes the following form:

$$q_{\lambda H} = \frac{1}{20(\sigma - 1/2)}\tag{1}$$

$$p_l = 1 - \sigma \tag{2}$$

with $p_h = 1 - p_l$ and $q_{HH} = 1 - q_{\lambda H}$. Notice that $q_{\lambda H} \leq 1$ in the scenario of high signal precision, with $q_{\lambda H} = 1$ holding if and only if $\sigma = 11/20$.

In this mixed-strategy Nash equilibrium, as σ goes to one, p_h goes to one. Player 2 is then almost sure to observe the signal b, leading to a choice of H (since $q_{\lambda\lambda} = 0$). Thus, as the signal becomes arbitrarily precise, the outcome of this mixed-strategy Nash equilibrium converges to the outcome of the unique Nash equilibrium of the simultaneous-move game. In other words, the outcome of this equilibrium converges to the "Cournot" outcome.

The second mixed-strategy Nash equilibrium arises when player 2 plays the strategy HH with zero probability. Let $q_{\lambda\lambda}$ denote the probability that player 2 selects $\lambda\lambda$. The equilibrium takes the following form:

$$q_{\lambda\lambda} = \frac{20\sigma - 11}{20(\sigma - 1/2)} \tag{3}$$

$$p_l = \sigma \tag{4}$$

with $p_h = 1 - p_l$ and $q_{\lambda H} = 1 - q_{\lambda \lambda}$. Notice that $q_{\lambda \lambda} \ge 0$ in the scenario of high signal precision, with $q_{\lambda \lambda} = 0$ holding if and only if $\sigma = 11/20$.

In this mixed-strategy Nash equilibrium, as σ goes to one, p_l goes to one. Player 2 is then almost sure to observe the signal a, leading to a choice of λ (since $q_{HH} = 0$). Thus, as the signal becomes arbitrarily precise, the outcome of this mixed-strategy Nash equilibrium converges to the outcome of the unique SPE of the sequential-move game. In other words, the outcome of this equilibrium converges to the "Stackelberg" outcome.¹⁰

2.2 Who's-on-first game

We now embed the commitment example into the who's-on-first game. The game begins with a move by Nature, which selects either player 1 or player 2 to move first. The commonly known probability that Nature selects player 1 to move first is denoted as η , where $\eta \in (0, 1)$. Neither player observes Nature's selection.

The player selected to go first receives a signal s, where $s \in \{a, b\}$ with b > a. For the first mover, Nature selects each signal with probability 1/2. If player 1 (2) is selected as the first mover, then her chosen action is imperfectly observed by player 2 (1). Specifically, the second mover observes the signal s, where $s \in \{a, b\}$ is now generated by a conditional distribution. If player 1 is selected as the first mover, then the signal received by player 2 is generated as follows:

$$Prob(s = a|l) = \sigma = Prob(s = b|h).$$

Similarly, if player 2 is selected as the first mover, the signal received by player 1 satisfies:

$$Prob(s = a|\lambda) = \sigma = Prob(s = b|H).$$

As in the noisy-leader game, we assume $\sigma \in (1/2, 1)$ with a higher value for σ again corresponding to a more precise signal.

Upon observing a realization of the signal, a player must choose an action. The player does not know, however, whether the observed signal realization was generated by the unconditional or the conditional distribution. In other words, a player does not know whether she is the first or second mover. Thus, for each $s \in \{a, b\}$, a (pure) strategy for player 1 specifies an action in $\{l, h\}$ and a (pure) strategy for player 2 specifies an action in $\{\lambda, H\}$. Mixed strategies can be defined in the usual way. A pair of strategies then induces a probability distribution over outcomes for the game. Payoffs thus can be computed in standard fashion.

 $^{^{10}}$ This corresponds to the equilibrium outcome selected by van Damme and Hurkens (1997).

As in the noisy-leader game, player 2 again has four pure strategies: $\lambda\lambda, \lambda H, H\lambda$ and HH, where the first (second) symbol in any pair indicates the action that is taken following the observation of the signal value a (b). In the who's-on-first game, however, player 1 also has four pure strategies: ll, lh, hl and hh, where again the first (second) symbol in any pair indicates the action that is taken following the observation of the signal value a (b).

To illustrate the calculations of payoffs, suppose that player 1 uses strategy lh and player 2 uses strategy $H\lambda$. Player 1's expected payoff is then calculated as follows:

$$\eta(1/2)(1-\sigma)10 + \eta(1/2)(\sigma 11 + 1 - \sigma) + (1-\eta)(1/2)\sigma + (1-\eta)(1/2)(\sigma 10 + (1-\sigma)11).$$

The first term describes the situation in which player 1 is selected to go first, observes the signal a and thus chooses l. Player 2 then observes the signal a with probability σ , which induces the choice of H and leads to a payoff of zero for player 1 since $u_1(l, H) = 0$. With probability $1 - \sigma$, player 2 observes the signal b, resulting in the choice of λ and thus a payoff of 10 for player 1 since $u_1(l, \lambda) = 10$. The second term likewise captures the situation in which player 1 is selected to go first but observes the signal b and thus chooses h. Finally, the third and fourth terms correspond to the situations in which player 2 is selected to go first and b, respectively. Player 2's expected payoff under this strategy pair can be similarly calculated.

Making similar calculations for both players for each of the 16 strategy pairs, we can construct a 4x4 strategic form to capture the who's-on-first game with the commitment example embedded. After simplifying expressions, we present this strategic form in Table 3.

	$\lambda\lambda$	λH	$H\lambda$	HH
11	10, B	$5 + 10\eta \left(\sigma - \frac{1}{2}\right),$	$5 - 10\eta \left(\sigma - \frac{1}{2}\right),$	0.0
		$B\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	$B\left[\frac{1}{2} - \eta\left(\sigma - \frac{1}{2}\right)\right]$	0,0
lh	$11 - \sigma + \eta \left(\sigma - \frac{1}{2} \right),$	$\frac{11}{2}\sigma B$	$\frac{11}{2}$,	$\sigma - \eta \left(\sigma - \frac{1}{2} \right),$
	$B\left(\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right)$	2,012	$B\left(\sigma - 2\eta\left(\sigma - \frac{1}{2}\right)\right)$	$B\left(\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right)$
hl	$10 + \sigma - \eta \left(\sigma - \frac{1}{2} \right),$	$\frac{11}{2}$,	$\frac{11}{1-\sigma}B$	$1 - \sigma + \eta \left(\sigma - \frac{1}{2} \right),$
	$B\left(1 - \sigma + \eta\left(\sigma - \frac{1}{2}\right)\right)$	$B\left(1-\sigma+2\eta\left(\sigma-\frac{1}{2}\right)\right)$	2,(1,0)D	$B\left(1-\sigma+\eta\left(\sigma-\frac{1}{2}\right)\right)$
hh	11, 0	$6 - 10\eta \left(\sigma - \frac{1}{2}\right),$	$6+10\eta\left(\sigma-\frac{1}{2}\right),$	1 <i>B</i>
		$B\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	$B\left[\frac{1}{2} - \eta\left(\sigma - \frac{1}{2}\right)\right]$	т, D

Table 3: Who's-on-first 4x4 strategic-form game

Similar to the noisy-leader game, certain strategies are strictly dominated and can be eliminated on that basis. We make two points in this regard. First, we may use $\sigma > 1/2$ to establish that player 2's strategy $H\lambda$ is strictly dominated by the strategy λH . Importantly, this conclusion thus holds without further restrictions on model parameters. We can thus focus on the 4x3 strategic form as presented in Table 4.

	$\lambda\lambda$	λH	HH
11	10 <i>B</i>	$5 + 10\eta \left(\sigma - \frac{1}{2}\right),$	0.0
	10, D	$B\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	0,0
lh	$11 - \sigma + \eta \left(\sigma - \frac{1}{2} \right),$	$\frac{11}{\sigma}\sigma B$	$\sigma - \eta \left(\sigma - \frac{1}{2} \right),$
	$B\left(\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right)$	$_{2}, 0D$	$B\left(\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right)$
hl	$10 + \sigma - \eta \left(\sigma - \frac{1}{2}\right),$	$\frac{11}{2}$,	$1 - \sigma + \eta \left(\sigma - \frac{1}{2} \right),$
	$B\left(1-\sigma+\eta\left(\sigma-\frac{1}{2}\right)\right)$	$B\left(1-\sigma+2\eta\left(\sigma-\frac{1}{2}\right)\right)$	$B\left(1 - \sigma + \eta\left(\sigma - \frac{1}{2}\right)\right)$
hh	11 0	$6 - 10\eta \left(\sigma - \frac{1}{2}\right),$	1 B
	11,0	$B\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	1, D

Table 4: Who's-on-first 4x3 strategic-form game

Second, and as in the noisy-leader game, under a condition on model parameters, we can make further progress by eliminating strictly dominated strategies. For the who's-on-first game, the corresponding condition is

$$20\eta(\sigma - 1/2) < 1\tag{5}$$

Notice that, if we set $\eta = 1$, then (5) reduces to the low-precision condition used above for the noisy-leader game. The generalized condition stated here, however, allows for general values of $\eta \in (0, 1)$ and is more likely to hold when the signal precision is low and/or player 1 is less likely to go first. Note that (5) is sure to hold if $\sigma \in (1/2, 11/20]$, but as η drops further below 1 it holds also for higher values of σ .

Under the parameter condition in (5), direct calculations confirm that player 1's strategy hh strictly dominates all other strategies for player 1 in the original 4x4 strategic form game. After eliminating all strategies for player 1 other than hh, it is direct that strategy HH strictly dominates all other strategies for player 2. We thus have the following conclusion:

Observation 1: Consider the who's-on-first game for the commitment example. Assume (5) holds. After two rounds of eliminating strictly dominated strategies for the associated

4x4 strategic form game, the unique remaining strategy pair is (hh, HH), yielding the payoff (1, B).

Thus, if (5) holds, then simple dominance arguments deliver the same outcome and payoffs as found in the PSNE of the simultaneous-move game. Under (5), a strategic advantage to commitment thus does not arise in the who's-on-first game for the commitment example.

