Commitment with Private Resolve*

Kyle Bagwell†

January 22, 2018

Abstract

This paper considers a model with a leader and follower in which the leader is privately informed about its resolve to follow through on a proposed course of action. The leader's initial or promised action, referred to as the "leader action," may then play both commitment and signaling roles. After putting the game in a reduced form where the leader's payoffs are expressed as a function of the leader action and the follower's belief, we show that a Nash pooling equilibrium exists in which, whether the leader is resolute or flexible, the leader action is set equal to the Nash action that the leader would choose in a pure-strategy equilibrium of the associated simultaneous-move game. Motivated by the D1 refinement for signaling games, we define a refinement for the reduced-form game and show that the Nash pooling equilibrium outcome is the unique refined pooling equilibrium outcome. We show further that in any refined equilibrium the final actions are the Nash actions of the simultaneous-move game, with each player thus earning Nash payoffs. These findings hold even when the leader is almost certain to have resolve. The arguments are first developed in a strategic setting motivated by the Cournot quantity game. We then identify sufficient conditions on general reduced-form payoff functions and provide applications.

*First posted draft: November 2, 2017. Very helpful comments from V. Bhaskar, Steve Callander, Isaias Chaves Villamizar, Kyungmin Kim, Muriel Niederle, John Roberts, Andy Skrzypacz, Takuo Sugaya, Claire Xue, Robert Wilson and participants at the Stanford Political Economy Theory Lunch seminar and the University of Texas Microtheory seminar are gratefully acknowledged.

†Department of Economics at Stanford University; and NBER.

1 Introduction

The strategic value of commitment is a fundamental and celebrated insight from game theory that has been widely used in economics and other disciplines. The idea is subtle but powerful: by moving first with an unalterable and observable action, a player may be able to change the behavior of the second mover in a way that is advantageous to the first mover. In the classic quantity-game setting, the simultaneous-move game admits a unique pure-strategy (Cournot) equilibrium, and the first mover in the associated sequential-move game gains precisely because of its ability to commit itself to an action that is not a best-response to the induced follower action. Stackelberg (1934) offers a formalization of this idea, and Schelling's (1960) classic work explores many interesting implications.

A direct interpretation of the standard sequential-move game is that the leader takes an unalterable action that is observed by the follower, but this game can also be interpreted as capturing situations in which the leader makes a credible and observed promise about a future action. Under either interpretation, the follower observes the action or promise and then selects a best response. For many applications, however, the standard representation of the first-mover advantage may be questioned in two important respects. First, it may be more realistic to allow that the leader's initial action or promise is of uncertain significance for the follower. The follower may reasonably wonder if the leader's initial action might be later modified; likewise, the follower may reasonably wonder if the leader will follow through on a promise about a future action. In short, the follower may be uncertain of the leader's resolve to follow through on a proposed course of action. Second, the leader may have private information about its resolve when taking an initial action or issuing a promise. The level of the initial or promised action may then also play a signaling role as the follower attempts to gauge the relevance of the leader's initial move.

With these considerations in mind, we consider here a game with private resolve. The incomplete-information game works as follows. Nature determines whether the leader has resolve or is flexible, where Nature picks each type with positive probability. The leader is privately informed of Nature's choice and then selects a "leader action" corresponding to an initial action or promise. The follower observes the leader action, forms a belief about the leader's type, and then takes its own action. If the leader has resolve, then the leader is committed to the leader action. The payoffs to the leader and follower are then determined based on the leader action and the follower's action. If instead the leader is flexible, then the leader is unconstrained by the leader action and optimally selects its final action as a best response to the follower's anticipated and simultaneous action. Given its beliefs, the follower thus optimally selects its action as a best response against the leader's uncertain behavior, knowing that a resolute leader will stick with the leader action while a flexible leader will simultaneously select its best response.

A key feature of the game with private resolve is that the leader action may affect the follower's behavior through two channels. First, as in the standard sequential-move game, the follower's behavior is directly influenced by the leader action if the leader has resolve. Second, as in a signaling game, the leader action may signal the leader's type and thus affect the probability with which the follower believes the leader to have resolve.[1] Given the central role of beliefs, it is perhaps not surprising that multiple equilibria exist in the game. Following the literature on signaling games, we thus use a refinement motivated by the D1 refinement to select among equilibria, and we refer to any equilibrium (outcome) that satisfies this refinement as a refined equilibrium (outcome).[2]

The game can be interpreted in several ways. For example, in the classical quantity-game setting, the Cournot and Stackelberg outcomes correspond to limiting cases where Nature is known to give the leader resolve with probability zero and one, respectively. In the game with private resolve, we add private information and focus on the middle ground, where this probability resides between zero and one. One interpretation is that the leader chooses an initial quantity that is observed by the follower, as in the standard Stackelberg game, but with the added wrinkle that the leader is privately informed about whether the quantity intended for the relevant market can be costlessly re-set at the time that the follower makes its quantity choice. An alternative interpretation is that the "leader quantity" represents instead a promise or pre-announcement as to a planned output. The leader may have superior information in this case as to whether it is bound by its promise, whether for financial, legal or moral reasons. A large theoretical and empirical literature addresses the potential strategic aspects of firms' preannouncements about future production or capacity expansion plans.[3]

Another interpretation is that the leader is a government official who seeks to shape market expectations with a leader action that corresponds to a promise. An example arises in the literature following Barro and Gordon (1983) and Backus and Driffill (1985) that addresses a monetary-policy commitment issue. In this model, the central bank would like to commit to zero inflation but has an incentive to surprise the public with positive inflation in order to expand aggregate output. Since the public has rational expectations, the no-commitment solution entails inflation that is fully anticipated, resulting in a Pareto-inferior outcome in comparison to the commitment solution of zero inflation. The analysis

[1]Another way of describing the game is that the leader is privately informed when taking its leader action as to whether this action is an irrevocable commitment or simply "cheap talk." See Crawford and Sobel (1982) for the pioneering analysis of cheap-talk games.

[2]As explained below, we apply the refinement to a reduced-form game. We show in Section 5.3 that the refinement's restrictions are much stronger than needed for our findings. See Cho and Kreps (1987) and Fudenberg and Tirole (1991) for further discussion of the D1 refinement in signaling games.

[3]See Bayus, et al (2001), Christensen and Caves (1997), Corona and Nan (2013), Doyle and Snyder (1999) and Gilbert and Lieberman (1987). For discussions of the relationships between quantity and capacity choices, see Kreps and Scheinkman (1983), Maggi (1996) and Poddar and Sasaki (2002).

developed here offers a framework within which to explore commitment games where the leader has private information about its capacity to keep its promise.[4]

Motivated by these and other applications, we analyze equilibrium and refined equilibrium behavior in a game with private resolve. To fix ideas, we begin with a strategic setting motivated by the classic quantity-game setting, wherein Firm 1, the leader, is either resolute or flexible and selects a leader quantity. The follower in this setting is Firm 2, which forms a belief about Firm 1's type and selects its best-response quantity under uncertainty. In line with the main features of standard treatments of quantity games, we assume that the best-response functions are decreasing and that a higher quantity by one firm lowers the profit of the other. We also fully model the "period-2 equilibrium quantity" choices that arise after the leader quantity is selected. The period-2 equilibrium quantity choices are the best-response choices for the flexible Firm 1 and Firm 2, when the resolute Firm 1 is constrained to maintain its leader quantity and Firm 2 best responds against an uncertain output from Firm 1. The period-2 equilibrium quantities are functions of the leader quantity and Firm 2's belief after observing the leader quantity. They are the final quantities that would be chosen by the flexible Firm 1 and Firm 2 in the continuation of any sequential equilibrium (Kreps and Wilson, 1982a) of the game.

To characterize equilibrium behavior, we embed the period-2 equilibrium quantities into Firm 1's payoff function and put the game in "reduced form." In the reduced-form game, payoffs are expressed directly as a function of the leader quantity and the follower's belief. A pure-strategy equilibrium for the reduced-form game is then comprised of an optimal leader-quantity choice for each type of Firm 1 along with a belief function for Firm 2, where beliefs are consistent with Bayes' rule along the equilibrium path. We show that a "Nash pooling equilibrium" exists in which, whether Firm 1 is resolute or flexible, it sets its leader quantity equal to the (Cournot-) Nash quantity that it would choose in a pure-strategy equilibrium of the associated simultaneous-move game.[5] We also characterize the set of separating equilibria and a class of "generalized-Stackelberg pooling equilibria." The separating equilibria are only superficially distinct from Nash pooling equilibria and lead to the same final Nash outputs and payoffs. The Nash pooling equilibrium outcome is also the only pooling equilibrium outcome that exists robustly.

[4]In a "zero lower bound" world, the central bank may instead wish to use "forward guidance" over the path of future interest rates to *raise* the expected level of future inflation. As Rogoff (2017, pp. 53-4) notes, however, a credible promise associated with forward guidance may be difficult to achieve, "given 1) the turnover in central bank governing boards, and 2) the central bank has an incentive not to keep its promise if the economy does indeed recover." Such considerations may limit a central bank's ability to make credible promises, whether the promise concerns lowering or raising the future rate of inflation.

[5]In the sequel, when we refer to "Nash quantities" and the "Nash equilibrium," it is understood that we are referring to the actions that would be chosen in the pure-strategy Nash equilibrium of the associated simultaneous-move game. In terms of the quantity-game setting, therefore, Nash quantities are the standard Cournot-Nash quantities.

Motivated by the D1 refinement for signaling games, we define a refinement for the reduced-form game and show that the Nash pooling equilibrium outcome is the unique refined pooling equilibrium outcome. We conclude that refined equilibria exist and that in any refined equilibrium the final outputs are the Nash outputs, with each firm thus earning its Nash profits. In other words, under the refinement for the reduced-form game, the final actions and payoffs are the same as would have occurred had the interaction been modeled as a simultaneous-move game.[6] From this perspective, our results indicate a sense in which the strategic advantage of commitment may be lost when the leader has private information about its resolve. We note that the value of commitment may be lost in this regard even when the leader is almost certain to have resolve.

The power of the refinement builds from a simple intuition. Consider first the reason that the Nash pooling equilibrium is refined. If Firm 1 were to deviate to a leader quantity different from the Nash quantity, then Firm 2 might plausibly evaluate which type of Firm 1 would be most likely to gain from the deviation. An important feature of the Nash pooling equilibrium is that both types of Firm 1 make the same Nash payoff in equilibrium; however, if Firm 2 were to form a belief after observing a deviation under which the resolute Firm 1 gains, then the flexible Firm 1 would be sure to gain as well. The reason is that the flexible Firm 1 has an added benefit from deviation: it can adjust its final quantity to be a best response to Firm 2's quantity. In terms of the refinement, we are thus on solid ground if we associate any deviation with a flexible Firm 1. This is enough to ensure that the Nash pooling equilibrium exists as a refined equilibrium. We note that this argument is quite general and, in particular, makes no reference to the relative magnitudes of the Stackelberg and Nash quantities for Firm 1.

Consider second the reason that no other pooling equilibrium is refined. For this strategic setting, we show that any other pooling equilibrium must be at a leader quantity that exceeds the Nash quantity (i.e., that is in the direction of Firm 1's Stackelberg quantity). At such an equilibrium, the flexible Firm 1 earns a greater equilibrium payoff than does the resolute Firm 1, whose quantity is not a best response. If Firm 1 were to deviate to a slightly lower leader quantity, then the period-2 equilibrium quantity of Firm 2 would remain unchanged if Firm 2's belief were to adjust to place just the right amount of additional weight on a resolute Firm 1. At this belief, the flexible Firm 1 does not gain from the deviation, since it continues to best respond against an unaltered output by Firm 2. The resolute Firm 1, however, *does* gain from the deviation under this belief, since the deviation enables it to place its quantity closer to its best-response level without altering Firm 2's behavior. Building from this insight, we show that under the refinement the deviation should be believed to have come from a resolute Firm 1. With this belief, the deviation would lead to a beneficial quantity reduction from Firm 2, and so the resolute

[6] As confirmed in Section 5.5, this result also holds when refined mixed-strategy equilibria are included.

Firm 1 would indeed deviate, ensuring that alternative pooling equilibria are not refined. We note that this argument makes reference to the relative magnitudes of the Stackelberg and Nash quantities and so requires modification in other strategic settings.

We develop our arguments in a strategic setting motivated by a standard quantity game. The quantity-game application is of direct interest, and to fix ideas we utilize the language of quantity competition when presenting our initial model. It is important to emphasize, though, that our description of this strategic setting draws only on two main features: decreasing best-response functions and a negative externality to one player when the other player selects a higher action. These strategic features arise as well in a range of other applications. For example, as is well known (Cournot, 1963, chapter ix; Sonnenschein, 1968), when marginal costs are zero, this strategic setting also arises in a price-setting game if firms sell complementary products that are of no use unless combined on a 1-to-1 basis to form a composite product. As another example, the strategic choices of the two firms could be investments in cost reduction, where some oligopoly game generates payoffs once the final investments and thus production costs are determined. For many post-investment oligopoly environments, investment best-response functions are negatively sloped and a higher investment by one firm lowers the profit of the other.[7]

At the same time, other important applications arise in different strategic settings. For example, when firms sell differentiated products and set prices, the price best-response functions may be increasing and a higher price by one firm may convey a positive externality to the other firm. While the price- and quantity-game settings differ in those respects, they share the property that the Stackelberg action of the leader exceeds the Nash action that it would take in a simultaneous-move game. By contrast, in the monetary-policy game, for example, the Stackelberg solution lies below the Nash action. From such examples, we are led to ask two questions. First, do related results hold for other strategic settings? Second, if so, what are the unifying features of strategic settings that underlie the characterizations of (refined) equilibrium behavior?

To address these questions, we approach the problem in a more abstract way and simply begin with a reduced-form payoff for Firm 1. The reduced-form payoff could reflect any of a variety of channels through which the leader action and associated belief affect the subsequent interaction between the leader and follower. Starting directly with a reduced-form payoff function for Firm 1, our approach is to look for general sufficient conditions on the reduced-form payoff functions that deliver our results. This approach

[7] For a standard Cournot setting, a firm is hurt when its rival invests more in cost reduction; furthermore, investment best-response functions are decreasing if the post-investment oligopoly game entails Cournot competition with linear demand and costs. Besley and Suzurmura (1992) and Reinganum (1983) show further that investment best-response functions are decreasing under subsequent Cournot competition for some popular non-linear demand structures. Bagwell and Staiger (1994) show that the described strategic features for the investment game are also satisfied when the investments are followed by price competition among firms selling differentiated products when demand and costs are linear.

clarifies the driving forces behind the analysis and facilitates future applications. To implement this approach, we find it useful to organize the discussion around two strategic settings, those for which the Stackelberg solution exceeds the Nash action and those for which the Stackelberg solution lies below the Nash action. We refer to the respective settings as the "Stackelberg-up" and "Stackelberg-down" cases, respectively.

We begin by identifying a general set of baseline assumptions under which the Nash pooling equilibrium is refined. As anticipated, the baseline assumptions do not rely on the distinction between the Stackelberg-up and -down cases. We then identify additional assumptions and a single-crossing-property assumption, which each take different forms for the Stackelberg-up and -down cases and which together ensure that any refined equilibrium in the reduced-form game with private resolve generates Nash payoffs. The sufficient conditions are easy to use for applications, once Firm 1's payoff is captured in reduced form. To illustrate the value of these results, we consider two applications, each of which is explored for simplicity in a quadratic-payoff formulation. The applications are a quantity-game setting with linear demand and costs, and the monetary-policy game. The first application fits in the Stackelberg-up category, whereas the monetary-policy game belongs to the Stackelberg-down category. For each application, we use our general sufficient conditions to show that refined equilibria exist and that in any refined equilibrium Firm 1 earns Nash payoffs, regardless of its type. We also discuss how are results could be used for other applications, including price-game settings.

The paper is organized as follows. Section 2 offers a review of the related literature. In Section 3, we set up the general game that is motivated by the strategic setting associated with quantity competition. We put the game into reduced form and characterize equilibrium behavior in Section 4. The analysis of refined equilibria is developed in Section 5. In Section 6, we start with general reduced-form payoff functions and develop sufficient conditions for our findings in both Stackelberg-up and -down settings. We also show that our results can be easily used in applications. Section 7 discusses an extension in which the leader is privately informed about the probability that it will be able to carry though on its commitment. Section 8 concludes. Remaining proofs are collected in the Appendix.

2 Related Literature

The paper contributes to several literatures. The main finding - that all refined equilibria in the reduced-form game with private resolve generate Nash payoffs - is broadly reminiscent of research on the value of commitment when the standard sequential-move game is perturbed to allow for imperfect observability. Bagwell (1995) considers a general "noisy-leader game" with a non-moving support assumption. Assuming that the follower's best-response correspondence is single-valued, he shows that the set of pure-strategy Nash equilibrium outcomes of the noisy-leader game coincides exactly with the set of pure-strategy Nash equilibrium outcomes in the associated simultaneous-move game, even when the noise in the signal is arbitrarily small.[8]

In comparison to this literature, the model considered here maintains the standard assumption of perfect observability and instead introduces private information regarding the significance of the initial action or the credibility of the promise. Despite this difference, the current paper shares with Bagwell (1995) a common perspective: the defining feature of the Stackelberg solution, namely, that the leader commits to an action that is not a best response, may generate deviation opportunities for the leader and thus raise potential robustness concerns for this solution when the game is appropriately perturbed.

The paper is also related at a broad level to the literature on reputation formation in finitely-repeated games. The pioneering papers in this literature are Kreps and Wilson (1982b) and Milgrom and Roberts (1982). A key message is that the introduction of a little incomplete information can generate rich behavior as compared to that which emerges in the complete-information benchmark. With regard to payoffs, Fudenberg and Levine (1989) show that a sufficiently patient "normal" long-lived player facing a sequence of short-lived players can effectively achieve the discounted expected payoff associated with the Stackelberg action, provided that there is at least a small probability of a type that is committed to the Stackelberg action. By contrast, we focus here on the addition of a possibly small amount of incomplete information to a sequential-move game so that the follower is uncertain about the significance of the leader's commitment. In this context, any refined equilibrium of the reduced-form game generates Nash payoffs, provided that there is at least a small probability of a non-Stackelberg (i.e., flexible) type.

In the literature on reputational bargaining, Abreu and Gul (2000) examine an infinite-horizon model with two-sided offers in which players may develop reputations as behavioral types that insistently demand fixed surplus shares, where each player has many potential behavioral types with each such type exogenously committed to a different share.[9] Our model is broadly related in that the leader may be resolute or flexible, where a resolute (flexible) leader is similar to a behavioral (normal) type. Kim's (2009) analysis of bargaining with one-sided offers in the context of the durable-goods monopoly problem is more closely related. In one version of his model, the seller is either a normal type or a

[8]For a 2×2 example, Bagwell also shows that mixed-strategy equilibria exist for the noisy-leader game, where one such equilibrium converges to the Stackelberg outcome as the noise in the signal goes to zero. van Damme and Hurkens (1997) show for general noisy-leader games that a mixed-strategy equilibrium exists that is close to the Stackelberg equilibrium when the noise is small. They also construct a refinement that selects this equilibrium. Oechssler and Schlag (2000) examine the noisy-leader game with a wide range of evolutionary and learning dynamics. They find that the pure-strategy Nash equilibrium is rarely eliminated and often uniquely selected. For other contributions, see Bhaskar (2009), Guth, et al (1998), Maggi (1999), Morgan and Vardy (2007, 2013) and Vardy (2004).

[9]See also Chatterjee and Samuelson (1987, 1988), Compte and Jehiel (2002) and Myerson (1991).

"rational commitment" type, where the seller's type is privately observed before any price offers are made and a rational commitment seller chooses the price to which a commitment is made.[10] Like Kim's rational commitment seller type, the resolute leader in our model privately observes its type before selecting the leader action to which it is committed. Our finding that the Nash pooling equilibrium exists robustly has an interesting counterpart in Kim's model. In his continuous-time, gap-case model, the no-commitment solution is given by the Coase conjecture, and the corresponding Coasian equilibrium exists robustly.

Our analysis differs from Kim's (2009) in three main ways. First, the aim and scope is different. We do not consider the durable-goods monopoly problem but instead examine a family of two-period games. This approach explicitly links our analysis to standard Stackelberg settings and facilitates applications to oligopoly theory and monetary policy. Second, the formal models have different features. For example, the flexible leader is never committed to its leader action in our model whereas the initial price offer of the normal seller in Kim's model may be accepted and thus may have direct payoff relevance.[11] Third, the refinement and resulting selected outcomes differ. The D1-based refinement used here achieves its power through a single-crossing property that builds directly from the defining feature of the Stackelberg solution and ensures that any refined equilibrium delivers Nash payoffs. Kim is able to refine the equilibrium set for his game by selecting those equilibria that remain when the game is perturbed to allow for a small probability of behavioral types and the probability of those types is then taken to zero. He shows that this refinement is always satisfied by the no-commitment (Coasian) equilibrium but is also satisfied by a pooling equilibrium with a higher initial price when the probability of a rational commitment type exceeds a cutoff value.

The paper is also related to recent work by Sanktjohanser (2017) and Dai (2017).[12] Sanktjohanser extends the Abreu-Gul model to allow that players make simultaneous offers after observing whether they are rational or "stubborn," where a stubborn type chooses the offer to which a commitment is made.[13] As in the discussion in the previous paragaph, the current paper has a different aim and scope and also differs formally, since here the flexible leader is never committed to its leader action. As well, the simultaneous-offer structure considered by Sanktjohanser introduces payoff discontinuities for the stub-

[10] See also Inderst (1995) for a related model with one-sided offers and an exogenous commitment price.

[11] The equilibrium set characterizations also differ. Kim establishes that the no-commitment (Coasian) equilibrium is the unique equilibrium when the prior probability of the rational-commitment type is below a cutoff value. By contrast, in our model, for any fixed prior probability of the resolute type, the counterpart Nash pooling equilibrium outcome is never unique without refinement. See Proposition 4.

[12] I thank V. Bhaskar for bringing Kim's (2009) and Sanktjohanser's (2017) papers to my attention, and I thank Kyungmin Kim for bringing Dai's (2017) work to my attention.

[13] See also Kambe (1995). He analyzes a bargaining model with two-sided offers in which players make offers to which they may become committed, where in one version of his model players privately observe their types before making offers.

born type that are not present in the model considered here. One point of contact is that Sanktjohanser refines equilibria using the D1 refinement (applied to a reduced-form game) and finds that every symmetric one offer equilibrium survives D1. The argument is based on the additional flexibility that the rational type enjoys and is similar to the argument presented here under which the Nash pooling equilibrium is refined for the reduced-form, leader-follower game.

Dai (2017) considers a model of pricing and search in which the seller is privately informed as to whether it is a commitment or non-commitment type, where the final price of the former (latter) type must equal (is unconstrained by) its advertised price. Consumers observe the seller's advertised price but are uncertain of the final price when deciding whether to incur a search cost and visit the seller. For this application, Dai fully characterizes the sets of pure-strategy separating and pooling equilibria. The current paper independently offers some related equilibrium characterizations and has two key distinguishing features. First, we analyze a strategic setting motivated by the standard quantity game and then proceed to consider a general-payoff setting that includes a range of applications. Second, for these settings, we identify a single-crossing property and thereby establish our main finding that all refined equilibria in the reduced-form game with private resolve generate Nash payoffs.[14]

The paper is also related to signaling models of electoral competition. Banks (1990) considers a model where the ideal policies of candidates are private information, winning candidates implement their ideal policies, and voters use each candidate's announced policy position as a signal of that candidate's ideal policy. Each candidate prefers to be believed to have the ideal policy of the median voter, but the winning candidate faces a lying or announcement cost when a discrepancy arises between the candidate's announced and ideal (implemented) policies.[15] Callander and Wilkie (2005) extend Banks' model to allow that the candidates also have private information about their costs of lying, where the announced position for a candidate with no such cost amounts to cheap talk. Kartik and McAfee (2007) explore related themes in a model in which a fraction of (non-strategic) candidates are exogenously committed to a campaign platform and are preferred by voters. In Callander's (2008) model, a candidate's campaign policy position is endogenously selected and commits the candidate to a post-election policy; in addition, the winning candidate chooses an effort level in the post-election phase, where the level of effort is related to the candidate's motivation type.

The model developed here differs in basic ways from those in the signaling literature

[14] Dai focuses on characterizing the full set of pure-strategy equilibria for her application, but she does argue that the Cho-Kreps (1987) intuitive criterion eliminates a subset of pooling equilibria (those located along one boundary of the equilibrium set) in an extended model with ex ante heterogeneous consumers.

[15] When lying costs are sufficiently great, the refined equilibria are such that moderate (extreme) types pool (separate). Bernheim (1994) develops a related characterization for a model of social conformity.

on electoral competition and features an outcome, the Nash pooling equilibrium outcome, which does not have a clear analog in the findings of that literature.[16] Even so, the model can be interpreted as also examining the relationship between promises and final actions, where under this interpretation the resolute leader is a promise-keeper, the flexible leader is unconstrained by past promises, and the leader is privately informed as to whether it is resolute or flexible. From this perspective, the model is broadly similar to those in Callander and Wilkie (2005) and Kartik and McAfee (2007) in that the flexible leader engages in cheap talk and is free to select a distinct final action; however, the preferred final action for the flexible leader here is endogenously determined in equilibrium as a best response to the behavior of the follower rather than being given exogenously (Banks, 2000; Callander and Wilkie, 2005; and Kartik and McAfee, 2007) or determined from a single-agent optimization problem given characteristics (Callander, 2008). As a consequence, the reduced-form payoffs in our model exhibit important, distinct properties.[17]

The paper is also broadly related to an oligopoly literature that endogenously determines the order of moves among firms. Hamilton and Slutsky (1990) provide an early contribution of this kind. Using the D1 refinement, Mailath (1993) explores the endogenous order of moves in a quantity-game setting with linear demand and costs when one firm is privately informed about market demand. The informed player chooses whether to move first or simultaneously, and this choice is an additional signal about the informed player's information. A distinguishing feature of the model considered here is that the follower does not directly observe whether the leader's "real" move is before or simultaneous with that of the follower. The order of moves in this sense is itself the source of private information, and the value of commitment is examined from this perspective.

The paper is related as well to papers in which players publicly select actions while aware of the possibility that they may subsequently become committed to those actions. In the literature on bargaining, Crawford (1983), Kambe (1999) and Wolitsky (2011) develop models of this general nature. In revision games, as introduced by Kamada and Kandori (2011), players prepare strategies in advance and then have stochastic opportunities to revise their actions. Relative to these papers, a distinguishing feature of the current paper is that the leader takes an initial action or makes a promise about a future action after becoming privately informed as to whether it is committed to that action or promise.[18]

[16] Some basic differences are: the model considered here studies the strategic value of leader commitment separate from competitive behavior associated with multiple leaders, and it also includes settings (such as the quantity-game setting) in which the follower has a rich action space rather than a binary vote.

[17] In particular, when the leader action (the signal) is set at the Nash level, the follower is not interested in the leader's type, since the leader's final action then equals its "promised" Nash action whether the leader is resolute or flexible. Also, the leader's preference over beliefs depends on the level of the leader action and specifically on whether the leader action is above, below or equal to the Nash action.

[18] For comparison, we can consider the complete-information game in which the follower and leader are both initially uncertain of whether the leader will have the opportunity to modify the leader action. In

Finally, the paper is related to research on games with uncertainty or private information about timing. Kreps and Ramey (1987) illustrate the tension that can arise between sequential rationality and structurally consistent beliefs (Kreps and Wilson, 1982a) when players are uncertain about which player moves first. Kamada and Moroni (2017) study dynamic games in which the timing of players' moves is private information, players choose whether to incur a small cost and disclose their moves, and payoffs are determined by a component game once all moves are made. They characterize conditions leading to a unique perfect bayesian equilibrium outcome when the component games is a coordination game and when it is a game of opposing interests. Our focus and results are quite different. We explore a two-period set up in which the identity of the leader is commonly known, the leader is privately informed as to whether it is resolute (moves once) or flexible (moves twice), and the leader action is publicly observed and serves as a potential signal. We also show that traditional Stackelberg behavior is absent in the refined equilibrium.

3 Set up and general structure

In this section, we present some basic structural assumptions and characterize the Nash equilibrium of the benchmark simultaneous-move game. The assumptions are motivated by the Cournot game of quantity competition. We then define the game with private resolve. We next characterize the "period-2 equilibrium quantities" that are functions of the observed "leader quantity" and Firm 2's resulting belief. We provide existence and comparative-statics results for the period-2 equilibrium quantities. These properties are used in the equilibrium analysis of the leader quantity in subsequent sections.

3.1 Nash benchmark

We begin with the benchmark of a simultaneous-move game, in which Firms 1 and 2 simultaneously select their quantities, q_1 and q_2, to maximize their payoffs, $\pi_1(q_1, q_2)$ and $\pi_2(q_1, q_2)$, respectively. Firm i, $i = 1, 2$, chooses its quantity q_i from its quantity space $Q_i \equiv [0, \overline{q}_i) \subset \Re$, where further discussion of $\overline{q}_i > 0$ is provided below.[19]

Throughout, we assume that $\pi_1(q_1, q_2)$ and $\pi_2(q_1, q_2)$ are twice-continuously differentiable over $Q_1 \times Q_2$, where, for all $(q_1, q_2) \in Q_1 \times Q_2$,

$$\frac{\partial^2 \pi_1(q_1, q_2)}{\partial {q_1}^2} < 0 \text{ and } \frac{\partial^2 \pi_2(q_1, q_2)}{\partial {q_2}^2} < 0. \tag{1}$$

that game, the leader would select the leader action at a level different from its Nash action. Once we assume that the leader is privately informed about its resolve, however, the refinement as applied to the reduced-form game directs attention to the Nash outcome.

[19] An alternative approach would be to specify that $Q_i = \Re$ and to define the best-response function for each firm as the maximum of zero and the quantity that satisfies the corresponding first-order condition.

Motivated by the Cournot game, we assume further that, for all $(q_1, q_2) \in Q_1 \times Q_2$,

$$\frac{\partial^2 \pi_1(q_1, q_2)}{\partial q_1 \partial q_2} < 0 \text{ and } \frac{\partial^2 \pi_2(q_1, q_2)}{\partial q_1 \partial q_2} < 0. \tag{2}$$

We now define the best-response functions for the respective firms. Firm 1's *best-response function*, $q_1^{br}(q_2)$, is defined as the solution to the following first-order condition:

$$\frac{\partial \pi_1(q_1, q_2)}{\partial q_1} = 0 \text{ at } q_1 = q_1^{br}(q_2). \tag{3}$$

Similarly, Firm 2's *best-response function*, $q_2^{br}(q_1)$, is defined as the solution to the following first-order condition:

$$\frac{\partial \pi_2(q_1, q_2)}{\partial q_2} = 0 \text{ at } q_2 = q_2^{br}(q_1). \tag{4}$$

Observe that the second-order conditions hold by (1) and that the best-response functions are decreasing by (2), as in the standard case.

The *Nash output vector* (q_1^N, q_2^N) is defined as the simultaneous solution to (3) and (4): $q_1^N = q_1^{br}(q_2^N)$ and $q_2^N = q_2^{br}(q_1^N)$. For $i, j = 1, 2$ and $i \neq j$, we now define $\overline{q}_i$ by $q_j^{br}(\overline{q}_i) = 0$. Thus, for example, $\overline{q}_1$ is the output level for Firm 1 at which Firm 2's best-response output is zero. We assume that

$$q_1^{br}(0) < \overline{q}_1 \text{ and } q_2^{br}(0) < \overline{q}_2, \tag{5}$$

so that the monopoly output for any one firm is less than the upper bound of feasible outputs for that firm at which the best response for the other firm is zero. Given that the best-response functions are continuous, (5) ensures the existence of a Nash output vector, (q_1^N, q_2^N), satisfying $q_i^N \in (0, \overline{q}_i)$. To ensure that the Nash output vector is unique, we assume further that the reaction curves are stable; that is, at any (q_1, q_2) satisfying $q_1 = q_1^{br}(q_2)$ and $q_2 = q_2^{br}(q_1)$, we assume that

$$\frac{\partial q_1^{br}(q_2)}{\partial q_2} = -\frac{\frac{\partial^2 \pi_1(q_1,q_2)}{\partial q_1 \partial q_2}}{\frac{\partial^2 \pi_1(q_1,q_2)}{\partial {q_1}^2}} > -\frac{\frac{\partial^2 \pi_2(q_1,q_2)}{\partial {q_2}^2}}{\frac{\partial^2 \pi_2(q_1,q_2)}{\partial q_1 \partial q_2}} = \frac{1}{\frac{\partial q_2^{br}(q_1)}{\partial q_1}}. \tag{6}$$

In a graph with q_1 on the y axis and q_2 on the x axis, (6) ensures that at any point of intersection of the two best-response functions, $q_1^{br}(q_2)$ is flatter (less negative) than $q_2^{br}(q_1)$. Using (6), we may conclude that a unique Nash output vector (q_1^N, q_2^N) exists.

Summarizing, our assumptions generate continuous, decreasing best-response functions which intersect at a single point, (q_1^N, q_2^N), where $q_i^N \in (0, \overline{q}_i)$ for $i = 1, 2$. Equivalently, the simultaneous-move game admits a unique pure-strategy Nash equilibrium,

(q_1^N, q_2^N), where $q_i^N \in (0, \overline{q}_i)$ for $i = 1, 2$.[20] The best-response functions also exhibit a stability property, which ensures that $q_1^{br}(q_2)$ lies below (above) $q_2^{br}(q_1)$ for $q_1 > q_1^N$ ($q_1 < q_1^N$) for a graph with q_1 on the y axis and q_2 on the x axis. We maintain this structure throughout our analysis of this model.

Finally, having defined the Nash outputs, we may now also define the *Nash profits*:

$$\pi_1^N \equiv \pi_1(q_1^N, q_2^N) \text{ and } \pi_2^N \equiv \pi_2(q_1^N, q_2^N).$$

These payoffs represent the benchmark against which we will compare when considering commitment with private resolve.