Our next step is to consider the set of PSNE. As noted above, for player 2, the strategy $H\lambda$ is strictly dominated by the strategy λH . We can thus search for PSNE in the 4x3 strategic form in Table 4. This search leads to the following conclusion:

Observation 2: Consider the who's-on-first game for the commitment example.

(i) If $1 \neq 20\eta(\sigma - 1/2)$, then there exists a unique PSNE, in which the strategy pair is (hh, HH).

(ii) If $1 = 20\eta(\sigma - 1/2)$, then there exist exactly three PSNE, in which the respective strategy pairs are $(hh, HH), (lh, \lambda H)$ and $(hl, \lambda H)$.

In the generic case described by part (i) of the observation, the unique (and strict) PSNE outcome is again that player 1 selects the action h and player 2 selects the action H for all possible signals and regardless of which player is selected to go first. For PSNE, we thus conclude for generic payoffs that a strategic advantage to commitment does not arise in the who's-on-first game for the commitment example. This result holds without further restrictions on model parameters (i.e., it does not require (5)) and generalizes Bagwell's (1995) findings to the who's-on-first game.

The remaining task is to identify the set of mixed-strategy Nash equilibria. Again, we may focus on the 4x3 strategic form provided in Table 4. Using observation 1 and dispensing with the non-generic case for which $1 = 20\eta(\sigma - 1/2)$, we may limit our analysis to the parameter constellation in which

$$20\eta(\sigma - 1/2) > 1,$$
 (6)

corresponding to a high probability that player 1 moves first and/or a high signal precision. Note that (6) requires $\eta > 1/10$.

Our approach is to consider all possible supports for the two players' mixed strategies. For example, a player mixing over exactly two strategies must be indifferent between those strategies and not gain from playing any other strategy with positive probability. The first requirement is an equality constraint that structures the mixing probabilities of the other player, and the second requirement is a set of inequality constraints. Additionally, no mixing probabilities so determined can be negative or greater than one. Proceeding on a case-bycase basis, we can use these requirements to rule out most candidates. In the end, we can confirm the existence of exactly two mixed-strategy Nash equilibria.¹¹

To describe these mixed-strategy Nash equilibria, we must enrich the notation with which we describe player 1's mixed strategies. Let p_{xy} denote the probability that player 1 selects the strategy xy, where $x \in \{l, h\}$ and $y \in \{l, h\}$. To be clear, let it also be understood that when we refer to "mixed-strategy" Nash equilibria, we are referring to Nash equilibria that are not PSNE. We may now report the following finding:

Observation 3: Consider the who's-on-first game for the commitment example. Assume (6) holds. There exist exactly two mixed-strategy Nash equilibria:

(i) In one mixed-strategy Nash equilibrium, player 1 mixes over lh and hh only (i.e., $p_{lh} \in (0,1)$, $p_{hh} \in (0,1)$ and $p_{lh} + p_{hh} = 1$) while player 2 mixes over λH and HH only (i.e., $q_{\lambda H} \in (0,1)$, $q_{HH} \in (0,1)$ and $q_{\lambda H} + q_{HH} = 1$). The mixed-strategy probabilities are

$$q_{\lambda H} = \frac{1 - \sigma + \eta(\sigma - 1/2)}{(\sigma - 1/2)(11\eta - 1)}$$
(7)

$$p_{hh} = 2\eta(\sigma - 1/2) \tag{8}$$

(ii) In the other mixed-strategy Nash equilibrium, player 1 mixes over ll and hl only (i.e., $p_{ll} \in (0,1)$, $p_{hl} \in (0,1)$ and $p_{ll} + p_{hl} = 1$) while player 2 mixes over $\lambda\lambda$ and λH only (i.e., $q_{\lambda\lambda} \in (0,1)$, $q_{\lambda H} \in (0,1)$ and $q_{\lambda\lambda} + q_{\lambda H} = 1$). The mixed-strategy probabilities are

$$q_{\lambda H} = \frac{\sigma(1-\eta) + \eta/2}{(\sigma - 1/2)(1+9\eta)}$$
(9)

$$p_{ll} = 2\eta(\sigma - 1/2) \tag{10}$$

¹¹For example, it is straightforward to see that a mixed-strategy Nash equilibrium does not exist in which player 1 mixes only over the strategies lh and hl. This follows because player 2's unique best response is then the strategy λH , which would then induce player 1 to select the strategy ll instead. We can similarly rule out that player 2 mixes over only the strategies $\lambda\lambda$ and HH. To see this, we observe that player 1's unique best response is then the strategy hh, which would in turn induce player 2 to select the strategy HHinstead. We may also observe that if player 2 mixes only over λH and HH (with $q_{HH} > 0$ thus implied), then player 1 must place zero probability on strategy hl. This is because player 1 is indifferent between the strategies hl and lh when player 2 uses the strategy λH and strictly prefers lh when player 2 uses HH. To illustrate further, we continue with the case where player 2 mixes only over λH and HH. Player 1 then cannot mix over each of her remaining three strategies, since the resulting two (indifference) equalities impose inconsistent requirements on player 2's mixed strategy given $q_{\lambda\lambda} = 0$ and $\sigma > 1/2$. Also, there then does not exist a mixed-strategy Nash equilibrium in which player 1 mixes only over ll and hh, since given $\sigma > 1/2$ there does not exist a feasible value for $q_{\lambda H}$ under which player 1 is indifferent to ll and hh and weakly prefers those strategies to lh. By contrast, we find for this case that there do exist feasible values for $q_{\lambda H}$ (with $q_{HH} = 1 - q_{\lambda H}$) under which player 1 would be willing to mix only over the strategies *ll* and *lh* and the strategies lh and hh, respectively. The analysis then moves forward to examine whether player 2 in fact would be willing to mix only over λH and HH, if player 1 were to mix in one of these two ways. Other cases are handled similarly. The proof process is tedious but conceptually straightforward. We therefore do not provide details for each candidate case here and instead just report the main findings.

To interpret these mixed-strategy Nash equilibria, we note that, as expected, they return the mixed-strategy Nash equilibria of the noisy-leader game when $\eta = 1$ is imposed.¹² To see this in the context of the first mixed-strategy Nash equilibrium, we note immediately that (7) then returns the expression for $q_{\lambda H}$ in (1). Further, if $\eta = 1$, then player 1 moves first with certainty, observes the signal *a* with probability 1/2, and thus chooses *l* with probability $(1/2)p_{lh} = 1 - \sigma$, returning thereby the value for p_l as given in (2). Similarly, when $\eta = 1$ is imposed, (9) and (10) respectively return (3) and (4).

The two mixed-strategy Nash equilibria identified in (7)-(10) thus also share limit properties with their counterpart mixed-strategy Nash equilibria for the noisy-leader game. In particular, as η and σ both go to 1, the outcome of the mixed-strategy Nash equilibrium captured in (7)-(8) converges to the "Cournot" outcome in which player 1 selects action hand player 2 selects action H, and likewise the outcome of the mixed-strategy Nash equilibrium captured in (9)-(10) converges to the "Stackelberg" outcome in which player 1 selects action l and player 2 selects action λ .

Finally, while the mixed-strategy Nash equilibria have natural limit properties, we emphasize that they characterize equilibrium behavior away from limits as well. Players then face more uncertainty regarding move order and opponent behavior. For example, with positive probability, a player may move second and yet not best respond to the other player's action. An interesting question then concerns how expected payoffs vary in a given equilibrium as model parameters vary. We consider this question next in the context of variation in the parameter η .

2.3 Comparative statics and the first-mover advantage

For each Nash equilibrium of the who's-on-first game, we now consider how payoffs vary with respect to η . We then consider whether player 1 enjoys a first-mover advantage in a given equilibrium. We explore a marginal notion of a first-mover advantage. Specifically, we say that an equilibrium exhibits a first-mover advantage (disadvantage) at $\eta \in (0, 1)$ if player 1's expected payoff strictly increases (decreases) following a slight rise in η .¹³ If an equilibrium exhibits a first-mover advantage (disadvantage) for all $\eta \in (0, 1)$, then we will simply say that it exhibits a first-mover advantage (disadvantage).

To begin, we consider PSNE. As noted in Observation 2, for the generic case in which

¹²Note in (7) that $11\eta > 1$ is implied by (6).

¹³For the commitment example, we focus on the first-mover advantage for player 1, since player 1 is then the only player who gains from moving first in the associated sequential-move game.

 $1 \neq 20\eta(\sigma - 1/2)$, a unique (and strict) PSNE exists. The unique PSNE strategy pair is (hh, HH), and the resulting payoff pair is (1, B). Since the payoff pair is independent of η , this equilibrium does not generate a first-mover advantage or disadvantage at any $\eta \in (0, 1)$.

Consider next the mixed-strategy Nash equilibrium described in Observation 3 by (7) and (8). This is the equilibrium whose outcome converges to the Cournot outcome as η and σ each go to one. To conduct this comparative statics exercise, we must represent each player's expected payoff in the equilibrium. Since player 1 mixes over *lh* and *hh*, she is indifferent between the two. Thus, player 1's expected payoff equals the expected payoff that she enjoys from using *hh*. With player 2 mixing over λH and *HH* only, player 1's expected payoff in this equilibrium can be calculated as

$$q_{\lambda H} 10[(1+\eta)/2 - \eta\sigma] + 1 \tag{11}$$

Given that $q_{\lambda H}$ as defined in (7) is strictly decreasing in η , we may differentiate (11) with respect to η and confirm that player 1's expected payoff strictly *decreases* with respect to η . Player 1 thus experiences a first-mover disadvantage in this equilibrium.

Consider now player 2. Since player 2 mixes over λH and HH, she is indifferent between the two; hence, player 2's expected payoff equals the expected payoff that she receives from selecting HH. Given that player 1 mixes over lh and hh only, we may calculate player 2's expected payoff as

$$p_{lh}B[\eta/2 + (1-\eta)\sigma - 1] + B \tag{12}$$

Using (8), we see that p_{lh} is strictly decreasing in η . Differentiating player 2's expected payoff as given in (12) with respect to η , we find that player 2's expected payoff is convex in η and minimized when $\eta = 1/2$.