3.2 The Game with Private Resolve

We now define *the game with private resolve.* This game is an incomplete-information game in which Nature selects whether Firm 1 is resolute or flexible, where the probability that Firm 1 has resolve is $r \in (0, 1)$. Firm 1 privately observes Nature's choice. Firm 1 then selects its "leader quantity," q_1^L, in the first period. Firm 2 observes q_1^L and forms a belief $b = b(q_1^L)$ as to the likelihood that Firm 1 is resolute. In the second period, the flexible Firm 1 and Firm 2 simultaneously select their final quantities, q_1 and q_2. The resolute Firm 1 has no choice in the second period (i.e., $q_1 \equiv q_1^L$ when Firm 1 is resolute), while the flexible Firm 1's choice of q_1 is unconstrained by its previous choice of q_1^L.

Formally, let $Q_i \equiv [0, \overline{q}_i)$ be the quantity space for Firm i, where $i = 1, 2$. Let Firm 1's type t be R, indicating resolute, or F, indicating flexible. Nature selects R with probability r. Firm 1's strategy for the game is then a leader quantity, $q_1^L : \{R, F\} \to Q_1$, and a final or "period-2 quantity," $q_1 : Q_1 \to Q_1$, for the flexible Firm 1 where $q_1(q_1^L)$ can depend on q_1^L but is not constrained by q_1^L. Firm 2's strategy is a final or "period-2 quantity," $q_2 : Q_1 \to Q_2$. Firm 2's belief function is $b : Q_1 \to [0, 1]$, where $b = b(q_1^L)$.

For a given leader quantity q_1^L selected by Firm 1, the respective payoffs to the resolute and flexible Firm 1 are given as

$$\pi_1(q_1^L, q_2(q_1^L)) \text{ and } \pi_1(q_1(q_1^L), q_2(q_1^L)),$$

where $q_1(q_1^L)$ and $q_2(q_1^L)$ are the resulting period-2 quantity choices for the flexible Firm 1 and Firm 2, respectively. Firm 2's expected payoff is then

$$b(q_1^L) \cdot \pi_2(q_1^L, q_2(q_1^L)) + (1 - b(q_1^L)) \cdot \pi_2(q_1(q_1^L), q_2(q_1^L))$$

[20] Observe also that, given the concavity of the profit functions as captured in (1), the simultaneous-move game does not admit a mixed-strategy equilibrium.

Notice that Firm 1's private information concerns its period-2 strategy space and not its payoff function.

3.3 Period-2 equilibrium quantities

For a given *leader quantity,* q_1^L, and a belief, b, we now represent and characterize the period-2 interaction between the firms. In the period-2 subform indexed by the pair (q_1^L, b), Firm 1's quantity is already determined as q_1^L if resolute, Firm 1 selects $q_1^*(q_1^L, b)$ if flexible, and Firm 2 selects $q_2^*(q_1^L, b)$, where we henceforth refer to $q_1^*(q_1^L, b)$ and $q_2^*(q_1^L, b)$ as the *period-2 equilibrium quantities.*

We refer to $q_1^*(q_1^L, b)$ and $q_2^*(q_1^L, b)$ as the period-2 equilibrium quantities, since they are the final quantities that would be chosen by the flexible Firm 1 and Firm 2 in the continuation of any sequential equilibrium (Kreps and Wilson, 1982a) of the game.[21] Formally, the period-2 equilibrium quantities satisfy mutual best-response conditions and are thus defined as the solution to a standard system of first-order conditions. Given q_1^L and b, $q_1^*(q_1^L, b)$ solves

$$\max_{q_1} \pi_1(q_1, q_2^*(q_1^L, b))$$

with first-order condition

$$\frac{\partial \pi_1(q_1, q_2^*(q_1^L, b))}{\partial q_1} = 0 \text{ at } q_1 = q_1^*(q_1^L, b). \tag{7}$$

Similarly, given q_1^L and b, $q_2^*(q_1^L, b)$ solves

$$\max_{q_2} \{b \cdot \pi_2(q_1^L, q_2) + (1-b) \cdot \pi_2(q_1^*(q_1^L, b), q_2)\}$$

with first-order condition

$$b \cdot \frac{\partial \pi_2(q_1^L, q_2)}{\partial q_2} + (1-b) \cdot \frac{\partial \pi_2(q_1^*(q_1^L, b), q_2)}{\partial q_2} = 0 \text{ at } q_2 = q_2^*(q_1^L, b). \tag{8}$$

We now make four observations. The first observation concerns the definition of the Nash output vector. Recall that the Nash output vector (q_1^N, q_2^N) is the simultaneous solution to (3) and (4): $q_1^N = q_1^{br}(q_2^N)$ and $q_2^N = q_2^{br}(q_1^N)$. Clearly, the Nash output vector equivalently can be defined by

$$(q_1^N, q_2^N) \equiv (q_1^*(q_1^L, 0), q_2^*(q_1^L, 0)),$$

[21] Given the concavity of the profit functions as captured in (1), the flexible Firm 1 and Firm 2 do not randomize in equilibirum with respect to their period-2 quantity choices.

and thus corresponds to the situation in which Firm 1 is flexible and Firm 2 believes that Firm 1 is flexible. Second, for any $b \in [0,1]$, the period-2 equilibrium quantity for the flexible Firm 1 is its best response to Firm 2's period-2 equilibrium quantity:

$$q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b)). \tag{9}$$

By contrast, the period-2 equilibrium quantity for Firm 2 need not be a best response to the flexible Firm 1's period-2 equilibrium quantity, since Firm 2 must allow as well that Firm 1 could be a resolute type whose output choice is q_1^L. Third, for the case in which $b = 1$ so that Firm 2 holds a belief that Firm 1 has resolve, the period-2 equilibrium quantities are given as

$$q_1^*(q_1^L, 1) = q_1^{br}(q_2^{br}(q_1^L)) \text{ and } q_2^*(q_1^L, 1) = q_2^{br}(q_1^L), \tag{10}$$

since in that case Firm 2 best responds to q_1^L and as (9) indicates the flexible Firm 1 best responds to Firm 2's choice.

Finally, our fourth observation is that, for all $(q_1^L, b) \in Q_1 \times [0,1]$, the period-2 equilibrium quantities, $(q_1^*(q_1^L, b), q_2^*(q_1^L, b))$, satisfying (7) and (8) exist and satisfy $q_1^*(q_1^L, b) \in (0, \overline{q}_i)$ for $i = 1, 2$. We have already covered the cases where $b \in \{0, 1\}$. To see the argument for $b \in (0,1)$, suppose $q_1^L > q_1^N$. Our approach is to consider different values for q_2 with q_1 then set as $q_1 = q_1^{br}(q_2)$ as suggested by (9). Start with $q_2 = q_2^N$ in which case $q_1 = q_1^N$. It is clear that $q_2 = q_2^N$ is too high to satisfy (8) when $q_1^L > q_1^N$ and $q_1^*(q_1^L, b)$ is replaced with $q_1 = q_1^N$. Next, consider the opposite extreme with $q_2 = q_2^{br}(q_1^L)$ in which case $q_1 = q_1^{br}(q_2)$. It is direct to confirm that q_1 so defined satisfies $q_1 \in (q_1^N, q_1^L)$. It is now clear that $q_2 = q_2^{br}(q_1^L)$ is too low to satisfy (8) when $q_1^L > q_1^N$ and $q_1^*(q_1^L, b)$ is replaced with $q_1 = q_1^{br}(q_2)$. We can now conclude that there exists $q_2 \in (q_2^{br}(q_1^L), q_2^N)$ with q_1 defined by $q_1 = q_1^{br}(q_2)$ such that $q_1 = q_1^*(q_1^L, b)$ and $q_2 = q_2^*(q_1^L, b)$ satisfy (7) and (8) and thus constitute period-2 equilibrium quantities. Other cases can be handled similarly.[22] We assume further that the period-2 equilibrium is unique and stable for all $(q_1^L, b) \in Q_1 \times [0,1]$.[23]

Building from these observations, we now provide a lemma that provides helpful structure for some of the analysis that follows.

Lemma 1 *Fix* $b \in (0,1)$.

If $q_1^L > q_1^N$, *then* $q_1^L > q_1^*(q_1^L, b) > q_1^N$ *and* $q_2^N > q_2^*(q_1^L, b) > q_2^{br}(q_1^L)$.

If $q_1^L < q_1^N$, *then* $q_1^L < q_1^*(q_1^L, b) < q_1^N$ *and* $q_2^N < q_2^*(q_1^L, b) < q_2^{br}(q_1^L)$.

The proof of Lemma 1 is in Appendix A.

[22] The case in which $q_1^L = q_1^N$ is addressed in further detail in Lemma 3.

[23] Stability ensures that the Jacobian determinant associated with the first-order conditions (7) and (8) is positive. We use this property in deriving the comparative statics properties below in (11) and (12).

As Figure 1 illustrates for a setting with $b \in (0, 1)$, a key point is that the period-2 equilibrium quantities are pulled away from the Nash output vector and in the direction suggested by the leader output. Hence, if $q_1^L > q_1^N$, then Firm 2 recognizes that a resolute Firm 1 will produce the high output q_1^L and is thus encouraged to lower its output choice. Firm 2's lower output choice in turn encourages the flexible Firm 1 to increase its output as a best response, which further encourages a reduction in Firm 2's output. In the end, Firm 2's output choice, $q_2^*(q_1^L, b)$, lies between its best responses to the output of the resolute Firm 1, q_1^L, and the output of the flexible Firm 1, $q_1^*(q_1^L, b)$. With $q_1^*(q_1^L, b) > q_1^N$, it then follows that $q_2^N > q_2^*(q_1^L, b) > q_2^{br}(q_1^L)$, as Lemma 1 indicates. Notice also that a flexible Firm 1 selects its best-response output below that of a resolute Firm 1, $q_1^*(q_1^L, b) < q_1^L$, when $q_1^L > q_1^N$. The defining feature of the standard Stackelberg model - that the leader commits to a quantity above its best response to the follower's output - thus carries over to this setting as well when $q_1^L > q_1^N$.

We next extend Lemma 1 in a straightforward way to include the case in which $b = 1$ so that Firm 2 holds a belief that Firm 1 has resolve.

Lemma 2 *Fix $b = 1$.*
If $q_1^L > q_1^N$, then $q_1^L > q_1^*(q_1^L, 1) > q_1^N$ and $q_2^N > q_2^*(q_1^L, 1) = q_2^{br}(q_1^L)$.
If $q_1^L < q_1^N$, then $q_1^L < q_1^*(q_1^L, 1) < q_1^N$ and $q_2^N < q_2^*(q_1^L, 1) = q_2^{br}(q_1^L)$.

Lemma 2 is proved in Appendix A. Notice that Lemma 2 takes the same form as Lemma 1, with the exception that $q_2^*(q_1^L, 1) = q_2^{br}(q_1^L)$ when $b = 1$.

Lemmas 1 and 2 do not consider the case in which $q_1^L = q_1^N$. As we confirm in the following lemma, for any $b \in [0, 1]$, the Nash output vector is achieved as period-2 equilibrium quantities when $q_1^L = q_1^N$.

Lemma 3 *For any $b \in [0, 1]$, $(q_1^*(q_1^N, b), q_2^*(q_1^N, b)) = (q_1^N, q_2^N)$.*

Lemma 3 is proved in Appendix A. Intuitively, if Firm 1 selects its Nash output as its leader quantity, $q_1^N = q_1^L$, then the relevance of the distinction between a resolute and flexible Firm 1 is removed and the Nash outcome is thus assured. As we confirm below after defining our equilibrium concept, this finding provides a natural lower bound for Firm 1's equilibrium profit.

We now complete our characterization of the period-2 Nash quantities by providing some comparative statics. It is direct to confirm the following properties: for $b \in (0, 1)$,

$$\frac{\partial q_1^*(q_1^L, b)}{\partial q_1^L} > 0 > \frac{\partial q_2^*(q_1^L, b)}{\partial q_1^L}, \tag{11}$$

and

$$\begin{aligned}
\frac{\partial q_1^*(q_1^L,b)}{\partial b} &> 0 > \frac{\partial q_2^*(q_1^L,b)}{\partial b} \text{ if } q_1^L > q_1^*(q_1^L,b) \\
\frac{\partial q_1^*(q_1^L,b)}{\partial b} &= 0 = \frac{\partial q_2^*(q_1^L,b)}{\partial b} \text{ if } q_1^L = q_1^*(q_1^L,b) \\
\frac{\partial q_1^*(q_1^L,b)}{\partial b} &< 0 < \frac{\partial q_2^*(q_1^L,b)}{\partial b} \text{ if } q_1^L < q_1^*(q_1^L,b).
\end{aligned} \tag{12}$$

Thus, for given beliefs $b \in (0,1)$, an increase in q_1^L raises the expected output by Firm 1, leading Firm 2 to respond with a lower value of $q_2^*(q_1^L,b)$. This response in turn induces the flexible Firm 1 to expand its output, $q_1^*(q_1^L,b)$. Similarly, for a given q_1^L, an increase in $b \in (0,1)$ causes Firm 2 to weigh more (less) heavily q_1^L ($q_1^*(q_1^L,b)$) when forming an expectation of Firm 1's output, so that Firm 2 reduces (raises) its output $q_2^*(q_1^L,b)$ if $q_1^L > q_1^*(q_1^L,b)$ ($q_1^L < q_1^*(q_1^L,b)$). The flexible Firm 1 anticipates this reaction by Firm 2 and adjusts its output $q_1^*(q_1^L,b)$ in the opposite direction. Of course, if $q_1^L = q_1^*(q_1^L,b)$, then a change in b has no effect on $q_2^*(q_1^L,b)$ nor therefore on $q_1^*(q_1^L,b)$.

4 Equilibrium analysis of the reduced-form game

In this section, we embed our analysis of the period-2 equilibrium quantities directly into the payoff functions. We refer to the resulting game as a "reduced-form game," and we define an equilibrium for this game. We then characterize pooling and separating equilibria, and we show that the equilibrium value for the resolute Firm 1's leader quantity is never below its Nash output for the simultaneous-move game.[24]

4.1 The reduced-form game

In any sequential equilibrium of the game with private resolve, and for any $q_1^L \in Q_1$, the period-2 quantities of the flexible Firm 1 and Firm 2, respectively, satisfy $q_1(q_1^L) = q_1^*(q_1^L,b)$ and $q_2(q_1^L) = q_2^*(q_1^L,b)$ when Firm 2 forms the belief $b = b(q_1^L)$ after observing q_1^L. Building from this property, we now embed the final outputs - $q_1 = q_1^L$ for the resolute Firm 1, $q_1 = q_1^*(q_1^L,b)$ for the flexible Firm 1, and $q_2 = q_2^*(q_1^L,b)$ for Firm 2 - directly into the payoff structure of a reduced-form game. In the reduced-form game, a strategy for Firm 1 is then just a leader-quantity strategy defined by a mapping from the type space, $\{R,F\}$, to the output space, Q_1, while Firm 2 simply forms a belief.

Definition 1 *In the reduced-form game, Firm 1 observes its type t and selects its leader*

[24] As noted in the Introduction, some results in this section have counterparts in Dai's (2017) work, although for a distinct application.

quantity, q_1^L, which induces a belief $b = b(q_1^L)$ by Firm 2 and leads to the following payoffs:

$$\pi_1(q_1^L, q_2^*(q_1^L, b)) \text{ for the resolute Firm 1}$$
$$\pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b)) \text{ for the flexible Firm 1}$$
$$b \cdot \pi_2(q_1^L, q_2^*(q_1^L, b)) + (1-b) \cdot \pi_2(q_1^*(q_1^L, b), q_2^*(q_1^L, b)) \text{ for Firm 2}$$

Recall that $q_1^*(q_1^L, b)$ and $q_2^*(q_1^L, b)$ as used in this definition uniquely satisfy the system of first-order conditions in (7) and (8).[25]

4.2 Equilibrium concept and a lower bound on profits

We are now ready to define a sequential equilibrium for the reduced-form game.

Definition 2 *An equilibrium is a triplet $\{q_1^L(R), q_1^L(F), b(q_1^L)\}$ such that*

$$\begin{aligned}
q_1^L(R) &\in \arg\max_{q_1^L \in Q_1} \pi_1(q_1^L, q_2^*(q_1^L, b(q_1^L))) \\
q_1^L(F) &\in \arg\max_{q_1^L \in Q_1} \pi_1(q_1^*(q_1^L, b(q_1^L)), q_2^*(q_1^L, b(q_1^L))) \\
\textit{If } q_1^L(R) &= q_1^L(F), \ \textit{then } b(q_1^L(R)) = r \\
\textit{If } q_1^L(R) &\neq q_1^L(F), \ \textit{then } b(q_1^L(R)) = 1 > 0 = b(q_1^L(F))
\end{aligned}$$

Thus, whatever its type, Firm 1 selects the leader quantity that maximizes its profit given the belief function of Firm 2. In turn, Firm 2's beliefs must obey Bayes' rule whenever possible. Notice that we focus on pure-strategy equilibria.[26]

A *pooling equilibrium* occurs when $q_1^L(R) = q_1^L(F)$, and in this case the belief function must satisfy $b(q_1^L(R)) = r$. A *separating equilibrium* occurs when $q_1^L(R) \neq q_1^L(F)$, and in this case $b(q_1^L(R)) = 1 > 0 = b(q_1^L(F))$. For leader quantities that are off the equilibrium path (i.e., for $q_1^L \notin \{q_1^L(R), q_1^L(F)\}$), the belief function is unrestricted. As is familiar from signaling games, this freedom in specifying beliefs for off-the-equilibrium-path actions is a source of multiple equilibria.

For the reduced-form game, an *equilibrium outcome* is defined by the pair $\{q_1^L(R), q_1^L(F)\}$. The embedded period-2 equilibrium quantities can be easily recovered from this pair. If

[25] The game with private resolve is not a standard signaling game, since a flexible Firm 1 moves twice. By constructing the reduced-form game, we represent the interaction in a signaling-game format but with the special feature that Firm 2 has no response beyond the formation of its belief. We define equilibrium and the refinement (in Section 5) relative to the reduced-form game. Other work that studies reduced-form signaling games includes, for example, Bagwell (2007), Bernheim (1994), Bernheim and Severinov (2003) and Kartik and Frankel (2017). Reduced-form payoffs are also commonly used to study signaling in financial markets, where the market is the "receiver" and forms a belief (see, e.g., Ross 1977).

[26] As we discuss in Section 5.5, our key findings continue to hold when mixed strategies are allowed.

$q_1^L(R) = q_1^L(F)$, then the period-2 equilibrium quantities are $q_1^*(q_1^L(R), r)$ and $q_2^*(q_1^L(R), r)$; and if $q_1^L(R) \neq q_1^L(F)$, then the period-2 equilibrium quantities are $q_1^*(q_1^L(R), 1)$ and $q_2^*(q_1^L(R), 1)$ following $q_1^L = q_1^L(R)$ and likewise $q_1^*(q_1^L(F), 0)$ and $q_2^*(q_1^L(F), 0)$ following $q_1^L = q_1^L(F)$.

The *equilibrium payoffs* for the two types of Firm 1 are now given as

$$\begin{aligned} \Pi_1(R) &\equiv \pi_1(q_1^L(R), q_2^*(q_1^L(R), b(q_1^L(R))) \\ \Pi_1(F) &\equiv \pi_1(q_1^*(q_1^L(F), b(q_1^L(F))), q_2^*(q_1^L(F), b(q_1^L(F))) \end{aligned} \tag{13}$$

Thus, in a separating equilibrium, the equilibrium payoffs for the resolute and flexible types of Firm 1 are, respectively,

$$\pi_1(q_1^L(R), q_2^*(q_1^L(R), 1)) \text{ and } \pi_1(q_1^*(q_1^L(F), 0), q_2^*(q_1^L(F), 0)).$$

Similarly, in a pooling equilibrium, the equilibrium payoffs for the resolute and flexible types of Firm 1 are respectively given as

$$\pi_1(q_1^L(R), q_2^*(q_1^L(R), r)) \text{ and } \pi_1(q_1^*(q_1^L(F), r), q_2^*(q_1^L(F), r)),$$

where of course $q_1^L(R) = q_1^L(F)$ in a pooling equilibrium.

We may now state the following corollary to Lemma 3:

Corollary 1 *In any equilibrium,* $\Pi_1(R) \geq \pi_1^N$ *and* $\Pi_1(F) \geq \pi_1^N$.

Using Lemma 3, the simple observation in Corollary 1 is that, regardless of its type, Firm 1 always has the option of achieving its Nash payoff by setting a leadership quantity equal to its Nash output, $q_1^L = q_1^N$.

With Corollary 1 in hand, we next characterize equilibrium payoffs and behavior in pooling and separating equilibria.

4.3 Pooling equilibria

Our first result for pooling equilibria characterizes the relationship between $\Pi_1(F)$ and $\Pi_1(R)$, depending on whether pooling occurs at Firm 1's Nash output or at another leader quantity value.

Proposition 1 *Fix a pooling equilibrium. If* $q_1^L(R) \neq q_1^N$*, then* $\Pi_1(F) > \Pi_1(R)$*; and if* $q_1^L(R) = q_1^N$*, then* $\Pi_1(F) = \Pi_1(R)$.

The proof of Proposition 1 is in Appendix A.

Intuitively, even though $q_1^L(R) = q_1^L(F)$ and thus $q_2^*(q_1^L(R), r) = q_2^*(q_1^L(F), r)$ in a pooling equilibrium, a flexible Firm 1 does at least as well as a resolute Firm 1 in a pooling equilibrium, since the flexible Firm 1 can best respond against Firm 2's output. The only case in which the flexible Firm 1 does not do better than the resolute Firm 1 is when pooling occurs at Firm 1's Nash output, as then by Lemma 3 both types of Firm 1 best respond against the resulting Nash output of Firm 2.

We next show that a pooling equilibrium exists in which the both types of Firm 1 select the Nash output as the leader quantity.

Proposition 2 *There exists a pooling equilibrium in which* $q_1^L(R) = q_1^L(F) = q_1^N$.

The proof of Proposition 2 is in Appendix A.

The key idea here is that, if pooling occurs at the Nash output and if any deviant leader quantity elicits the belief that Firm 1 is flexible, then Firm 2's period-2 quantity is its Nash output, independent of the leader quantity that Firm 1 selects. Below, we refer to the described beliefs - where $b(q_1^L) = 0$ for any deviant q_1^L - as "punishing" beliefs, since Firm 2 then believes that an observed deviant leader quantity comes from a flexible Firm 1 and will have no direct bearing on Firm 1's final output selection.[27]

We define the *Nash pooling equilibrium* as the pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$ and $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$. Notice that the Nash pooling equilibrium uses the punishing belief specification (i.e., $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$). The Nash pooling equilibrium is thus the pooling equilibrium constructed in the proof of Proposition 2, and so Proposition 2 establishes the existence of the Nash pooling equilibrium. Of course, other belief specifications can be used to support a pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$. Allowing for alternative belief specifications, we define the *Nash pooling equilibrium outcome* by $\{q_1^L(R), q_1^L(F)\} = \{q_1^N, q_1^N\}$.

Let us now say that a pooling equilibrium outcome $\{q_1^L(R), q_1^L(F)\} = \{q_1^L, q_1^L\}$ is *r-robust* if a pooling equilibrium with $q_1^L(R) = q_1^L(F) = q_1^L$ exists for all $r \in (0, 1)$. As Proposition 2 confirms, the Nash pooling equilibrium exists for any value of $r \in (0, 1)$, and so the Nash pooling equilibrium outcome is thus r-robust. The following proposition establishes that no other pooling equilibrium outcome is r-robust.

Proposition 3 *The unique pooling equilibrium outcome that is r-robust is the Nash pooling equilibrium outcome.*

The proof of Proposition 3 is in Appendix A.

Intuitively, for any fixed $q_1^L \neq q_1^N$, a pooling equilibrium at q_1^L must fail to exist as r approaches zero, since a resolute Firm 1 is then sure to gain by deviating from $q_1^L(R) = q_1^L$

[27] As discussed further in the Introduction, Kim (2009) finds a related no-commitment equilibrium for his continuous-time, gap-case model of the durable-goods monopoly problem.

to q_1^N. The only pooling equilibrium outcome that is robust to arbitrarily low specifications for r is therefore the Nash pooling equilibrium outcome.

Our preceding discussion indicates that there is limited scope for the existence of a pooling equilibrium outcome that is robust to arbitrarily low values for r. While an r-robust pooling equilibrium outcome has attractive features, if we are interested in whether traditional Stackelberg arguments are robust to the introduction of a small degree of private information concerning Firm 1's resolve, then the fact that pooling equilibria associated with a given leader quantity fail to exist when Firm 1 is almost certain to lack resolve (i.e., when r is arbitrarily close to zero) may not be of great concern. For this question, we would be more interested to understand behavior as we move from the complete-information setting corresponding to $r = 1$ to a private-information setting in which r is close to unity.

Motivated by these considerations, we now construct a class of pooling equilibria that is parameterized by r. Let us define the *generalized Stackelberg quantity*, $q_1^{gs}(r)$, as the solution to the following maximization program:

$$\max_{q_1^L} \pi_1(q_1^L, q_2^*(q_1^L, r)).$$

The generalized Stackelberg solution corresponds to the traditional Stackelberg solution when $r = 1$ and to the Nash output, q_1^N, when $r = 0$. The first-order condition for the maximization problem is given as

$$\frac{\partial \pi_1(q_1^L, q_2^*(q_1^L, r))}{\partial q_1} + \frac{\partial \pi_1(q_1^L, q_2^*(q_1^L, r))}{\partial q_2} \cdot \frac{\partial q_2^*(q_1^L, r)}{\partial q_1^L} = 0 \text{ at } q_1^L = q_1^{gs}(r). \tag{14}$$

We now assume that the associated second-order condition holds with strict inequality for any $r \in (0, 1)$. Our previous assumptions suffice to ensure that the second-order condition holds as well in the limiting case of $r = 0$, and to ensure that the standard Stackelberg solution is well-defined we assume further that the second-order condition holds with strict inequality, too, for the limiting case where $r = 1$.

To characterize the generalized Stackelberg solution, we now assume further that a higher quantity by Firm 2 results in a negative externality for Firm 1's profit:

$$\frac{\partial \pi_1(q_1, q_2)}{\partial q_2} < 0. \tag{15}$$

This externality assumption is standard in models of Cournot competition, and captures there the simple idea that higher rival output lowers the market-clearing price. When we combine (3) and (15), we see further that

$$\frac{d\pi_1(q_1^{br}(q_2), q_2)}{dq_2} < 0, \tag{16}$$

which indicates that Firm 1's profit rises as q_2 is reduced and q_1 is adjusted along Firm 1's best-response function.

Returning to our characterization of the generalized Stackelberg solution, if we combine (15) with our earlier finding from (11) that $\frac{\partial q_2^*(q_1^L,r)}{\partial q_1^L} < 0$, then we see from (14) that $\frac{\partial \pi_1(q_1^L, q_2^*(q_1^L,r))}{\partial q_1} < 0$ at $q_1^L = q_1^{gs}(r)$. Using (7) with $b = r$, we conclude that

$$q_1^{gs}(r) > q_1^*(q_1^{gs}(r), r) = q_1^{br}(q_2^*(q_1^{gs}(r), r)), \tag{17}$$

where the equality in (17) uses (9). Using Lemmas 1 and 3, we conclude from (17) that

$$q_1^{gs}(r) > q_1^N. \tag{18}$$

Thus, the generalized Stackelberg solution exceeds the Nash quantity.

We can now state the following proposition:

Proposition 4 *There exists a pooling equilibrium in which* $q_1^L(R) = q_1^L(F) = q_1^{gs}(r)$.

The proof of Proposition 4 is in Appendix A.

Proposition 4 establishes that the generalized Stackelberg solution can be captured in the reduced-form game as a pooling equilibrium. Furthermore, as r rises toward unity, the corresponding equilibrium behavior and payoffs for the resolute Firm 1 and Firm 2 approach those achieved in the traditional Stackelberg solution of the complete-information game. From this perspective, the introduction of a slight degree of private information concerning Firm 1's resolve does not require a discontinuous change in behavior or payoffs. Below, we return to this issue when using a refined equilibrium concept for the reduced-form game. As we will see, this concept does not permit the "punishing" beliefs of $b(q_1^L) = 0$ for $q_1^L \neq q_1^{gs}(r)$ which are used in the proof of Proposition 4 to support the pooling equilibrium.

4.4 Separating equilibria

In a separating equilibrium, the flexible Firm 1 is revealed, and so $b(q_1^L(F)) = 0$. The Nash output vector thus obtains when Firm 1 is the flexible type: $(q_1^N, q_2^N) \equiv (q_1^*(q_1^L, 0), q_2^*(q_1^L, 0))$. The resolute type of Firm 1 is also revealed in a separating equilibrium, and so $b(q_1^L(R)) = 1$. When Firm 1 is the resolute type, the final quantities are thus $(q_1^L(R), q_2^*(q_1^L(R), 1)) =$

$(q_1^L(R), q_2^{br}(q_1^L(R)))$. Given these observations, in any separating equilibrium, Firm 1's payoffs are $\Pi_1(F) = \pi_1^N$ and $\Pi_1(R) = \pi_1(q_1^L(R), q_2^*(q_1^L(R), 1))$.

For a separating equilibrium to obtain, the flexible Firm 1 must prefer to select $q_1^L(F)$ and induce the belief $b(q_1^L(F)) = 0$ than to mimic $q_1^L(R)$ and induce the belief $b(q_1^L(F)) = 1$. The no-mimic constraint thus takes the form

$$\pi_1^N = \pi_1(q_1^N, q_2^N) = \pi_1(q_1^*(q_1^L(F), 0), q_2^*(q_1^L(F), 0)) \geq \pi_1(q_1^*(q_1^L(R), 1), q_2^*(q_1^L(R), 1)).$$

Since Firm 1 is best responding on both sides of the inequality, with $q_1^*(q_1^L(F), 0) = q_1^{br}(q_2^N)$ and $q_1^*(q_1^L(R), 1) = q_1^{br}(q_2^*(q_1^L(R), 1))$, we may use (16) to conclude that $q_2^N = q_2^*(q_1^L(F), 0) \leq q_2^*(q_1^L(R), 1)$. Next, recall that q_2^N is a best response for Firm 2 to q_1^N, and also that $q_2^*(q_1^L(R), 1)$ is a best response for Firm 2 to $q_1^L(R)$. Given that Firm 2's best-response function is decreasing, it thus follows from $q_2^N \leq q_2^*(q_1^L(R), 1)$ that the no-mimic constraint can hold only if $q_1^N \geq q_1^L(R)$.

Suppose then that $q_1^N > q_1^L(R)$. The resolute Firm 1's equilibrium payoff would then be below its Nash equilibrium payoff,

$$\Pi_1(R) = \pi_1(q_1^L(R), q_2^*(q_1^L(R), 1)) < \pi_1(q_1^N, q_2^N) = \pi_1^N,$$

since in each case Firm 2 best responds and under our assumptions Firm 1's Stackelberg payoff function, $\pi_1(q_1^L, q_2^*(q_1^L, 1))$, is strictly concave with $q_1^{gs}(1) > q_1^N$ following from (18). We then have a contradiction with Corollary 1. Thus, $q_1^L(R) = q_1^N$ in a separating equilibrium.

Given that a separating equilibrium is possible only if $q_1^L(R) = q_1^N$, the resolute Firm 1 must earn its Nash profit, π_1^N, in any separating equilibrium. Recalling that we argue above that $\Pi_1(F) = \pi_1^N$, we may summarize our findings to this point on separating equilibria with the following proposition:

Proposition 5 *In any separating equilibrium,* $q_1^L(R) = q_1^N$ *and* $\Pi_1(R) = \Pi_1(F) = \pi_1^N$.

Thus, in any separating equilibrium, we have that $q_1^L(R) = q_1^N \neq q_1^L(F)$, where the alternative leader quantity selected by the flexible Firm 1 is incidental. Whether Firm 1 is flexible or resolute, its final output choice is q_1^N, against which Firm 2 best responds by choosing q_2^N. For all practical purposes, the separating equilibria identified in Proposition 5 are thus equivalent to the Nash pooling equilibrium identified in Proposition 2.

4.5 A lower bound on the $q_1^L(R)$

Our results above take two forms. First, Corollary 1 considers the full class of equilibria and establishes a lower bound for $\Pi_1(R)$ and $\Pi_1(F)$. Second, we consider pooling and separating equilibria, respectively, and provide characterization results for each of these two classes of equilibria. Proposition 5 offers a general characterization for any separating equilibrium, and Proposition 1 establishes that $\Pi_1(F) \geq \Pi_1(R)$ in any pooling equilibrium. While we construct specific pooling equilibria in Propositions 2 and 4, we do not examine how the full set of pooling equilibria varies with r.

Against this backdrop, we now is to consider the full class of equilibria and establish a lower bound for the resolute Firm 1's equilibrium leader quantity, q_1^L. In particular, using Lemma 1 and (16), we establish that in equilibrium the resolute Firm 1 never chooses a leader quantity below its Nash output.

Proposition 6 *In any equilibrium, $q_1^L(R) \geq q_1^N$.*

The proof of Proposition 6 is in Appendix A. The proof complements Proposition 5 and shows that a resolute Firm 1 would deviate and set its leader quantity equal to q_1^N if a pooling equilibrium were posited in which $q_1^L(R) < q_1^N$.

5 Refined Equilibria

Having now offered equilibrium characterization results, we consider in this section the equilibria that survive as "refined" equilibria. To this end, we define a refinement for our reduced-form game that is motivated by the D1 refinement for signaling games.[28] We show that the Nash pooling equilibrium as constructed in Proposition 2 is refined. We then show further that no other pooling equilibrium outcome is refined. The end result is thus that refined equilibria exist, and in any refined equilibrium the final outputs are the Nash outputs, (q_1^N, q_2^N), with each firm thus earning its Nash profits. From this perspective, our results indicate a sense in which the strategic advantage of commitment may be lost when the leader has private information about its resolve. We also consider alternative refinement criteria, mixed-strategy equilibria, and the ways in which the payoff functions considered here differ from those in standard signaling models.