We turn next to the mixed-strategy Nash equilibrium described in Observation 3 by (9) and (10). Recall that as η and σ each go to one, the associated equilibrium outcome converges to the Stackelberg outcome. Player 1 mixes over ll and hl, and her expected payoff thus equals the expected payoff that she enjoys from using ll. Since player 2 mixes over $\lambda\lambda$ and λH only, player 1's expected payoff in this equilibrium can be calculated as

$$q_{\lambda H}[10\eta(\sigma - 1/2) - 5] + 10 \tag{13}$$

Since $q_{\lambda H}$ as defined in (9) is strictly decreasing in η , we may differentiate (11) with respect to η and confirm that player 1's expected payoff strictly increases with respect to η . In this equilibrium, therefore, layer 1 experiences a first-mover advantage. Finally, since player 2 mixes over $\lambda\lambda$ and λH , her expected payoff equals the expected payoff that she receives from selecting $\lambda\lambda$. With player 1 mixing over *ll* and *hl* only, we may calculate player 2's expected payoff as

$$p_{hl}B[-\sigma + \eta(\sigma - 1/2)] + B \tag{14}$$

Using (10), we see that p_{h1} is strictly decreasing in η . Using (14) to differentiate player 2's expected payoff with respect to η , we find that player 2's expected payoff strictly increases with respect to η .

Focusing on player 1, we summarize our findings for this subsection with the following proposition:

Proposition 1. Consider the who's-on-first game for the commitment example. The unique PSNE defined in Observation 2(i) does not generate a first-mover advantage or disadvantage at any $\eta \in (0, 1)$. Under (6), the mixed-strategy Nash equilibrium described in Observation 3 by (7) and (8) exhibits a first-mover disadvantage, and the mixed-strategy Nash equilibrium described in Observation 3 by (9) and (10) exhibits a first-mover advantage.

Thus, the PSNE is (generically) unique and strict and never generates a first-mover advantage or disadvantage. On the other hand, for mixed-strategy Nash equilibria, the presence of a first-mover advantage varies across equilibria: a first-mover advantage exists in one mixedstrategy Nash equilibrium, but a first-mover disadvantage exists in the other. For mixedstrategy Nash equilibria, therefore, whether player 1 gains when she becomes more likely to be the first mover depends on the equilibrium that is selected. Finally, our findings also identify circumstances under which player 2 gains when the probability that player 1 is the first mover increases. As η rises, player 2 enjoys a higher payoff in both mixed-strategy Nash equilibria, although for one equilibrium the gain is enjoyed if and only if $\eta > 1/2$.¹⁴

2.4 Newcomb's paradox

The commitment example also may be used to provide a new perspective on Newcomb's paradox. As described in the Introduction, Newcomb's paradox refers to a thought experiment

¹⁴Our focus here is on how equilibrium payoffs vary with respect to η , but we can also compute how equilibrium payoffs vary with respect to σ . PSNE payoffs do not vary with σ . In the mixed-strategy Nash equilibrium described in Observation 3 by (7) and (8), we find that player 1's equilibrium payoff strictly decreases in σ , and player 2's equilibrium payoff strictly increases in σ . By contrast, in the mixed-strategy Nash equilibrium described in Observation 3 by (9) and (10), we find that player 1's equilibrium payoff strictly increases in σ , and player 2's equilibrium payoff strictly increases in σ for all σ satisfying (6) if $\eta \geq 4/9$.

under which an agent is presented with two boxes, box 1 and box 2, and can choose either to open box 2 only or to open both boxes. Box 1 is known to possess one dollar. A being decides whether to put 10 dollars in box 2 or to leave box 2 empty. If the being anticipates that the agent will open both boxes, then the being prefers to leave box 2 empty. If, however, the being anticipates that the agent will open box 2 only, then the being prefers to place 10 dollars in box 2. The being predicts the agent's behavior with high precision.

What should the agent do in this situation? The problem is sometimes referred to as highlighting a tension between dominance logic under which, assuming the being has already made a choice, the agent should open both boxes, and expected utility logic under which, given the being's excellent skills at prediction, the agent would seem to enjoy the highest expected payoff by opening box 2 only.

The perspective offered here is made within the context of a fully specified game-theoretic model. The agent is captured as player 1, the being is represented as player 2, the agent either opens box 2 (action l) or both boxes (action h), and the being either puts 10 dollars in box 2 (action λ) or leaves box 2 empty (action H). We thus represent not just the agent but also the being as a player in the game. The agent has free will. When preparing to move, the being forms beliefs as to the agent's behavior based on the observed signal and the equilibrium hypothesis. We can capture the possibility of a very precise signal of the agent's behavior by allowing that η and σ are close to one. The model also embodies dominance considerations that encourage the agent to open both boxes. These considerations are naturally stronger when η is smaller and/or that the signal precision σ is weaker.

Our model departs from conventional formulations (see, e.g., Gardner, 2001) in at least three respects. First, we do not insist that the being moves first; rather, we allow that the being may or may not move first. We thereby use move uncertainty as a way to capture both the dominance and expected utility logic emphasized in earlier treatments. Second, traditional formulations assume further that the being leaves no money in box 2 if the being believes that the agent randomizes her choice. We depart from this approach here by assuming that the being, if moving second, receives an imperfect signal of the agent's *action* rather than of the agent's strategy. Finally, we model the being as a player in a game, and we use the Nash equilibrium solution concept. This implies that the being perfectly anticipates the agent's strategy *in equilibrium*. It thus does not imply that the being would directly observe the agent's strategy were the agent to deviate to an alternative strategy.¹⁵

¹⁵Weber (2016) also examines Newcomb's paradox in a game-theoretic model. His dynamic game allows for an exogenous probability under which the being can revise her action. Weber also finds that the agent's equilibrium behavior is sensitive to model parameters. Our work differs in several ways. We do not allow for revision opportunities and explore a different game, leading, for example, to different findings for PSNE in that the agent always opens both boxes. We also study mixed-strategy Nash equilibria in detail, wherein

Our findings provide a different perspective on the paradox. As Observation 1 confirms, when (5) holds so that the likelihood that the agent goes first is sufficiently low and/or the signal is sufficiently imprecise, then two rounds of elimination of strictly dominated strategies results in a unique remaining strategy pair, (hh, HH), under which the agent always opens both boxes and the being always leaves box B empty. Furthermore, as Observation 2 indicates, for generic settings, this same strategy pair forms the unique PSNE. Thus, even if the agent is extremely likely to move first and the being would then receive an extremely precise signal of the agent's choice, the agent can get "stuck" in an equilibrium in which she opens both boxes and is anticipated to do so. The agent's dilemma in this case is that, were she to decide instead to open only box 2, the being - as a believer in Nash equilibrium - would maintain faith in the equilibrium hypothesis and thus leave box 2 empty.¹⁶ This equilibrium is also a strict Nash equilibrium.

Finally, as Observation 3 confirms, when (5) fails, two mixed-strategy Nash equilibria also exist. As indicted in part ii of Observation 3, when the agent is almost sure to go first and the signal precision is very high, we can approximately capture as a mixed-strategy Nash equilibrium outcome the outcome under which the agent opens only box 2 and the being places 10 dollars in box 2. Thus, there exist parameter ranges and a mixed-strategy Nash equilibrium that provides support for this outcome as well. But as part i of Observation 3 shows, for this same range of parameters, we also approximately capture as a mixed-strategy Nash equilibrium outcome the outcome under which the agent opens both boxes and the being being leaves box 2 empty.

3 The Battle-of-the-Sexes Example

We now consider an alternative underlying 2x2 strategic environment that features a strategic advantage to commitment. Specifically, the two players interact in a Battle-of-the-Sexes (BoS) setting. We begin by characterizing equilibrium behavior for the associated benchmark scenarios of a simultaneous-move game, a sequential-move game and the noisy-leader game.

both players randomize. By contrast, Weber assumes that the being observes the agent's *strategy* with some probability and punishes a randomizing agent by leaving box B empty, ensuring thereby that the agent doesn't randomize. With the who's-on-first game analyzed here, we also pick up other benchmark games as limiting cases, and we offer a new examination of first-mover advantages in given Nash equilibria by varying the probability η that player 1 moves first. Weber examines the subgame perfect Nash equilibria for his dynamic game. For the who's-on-first game considered here, the only proper subgame is the full game, and so the sets of Nash and subgame perfect Nash equilibria coincide. Finally, we also study equilibrium behavior when a BoS game is embedded.

¹⁶The being could relinquish this faith if a signal value were observed that is impossible given the equilibrium hypothesis; however, under the non-moving support assumption made here, a signal value may be very unlikely but is never impossible.

We then analyze equilibrium behavior in the associated who's-on-first game. We conclude the section by examining the first-mover advantage.

3.1 Benchmark games

In the BoS example, there are two players who operate in a symmetric setting. Each player has two actions. We again denote player 1's two actions as l and h, and we likewise denote player 2's two actions as λ and H. As before, $u_i(x, y)$ denotes the payoff to player i, i = 1, 2, when player 1 (2) selects action x(y), where $x \in \{l, h\}$ and $y \in \{\lambda, H\}$. The payoffs take a symmetric form and are given as follows:

$$u_1(l,\lambda) = R > r = u_1(h,H) > 0 = u_1(l,H) = u_1(h,\lambda)$$
$$u_2(h,H) = R > r = u_2(l,\lambda) > 0 = u_2(l,H) = u_2(h,\lambda).$$

As is well known, the BoS game features a setting in which the players gain from matching their actions (l with λ , and h with H) but have different opinions about which of the two matched action pairs is best. We assume R > r > 0, so that player 1 prefers the pair (l, λ) while player 2 prefers the pair (h, H). As discussed below, a player gains a strategic advantage to commitment in this setting since the player can thereby ensure matching on her favorite matched action pair. These payoffs are illustrated in Table 5.