5.1 Definition

To define the refinement, we begin by considering the gain from a deviation from an equilibrium. For a deviation $q_1^L \notin \{q_1^L(R), q_1^L(F)\}$ and associated belief $b = b(q_1^L)$, the *gain from deviation* for Firm 1 of types R and F, respectively, is

$$\begin{aligned} \Delta^R(q_1^L, b) &\equiv \pi_1(q_1^L, q_2^*(q_1^L, b)) - \Pi_1(R) \\ \Delta^F(q_1^L, b) &\equiv \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b)) - \Pi_1(F) \end{aligned} \tag{19}$$

[28] See Cho and Kreps (1987) and Fudenberg and Tirole (1991, Chapter 11) for further discussion of the D1 refinement in signaling games.

where $\Pi_1(R)$ and $\Pi_1(F)$ are the respective payoffs to the resolute and flexible types of Firm 1 in the given equilibrium. As (19) indicates, in the reduced-form game, a deviant leader quantity q_1^L and associated belief b induces the period-2 equilibrium quantities associated with that leader quantity and belief, $q_1^*(q_1^L, b)$ and $q_2^*(q_1^L, b)$.

It is not immediately clear whether $\Delta^R(q_1^L, b)$ is larger or smaller than $\Delta^F(q_1^L, b)$. On the one hand, as confirmed in Propositions 1 and 5, the flexible Firm 1 earns at least as high of an equilibrium payoff as the resolute Firm 1. Specifically, for pooling equilibria, $\Pi_1(F) > \Pi_1(R)$ if $q_1^L(R) \neq q_1^N$, and $\Pi_1(F) = \Pi_1(R)$ if $q_1^L(R) = q_1^N$; and for separating equilibria, $\Pi_1(F) = \Pi_1(R)$. The key point is that the flexible Firm 1 has an advantage, since it can move away from its leader output and best respond against the induced period-2 equilibrium quantity from Firm 2. A similar logic also applies, however, when a deviant leader quantity is selected. Given a deviation $q_1^L \notin \{q_1^L(R), q_1^L(F)\}$ and an associated belief $b = b(q_1^L)$, if $q_1^*(q_1^L, b) \neq q_1^L$, then the flexible Firm 1 also enjoys a greater deviation payoff. We return to the relationship between $\Delta^R(q_1^L, b)$ and $\Delta^F(q_1^L, b)$ in greater detail below, since this relationship is fundamental to the refinement that we now define.

To that end, let us define the respective belief values for which the gain from deviation is positive or zero. For the flexible type of Firm 1, these sets are defined as:

$$\begin{aligned} D^F(q_1^L) &\equiv \{b \in [0,1] | \Delta^F(q_1^L, b) > 0\} \\ D_0^F(q_1^L) &\equiv \{b \in [0,1] | \Delta^F(q_1^L, b) = 0\} \end{aligned} \tag{20}$$

For the resolute type of Firm 1, we may similarly define

$$\begin{aligned} D^R(q_1^L) &\equiv \{b \in [0,1] | \Delta^R(q_1^L, b) > 0\} \\ D_0^R(q_1^L) &\equiv \{b \in [0,1] | \Delta^R(q_1^L, b) = 0\} \end{aligned} \tag{21}$$

Motivated by the D1 refinement for signaling games, we are now ready to define a refined equilibrium for the reduced-form game.

Definition 3 *A refined equilibrium is an equilibrium such that, for $q_1^L \notin \{q_1^L(R), q_1^L(F)\}$,*

$$\begin{aligned} &\text{if } D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L) \text{ and } D^R(q_1^L) \neq \varnothing, \text{ then } b(q_1^L) = 1; \text{ and} \\ &\text{if } D^R(q_1^L) \cup D_0^R(q_1^L) \sqsubseteq D^F(q_1^L) \text{ and } D^F(q_1^L) \neq \varnothing, \text{ then } b(q_1^L) = 0. \end{aligned} \tag{22}$$

To understand the idea behind the refinement, suppose that a deviant quantity $q_1^L \notin \{q_1^L(R), q_1^L(F)\}$ is observed that satisfies $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$ and $D^R(q_1^L) \neq \varnothing$. This means that the set of beliefs b for Firm 2 under which the resolute Firm 1 enjoys a gain from the deviation includes all of the beliefs under which the flexible Firm 1 weakly gains from the deviation and that there indeed exists some belief under which the resolute

Firm 1 would enjoy a gain from the deviation.[29] It then seems natural for Firm 2 to believe that such a deviation is more likely to come from the resolute Firm 1 than the flexible Firm 1. The refinement takes this logic to its limit and holds that the deviation is then infinitely more likely to have come from a resolute Firm 1. While the D1 refinement is commonly used in the signaling-game literature, the associated restriction on beliefs may seem strong. After developing our findings, we return to this issue in Section 5.3 and argue that our findings also hold when a milder restriction on beliefs is imposed.[30]

With our notion of a refined equilibrium now defined, we consider in turn two questions. First, is the Nash pooling equilibrium as constructed in Proposition 2 a refined equilibrium? Second, can any other pooling equilibrium outcome be supported by a refined equilibrium? In organizing our analysis in this way, we focus on pooling equilibria and ignore separating equilibria. In terms of predicted outcomes, however, this focus is inconsequential. We recall from Proposition 5 all separating equilibrium induce the same equilibrium final quantities as does the Nash pooling equilibrium. Thus, the distinction between separating equilibria and the Nash pooling equilibrium is in this respect only cosmetic. We return to this issue at the end of Section 5.4 and describe there the implications of our results for the full set of refined equilibria.

5.2 Is the Nash pooling equilibrium refined?

Consider then the Nash pooling equilibrium analyzed in Proposition 2 and defined by $q_1^L(R) = q_1^L(F) = q_1^N$, $b(q_1^N) = r$, and $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$. We establish now that the Nash pooling equilibrium is a refined equilibrium. The key point is that this equilibrium has a special feature: Firm 1 makes the same equilibrium payoff whether it is resolute or

[29] For standard signaling games, the D1 refinement defines the "D" and "D_0" sets with respect to the mixed-strategy best responses of the receiver for different possible beliefs. The use of mixed-strategy best responses may be replaced with best responses when the receiver's best response is a singleton. With our analysis of the reduced-form game, we embed the period-2 best responses of the flexible Firm 1 and Firm 2 into the payoffs of the reduced-form game and define the "D" and "D_0" sets in terms of belief sets. We then apply the proposed refinement to the reduced-form game. Recall that the period-2 best responses of Firms 1 and 2 are uniquely determined given the leader quantity and belief (see footnote 21).

[30] The definition of a refined equilibrium used here incorporates elements of backward and forward induction. Backward induction is reflected in the use of the sequential equilibrium concept, and a form of forward induction is captured in the belief restrictions in (22). Kohlberg and Mertens (1986) define stable sets of Nash equilibria for finite games and associate forward induction with the elimination of never weak best response strategies in that context. We do not formally explore the relationship between the refinement used here and stable sets of Nash equilibria, but we do show in the Supplementary Appendix that the belief restriction given in (22) can also be motivated by the following experiment. Consider the game with private resolve as defined in Section 3.2 and fix a sequential equilibrium outcome. Consider the set of all sequential equilibria for that game giving rise to this outcome, and suppose that there exists a deviant leader quantity that is not selected by any sequential equilibrium supporting this outcome. If we eliminate strategies that are never weak best responses to any sequential equilibrium in this set, then any sequential equilibrium of the resulting game must sastisfy the belief restriction given in (22).

flexible, $\Pi_1(F) = \Pi_1(R)$. Since the flexible Firm 1 has a greater benefit from deviating, a belief function that assigns deviations to the flexible type - i.e., which specifies $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$ - is sure to satisfy the additional condition (22) for a refined equilibrium.

Formally, consider any deviant leader quantity $q_1^L \neq q_1^N$. Clearly, if $b = 0$, then Firm 2's period-2 quantity is unchanged, $q_2^*(q_1^L, 0) = q_2^N$. Since q_1^N is a best response to q_2^N and $q_1^L \neq q_1^N$ is thus not a best response to q_2^N, the resolute Firm 1 is sure to lose from the deviation under this belief: $\Delta^R(q_1^L, 0) < 0$. By contrast, following a deviant leader quantity $q_1^L \neq q_1^N$ that induces a belief of $b = 0$, the flexible Firm 1 still sets its period-2 quantity at q_1^N. We thus have that $\Delta^F(q_1^L, 0) = 0$. We conclude that

$$b = 0 \notin D^R(q_1^L) \cup D_0^R(q_1^L) \text{ and } b = 0 \in D_0^F(q_1^L). \tag{23}$$

By (23), $b = 0 \in D^F(q_1^L) \cup D_0^F(q_1^L)$ and yet $b = 0 \notin D^R(q_1^L)$. This is enough to conclude that our specification of $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$ satisfies the refinement condition (22).

We thus may now report our finding with regard to our first question:

Proposition 7 *The Nash pooling equilibrium is a refined equilibrium.*

Before proceeding to our second question, we make four additional points. First, while our arguments above apply to any $q_1^L \neq q_1^N$, the specification of beliefs for $q_1^L < q_1^N$ is not significant, since neither the resolute nor flexible Firm 1 can gain from a deviation to such a leader quantity. Second, our arguments concerning the Nash pooling equilibrium and its existence as a refined equilibrium are quite general and do not require the additional assumptions made above (such as (15) and the conditions for the existence of a generalized Stackelberg solution) that we imposed to establish Proposition 4. Third, the plausibility of the belief specification is further enhanced by noting that, if we find $b > 0$ such that the resolute Firm 1 weakly gains from deviating to $q_1^L \neq q_1^N$, then it must be true that the flexible Firm 1 gains from the deviation.[31]

To establish this third point, let us consider any deviant leader quantity $q_1^L \neq q_1^N$ and suppose that we find $b > 0$ such that $\Delta^R(q_1^L, b) \geq 0$. It then follows that

$$\begin{aligned} 0 &\leq \Delta^R(q_1^L, b) = \pi_1(q_1^L, q_2^*(q_1^L, b)) - \Pi_1(R) \\ &= \pi_1(q_1^L, q_2^*(q_1^L, b)) - \Pi_1(F) \\ &< \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b)) - \Pi_1(F) = \Delta^F(q_1^L, b), \end{aligned}$$

where the second equality utilizes $\Pi_1(F) = \Pi_1(R)$ at the Nash pooling equilibrium and the strict inequality follows since $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b))$ by (9) and $q_1^L \neq q_1^*(q_1^L, b)$ by

[31] Indeed, the refinement thus *requires* that $b(q_1^L) = 0$ for $q_1^L \neq q_1^N$ if we can find $b > 0$ such that $\Delta^R(q_1^L, b) \geq 0$. The Nash pooling equilibrium obviously satisfies this requirement. In Proposition 9, we build on this point and establish that, in any refined Nash pooling equilibrium, $b(q_1^L) = 0$ for any $q_1^L > q_1^N$.

Lemmas 1 and 2. Thus, given $q_1^L \neq q_1^N$ and for $b > 0$, if $b \in D^R(q_1^L) \cup D_0^R(q_1^L)$, then $b \in D^F(q_1^L)$. Hence, if a belief $b > 0$ is such that the resolute Firm 1 weakly gains from the deviation to q_1^L, then the flexible Firm 1 is sure to gain from this deviation.

Fourth, it is interesting to contrast the result here with a common finding for signaling games that refinements direct attention to the least-cost separating equilibrium.[32] In the model analyzed here, starting at the Nash pooling equilibrium, the flexible Firm 1 gains from a deviation under a wider range of beliefs than does the resolute Firm 1. The resolute Firm 1 is thus unable starting at $q_1^L(R) = q_1^L(F) = q_1^N$ to invoke a single-crossing property and reveal itself under the refinement through a deviation.[33] The model considered here is also distinct from the standard signaling model in that the belief that the sender prefers to elicit varies with the level of the signal; specifically, both types of Firm 1 prefer to be perceived as a resolute type when $q_1^L > q_1^N$ and are indifferent about Firm 2's belief when $q_1^L = q_1^N$.[34] We develop these distinctions in more detail in Section 5.6.

5.3 Are any other pooling equilibrium outcomes refined?

We now turn to our second question and consider the existence of other refined pooling equilibrium outcomes. We know that other pooling equilibria may exist. For example, as Proposition 4 reveals, under additional assumptions, there exists a pooling equilibrium at the generalized Stackelberg leader output: $q_1^L(R) = q_1^L(F) = q_1^{gs}(r)$ with $b(q_1^{gs}(r)) = r$ and $b(q_1^L) = 0$ for all $q_1^L \neq q_1^{gs}(r)$. The question we now ask is whether any pooling equilibrium in which $q_1^L(R) = q_1^L(F) \neq q_1^N$ can be a refined equilibrium. We establish here that the answer is negative; in other words, the Nash pooling equilibrium outcome is the only refined pooling equilibrium outcome.

To begin, we assume the existence of a pooling equilibrium in which $q_1^L(R) = q_1^L(F) \neq q_1^N$. Next, we recall from Proposition 6 that $q_1^L(R) \geq q_1^N$. Thus, we may assume the existence of a pooling equilibrium in which $q_1^L(R) = q_1^L(F) > q_1^N$. By Lemmas 1 and 2, we thus have that the pooling leader quantity $q_1^L(R)$ satisfies $q_1^L(R) > q_1^*(q_1^L(R), b) > q_1^N$

[32] See Cho and Sobel (1990) for a general statement of sufficient conditions under which the D1 refinement for signaling games selects the least-cost separating equilibrium.

[33] When $q_1^L(R) = q_1^L(F) = q_1^N$, the flexible Firm 1 is indifferent about revealing itself through a deviation. It is significant that we consider here deviations from the Nash pooling equilibrium. In Section 5.3, we consider pooling equilibria in which $q_1^L(R) = q_1^L(F) > q_1^N$ and show that the resolute Firm 1 can then invoke a single-crossing property and reveal itself under the refinement through a deviation.

[34] Both types of Firm 1 prefer to be perceived as a flexible type when $q_1^L < q_1^N$, but as shown in Proposition 6 the resolute Firm 1 would never choose such a leader quantity in equilibrium. Bernheim and Sererinov (2003) consider a signaling model in which the sender's preferred belief varies with its type. The model considered here is different in that both types of the sender (i.e., Firm 1) agree on the ranking of different beliefs but the ranking itself depends on the magnitude of the signal (i.e., the sign of $q_1^L - q_1^N$).

for any $b > 0$. We may now use (11) and (12) to conclude that, for all $b \in (0, 1)$,

$$0 > \frac{\partial q_2^*(q_1^L(R), b)}{\partial q_1^L} \text{ and } 0 > \frac{\partial q_2^*(q_1^L(R), b)}{\partial b}, \tag{24}$$

where the relationships in this and the preceding sentence hold in particular when $b = r$.

Let us now consider a small deviation from $q_1^L(R)$; specifically, consider the deviant leader quantity $q_1^L = q_1^L(R) - \varepsilon$, where $\varepsilon > 0$ is a small number, ensuring that $q_1^L > q_1^N$ and thus that $q_1^L > q_1^*(q_1^L, b) > q_1^N$ for any $b > 0$ and that (24) continues to hold for all $b \in (0, 1)$ when $q_1^L(R)$ is replaced by q_1^L. Suppose that we define a corresponding belief value b' such that

$$q_2^*(q_1^L, b') = q_2^*(q_1^L(R), r), \tag{25}$$

where $b' \in (0, 1)$ is thus close to r and satisfies $b' > r$ by (24) when evaluated at $b = r$.

Observe first that, with the period-2 equilibrium quantity of Firm 2 fixed by (25), the flexible Firm 1 is indifferent about deviating to q_1^L when the belief b' is elicited: $\Delta^F(q_1^L, b') = 0$. This observation reflects the fact that the final output for the flexible Firm 1 is unconstrained by the leader quantity selected by this firm. We conclude that

$$b' \in D_0^F(q_1^L). \tag{26}$$

As we establish just below, however, the resolute Firm 1 gains from deviating to q_1^L when the belief b' is elicited: $\Delta^R(q_1^L, b') > 0$. We thus have that

$$b' \in D^R(q_1^L). \tag{27}$$

To understand why the resolute Firm 1 gains from deviating to q_1^L when the belief b' is elicited, observe that, at the fixed period-2 equilibrium quantity for Firm 2 given by (25), the leader quantity $q_1^L(R)$ is above Firm 1's best-response value:

$$q_1^L(R) > q_1^*(q_1^L(R), r) = q_1^{br}(q_2^*(q_1^L(R), r)) = q_1^{br}(q_2^*(q_1^L, b')).$$

Thus, the resolute Firm 1's profit rises with the deviation, since the deviation doesn't alter Firm 2's period-2 equilibrium quantity and yet enables the resolute Firm 1 to position its final output closer to its best-response value. Hence, the defining feature of the Stackelberg solution, namely, that the leader select an action that is not a best response, explains why the resolute Firm 1 gains from this deviation. We conclude that (27) indeed holds.