	λ	H
l	R, r	0, 0
h	0, 0	r, R

Table 5: BoS example

We begin with the simultaneous-move game. The simultaneous-move game has exactly two PSNE, namely, (l, λ) and (h, H). The associated payoff pairs are (R, r) and (r, R), respectively. Thus, player 1 prefers the first PSNE, and player 2 prefers the second PSNE. The simultaneous-move game also has a unique mixed-strategy Nash equilibrium, in which player 1 selects l with probability R/(R+r) while player 2 selects λ with probability r/(R+r). In the unique mixed-strategy Nash equilibrium, each player enjoys an expected payoff of rR/(R+r). Thus, each player earns an expected payoff in the mixed-strategy Nash equilibrium that is less than r, the lowest payoff that a player makes in a PSNE.

We turn now to the *sequential-move game* in which player 1 moves first, player 2 moves second, and player 2 perfectly observes player 1's choice. In the sequential-move game, the unique SPE outcome entails a choice of l by player 1 and a choice of λ by player 2. The resulting payoff pair is (R, r). Thus, in the unique SPE of the sequential-move game, player 1 achieves a strategic advantage by committing to an action that leads to her preferred PSNE outcome in the simultaneous-move game.

The final benchmark game is the *noisy-leader game*. For this game, player 1 has two pure strategies, l and h. As before, once player 1 selects an action, a signal s is generated where s either takes value a or b with a < b. The signal s is again generated according to the following probability distribution:

$$Prob(s = a|l) = \sigma = Prob(s = b|h)$$

where $\sigma \in (1/2, 1)$. Upon observing the signal realization, player 2 chooses between two actions, λ and H. Player 2 thus has four pure strategies: $\lambda\lambda, \lambda H, H\lambda$ and HH, where the first (second) symbol in any pair again indicates the action that is taken following the observation of the signal value a (b).

The noisy-leader game thus has a 2x4 strategic form. Using $\sigma > 1/2$, we find for this game as well that for player 2 the strategy $H\lambda$ is strictly dominated by the strategy λH . We thus analyze the 2x3 strategic-form game as depicted in Table 6.

	$\lambda\lambda$	λH	HH
l	R, r	$\sigma R, \sigma r$	0, 0
h	0, 0	$\sigma r, \sigma R$	r, R

Table 6: Noisy-leader 2x3 strategic-form game

The set of PSNE is easily characterized.¹⁷ Examining the 2x3 strategic-form game in Table 6, it is apparent that the noisy-leader game has exactly two PSNE. The PSNE are given by $(l, \lambda\lambda)$ and (h, HH). They lead to the payoff pairs (R, r) and (r, R), respectively. Thus, the set of PSNE outcomes for the noisy-leader game again coincides with the set of PSNE outcomes for the simultaneous-move game, aligning with Bagwell's (1995) finding. The PSNE are strict Nash equilibria.

Finally, we identify the set of mixed-strategy Nash equilibria. Considering each of the possible supports for player 2's mixed strategies, we find exactly one mixed-strategy Nash equilibrium. To represent this equilibrium, we again let p_l (p_h) denote the probability that player 1 selects l (h). Likewise, we again let $q_{\lambda H}$ denote the probability that player 2 selects λH , q_{HH} denote the probability that player 2 selects HH, and $q_{\lambda\lambda}$ denote the probability that player 2 selects $\lambda\lambda$. The equilibrium takes the following form:

¹⁷Unlike the commitment example, when the BoS example is embedded into the noisy-leader game, we do not find parameter ranges such that a unique solution can be obtained by successive elimination of strictly dominated strategies.

$$q_{\lambda H} = \frac{r}{r + \sigma(R - r)} \tag{15}$$

$$p_l = \frac{R(1-\sigma)}{R(1-\sigma) + \sigma r} \tag{16}$$

with $p_h = 1 - p_l$ and $q_{HH} = 1 - q_{\lambda H}$. The payoffs for this mixed-strategy Nash equilibrium are easily computed to be

$$\left(\frac{\sigma rR}{r+\sigma(R-r)}, \frac{\sigma rR}{R(1-\sigma)+\sigma r}\right),$$
(17)

from which we can see that both players earn lower payoffs than they enjoy in the PSNE that delivers the payoff pair (r, R), corresponding to player 2's preferred PSNE.

Using (15)-(17), we can also examine limiting behavior in this mixed-strategy Nash equilibrium. First, as σ goes to 1/2, p_l goes to R/(R+r) and $q_{\lambda H}$ goes to 2r/(r+R). Since player 2 then observes the signal *a* with a probability that approaches 1/2, and since $q_{\lambda\lambda} = 0$, the probability that player 2 selects λ goes to r/(r+R). The equilibrium payoff pair likewise goes to (rR/(r+R), rR/(r+R)). It follows that, as σ goes to 1/2, the mixed-strategy Nash equilibrium outcome and payoffs for the noisy-leader game converge to those of the mixed-strategy Nash equilibrium for the simultaneous-move game.

Second, as σ goes to one, p_h goes to one. Player 2 is then almost sure to observe the signal b, leading to a choice of H (since $q_{\lambda\lambda} = 0$). Thus, and perhaps surprisingly, as the signal becomes arbitrarily precise, the outcome of this mixed-strategy Nash equilibrium converges to player 2's preferred PSNE outcome of the simultaneous-move game. This outcome is not the SPE outcome of the sequential-move game (in which player 1 moves first). The SPE outcome of the sequential-move game, however, coincides with the outcome of one of the PSNE of the noisy-leader game.

3.2 Who's-on-first game

We now embed the BoS example into the who's-on-first game. The game begins when Nature selects either player 1 or player 2 to move first, where Nature picks player 1 to move first with probability η . As before, Nature's choice is not observed by either player. Since the players' payoffs are symmetric in the BoS example, we start with the case where $\eta \in [1/2, 1)$. Once we have results for this case, we can capture the remaining case where $\eta \in (0, 1/2)$ by simply reversing the players' names. We utilize this approach in the next subsection while examining the first-mover advantage.

As before, the player selected to go first observes a signal s, where $s \in \{a, b\}$ with b > a.

For the first mover, each signal occurs with probability 1/2. The second mover imperfectly observes the choice of the first mover; in particular, for the second mover, the signal $s \in \{a, b\}$ is generated by a conditional distribution. Once again, if player 1 is selected as the first mover, then the signal received by player 2 is generated as follows:

$$Prob(s = a|l) = \sigma = Prob(s = b|h).$$

Likewise, if player 2 is the first mover, then the signal received by player 1 satisfies:

$$Prob(s = a|\lambda) = \sigma = Prob(s = b|H).$$

Recall that $\sigma \in (1/2, 1)$.

In the who's-on-first game, upon observing a realization of the signal, a player must choose an action without knowing whether she is the first or second mover. A pure strategy for player 1 thus specifies an action in $\{l, h\}$ for each $s \in \{a, b\}$, and a pure strategy for player 2 assigns an action in $\{\lambda, H\}$ to each $s \in \{a, b\}$. Mixed strategies are defined in the usual way. Given a pair of strategies, a probability distribution over outcomes is induced, and payoffs are computed accordingly.

Once payoffs are calculated for each pair of pure strategies, we can construct a 4x4 strategic-form representation of the who's-on-first game with the BoS example embedded. We represent the strategic-form game in Table 7.

	$\lambda\lambda$	λH	$H\lambda$	HH
11	B r	$R\left[\frac{1}{2}+\eta\left(\sigma-\frac{1}{2}\right) ight],$	$R\left[\frac{1}{2} - \eta\left(\sigma - \frac{1}{2}\right)\right],$	0.0
	10,7	$r\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	$r\left[\frac{1}{2} - \eta\left(\sigma - \frac{1}{2}\right)\right]$	0,0
lh	$R\left[\sigma-\eta\left(\sigma-\frac{1}{2}\right) ight],$	$\frac{1}{2}\sigma(r+R),$	$\frac{1}{2}\left[\sigma - 2\eta\left(\sigma - \frac{1}{2}\right)\right](r+R),$	$r\left[\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right],$
	$r\left[\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right]$	$\frac{1}{2}\sigma(r+R)$	$\frac{1}{2}\left[\sigma - 2\eta\left(\sigma - \frac{1}{2}\right)\right](r+R)$	$R\left[\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right]$
hl	$R\left[1-\sigma+\eta\left(\sigma-\frac{1}{2}\right)\right],$	$\frac{1}{2}\left[1-\sigma+2\eta\left(\sigma-\frac{1}{2}\right)\right](r+R),$	$\frac{1}{2}(1-\sigma)(r+R),$	$r\left[1-\sigma+\eta\left(\sigma-\frac{1}{2}\right) ight],$
	$r\left[1-\sigma+\eta\left(\sigma-\frac{1}{2} ight) ight]$	$\frac{1}{2}\left[1-\sigma+2\eta\left(\sigma-\frac{1}{2}\right)\right](r+R)$	$\frac{1}{2}(1-\sigma)(r+R)$	$R\left[1 - \sigma + \eta\left(\sigma - \frac{1}{2}\right)\right]$
hh	0.0	$r\left[\frac{1}{2}+\eta\left(\sigma-\frac{1}{2}\right) ight],$	$r\left[\frac{1}{2}-\eta\left(\sigma-\frac{1}{2}\right) ight],$	r B
	0,0	$R\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right) ight]$	$R\left[\frac{1}{2} - \eta\left(\sigma - \frac{1}{2} ight) ight]$	7,10

Table 7: Who's-on-first 4x4 strategic-form game

As in the noisy-leader game, certain strategies are strictly dominated and thus can be eliminated. Using $\sigma > 1/2$, we find that player 2's strategy $H\lambda$ is strictly dominated by the strategy λH , and we likewise find that player 1's strategy hl is strictly dominated by the strategy lh. We can thus focus on the 3x3 strategic-form game as represented in Table 8.

	$\lambda\lambda$	λH	HH
ll	R,r	$R\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right],$	0, 0
		$r\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	
lh	$R\left[\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right],$	$\frac{1}{2}\sigma(r+R),$	$r\left[\sigma-\eta\left(\sigma-\frac{1}{2}\right) ight],$
	$r\left[\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right]$	$\frac{1}{2}\sigma(r+R)$	$R\left[\sigma - \eta\left(\sigma - \frac{1}{2}\right)\right]$
hh	0, 0	$r\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right],$	r B
		$R\left[\frac{1}{2} + \eta\left(\sigma - \frac{1}{2}\right)\right]$	7,10

Table 8: Who's-on-first 3x3 strategic-form game

Employing the 3x3 strategic-form game, the set of PSNE is easily characterized.

Observation 4: Consider the who's-on-first game for the BoS example. There exist exactly two PSNE, in which the respective strategy pairs are $(ll, \lambda\lambda)$ and (hh, HH).

As expected, the PSNE outcomes of the who's-on-first game coincide with those of the simultaneous-move and noisy-leader games, confirming that a strategic advantage to commitment does not arise in the PSNE of the who's-on-first game for the BoS example.

Using the 3x3 strategic-form game provided in Table 8, our next task is to identify the set of mixed-strategy Nash equilibria. To organize the analysis, we henceforth maintain the assumption that $R \ge 2r$. With this assumption, we pick up, for example, the common BoS specification under which R = 2 and r = 1. Using this assumption and that $\eta \in [1/2, 1)$, we find that player 1's strategies are strictly ranked for any fixed strategy by player 2. In particular, player 1 then strictly prefers strategy ll to strategy lh when player 2 uses the strategy λH .¹⁸ Player 2's strategies are likewise strictly ranked for any given strategy by player 1, with one exception: player 2's ranking of strategies λH and HH when player 1's strategy is lh depends on η and σ . While player 2's ranking over λH and HH is unclear when player 1 uses the strategy lh, we can determine that both of these strategies are then strictly better for player 2 than the strategy $\lambda \lambda$.¹⁹

¹⁸If $\sigma \leq R/(R+r)$, then this ranking holds for any $\eta \in (0,1)$.

¹⁹Let \succ_i denote a strict preference relation for player *i* defined over player *i* strategies, for a fixed strategy by the other player. We may then represent the described strict rankings as follows. For player 1, $ll \succ_1 lh \succ_1 hh$ given $\lambda\lambda$ or λH , and $hh \succ_1 lh \succ_1 ll$ given HH. For player 2, $\lambda\lambda \succ_2 \lambda H \succ_2 HH$ given ll; $HH \succ_2 \lambda H \succ_2 \lambda\lambda$ given hh; and $\lambda H \succ_2 \lambda\lambda$ and $HH \succ_2 \lambda\lambda$ given lh.

Finally, with further calculations, we find that player 2 strictly prefers λH to HH when player 1 uses lh if and only

$$\eta > \hat{\eta}(\sigma) \equiv \frac{\sigma(R-r)/2}{(\sigma - 1/2)R} \tag{18}$$

for $\eta \in [1/2, 1)$ and $\sigma \in (1/2, 1)$. The function $\hat{\eta}(\sigma)$ is illustrated in Figure 1. As depicted, $\hat{\eta}(\sigma)=1$ when $\sigma = R/(R+r)$, with $\hat{\eta}(\sigma) < 1$ thus holding for $\sigma \in (R/(R+r), 1)$. Observe further that, when $\sigma \in (1/2, R/(R+r)]$, we have $\hat{\eta}(\sigma) \ge 1$ so that $\eta < \hat{\eta}(\sigma)$ is sure to hold.



Figure 1: The function $\hat{\eta}(\sigma)$.

The rankings embodied in Table 8 lead to several useful observations that narrow the set of mixed-strategy Nash equilibrium candidates. First, there does not exist a mixed-strategy Nash equilibrium in which player 2 mixes over $\lambda\lambda$ and λH only. This follows under our assumptions since ll is then the unique best response for player 1 against each of $\lambda\lambda$ and λH , which in turn induces player 2 to select $\lambda\lambda$ with probability one. Second, if player 1 mixes over lh and hh only, then player 2 must select the strategy $\lambda\lambda$ with probability zero. This implication holds since $\lambda\lambda$ is the worst strategy for player 2 against each of lh and hh. Finally, if player 2 mixes over $\lambda\lambda$ and HH only, then player 1 must play strategy lhwith positive probability. The argument builds from the following logic: given that player 2 selects strategy λH with probability zero, player 1 can be indifferent between strategies lland hh if and only if player 2 selects strategy $\lambda\lambda$ with probability r/(R+r); however, given this mixing probability for player 2, player 1 would do better by selecting the strategy lh. To find the set of mixed-strategy Nash equilibria for the 3x3 strategic-form game, we proceed on a case-by-case basis and consider all possible supports for the mixed strategies of players 1 and 2. The observations just mentioned narrow the set of possibilities, but many cases remain. We eliminate further candidates when they violate one of three requirements that a mixed-strategy Nash equilibrium must satisfy. First, a player must be indifferent over strategies that receive positive probability, with the associated equality constraint structuring the mixing probabilities used by the other player. Second, if a player mixes over a subset of strategies, then an inequality constraint must be satisfied such that the player does not strictly gain from choosing a strategy outside the support. Finally, the mixed-strategy probabilities so determined must be feasible. The proof process is thus conceptually straightforward but tedious. We simply report here the findings. To do so, we follow the notation conventions previously used. Thus, for example, p_{lh} represents the probability that player 1 uses the strategy lh, and $q_{\lambda H}$ denotes the probability that player 2 uses the strategy λH .

We may now report the following finding:

Observation 5: Consider the who's-on-first game for the BoS example. Assume $9.8r \ge R \ge 2r$ and that $\eta \ne \hat{\eta}(\sigma)$ where $\hat{\eta}(\sigma)$ is defined in (18). Then there exists exactly one mixed-strategy Nash equilibrium:

(i) Suppose $\eta > \hat{\eta}(\sigma)$. Then there exists a unique mixed-strategy Nash equilibrium, in which player 1 mixes over *lh* and *hh* only (i.e., $p_{lh} \in (0,1)$, $p_{hh} \in (0,1)$ and $p_{lh} + p_{hh} = 1$) while player 2 mixes over λH and HH only (i.e., $q_{\lambda H} \in (0,1)$, $q_{HH} \in (0,1)$ and $q_{\lambda H} + q_{HH} = 1$). The mixed-strategy probabilities are

$$q_{\lambda H} = \frac{r[1 - \sigma + \eta(\sigma - 1/2)]}{r/2 + (R - r)\sigma/2}$$
(19)

$$p_{lh} = \frac{R[1/2 - \eta(\sigma - 1/2)]}{R/2 - (R - r)\sigma/2}$$
(20)

(ii) Suppose $\eta < \hat{\eta}(\sigma)$. Then there exists a unique mixed-strategy Nash equilibrium, in which player 1 mixes over *ll* and *lh* only (i.e., $p_{ll} \in (0,1)$, $p_{lh} \in (0,1)$ and $p_{ll} + p_{lh} = 1$) while player 2 mixes over λH and HH only (i.e., $q_{\lambda H} \in (0,1)$, $q_{HH} \in (0,1)$ and $q_{\lambda H} + q_{HH} = 1$). The mixed-strategy probabilities are

$$q_{\lambda H} = \frac{r[\sigma - \eta(\sigma - 1/2)]}{r[\sigma - \eta(\sigma - 1/2)] + R[1/2 + \eta(\sigma - 1/2)] - \sigma(R + r)/2}$$
(21)

$$p_{lh} = \frac{r[1/2 + \eta(\sigma - 1/2)]}{r[1/2 + \eta(\sigma - 1/2)] + \sigma(R - r)/2 - R\eta(\sigma - 1/2)}$$
(22)

Observation 5 imposes the assumptions that $9.8r \ge R \ge 2r$. We maintain these assumptions henceforth. The assumption $R \ge 2r$ is already discussed above. The new assumption that $9.8r \ge R$ allows for a wide range of values for R and r. When we consider the case in which each player selects each of her three strategies with positive probability, we find that the required mixing probabilities are infeasible if this assumption is imposed.²⁰

The assumption that $\eta \neq \hat{\eta}(\sigma)$ eliminates only the non-generic parameter range under which $\eta = \hat{\eta}(\sigma)$. This range defines a boundary between the mixed strategies described in parts (i) and (ii) of Observation 5. On this boundary, there exist a continuum of mixedstrategy Nash equilibria in which $p_{lh} = 1$ and player 2 mixes over λH and HH only. Included in this set are the strategies featured in parts (i) and (ii) of Observation 5, when we set $p_{lh} = 1$. Player 2's equilibrium payoff is continuous as this boundary is crossed, since p_{lh} goes to one as η approaches $\hat{\eta}(\sigma)$ in both (20) and (22) and since player 2 is then indifferent between λH and HH when $\eta = \hat{\eta}(\sigma)$. Player 2's equilibrium strategy jumps as this boundary is crossed, however, since $q_{\lambda H}$ is strictly higher when $\eta = \hat{\eta}(\sigma)$ under (21) than under (19) given that $\sigma > R/(R+r)$ and thus $\hat{\eta}(\sigma) < 1$. Since player 1 strictly prefers a higher value for $q_{\lambda H}$ when $\eta = \hat{\eta}(\sigma)$, $p_{lh} = 1$, and player 2 mixes over λH and HH only, player 1's equilibrium payoff features a downward jump as the $\eta = \hat{\eta}(\sigma)$ boundary is crossed. Put differently, player 1's equilibrium payoff is strictly higher as η approaches $\hat{\eta}(\sigma)$ from below than above.

The parameter regions corresponding to parts (i) and (ii) of Observation 5, respectively, are captured in Figure 1. We note the parameter region described in part (i) requires $\sigma > R/(R+r)$ and $\eta > 1/2$, where the latter inequality follows given $\sigma < 1$ since $(R-r)/R \ge 1/2$ under our assumption that $R \ge 2r$. As previously noted, the parameter region described in part (ii) is implied if $\sigma \in (1/2, R/(R+r)]$. Referring to Figure 1, we also see that a parameter pair (η, σ) may belong to this parameter region even if $\sigma > R/(R+r)$.