Holding the deviant leader quantity q_1^L fixed, we now characterize the set of beliefs b

such that $b \in D^F(q_1^L) \cup D_0^F(q_1^L)$. To this end, we observe that, for $b \in (0,1)$,

$$\begin{aligned}\frac{d\Delta^F(q_1^L,b)}{db} &= \frac{\partial \pi_1(q_1^*(q_1^L,b),q_2^*(q_1^L,b))}{\partial q_1} \cdot \frac{\partial q_1^*(q_1^L,b)}{\partial b} \\ &\quad + \frac{\partial \pi_1(q_1^*(q_1^L,b),q_2^*(q_1^L,b))}{\partial q_2} \cdot \frac{\partial q_2^*(q_1^L,b)}{\partial b} \\ &= \frac{\partial \pi_1(q_1^*(q_1^L,b),q_2^*(q_1^L,b))}{\partial q_2} \cdot \frac{\partial q_2^*(q_1^L,b)}{\partial b} > 0,\end{aligned} \tag{28}$$

where the first equality follows directly from (19), the second equality follows from the envelope theorem as captured here in (7), and the inequality follows from (15), $q_1^L = q_1^L(R) - \varepsilon$ and (24). Given $\Delta^F(q_1^L, b') = 0$, we may now conclude from (28) that

$$D^F(q_1^L) \cup D_0^F(q_1^L) = \{b | b \geq b'\}. \tag{29}$$

The key intuition is straightforward. Since by construction the flexible Firm 1 is indifferent to deviating to q_1^L when the belief b' is elicited that holds fixed Firm 2's period-2 equilibrium quantity, the flexible Firm 1 gains from deviating to q_1^L if and only if the elicited belief b induces a reduction in Firm 2's period-2 equilibrium quantity. Accordingly, the flexible Firm 1 weakly gains from deviation to q_1^L if and only if $b \geq b'$.

Let us next consider whether $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$. Recall from (27) that $b' \in D^R(q_1^L)$. To consider $b > b'$, we use (19) and find that, for any $b \in (0,1)$,

$$\frac{d\Delta^R(q_1^L,b)}{db} = \frac{\partial \pi_1(q_1^L,q_2^*(q_1^L,b))}{\partial q_2} \cdot \frac{\partial q_2^*(q_1^L,b)}{\partial b} > 0, \tag{30}$$

where the inequality follows from (15), $q_1^L = q_1^L(R) - \varepsilon$ and (24). Given $\Delta^R(q_1^L, b') > 0$, we may now conclude from (30) that

$$\{b | b \geq b'\} \subset D^R(q_1^L), \tag{31}$$

where the inclusion is strict since by continuity $b' - \eta \in D^R(q_1^L)$ for $\eta > 0$ sufficiently small. Referring to (29) and (31), we conclude that

$$D^F(q_1^L) \cup D_0^F(q_1^L) \subset D^R(q_1^L). \tag{32}$$

With (27) and (32) in hand, we now consider the implications of our analysis for refined beliefs. According to (22), our refinement requires that Firm 2 believes that the deviation to q_1^L is undertaken by the resolute Firm 1: $b(q_1^L) = 1$. Notice that this belief function is different than the punishing beliefs we used to support the pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^{gs}(r)$, as in that case we specified that Firm 2 believes that any deviation is undertaken by the flexible Firm 1. The key issue is now whether a pooling equilibrium other than the Nash pooling equilibrium can be enforced when $q_1^L = q_1^L(R) - \varepsilon$ generates the belief that the deviation is undertaken by the resolute Firm 1, $b(q_1^L) = 1$. But we already know from (31) that $1 \in D^R(q_1^L)$. Equivalently, we know that $\Delta^R(q_1^L, 1) = \pi_1(q_1^L, q_2^*(q_1^L, 1)) - \Pi_1(R) > 0$, which is to say that the resolute Firm 1 would gain from a deviation to q_1^L if the belief $b(q_1^L) = 1$ were elicited. Intuitively, the resolute Firm 1 gains in this case, since it is then able to position its final output closer to its best-response value *and* induce a reduction in Firm 2's period-2 equilibrium quantity. We conclude that the putative pooling equilibrium fails to be refined.

We summarize our argument with the following proposition:

Proposition 8 *The Nash pooling equilibrium outcome is the unique refined pooling equilibrium outcome.*

Since any separating equilibrium leads to the same final outputs as does the Nash pooling equilibrium, we have as well the following corollary:

Corollary 2 *Refined equilibria exist, and in any refined equilibrium the final outputs are the Nash outputs,* (q_1^N, q_2^N)*, with each firm thus earning its Nash profits,* $\Pi_1(R) = \Pi_1(F) = \pi_1^N$ *and* π_2^N*.*

We now make four additional points. First, our uniqueness results utilize additional structure in comparison to our finding in Proposition 7 that the Nash pooling equilibrium is refined. In particular, we use (15) in the proof of Proposition 8. Motivated in part by this observation, we consider in the next section a general-payoff structure and thereby identify the driving forces behind our results.

Second, the results above indicate that, starting at a pooling equilibrium in which $q_1^L(R) = q_1^L(F) > q_1^N$, a single-crossing property holds for the model in that a reduction in the leader quantity and increase in the belief that together leave the flexible Firm 1 indifferent must generate a gain for the resolute Firm 1. By contrast, and in line with the discussion above, when we start at the Nash pooling equilibrium where $q_1^L(R) = q_1^L(F) = q_1^N$, this single-crossing property fails to hold. Indeed, starting at the Nash pooling equilibrium, the flexible Firm 1 is indifferent to a deviant leader quantity if and only if the deviation leads to the belief $b = 0$, but the resolute Firm 1 loses from a deviation when the belief $b = 0$ is induced.

Third, the refinement requires only that the resolute Firm 1 gains whenever a flexible Firm 1 weakly gains. A different consideration might be to see which type of Firm 1 gains *more* from deviating. While this goes beyond the refinement that we employ, it is interesting to observe that under our assumptions the resolute Firm 1 also gains more from deviating to q_1^L for any $b \geq b'$. Formally, as confirmed in Appendix A, for $b \in (0,1)$ and given $q_1^L = q_1^L(R) - \varepsilon > q_1^N$, we have that

$$\frac{d\Delta^R(q_1^L, b)}{db} - \frac{d\Delta^F(q_1^L, b)}{db} > 0. \tag{33}$$

As shown in Figure 2, it follows for any $b \geq b'$ that the resolute Firm 1 gains more from deviating to q_1^L than does the flexible Firm 1. This perhaps provides some additional intuitive support for the specification that $b(q_1^L) = 1$, which in turns ensures that no pooling equilibrium outcome other than the Nash pooling equilibrium outcome is refined.

Fourth, we note that the refinement is in any case stronger than necessary to eliminate pooling equilibria with $q_1^L(R) = q_1^L(F) > q_1^N$. We establish this point in two ways. First, we maintain focus on the "local" deviation to $q_1^L = q_1^L(R) - \varepsilon$ and argue that the belief restriction that $b(q_1^L) = 1$ goes well beyond what is necessary to induce the resolute Firm 1 to deviate to q_1^L. Second, we briefly expand our discussion to consider beliefs at "non-local" deviations, and we argue that the single-crossing property described above ensures that all pooling equilibria with $q_1^L(R) = q_1^L(F) > q_1^N$ are eliminated under a belief restriction motivated by the intuitive criterion of Cho and Kreps (1987).[35]

We begin by considering alternative belief restrictions given the local deviation, $q_1^L = q_1^L(R) - \varepsilon$. As can be inferred from Figure 2, there exists some $b'' \in (0, b')$ such that the resolute Firm 1 is sure to undertake the deviation to q_1^L if $b(q_1^L) > b''$.[36] It is thus interesting to consider whether other refinements might be sufficient to ensure that $b(q_1^L) > b''$. The intuitive criterion does not deliver this belief restriction. Starting at a pooling equilibrium with $q_1^L(R) = q_1^L(F) > q_1^N$, both types of Firm 1 would gain from a slight downward deviation to q_1^L if Firm 2 were to have the belief $b(q_1^L) = 1$ and thus reduce its period-2 equilibrium quantity. Hence, the deviant leader quantity q_1^L is not equilibrium dominated in the reduced-form game for either type of Firm 1. Motivated by the divinity refinement of Banks and Sobel (1987) for signaling games, we might also consider a refinement under which $b(q_1^L) \geq r$ whenever such a deviation q_1^L from a pooling equilibrium is observed and (32) is satisfied. This refinement is sufficient to eliminate pooling equilibria at leader quantities that exceed the generalized-Stackelberg quantity, $q_1^{gs}(r)$, since the resolute Firm 1 then gains from a slight downward deviation to q_1^L even if $b(q_1^L) = r$.

What would a divinity-based refinement imply if we were to start at a pooling equilibrium with a leader quantity that lies at or below the generalized-Stackelberg quantity so that $q_1^L(R) = q_1^L(F) \in (q_1^N, q_1^{gs}(r)]$? Interestingly, for this case, the divinity-based refinement is not quite sufficient to ensure a deviation by the resolute Firm 1. For the

[35] We give primary focus to the D1-based refinement since the associated local deviations are easily explained and would be utilized in extended models with richer type spaces. Section 7 offers one example.

[36] For the given $q_1^L = q_1^L(R) - \varepsilon$, b'' is defined by $\Delta^R(q_1^L, b) = 0$. Given $q_1^L > q_1^N$, we can show that the pooling equilibrium exists only if $\Delta^R(q_1^L, 0) < 0$. It follows that $b'' > 0$.

deviation $q_1^L = q_1^L(R) - \varepsilon$ with $\varepsilon > 0$ and small, the key point is that in this case b'' satisfies $b'' \in (r, b')$, and so a belief $b(q_1^L)$ such that $b(q_1^L) \in [r, b'']$ would satisfy the divinty-based refinement without inducing a deviation. Figure 3 illustrates the subcase in which $q_1^L(R) = q_1^L(F) \in (q_1^N, q_1^{gs}(r))$. The indifference curve for the resolute Firm 1 is then negatively sloped through the point $(q_1^L(R), r)$ that corresponds to the equilibrium leader action and associated belief, and $b'' \in (r, b')$ thus clearly follows.[37]

At the same time, given that the deviation $q_1^L = q_1^L = q_1^L(R) - \varepsilon$ satisfies (32), a belief function of this kind might be challenged as giving insufficient attention to the resolute Firm 1's greater potential to gain from the deviation. Consider the following alternative. Given a pooling equilibrium $q_1^L(R) = q_1^L(F) > q_1^N$, suppose we require that there exists $\phi > 0$ such that, for any deviation q_1^L satisfying (32), $b(q_1^L) \geq r + \phi$.[38] The simple idea here is that, if a deviation from a pooling equilibrium is observed that generates a gain for the resolute Firm 1 whenever it generates a weak gain for the flexible Firm 1, then we might reasonably require that the belief function rise above the prior by a certain amount ϕ when the deviation is observed, where $\phi > 0$ is defined independently of the specific deviation but can be arbitrarily small. A pooling equilibrium would then fail to satisfy this "milder restriction" if, for all $\phi > 0$, the resolute or flexible Firm 1 can find a deviation under which it enjoys a gain from deviation. We can easily see from Figure 3 for the subcase $q_1^L(R) = q_1^L(F) \in (q_1^N, q_1^{gs}(r))$ that, given any $\phi > 0$, there will always exist an $\varepsilon > 0$ sufficiently small that the resolute Firm 1 gains from deviating to $q_1^L = q_1^L(R) - \varepsilon$ and obtaining the belief $b(q_1^L) \geq r + \phi$. A pooling equilibrium at $q_1^L(R) = q_1^L(F) = q_1^{gs}(r) > q_1^N$ would fail by a similar argument, and as noted above a pooling equilibrium with $q_1^L(R) = q_1^L(F) > q_1^{gs}(r)$ would fail even if $\phi = 0$. Thus, the milder restriction suffices to eliminate all pooling equilibria with $q_1^L(R) = q_1^L(F) > q_1^N$.[39]

We next examine non-local deviations and consider the implications of a belief restriction motivated by the intuitive criterion. To this end, we fix a pooling equilibrium with $q_1^L(R) = q_1^L(F) > q_1^N$ and define $\tilde{q}_1^L < q_1^L(R)$ as the leader quantity that satisfies

[37]The subcase in which $q_1^L(R) = q_1^L(F) = q_1^{gs}(r)$ is similar, except that the slope of the indifference curve for the resolute Firm 1 then has a slope of zero at the point $(q_1^L(R), r)$ and takes a negative (positive) slope for leader quantities just below (above) $q_1^L(R)$. We can also refer to Figure 3 to interpret the finding reported in the previous paragraph for the situation in which $q_1^L(R) = q_1^L(F) > q_1^{gs}(r)$. The indifference curve for the resolute Firm 1 would then take a positive slope through the point $(q_1^L(R), r)$, and as noted in the previous paragraph the resolute Firm 1 would then deviate to $q_1^L = q_1^L(R) - \varepsilon$ even if $b(q_1^L) = r$.

[38]Notice that (32) holds with strict inclusion, and so for simplicity we define the restriction described here in a similar way. Notice also that we allow that ϕ may take different values depending on the pooling equilibrium under consideration. Alternatively, we could have defined a value for ϕ that applies uniformly for any pooling equilibrium or for any pooling equilibrium associated with a given pooling equilibirum outcome. The point made here holds under any of these formalizations.

[39]We can also eliminate all pooling equilibria with $q_1^L(R) = q_1^L(F) > q_1^N$ under alternative mild restrictions; in particular, it is not necessary that the belief function is discontinuous. A similar argument holds if the belief function is continuous and rises from r in Figure 3 at a slope that exceeds the slope of the resolute Firm 1's indifference curve.

$q_2^*(q_1^L(R), r) = q_2^*(\tilde{q}_1^L, 1)$. As (10) indicates, $q_2^*(\tilde{q}_1^L, 1) = q_2^{br}(\tilde{q}_1^L)$; thus, $\tilde{q}_1^L$ is simply the quantity for Firm 1 to which $q_2^*(q_1^L(R), r)$ is a best response. With $q_2^*(q_1^L(R), r) < q_2^N$ ensured by Lemma 1, it follows that $\tilde{q}_1^L \in (q_1^N, q_1^L(R))$.

Consider now a deviation to the leader quantity $\tilde{q}_1^L - \varepsilon$ where $\varepsilon > 0$ is small so that $\tilde{q}_1^L - \varepsilon > q_1^N$. Recalling that the flexible Firm 1 is indifferent to a deviation that leaves Firm 2's period-2 equilibrium quantity unaltered, we have that $\Delta^F(\tilde{q}_1^L, 1) = 0$. From here, it is a simple matter to utilize preceding arguments and conclude that the deviation $\tilde{q}_1^L - \varepsilon$ is equilibrium dominated in the reduced-form game for the flexible Firm 1: $\Delta^F(\tilde{q}_1^L - \varepsilon, b) < 0$ for all $b \in [0, 1]$. By contrast, for the resolute Firm 1, the deviant leader quantity $\tilde{q}_1^L$ is not equilibrium dominated in the reduced-form game since $\Delta^R(\tilde{q}_1^L, 1) > 0$. In line with preceding arguments, the key point is that the deviation to $\tilde{q}_1^L$ in combination with the belief $b(\tilde{q}_1^L) = 1$ doesn't alter Firm 2's period-2 equilibrium quantity and yet enables the resolute Firm 1 to position its final output closer to its best-response value. By continuity, for sufficiently small $\varepsilon > 0$, the deviant leader quantity $\tilde{q}_1^L - \varepsilon$ is likewise not equilibrium dominated for the resolute Firm 1. Motivated by the intuitive criterion, we may thus impose $b(\tilde{q}_1^L - \varepsilon) = 1$, which induces the resolute Firm 1 to undertake the deviation. Hence, by considering non-local deviations, we can also use a belief restriction motivated by the intuitive criterion to eliminate pooling equilibria with $q_1^L(R) = q_1^L(F) > q_1^N$.

5.4 The set of refined equilibria

We now further characterize the set of refined equilibria as formally defined in Section 5.1. As shown in Proposition 8 and Corollary 2, in any refined equilibrium, the final outputs and profits are uniquely determined as Nash outputs and profits. For completeness, we address two remaining issues. First, does the refinement also require that the punishing belief specification be used to support the Nash pooling equilibrium outcome? Second, are separating equilibria also refined, and to what extent does the refinement require that separating equilibria employ the punishing belief specification?

We start with the first issue. Can the Nash pooling equilibrium outcome be supported by a refined equilibrium that specifies a belief function that differs from the punishing belief function? Building on arguments made above, we show in Appendix A that for any pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$, the equilibrium is refined only if $b(q_1^L) = 0$ for any $q_1^L > q_1^N$. Thus, a Nash pooling equilibrium is refined only if the punishing belief specification is used for $q_1^L > q_1^N$. The refinement has no bite, however, for deviations from the Nash pooling equilibrium quantity to $q_1^L < q_1^N$. The reason is that neither type of Firm 1 can gain from such a deviation. We thus conclude as follows:

Proposition 9 *In any refined Nash pooling equilibrium, $b(q_1^L) = 0$ for any $q_1^L > q_1^N$.*

We consider next separating equilibria. By Proposition 5, $q_1^L(R) = q_1^N$ and $\Pi_1(R) = \Pi_1(F) = \pi_1^N$ in any separating equilibrium. Consider any value $q_1^L(F)$ such that $q_1^L(F) \neq q_1^N$. Since the equilibrium profit for both types of Firm 1 is the same as in any pooling equilibrium with $q_1^L(R) = q_1^L(F) = q_1^N$, we can draw directly on our findings above and conclude that the separating equilibrium is refined if punishing beliefs are used, so that $b(q_1^L) = 0$ for all $q_1^L \notin \{q_1^L(R) = q_1^N, q_1^L(F)\}$. We can also draw on the proof of Proposition 9 and conclude that a separating equilibrium is refined only if $b(q_1^L) = 0$ for any $q_1^L > q_1^N$ and such that $q_1^L \neq q_1^L(F)$. Once again, the refinement imposes no conditions on the belief function for any $q_1^L < q_1^N$ and such that $q_1^L \neq q_1^L(F)$. We thus have the following result:

Proposition 10 *Refined separating equilibria exist; and in any refined separating equilibrium, $b(q_1^L) = 0$ for any $q_1^L > q_1^N$ such that $q_1^L \neq q_1^L(F)$.*

Summarizing, the set of refined equilibria is characterized by pooling equilibria in which $q_1^L(R) = q_1^L(F) = q_1^N$ with $b(q_1^L) = 0$ for any $q_1^L > q_1^N$ and by separating equilibrium with $q_1^L(R) = q_1^N$ and $b(q_1^L) = 0$ for any $q_1^L > q_1^N$ such that $q_1^L \neq q_1^L(F)$. The refinement leaves some wiggle room for cosmetic differences associated with beliefs for deviations for which no type of Firm 1 would gain and the flexible Firm 1's leader quantity under separation. But at a substantive level the refinement pins down the relevant economic activity: for any refined equilibrium, the final quantities are the Nash outputs, (q_1^N, q_2^N), with each firm thus earning its Nash profits, $\Pi_1(R) = \Pi_1(F) = \pi_1^N$ and π_2^N.

5.5 Mixed-strategy equilibria

We focus on refined pure-strategy equilibria. It is natural to wonder whether the results would change in some significant way were refined mixed-strategy equilibria considered.

In the Supplementary Appendix, we extend the analysis to characterize mixed-strategy sequential equilibria for the game with private resolve as defined in Section 3.2. We restrict attention to mixed strategies in which the distribution over actions is discrete and has finite support.[40] A mixed strategy for a player thus indicates the probability that the player will select a given action in the support of the mixed strategy. Given the concavity of the profit functions as captured in (1), the flexible Firm 1 and Firm 2 do not randomize in equilibrium with respect to their period-2 quantity choices; therefore, any randomization must concern the selection of Firm 1's leader quantities. Thus, we may again embed the period-2 equilibrium quantities into the payoff functions and analyze the (refined) equilibria of the reduced-form game. The definitions of equilibria and refined equilibria extend in straightforward fashion when mixed strategies are included.

[40] Cho and Kreps (1987) adopt a similar approach for their analysis of the Spence signaling model.

As shown in the Supplementary Appendix, in any refined mixed-strategy equilibrium, the resolute Firm 1 adopts a pure strategy and sets its leader quantity equal to its Nash output, q_1^N, with probability one.[41] In refined mixed-strategy equilibria, the flexible Firm 1 either separates with probability one with randomly determined leader quantities that differ from q_1^N, or pools with the resolute Firm 1 at q_1^N with some positive probability and separates with complementary probability with randomly determined leader quantities that differ from q_1^N. In any case, and whether Firm 1 is resolute or flexible, Firm 2 expects that Firm 1's final output is q_1^N and thus best responds with a final output of q_2^N, ensuring that both types of Firm 1 earn π_1^N while Firm 2 earns π_2^N. Our main findings as summarized in Corollary 2 thus continue to hold when mixed strategies are allowed.

5.6 Comparison with standard signaling games

In a standard Spence-style signaling model with two types, the application of refinements leads to the least-cost separating equilibrium outcome.[42] Such models often feature two assumptions. First, for any fixed value of the signal, each type of sender gains when the receiver increases the belief probability that is associated with a high type of sender. This gain derives from the induced change in the receiver's response action. Second, a global single-crossing property holds and ensures that the high type of sender is more willing to send higher signals than is the low type. Our analysis of refined equilibria for the reduced-form game, by contrast, directs attention to the Nash pooling equilibrium outcome. The different results suggest that the payoff functions for the reduced-form game must differ in some systematic ways from those considered in standard signaling models.

Figure 4 provides a convenient illustration with which to identify some key distinctions. Consider first the preferences that Firm 1 holds with respect to Firm 2's beliefs. As illustrated in Figure 4, for a given leader quantity, both types of Firm 1 agree as to whether they prefer higher beliefs to lower beliefs; however, a feature of the model considered here is that their preferences regarding beliefs depend on the level of the leader quantity (i.e.,

[41] The argument proceeds via several observations, One is that Firm 2's period-2 equilibrium quantity must be constant following any leader quantity that the flexible Firm 1 selects with positive probability in a mixed-strategy equilibrium. A key next step is to consider various cases and show that the flexible Firm 1 earns its Nash profit, π_1^N, in any refined mixed-strategy equilibrium. For example, to rule out a refined equilibrium in which the flexible Firm 1 selects more than one leader action with positive probability and earns an equilibrium profit above π_1^N, we argue that it would then be necessary that the resolute Firm 1 also selects these leader quantities with positive probability. But given the strict concavity of the resolute Firm 1's profit and that the period-2 equilibrium output of Firm 2 is constant for these leader quantities, there can be at most two leader quantities that the flexible Firm 1 selects with positive probability, where the higher of the two is above the best-response value. From here, it is straightforward to use the refinement as above. With the flexible Firm 1's payoff equal to π_1^N in any refined mixed-strategy equilibrium, it is not difficult to show that the resolute Firm 1 must set its leader output at q_1^N.

[42] See Cho and Sobel (1990) for general characterizations of refined equilibria in signaling games.

on the value of the "signal"). We have the following taxonomy: when $q_1^L > q_1^N$, both types of Firm 1 prefer that Firm 2 has a higher belief b; when $q_1^L < q_1^N$, both types of Firm 1 prefer that Firm 2 has a lower belief b; and when $q_1^L = q_1^N$, both types of Firm 1 are indifferent about Firm 2's belief b.

As Figure 4 illustrates, the nature of the single-crossing property also varies as different values for the leader quantity are considered. With respect to this property, if we start at a point (q_1^L, b) with $b > 0$, the taxonomy is as follows: when $q_1^L > q_1^N$, the indifference curve for the flexible Firm 1 is steeper, ensuring that the resolute Firm 1 can benefit from a slight reduction in the leader quantity under a wider set of beliefs than can the flexible Firm 1; when $q_1^L < q_1^N$, the indifference curve for the resolute Firm 1 is steeper, ensuring that the resolute Firm 1 can benefit from a slight increase in the leader quantity under a wider set of beliefs than can the flexible Firm 1; and when $q_1^L = q_1^N$, the resolute Firm 1 is unable to find *any* alternative leader quantity and associated belief under which it gains while the flexible Firm 1 is indifferent or loses.[43]

Figure 4 illustrates other features as well. For example, one key feature is that the flexible Firm 1 is indifferent over different values of q_1^L when $b = 0$. The generalized Stackelberg solution, $q_1^{gs}(b)$, is also illustrated in Figure 4. In line with the analysis above, this function is depicted as exceeding the Nash output, q_1^N, when $b > 0$. For simplicity, this function is further represented as being an increasing function.[44] Finally, we notice as well that the indifference curves for both types of Firm 1 become steeper as q_1^L moves toward q_1^N, since changes in beliefs lead to smaller adjustments in Firm 2's second-period equilibrium quantity when the difference between the leader quantity and the Nash quantity diminishes.

The takeaway from Figure 4 is that the payoff functions considered here are significantly different from those adopted in standard signaling models. It is thus perhaps not surprising that our analysis directs attention to a different form of refined equilibrium play. Motivated by this discussion, we proceed in the next section to examine general properties of payoff functions that are sufficient for the findings reported above.

6 General payoffs and applications

In this section, we represent Firm 1's reduced-form payoffs as general functions of q_1^L and b, with the goal of identifying conditions for general payoff functions that suffice for our findings. This generalization clarifies the driving forces in the analysis and facilitates

[43] The indifference curve for the resolute Firm 1 is steeper when when $q_1^L < q_1^N$, since an increase in q_1^L and b that maintains Firm 2's second-period equilibrium quantity (and that thus holds fixed the payoff of the flexible Firm 1) enables the resolute Firm 1 to enjoy the additional benefit of moving its leader quantity closer to its best-response quantity.

[44] As confirmed in Section 6.2.1, this function is increasing when demand and cost functions are linear.

applications. To anchor the discussion, we interpret our assumptions at various points with reference to the quantity-game setting explored above. We organize the analysis around two cases: the "Stackelberg-up" case in which as in the quantity-game setting the generalized Stackelberg action exceeds the Nash action, and the "Stackelberg-down" case in which the Nash action exceeds the generalized Stackelberg action. After establishing general results for both cases, we illustrate the results with specific applications.

6.1 General payoffs and sufficient conditions

We now let $\pi^R(q_1^L, b)$ and $\pi^F(q_1^L, b)$ denote general payoff functions for the resolute and flexible Firm 1's, respectively, when the leader choice q_1^L is made and the belief $b = b(q_1^L)$ is induced. For example, in the quantity-game setting explored above, $\pi^R(q_1^L, b) = \pi_1(q_1^L, q_2^*(q_1^L, b))$ and $\pi^F(q_1^L, b) = \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b))$; thus, for that setting, the general payoff functions suppress the different channels through which q_1^L and b affect profits. In order to include this strategic setting and others, our approach in this section is to impose general assumptions directly on the reduced-form payoff functions $\pi^R(q_1^L, b)$ and $\pi^F(q_1^L, b)$. We refer to this environment as the *general-payoff setting.*[45]

To begin, we assume that the choice set for Firm 1 is defined by $Q_1 \equiv [0, \overline{q}_1)$ where $\overline{q}_1 > 0$, and we assume that $\pi^R(q_1^L, b)$ and $\pi^F(q_1^L, b)$ are twice-continuously differentiable over $Q_1 \times [0,1]$. To this, we add this following baseline assumptions on the payoff functions:

Definition 4 *For the general-payoff setting, our baseline assumptions are:*

1. For all $q_1^L \in Q_1$, $\pi^R(q_1^L, 0)$ is strictly concave in q_1^L with a unique maximizer $q_1^N \in (0, \overline{q}_1)$ which delivers the profit level $\pi_1^N \equiv \pi^R(q_1^N, 0)$.

2. For all $b \in [0,1]$, $\pi^F(q_1^N, b) = \pi^R(q_1^N, b) = \pi_1^N$.

3. For all $b \in [0,1]$, and for all $q_1^L \in Q_1$ such that $q_1^L \neq q_1^N$, $\pi^F(q_1^L, b) > \pi^R(q_1^L, b)$.

4. For all $q_1^L \in Q_1$, $\pi^F(q_1^L, 0) = \pi_1^N$.

While these assumptions may appear abstract, they are easily interpreted in the quantity-game setting considered above. To fix ideas in the general-payoff setting, we draw on our preceding analysis of the quantity-game setting and briefly interpret the baseline assumptions in that context.

For the quantity-game setting, the first baseline assumption simply defines the Nash output for Firm 1, with $b = 0$ ensuring that Firm 1's leader quantity has no impact on Firm 2's behavior. In line with Lemma 3, the second baseline assumption indicates that,

[45] To facilitate comparison between the analysis of the general-payoff setting and the previous analysis of the quantity-game setting, we ensure that any notation used in common across the two settings can be understood as an extension from the quantity-game setting to the general-payoff setting.

regardless of beliefs or Firm 1's type, the Nash profit is delivered if the leader quantity is set at the Nash output. The third baseline assumption highlights the advantage of being flexible: provided that the leader quantity differs from the Nash output, the leader output is not a best-response to Firm 2's output, ensuring that the flexible Firm 1 does better. The fourth baseline assumption indicates that the leader quantity selection for a flexible Firm 1 is irrelevant when Firm 2 believes that it is facing a flexible Firm 1, since then the leader quantity is irrelevant to both firms.

For the general-payoff setting, the reduced-form game can be described as follows. Nature selects whether Firm 1 has resolve or is flexible, where Nature selects "resolve" with probability $r \in (0,1)$ and where Nature privately informs Firm 1 of its choice. After learning its type, Firm 1 then selects a leader action, q_1^L. Firm 2 observes q_1^L and forms a belief $b = b(q_1^L)$ as to the likelihood that Firm 1 has resolve. Firm 1's payoffs are then determined as $\pi^R(q_1^L, b)$ and $\pi^F(q_1^L, b)$. For a given application, any subsequent actions that Firms 1 and 2 may take are embedded into the payoff functions.

An equilibrium for the reduced-form game in the general-payoff setting is exactly analogous to that presented previously for the quantity-game setting.

Definition 5 *For the general-payoff setting, an equilibrium is a triplet* $\{q_1^L(R), q_1^L(F), b(q_1^L)\}$ *such that*

$$\begin{aligned}
q_1^L(R) &\in \arg\max_{q_1^L} \pi_1^R(q_1^L, b(q_1^L)) \\
q_1^L(F) &\in \arg\max_{q_1^L} \pi_1^F(q_1^L, b(q_1^L)) \\
\textit{If } q_1^L(R) &= q_1^L(F), \textit{ then } b(q_1^L(R)) = r \\
\textit{If } q_1^L(R) &\neq q_1^L(F), \textit{ then } b(q_1^L(R)) = 1 > 0 = b(q_1^L(F))
\end{aligned}$$

The definitions of a pooling equilibrium and a separating equilibrium are unchanged; likewise, we continue to define a Nash pooling equilibrium, the Nash pooling equilibrium outcome, and an r-robust pooling equilibrium outcome exactly as above.

To relate our findings for the general-payoff setting to the propositions presented above, we note that the equilibrium payoffs for the general-payoff setting, $\pi^R(q_1^L(R), b(q_1^L(R)))$ and $\pi^F(q_1^L(F), b(q_1^L(F)))$, can be understood as an extension of the notation used previously in (13) as follows:

$$\begin{aligned}
\pi^R(q_1^L(R), b(q_1^L(R))) &\equiv \Pi_1(R) = \pi_1(q_1^L(R), q_2^*(q_1^L(R), b(q_1^L(R))) \\
\pi^F(q_1^L(F), b(q_1^L(F))) &\equiv \Pi_1(F) = \pi_1(q_1^L(F), q_2^*(q_1^L(F), b(q_1^L(F)))
\end{aligned} \tag{34}$$

Using this notational extension, we can speak of whether a proposition that characterizes equilibrium payoffs for the quantity-game setting also holds in the general-payoff setting.

Our definition of a refined equilibrium extends directly to the general-payoff setting if we define

$$\begin{aligned} \Delta^R(q_1^L, b) &\equiv \pi^R(q_1^L, b) - \Pi_1(R) \\ \Delta^F(q_1^L, b) &\equiv \pi^F(q_1^L, b) - \Pi_1(F), \end{aligned} \tag{35}$$

where the definitions of $\Pi_1(R)$ and $\Pi_1(F)$ are already extended to the general-payoff setting in (34).

In Appendix B, we establish our first result for the general-payoff setting.

Proposition 11 *For the general-payoff setting and under the baseline assumptions, Corollary 1 and Propositions 1, 2, 3 and 7 all hold.*

Thus, in equilibrium, Firm 1 earns at least its Nash profit, π_1^N, regardless of its type; and in a pooling equilibrium, the flexible Firm 1 makes more than (the same as) the resolute Firm 1 if $q_1^L \neq q_1^N$ ($q_1^L = q_1^N$). Further, the baseline assumptions alone are sufficient to ensure the existence of the Nash pooling equilibrium, that the Nash pooling equilibrium outcome is the unique pooling equilibrium outcome that is r-robust, and that the Nash pooling equilibrium is a refined equilibrium.

To construct other equilibria and analyze separating equilibria, we utilize some additional structure. At this point, it is useful to distinguish between two general strategic environments. In the "Stackelberg-up" case, Firm 1's generalized Stackelberg choice exceeds its Nash choice, provided that $b > 0$. The quantity-game setting fits into this category. The second general strategic environment is then the "Stackelberg-down" case, in which Firm 1's generalized Stackelberg choice lies below its Nash choice, again provided that $b > 0$. We define additional assumptions for both cases, and within each set of additional assumptions we again define the generalized Stackelberg quantity with the notation $q_1^{gs}(b)$ for any $b \in [0, 1]$ where we now explicitly extend the domain to include the limiting cases where $b \in \{0, 1\}$.

We first define the additional assumptions that we impose on the general payoff functions for the Stackelberg-up case.