As expected, if $\eta = 1$ is imposed, the mixed-strategy Nash equilibrium outcome as characterized in Observation 5 converges to the mixed-strategy Nash equilibrium outcome of the noisy-leader game. Consider first the scenario in which $\sigma \in (R/(R+r), 1)$, placing us in the setting described by part (i) of Observation 5. As η goes to 1, we find from (19) that $q_{\lambda H}$ converges to $r/[r + (R - r)\sigma]$, matching (15) from the noisy-leader game. Likewise, we find from (20) that p_{lh} converges to $2R(1 - \sigma)/[R(1 - \sigma) + r\sigma]$. In the limit, the probability

²⁰If player 2 mixes over each of λH , HH and $\lambda\lambda$, then player 1's mixing probabilities are determined by the corresponding indifference equations. We find that $p_{lh} = 2rR/[(R-r)^2(1-\eta)(\sigma-1/2)]$. Feasibility requires $p_{lh} \leq 1$. This inequality is sure to fail for all $\eta \in [1/2,1)$ and $\sigma \in (1/2,1)$ if it fails when we set $\eta = 1/2$ and $\sigma = 1$. We note that $p_{lh} \leq 1$ when $\eta = 1/2$ and $\sigma = 1$ if and only if $R^2 + r^2 - 10rR \geq 0$. From here, we can show that feasibility fails for all $\eta \in [1/2,1)$ and $\sigma \in (1/2,1)$ if $R \in [2r, 5r + 2r\sqrt{6}]$. It follows that feasibility fails if $R \in [2r, 9.8r]$.

that player 1 selects the action l thus equals 1/2 (i.e, the probability that player 1 observes the signal a) times the limiting value of p_{lh} , matching the value of p_l as given in (16) for the noisy-leader game. We can do an exactly analogous calculation in the second scenario given by $\sigma \in (1/2, R/(R+r))$. In this case, as η goes to one, we are in the setting described by part (ii) of Observation 5.²¹

We can also explore limit behavior when σ and η both go to one. We then find that $q_{\lambda H}$ goes to r/R while the probability the player 1 selects the action l goes to zero. Hence, just as in the noisy-leader game, as η and σ both go to one, the mixed-strategy equilibrium outcome converges to a choice of action h by player 1 and of action H by player 2. We thus again converge to the PSNE outcome favored by player 2.

When σ and η are both near 1/2, we are in the setting described in part (ii) of Observation 5. In the limit, and using (21) and (22), we find that $q_{\lambda H}$ and p_{lh} both go to 2r/(R+r). We can then calculate the expected payoff enjoyed by each player in the mixed-strategy Nash equilibrium of the who's-on-first game with the BoS example. We find that each player enjoys an expected payoff of rR/(r+R), which matches the expected payoff in the mixed-strategy Nash equilibrium of the simultaneous-move game. Thus, and as before, we can capture the Nash equilibria of the simultaneous-move game via the who's-on-first game when η and σ go to 1/2. In this way, the who's-on-first game nests an alternative extensiveform representation of the simultaneous-move game with the natural feature that each player regards herself as being equally likely to be the player with the first move.

As noted previously, the mixed-strategy Nash equilibria have natural limit properties but characterize equilibrium behavior more generally. Specifically, they characterize equilibrium behavior for general parameter regions for which players face uncertainty about move order and opponent behavior. Utilizing this property, we consider next how expected payoffs vary in a given equilibrium as the probability that player 1 moves first (i.e., η) varies.

3.3 Comparative statics and the first-mover advantage

Fix any Nash equilibrium of the who's-on-first game for the BoS example. We now examine how the associated equilibrium payoffs change as η is increased. As before, we say that an equilibrium exhibits a *first-mover advantage (disadvantage)* at η if player 1's expected payoff strictly increases (decreases) following a slight rise in η .

²¹The calculation of the probability that player 1 selects action l then proceeds as follows. After using (22) to calculate the limiting value of p_{lh} , we obtain the limiting value of p_{ll} via $p_{ll} + p_{lh} = 1$. We then calculate the probability of action l for player 1 as 1/2 plus 1/2 times the limiting value of p_{ll} , where the first (second) "1/2" refers to the probability that player 1 observes the signal a (b). The resulting value again matches p_l as given in (16) for the noisy-leader game.

Recall for the BoS example that we assume $\eta \ge 1/2$. This assumption is without loss of generality. As noted above, given the symmetry of the BoS setting, we can translate our results to the setting where $\eta < 1/2$ by reversing players' names.²² We thus maintain the assumption that $\eta \ge 1/2$ while developing our main findings (Proposition 2) for the subsection. At the end of the subsection, however, we provide findings (Proposition 3) in which we detail the implications of our findings for the general setting where $\eta \in (0,1)$.

We start with PSNE. As noted in Observation 4, there exist exactly two PSNE, and the associated strategy pairs are $(ll, \lambda\lambda)$ and (hh, HH). Clearly, the equilibrium strategies and thus payoffs are independent of η . We therefore conclude that the two PSNE do not exhibit a first-mover advantage or disadvantage. Both PSNE are strict Nash equilibria.

Next, we consider the mixed-strategy Nash equilibrium described in Observation 5 by (19)-(20) and (21)-(22), respectively. As captured in Observation 5, the equilibrium strategy takes a different form depending on whether $\eta > \hat{\eta}(\sigma)$ or $\eta < \hat{\eta}(\sigma)$. For each setting, we represent the expected equilibrium payoffs for both players.

Starting with the setting in which $\eta > \hat{\eta}(\sigma)$, we note that player 1 mixes over *lh* and *hh* and is thus indifferent between the two. Player 1's expected payoff therefore equals the expected payoff that she enjoys from using *hh*. With player 2 mixing over λH and *HH* only, player 1's expected payoff in this equilibrium can be calculated as

$$q_{\lambda H} r [\eta (\sigma - 1/2) - 1/2] + r.$$
(23)

As $q_{\lambda H}$ is defined in (19), we may readily calculate the derivative of $q_{\lambda H}$ with respect to η , finding that it is strictly positive. Using the corresponding expression and differentiating (23) with respect to η , we may confirm that player 1's expected payoff strictly increases with respect to η for $\eta > \hat{\eta}(\sigma)$. For the mixed-strategy Nash equilibrium when $\eta > \hat{\eta}(\sigma)$, player 1 thus enjoys a first-mover advantage.

As we discuss further below, it is again important to highlight that the notion of a firstmover advantage explored here is a marginal concept: we focus on how a player's payoff changes as η changes. Thus, while player 1 enjoys a first-mover advantage in the mixedstrategy Nash equilibrium when $\eta > \hat{\eta}(\sigma)$, we see from (23) that the *level* of player 1's equilibrium payoff is low and indeed strictly below r. Consistent with our findings for the noisy-leader game, as η and σ converge to one, player 1's equilibrium payoff converges to r.

Player 2's payoffs may be calculated similarly. Player 2 mixes over λH and HH and is thus indifferent between the two. Player 2's expected payoff therefore equals the expected payoff that she receives from selecting HH. Given that player 1 mixes over lh and hh only,

²²Correspondingly, and as discussed below, for a given equilibrium, the first-mover advantage for player 2 can be assessed by examining how her payoffs change as η falls.

player 2's expected payoff in this equilibrium can be calculated as

$$R - p_{lh}R[1 - \eta/2 - (1 - \eta)\sigma]$$
(24)

Employing (20), we may easily calculate the derivative of p_{lh} with respect to η , finding that it is strictly negative. Using the corresponding expression and differentiating (24) with respect to η , we find that player 2's expected payoff strictly increases with respect to η for $\eta > \hat{\eta}(\sigma)$.

Consider next the setting in which $\eta < \hat{\eta}(\sigma)$. According to part (ii) of Observation 5, in the associated mixed-strategy Nash equilibrium, player 1 mixes over ll and lh and is thus indifferent between the two. Player 1's expected payoff hence equals the expected payoff that she enjoys from using ll. With player 2 mixing over λH and HH only, player 1's expected payoff in this equilibrium can be calculated as

$$q_{\lambda H} R[\eta(\sigma - 1/2) + 1/2]$$
(25)

Using (21), we undertake straightforward calculations to determine the derivative of $q_{\lambda H}$ with respect to η , finding that it is strictly negative. Using the corresponding expression and differentiating (25) with respect to η , and using the assumption that $R \geq 2r$, we find that player 1's expected payoff strictly decreases with respect to η .²³ For the mixed-strategy Nash equilibrium when $\eta < \hat{\eta}(\sigma)$, we thus conclude that player 1 experiences a first-mover disadvantage.

Finally, since player 2 mixes over λH and HH, her expected payoff equals the expected payoff that she receives from selecting HH. With player 1 mixing over ll and lh only, we may calculate player 2's expected payoffs as

$$p_{lh}R[\sigma - \eta(\sigma - 1/2)] \tag{26}$$

Employing (22), we may easily calculate the derivative of p_{lh} with respect to η , finding that it is strictly positive. Using the corresponding expression and differentiating (26) with respect to η , and using the assumption that $R \geq 2r$, we find that player 2's expected payoff strictly increases with respect to η .²⁴

²³Briefly, we find that the derivative of player 1's expected payoff with respect to η can be written as $(\sigma - 1/2)AB$, where A depends on r, R, η and σ and is strictly positive and where B likewise depends on r, R, η and σ and is strictly decreasing with respect to R. Specifically, we find that $A = (Rr/2)[r(\sigma - \eta(\sigma - 1/2)) + R(1/2 + \eta(\sigma - 1/2)) - \sigma(R + r)/2]^{-2}$. The value of B when evaluated at $R \ge 2r$ is thus less than or equal to the value of B when evaluated at R = 2r. We then find that B < 0 when evaluated at R = 2r.

²⁴Briefly, we find that the derivative of player 2's expected payoff with respect to η can be written as $(\sigma - 1/2)CD$, where C depends on r, R, η and σ and is strictly positive and where D likewise depends on r, R, η and σ and is strictly increasing with respect to R. Specifically, we find that $C = (Rr)[R(\sigma - \eta(\sigma - 1/2)) + r(1/2 + \eta(\sigma - 1/2)) - \sigma(R + r)/2]^{-2}$. The value of D when evaluated at $R \ge 2r$ is thus greater than

We summarize our findings with the following proposition:

Proposition 2. Consider the who's-on-first game for the BoS example. The two PSNE defined in Observation 4 do not generate a first-mover advantage or disadvantage at any η . For $\eta > \hat{\eta}(\sigma)$, the mixed-strategy Nash equilibrium described in Observation 5 by (19) and (20) exhibits a first-mover advantage. For $\eta < \hat{\eta}(\sigma)$, the mixed-strategy Nash equilibrium described in Observation 5 by (21) and (22) exhibits a first-mover disadvantage.

For the case where $\eta < \hat{\eta}(\sigma)$, our maintained assumption $\eta \ge 1/2$ provides a lower bound for η . Also, and as noted above, since $\hat{\eta}(\sigma) \ge 1$ when $\sigma \le R/(R+r)$, the case where $\eta < \hat{\eta}(\sigma)$ assuredly applies when $\sigma \le R/(R+r)$.

Thus, for PSNE, no first-mover advantage or disadvantage is generated. Turning to the mixed-strategy Nash equilibrium, we find that a first-mover advantage exists when $\eta > \hat{\eta}(\sigma)$; however, a first-mover disadvantage is generated when $\eta < \hat{\eta}(\sigma)$. Within the context of the mixed-strategy Nash equilibrium, therefore, whether player 1 gains when she becomes more likely to be the first mover depends on whether $\eta > \hat{\eta}(\sigma)$ or $\eta < \hat{\eta}(\sigma)$. By contrast, our findings above indicate that, regardless of whether $\eta > \hat{\eta}(\sigma)$ or $\eta < \hat{\eta}(\sigma)$, player 2 gains when the probability that player 1 is the first mover increases. Finally, recall also that player 1's equilibrium payoff features a downward jump as the $\eta = \hat{\eta}(\sigma)$ boundary is crossed.

As noted, our findings above for the BoS example are stated under the maintained assumption that $\eta \in [1/2,1)$. Fortunately, we can extend our findings to the general setting where $\eta \in (0,1)$. The key point is that, if $\eta < 1/2$, then the prior probability $1-\eta$ that player 2 moves first lies in the interval (1/2, 1). We may thus apply our findings while thinking of player 2 as playing the role of player 1, so that a *decrease* in η corresponds to an increase in "player 1's" probability of moving first.

To operationalize this approach, it is helpful to consult Figure 2. The portion of the figure that lies above the line along which $\eta = 1/2$ corresponds to the setting analyzed above and in Proposition 2. The bottom portion of the figure then depicts an analogous representation when $\eta < 1/2$ that is directed toward player 2. Note that player 2's prior probability of moving first exceeds $\hat{\eta}(\sigma)$ if and only if $1 - \eta > \hat{\eta}(\sigma)$ or equivalently $1 - \hat{\eta}(\sigma) > \eta$. The bottom portion of Figure 2 thus depicts the boundary curve $1 - \hat{\eta}(\sigma)$.

Consider first the parameter region for which the signal precision is modest, as captured by $\sigma \in (1/2, R/(R+r)]$. For a signal precision in this range, $\hat{\eta}(\sigma) \ge 1$ and so $\eta < \hat{\eta}(\sigma)$ follows immediately. Thus, if $\eta \in [1/2, 1)$, we can apply Proposition 2 for the case where $\eta < \hat{\eta}(\sigma)$. Likewise, for the setting where $\eta \in (0, 1/2)$, we then have that $1 - \hat{\eta}(\sigma) \le 0$, and

or equal to the value of D when evaluated at R = 2r. We then find that D > 0 when evaluated at R = 2r.



Figure 2: The functions $\hat{\eta}(\sigma)$ and $1 - \hat{\eta}(\sigma)$.

so $\eta > 1 - \hat{\eta}(\sigma)$ follows immediately. If $\eta \in (0, 1/2)$, we can thus similarly apply Proposition 2 but with player 2 playing the role of player 1.

To develop these points in more detail, let us start in the region of modest signal precision and increase η slightly. If $\eta \in [1/2, 1)$, we may directly apply our findings above and conclude for the mixed-strategy Nash equilibrium that player 1's equilibrium payoff strictly falls while player 2's equilibrium payoff strictly increases. If instead $\eta \in (0, 1/2)$, then we may apply our previous findings with player 2 playing the role of player 1. Thus, for the mixed-strategy Nash equilibrium, a slightly lower value for η (i.e., a slightly higher chance that player 2 moves first) results in a strictly lower equilibrium payoff for player 2 and a strictly higher equilibrium payoff for player 1. Equivalently, If $\eta \in (0, 1/2)$, then in the mixed-strategy Nash equilibrium a slightly higher value for η induces a strictly higher equilibrium payoff for player 2 and a strictly lower equilibrium payoff for player 1. Hence, for any $\eta \in (0, 1)$, in the region of modest signal precision and for the mixed-strategy Nash equilibrium, an increase in η generates a strictly lower (higher) equilibrium payoff for player 1 (2).

Consider second the parameter region for which the signal precision is high, as captured by $\sigma \in (R/(R+r), 1)$. We then have $\hat{\eta}(\sigma) \in ((R-r)/R, 1)$. Two possibilities arise. The first possibility is that $\eta \in (1 - \hat{\eta}(\sigma), \hat{\eta}(\sigma))$. Note that this interval is non-empty and is centered around $\eta = 1/2$.²⁵ For all $\eta \in (1 - \hat{\eta}(\sigma), \hat{\eta}(\sigma))$, we may follow the logic just detailed and conclude for the mixed-strategy Nash equilibrium that a slightly higher value for η generates a strictly lower equilibrium payoff for player 1 and a strictly higher equilibrium payoff for player 2.

Given a high signal precision, the second possibility is that $\eta > \hat{\eta}(\sigma)$ or $\eta < 1 - \hat{\eta}(\sigma)$. If $\eta > \hat{\eta}(\sigma)$, then $\eta > 1/2$ follows. We may then directly invoke the findings described above. Thus, for the mixed-strategy Nash equilibrium, if $\eta > \hat{\eta}(\sigma)$, then a slightly higher value for η results in a strictly higher equilibrium payoff for player 1 and a strictly higher equilibrium payoff for player 2. By reversing player names, it is straightforward to extend the results to the analogous case where $\eta < 1 - \hat{\eta}(\sigma)$. Hence, for the mixed-strategy Nash equilibrium, if $\eta < 1 - \hat{\eta}(\sigma)$, then a slightly higher value for η results in a strictly lower equilibrium, if $\eta < 1 - \hat{\eta}(\sigma)$, then a slightly higher value for η results in a strictly lower equilibrium payoff for player 2.

Let π_1 and π_2 denote the equilibrium payoffs to players 1 and 2, respectively, in the mixed-strategy Nash equilibrium. We capture the information developed above in Figures 3 and 4, which illustrate the relationship between the equilibrium payoffs and η under both modest signal precision (Figure 3) and high signal precision (Figure 4)



Figure 3: Modest Signal Precision.

As illustrated, given modest signal precision, player 1's (2's) equilibrium payoff strictly

²⁵As illustrated in Figure 1, given $\sigma < 1$, we have that $\hat{\eta}(\sigma) > 1/2$ and thus $\hat{\eta}(\sigma) > 1 - \hat{\eta}(\sigma)$, where the inequalities hold even at $\sigma = 1$ if R > 2r.



Figure 4: High Signal Precision.

falls (rises) with η for $\eta \in (0, 1)$. By contrast, under high signal precision, player 1's equilibrium payoff strictly falls with η for $\eta \in (0, \hat{\eta}(\sigma))$ and strictly rises with η for $\eta \in (\hat{\eta}(\sigma), 1)$. As previously discussed, under high precision, π_1 also exhibits a downward jump as the $\eta = \hat{\eta}(\sigma)$ boundary is crossed. Under high signal precision, we may also confirm that π_1 is strictly lower as η approaches one than as η approaches $\hat{\eta}(\sigma)$ from below. As depicted, player 2's equilibrium payoff takes a similar pattern, except that the jump occurs at $\eta = 1 - \hat{\eta}(\sigma)$ rather than at $\eta = \hat{\eta}(\sigma)$.

Exploiting the symmetry of the BoS example, let us now explicitly expand our definition of a first-mover advantage to include player 2. Specifically, we say that an equilibrium exhibits a first-mover advantage (disadvantage) for player 1 at η if player 1's expected payoff strictly increases (decreases) following a slight rise in η and that an equilibrium exhibits a first-mover advantage (disadvantage) for player 2 at η if player 2's expected payoff strictly decreases (increases) following a slight rise in η .

We may now summarize our discussion with the following proposition:

Proposition 3. Consider the who's-on-first game for the BoS example, extended to allow for any $\eta \in (0, 1)$. In the mixed-strategy Nash equilibrium:

(1) (Modest signal precision) If $\sigma \in (1/2, R/(R+r)]$, then player 1's (2's) equilibrium payoff strictly falls (rises) with η for $\eta \in (0, 1)$. For all $\eta \in (0, 1)$, the equilibrium thus exhibits a first-mover disadvantage for players 1 and 2.

(2) (High signal precision) If $\sigma \in (R/(R+r), 1)$, then player 1's equilibrium payoff strictly falls with η for $\eta \in (0, \hat{\eta}(\sigma))$ and strictly rises with η for $\eta \in (\hat{\eta}(\sigma), 1)$, and player 2's equilibrium payoff strictly falls with η for $\eta \in (0, 1 - \hat{\eta}(\sigma))$ and strictly rises with η for $\eta \in (1 - \hat{\eta}(\sigma), 1)$. The equilibrium thus exhibits (i) a first-mover disadvantage for player 1 for all $\eta \in (0, \hat{\eta}(\sigma))$ and a first-mover advantage for player 1 for all $\eta \in (\hat{\eta}(\sigma), 1)$, and (ii) a first-mover advantage for player 2 for all $\eta \in (0, 1 - \hat{\eta}(\sigma))$ and a first-mover disadvantage for player 2 for all $\eta \in (1 - \hat{\eta}(\sigma), 1)$.

For the mixed-strategy Nash equilibrium in the BoS example, Proposition 3 thus indicates that a first-mover advantage arises for player 1 if and only if player 1 is sufficiently likely to move first $(\eta > \hat{\eta}(\sigma))$ and the precision of the signal is sufficiently high $(\sigma > R/(R+r))$. A similar statement applies for player 2.

We return now to discuss in greater detail the marginal nature of the first-mover advantage explored in this paper. While we define a first-mover advantage for a given player in terms of the change in that player's equilibrium payoff as that player's prior probability of moving first is increased, we note that this is not the only definition that could be explored. For example, it is also of interest to know if a player enjoys a greater equilibrium payoff when the player is almost certain to move first in comparison to other benchmarks, such as when the player is almost certain to move second or moves first half the time. In the past, these comparisons have been made by comparing equilibrium outcomes of different games. For a given Nash equilibrium, we can capture these comparisons here as well, by simply examining payoffs for different values of η .

Consider the who's-on-first game with the BoS example. The comparisons just described are immediate for PSNE, since equilibrium payoffs are then insensitive to η . For the mixedstrategy Nash equilibrium, it is also straightforward to make these comparisons if the signal precision is modest. As Figure 3 illustrates, a given player then does strictly better when she is almost sure to move second than when she is almost sure to move first, and enjoys an intermediate equilibrium payoff when she moves first with probability one half. Calculations also confirm that the same ranking holds for the mixed-strategy Nash equilibrium when the signal precision is high. Figure 4 reflects this ranking.²⁶

²⁶When the signal precision is high, we confirm the ranking as follows. First, to compare the endpoints for player 1, we simply note that π_1 when $\eta = 0$ equals π_2 when $\eta = 1$ These values are straightforward to compute. Next, to show that π_1 when $\eta = 1/2$ strictly exceeds π_1 when $\eta = 1$ (and thus likewise that π_2 when $\eta = 1/2$ strictly exceeds π_2 when $\eta = 0$), we use Proposition 2 and $\hat{\eta} > 1/2$ to conclude that π_1 when $\eta = 1/2$ strictly exceeds the limiting value for π_1 as η approaches $\hat{\eta}$ from below. As previously noted, direct calculations in turn confirm that this limiting value strictly exceeds π_1 when $\eta = 1$, given that $\sigma > R/(R+r)$ under high signal precision.

Whether a first-mover advantage exists can be sensitive to the definition used. Consider again the who's-on-first game with the BoS example. First, as Figure 4 illustrates, while under our definition player 1 enjoys a first-mover advantage at $\eta > \hat{\eta}$ when the signal precision is high, the level of player 1's equilibrium payoff is nevertheless then strictly lower for all $\eta > \hat{\eta}$ than for all $\eta < \hat{\eta}$. This possibility arises due to the downward jump in player 1's payoff at $\hat{\eta}$ in the mixed-strategy Nash equilibrium. From this perspective, the marginal definition used here may overstate the existence of a first-mover advantage. Second, consider the PSNE with the strategy pair $(ll, \lambda\lambda)$, which generates the payoff of R for player 1. Under the definition used here, player 1 does not enjoy a first-mover advantage in this equilibrium, since her equilibrium payoff is independent of η . The payoff level enjoyed by player 1 in this equilibrium, however, is the highest feasible payoff for player 1. From this perspective, the marginal definition used here may understate the existence of a first-mover advantage.

4 Conclusion

This paper examines a two-player game in which actions are imperfectly observed and move order is uncertain. A player who is called upon to move thus does not know whether the other player has already moved or will move subsequently. We assume, however, that a player does observe a signal before moving, and that, when the player moves second, the distribution of the signal is impacted by the other player's action. The game thus includes a channel through which the strategic advantage of commitment may be expressed. We refer to this game as the "who's-on-first game."

We embed two examples - the commitment example and the battle-of-the-sexes (BoS) example - into the who's-on-first game. In each example, each player has two possible actions and observes the realization of a binary signal prior to moving. The associated strategic-form game is thus a 4x4 game. For each example, we characterize the full set of pure- and mixed-strategy Nash equilibria.

An attractive feature of the who's-on-first game is that the strategic advantage of commitment can be studied within a single game by varying model parameters. In particular, we study the first-mover advantage by examining how, for each Nash equilibrium, player 1's payoff changes as she becomes more likely to move first. We find that no first-mover advantage arises in pure-strategy Nash equilibria; however, in mixed-strategy Nash equilibria, behavior and payoffs vary with the probability that player 1 moves first. We find that whether a first-mover advantage exists varies across mixed-strategy equilibria in the commitment example and across parameter regions in the BoS example. For the mixed-strategy Nash equilibrium in the BoS example, we find under our (marginal) definition that a firstmover advantage arises for a given player if and only if that player is sufficiently likely to move first and the precision of the signal is sufficiently high. Even in this scenario, however, the player would enjoy a strictly higher equilibrium payoff level were she to move first with a discretely lower probability, including all probabilities in (0, 1/2]. We also show that the who's-on-first game is versatile and picks up other benchmark games as limiting cases. Further, using the commitment example, we offer a new perspective on Newcomb's paradox.

Previous analyses of the first-mover advantage focus on the comparison of the Nash equilibrium outcomes in a simultaneous-move game with the subgame perfect equilibrium outcomes in a counterpart sequential-move game in which player 1 moves first and takes an action that is perfectly observed by player 2. This comparison can be picked up as a special case here, as η and σ both go to one. As the limit is taken, whether the standard first-mover advantage obtains then depends on the equilibrium that is selected. For the who's-on-first game with the commitment (BoS) example embedded, the traditional Stackelberg advantage is obtained in the limit for one of the two mixed-strategy (pure-strategy) Nash equilibria.

In future work, the basic approach taken here can be directly extended to other 2x2 examples. Conceptually, similar methods can also be used to characterize all Nash equilibria for the who's-on-first game for other strategic settings involving a finite number of actions for players. The main challenge in analyzing the who's-on-first game in this multi-action setting is that the full set of mixed-strategy Nash equilibria may be difficult to characterize, since the number of possible supports for mixed strategies grows quickly with the number of actions. Other directions for future research include games with a continuum of actions and richer signal structures.

5 References

- Bagwell, K. (1995), "Commitment and Observability in Games," Games and Economic Behavior, 8, 271-280.
- Bagwell, K. (2018), "Commitment with Private Resolve," January 22, working paper.
- Bhaskar, V. (2009), "Commitment and Observability in a Contracting Environment," Games and Economic Behavior, 66, 708-720.
- van Damme, E. and S. Hurkens (1996), "Commitment Robust Equilibria and Endogenous Timing," *Games and Economic Behavior*, 15, 290-311.
- van Damme, E. and S. Hurkens (1997), "Games with Imperfectly Observed Commitment," Games and Economic Behavior, 21, 282-308.

Doval, L. and J. C. Ely (2020), "Sequential Information Design," April 29, working paper.

- Gallice, A. and I. Monzon (2019), "Co-operation in Social Dilemmas Through Position Uncertainty," *Economic Journal*, 129, 2127-2154.
- Gardner. M. (1973), "Free Will Revisited, with a Mind-Bending Prediction Paradox by William Newcomb," *Scientific American*, 229(1), 104-109.
- Gardner, M. (2001), "Newcomb's Paradox," The Colossal Book of Mathematics, Norton, New York, NY, Chapter 44, 580-591.
- Guth, W., G. Kirchsteiger and K. Ritzberger (1998), "Imperfectly Observed Commitments in n-Player Games," *Games and Economic Behavior*, 23, 54-74.
- Hamilton, J. H. and S. M. Slutsky (1990), "Endogenous Timing in Duopoly Games: Stackelberg or Cournot Equilibria," *Games and Economic Behavior*, 2, 29-46.
- Hamilton, J. H. and S. M. Slutsky (1993), "Endogenizing the Order of Moves in Matrix Games," *Theory Decision*, 34, 47-62.
- Henkel, J. (2002), "The 1.5th Mover Advantage," *The RAND Journal of Economics*, 33(1), 156-170.
- Harsanyi, J. C. (1995), "A New Look at Strategic Rationality: At Equilibrium Selection, Risk Dominance and Payoff Dominance," *Games and Economic Behavior*, 8, 91-122.
- Harsanyi, J. C. and R. Selten. (1988), A General Theory of Equilibrium Selection in Games, MIT Press, Cambridge, MA.
- Kamada, Y. and M. Kandori (2020), "Revision Games," *Econometrica*, 88.4, 1599-1630.
- Kreps, D. M. and G. Ramey (1987), "Structural Consistency, Consistency, and Sequential Rationality," *Econometrica*, 55(6), 1331-1348.
- Maggi, G. (1999), "The Value of Commitment with Imperfect Observability and Private Information," *The Rand Journal of Economics*, 30, 555-574.
- Morgan, J. and F Vardy (2007), "The Value of Commitment in Contests and Tournaments when Observation is Costly," *Games and Economic Behavior*, 60, 326-338.
- Morgan, J. and F. Vardy (2013), "The Fragility of Commitment," Management Science, 59, 1344-1353.

- Nishihara, K. (1997), "A Resolution of N-person Prisoners' Dilemma," *Economic Theory*, 10(3), 531-40.
- Nozick, R. (1969), "Newcomb's Problem and Two Principles of Choice," in N. Rescher (Ed.), *Essays in Hone of Carl G. Hempel*, Springer Dordrecht, 114-146.
- Oechssler, J. and K. Schlag (2000), "Does Noise Undermine the First-Mover Advantage? An Evolutionary Analysis of Bagwell's Example," *International Game Theory Review*, 2, 83-96.
- Rosenthal, R. W. (1991), "A Note on Robustness of Equilibria with Respect to Commitment Opportunities," *Games and Economic Behavior*, 3, 237-243.
- Vardy, F. (2004), "Stackelberg Leader Games with Observation Costs," Games and Economic Behavior, 49, 374-400.
- Weber, T.A. (2016), "A Robust Resolution of Newcomb's Paradox," *Theory and Decision*, 81, 339-356.
- Wolpert, D. H. and G. Benford. (2013), "The Lesson of Newcomb's Paradox," *Synthese*, 190, 1637-1646.