Definition 6 *For the general-payoff setting in the Stackelberg-up case, our additional assumptions are:*

1. *For all $q_1^L \in Q_1$ and $b \in (0, 1]$, $\pi^F(q_1^L, b)$ is increasing in q_1^L.*

2. *For all $q_1^L \in (q_1^N, \overline{q}_1)$ and $b \in [0, 1]$, $\pi^F(q_1^L, b)$ and $\pi^R(q_1^L, b)$ are increasing in b.*

3. *For all $q_1^L \in Q_1$ and $b \in (0, 1]$, $\pi^R(q_1^L, b)$ is strictly concave in q_1^L and maximized at $q_1^{gs}(b) \in (q_1^N, \overline{q}_1)$.*

These additional assumptions are also readily interpreted in terms of the quantity-game setting considered above. The first additional assumption captures the idea that the flexible Firm 1 gains from a higher leader output when $b > 0$, since Firm 2 then reduces its output.[46] The second additional assumption captures a related idea: provided that $q_1^L > q_1^N$, a higher value for b increases the weight that Firm 2 attaches to the possibility that Firm will follow through and produce q_1^L, and so both the flexible and resolute Firm 1's gain from the resulting decrease in Firm 2's output. Finally, the third additional assumption defines the Stackelberg-up case and ensures that the generalized Stackelberg leader output is well defined.

We next define the additional assumptions that we impose on the general payoff functions for the Stackelberg-down case.

Definition 7 *For the general-payoff setting in the Stackelberg-down case, our additional assumptions are:*

1. *For all $q_1^L \in Q_1$ and $b \in (0, 1]$, $\pi^F(q_1^L, b)$ is decreasing in q_1^L.*
2. *For all $q_1^L \in [0, q_1^N)$ and $b \in [0, 1]$, $\pi^F(q_1^L, b)$ and $\pi^R(q_1^L, b)$ are increasing in b.*
3. *For all $q_1^L \in Q_1$ and $b \in (0, 1]$, $\pi^R(q_1^L, b)$ is strictly concave in q_1^L and maximized at $q_1^{gs}(b) \in [0, q_1^N)$.*

The additional assumptions for the Stackelberg-down case are similar to those for the Stackelberg-up case, except that now Firm 1 gains when believed to have selected a lower leader action. In this case, Firm 1 thus gains from an increase in b when the leader action is positioned below the Nash action. The third additional assumption defines the Stackelbeg-down case and ensures that the Stackelberg-down leader action is well defined.

In Appendix B, we confirm the following result:

Proposition 12 *For the general-payoff setting in the Stackelberg-up case and under the baseline and additional assumptions, Propositions 4, 5 and 6 all hold.*

Thus, when the additional assumptions are added to the baseline assumptions, we are able to construct an equilibrium in which Firm 1 pools at the generalized Stackelberg action, $q_1^{gs}(r)$. In addition, we can show that $q_1^L(R) \geq q_1^N$ in any equilibrium, where $q_1^L(R) = q_1^N$ with both types of Firm 1 making π_1^N in any separating equilibrium.

We also show in Appendix B that the results extend naturally for the Stackelberg-down case as well.

---

[46] For the quantity-game setting, recall that we define $\overline{q}_1$ as the lowest Firm 1 quantity that induces a best-response output of zero from Firm 2. If $b = 1$ and $q_1^L = \overline{q}_1$, then a higher value for q_1^L would have no impact on Firm 2's quantity nor therefore on the flexible Firm 1's payoff. We thus find it convenient to define $Q_1 \equiv [0, \overline{q}_1)$.

Proposition 13 *For the general-payoff setting in the Stackelberg-down case and under the baseline and additional assumptions, Propositions 4 and 5 both hold. Proposition 6 now holds with a reversed inequality: in any equilibrium,* $q_1^L(R) \leq q_1^N$.

Our last step is to establish conditions for the general-payoff setting under which Proposition 8 holds so that the Nash pooling equilibrium outcome is the only refined pooling equilibrium outcome. The extra condition that we require is a single-crossing property that is captured in the following assumption:

Definition 8 *For the general-payoff setting, we assume that the following single-crossing property holds:*

1. For the Stackelberg-up case where $q_1^{gs}(b) \in (q_1^N, \overline{q}_1)$ *for all* $b \in (0,1]$, *suppose* $q_1^L(F) \in (q_1^N, \overline{q}_1)$ *and let* $q_1^L = q_1^L(F) - \varepsilon > q_1^N$ *for* $\varepsilon > 0$. *If* $b' \in (r,1]$ *exists such that* $\pi^F(q_1^L(F), r) = \pi^F(q_1^L, b')$, *then* $\pi^R(q_1^L, b') > \pi^R(q_1^L(F), r)$.

2. For the Stackelberg-down case where $q_1^{gs}(b) \in [0, q_1^N)$ *for all* $b \in (0,1]$, *suppose* $q_1^L(F) \in [0, q_1^N)$ *and let* $q_1^L = q_1^L(F) + \varepsilon < q_1^N$ *for* $\varepsilon > 0$. *If* $b' \in (r,1]$ *exists such that* $\pi^F(q_1^L(F), r) = \pi^F(q_1^L, b')$, *then* $\pi^R(q_1^L, b') > \pi^R(q_1^L(F), r)$.

To interpret the single-crossing property, consider first the Stackelberg-up case and recall the quantity-game setting. For that setting, if $b = r$ and the leader quantity exceeds the Nash output and is dropped slightly, then Firm 2's Nash quantity rises. To keep Firm 2's quantity fixed, an offsetting increase in the belief is required, with $b' > r$ defining the required increase. A flexible Firm 1 then is indifferent to the described change, but a resolute Firm 1 gains, since the resolute Firm 1 gains from lowering its leader quantity to a level closer to its best-response level. The single-crossing-property assumption captures these relationships in terms of the general payoff functions. The Stackelberg-down case can be interpreted similarly for a situation in which, by choosing $q_1^L(F) < q_1^N$, the resolute Firm 1 commits to an action below its best-response.

In Appendix B, we prove our final general proposition.

Proposition 14 *For the general-payoff setting in the Stackelberg-up and Stackelberg-down cases and under the baseline and corresponding additional and single-crossing-property assumptions, Proposition 8 holds.*

Thus, for both the Stackelberg-up and Stackelberg-down cases, when the corresponding assumptions identified above all hold, the Nash pooling equilibrium outcome is the unique refined pooling equilibrium outcome. We note that the additional and single-crossing assumptions take different forms for the Stackelberg-up and Stackelberg-down cases, as indicated above.

With these findings in hand, we may now extend Corollary 2 to the general-payoff setting as follows:

Corollary 3 *For the general-payoff setting in the Stackelberg-up and Stackelberg-down cases and under the baseline and corresponding additional and single-crossing-property assumptions, refined equilibria exist, and in any refined equilibrium* $q_1^L(R) = q_1^N$ *and thus* $\Pi_1(R) = \Pi_1(F) = \pi_1^N$.

Thus, for both the Stackelberg-up and Stackelberg-down cases, when the corresponding assumptions identified above all hold, the resolute Firm 1 sets its leader action equal to its Nash action, and as a consequence Firm 1 earns its Nash payoff whether it is resolute or flexible.

6.2 Applications

The general-payoff setting enables us to identify driving forces in the analysis but is abstract and explicitly models only the leader action of Firm 1. While the flexible Firm 1's payoffs are treated, the final action selected by the flexible Firm 1 is suppressed in the general-payoff setting. Likewise, any action by Firm 2 is also suppressed. The quantity-game setting analyzed previously provides one fully specified model with which to interpret the analysis of the Stackelberg-up case. We now show how the propositions for the general-payoff setting can be easily used for applications with quadratic payoffs. Quadratic payoffs are convenient since they enable us to derive closed-form solutions for $\pi^F(q_1^L, b)$ and $\pi^R(q_1^L, b)$, but of course the results apply more generally.

To begin, we return to the quantity-model setting and now impose linear demand and costs. As previously noted, this game fits in the Stackelberg-up case. Next, we consider a monetary-policy model and show that this model fits into the Stackelberg-down case. For each application, we use our results for the general-payoff setting and show that the refined equilibrium generates Nash payoffs, even when the prior probability is high that the first-mover has resolve.

6.2.1 Stackelberg-up: A simple quantity-game setting

Consider the quantity-game setting as analyzed in Sections 3-5, but assume now that demand and costs take linear forms so that the profit functions are quadratic. Specifically, assume that the market price P is determined as $P(X) = \alpha - \beta X$, where $X \equiv q_1 + q_2$ is the aggregate quantity in the market when Firm 1 produces q_1 units and Firm 2 produces q_2 units. Assume further that the constant unit cost for production is c_1 for Firm 1 and c_2 for Firm 2. For $i, j = 1, 2$ with $i \neq j$, we impose the parameter restrictions that $\alpha > 0$, $\beta > 0$ and $c_i \geq 0$ where $\alpha/\beta > c_i$, $\alpha - 2c_i + c_j > 0$ and $\alpha + 2c_1 - 3c_2 > 0$.

The profit functions then take the following form:

$$\begin{aligned}\pi_1(q_1, q_2) &= [\alpha - \beta(q_1 + q_2) - c_1]q_1 \\ \pi_2(q_1, q_2) &= [\alpha - \beta(q_1 + q_2) - c_2]q_2.\end{aligned}$$

The quantity space for Firm i is $Q_i = [0, \overline{q}_i)$ where $\overline{q}_i = (\alpha - c_j)/\beta$ for $i, j = 1, 2$ with $i \neq j$. The Nash output levels and Nash profit levels can be calculated as $q_i^N = (\alpha - 2c_i + c_j)/(3\beta)$ and $\pi_i^N = (\alpha - 2c_i + c_j)^2/(9\beta)$, where $q_i^N \in (0, \overline{q}_i)$ and $i, j = 1, 2$ with $i \neq j$.

For a given leader quantity $q_1^L \in Q_i$ and belief $b \in [0, 1]$, the period-2 equilibrium quantities are denoted as $q_1^*(q_1^L, b)$ and $q_2^*(q_1^L, b)$, where $q_1^*(q_1^L, b)$ maximizes $\pi_1(q_1, q_2^*(q_1^L, b))$ and where $q_2^*(q_1^L, b)$ maximizes $b \cdot \pi_2(q_1^L, q_2) + (1 - b) \cdot \pi_2(q_1^*(q_1^L, b), q_2)$. The solutions are unique and take the following form:

$$\begin{aligned}q_1^*(q_1^L, b) &= \frac{\alpha - 2c_1 + c_2 + b\beta q_1^L}{\beta(3 + b)} \\ q_2^*(q_1^L, b) &= \frac{\alpha - 2c_2 + c_1 + b(\alpha - c_1 - 2\beta q_1^L)}{\beta(3 + b)},\end{aligned}$$

where $q_i^*(q_1^L, b) \in (0, \overline{q}_i)$. The flexible Firm 1's payoff is thus $\pi^F(q_1^L, b) = \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b))$, while the resolute Firm 1's payoff is $\pi^R(q_1^L, b) = \pi_1(q_1^L, q_2^*(q_1^L, b))$. For the simple linear model considered here, these payoffs take the following form:

$$\begin{aligned}\pi^F(q_1^L, b) &= (\frac{\alpha - 2c_1 + c_2 + b\beta q_1^L}{3 + b})^2 \frac{1}{\beta} \\ \pi^R(q_1^L, b) &= (\frac{2(\alpha - 2c_1 + c_2) - \beta(3 - b)q_1^L}{3 + b})q_1^L.\end{aligned}$$

With the payoffs now defined, we can easily check whether the baseline, additional and single-crossing-property assumptions hold.

Straightforward calculations confirm that the baseline assumptions hold. We also find that the generalized Stackelberg solution takes the form $q_1^{gs}(b) = 3q_1^N/(3 - b)$, from which it follows that $q_1^{gs}(b) > q_1^N$ for $b > 0$. Thus, the model belongs to the Stackelberg-up case. It is direct to verify that the model satisfies the additional assumptions for the Stackelberg-up case. We note that $q_1^{gs}(b)$ is increasing and that $q_1^{gs}(1) < \overline{q}_1$, since $\alpha + 2c_1 - 3c_2 > 0$. Finally, as we show in Appendix C, it is direct to confirm that the single-crossing-property for the Stackelberg-up case holds as well. Referring to Corollary 3, we conclude that, for this standard linear model of duopolistic quantity competition, refined equilibria exist, and in any refined equilibrium $q_1^L(R) = q_1^N$ and thus $\Pi_1(R) = \Pi_1(F) = \pi_1^N$.

More generally, the Stackelberg-up case can be associated with strategic settings in which best-response functions are decreasing and a higher action generates a negative cross-firm externality, as in the setting just analyzed, and also in which best-response functions are increasing and a higher action generates a positive cross-firm externality. As discussed in the Introduction, other examples of the former setting include certain price-game settings among firms selling complementary goods and also settings in which cost-reducing investments are selected prior to various oligopoly interactions. An example of the latter kind of setting occurs when two firms occupy opposite endpoints of the Hotelling line and set prices for their horizontally differentiated products. For an appropriately specified model, we could illustrate the application of our results for Stackelberg-up settings of this kind as well.

6.2.2 Stackelberg-down: A simple monetary-policy game

We consider next a monetary-policy game, in which a government wishes to commit to zero inflation but is tempted to surprise the market with inflation that exceeds expectations. As discussed below, in the Nash equilibrium of the game without commitment, the government selects a positive rate of inflation; by contrast, if the government has the ability to make a commitment, then it commits to zero inflation and enjoys a higher payoff. We now describe the monetary-policy game and embed this application into our general-payoff setting for the reduced-form game.

Barro and Gordon (1983) and Backus and Driffill (1985) offer formulations of the monetary-policy game. They assume the government has the following utility function: $u_g(x, x^e) = -\frac{1}{2}\alpha x^2 + \beta(x - x^e)$, where x and x^e are actual and expected inflation and where the public chooses the expected inflation rate so as to maximize $u_p(x, x^e) = -(x - x^e)^2$. In other words, the public resists being fooled about inflation. In the simultaneous-move game, the government chooses x to maximize $u_g(x, x^e)$ at the same time that the public selects x^e to maximize $u_p(x, x^e)$. The Nash equilibrium of the simultaneous-move game has $x = \beta/\alpha = x^e$. If instead the government could credibly commit to x before the public forms its expectations, then the resulting Stackelberg solution entails $x = 0 = x^e$ and thus higher payoffs for the government.

To put this game into the general-payoff setting as developed above, we imagine a government that announces a rate of inflation, q_1^L, while being privately informed as to whether it has resolve or is flexible. For a government with resolve, the announcement is fully credible, and the inflation rate of q_1^L is delivered. By contrast, for a flexible government, a rate of inflation is selected as the public forms its expectation, and that rate is not constrained by the announced rate, q_1^L. The public's "choice" of a rational belief about the expected rate of inflation is captured in our framework by the requirement that the belief function, $b(q_1^L)$, is consistent with Bayes' rule along the equilibrium path and that the public correctly forecasts the equilibrium choice of a flexible government.

Formally, we translate the monetary-policy setting into our framework as follows:

$$\begin{aligned}
\pi^R(q_1^L, b) &= -\frac{1}{2}\alpha(q_1^L)^2 + \beta(q_1^L - Eq_1(q_1^L, b)) \\
\pi^F(q_1^L, b) &= -\frac{1}{2}\alpha(q_1^*(q_1^L, b))^2 + \beta(q_1^*(q_1^L, b) - Eq_1(q_1^L, b))
\end{aligned}$$

where $Eq_1(q_1^L, b) \equiv b \cdot q_1^L + (1-b) \cdot q_1^*(q_1^L, b)$ and

$$q_1^*(q_1^L, b) = \arg\max_{q_1 \in Q_1} -\frac{1}{2}\alpha(q_1)^2 + \beta(q_1 - Eq_1(q_1^L, b)),$$

with $Q_1 \equiv [0, \overline{q}_1)$ and where $\overline{q}_1 > \beta/\alpha$.

It is direct to confirm that $q_1^*(q_1^L, b) = \beta/\alpha$; thus, $q_1^*(q_1^L, b)$ here is independent of $Eq_1(q_1^L, b)$ and hence q_1^L and b. We thus define $q_1^N = \beta/\alpha$ and $\pi_1^N = -\alpha(q_1^N)^2/2$, where $q_1^N \in (0, \overline{q}_1)$. Using $Eq_1(q_1^L, b) = b \cdot q_1^L + (1-b) \cdot q_1^N$, we represent the payoffs as follows:

$$\begin{aligned}
\pi^R(q_1^L, b) &= -\frac{1}{2}\alpha(q_1^L)^2 + \beta(1-b)(q_1^L - q_1^N) \\
\pi^F(q_1^L, b) &= -\frac{1}{2}\alpha(q_1^N)^2 + \beta b(q_1^N - q_1^L).
\end{aligned}$$

With the payoffs now defined, we can easily check whether the baseline, additional and single-crossing-property assumptions hold.

Straightforward calculations confirm that the baseline assumptions hold. The generalized Stackelberg solution, $q_1^{gs}(b)$, maximizes $\pi^R(q_1^L, b)$ and is given as $q_1^{gs}(b) = \beta(1-b)/\alpha$, where clearly $q_1^{gs}(b) < q_1^N$ for $b > 0$. We note that $q_1^{gs}(b)$ is decreasing with $q_1^{gs}(1) = 0$ as anticipated. Thus, $q_1^{gs}(b) \in [0, q_1^N)$ for all $b \in (0, 1]$. Building from these observations, it is direct to show that the monetary-policy game satisfies the additional assumptions for the Stackelberg-down case. Finally, as shown in Appendix C, it is direct to confirm that the single-crossing-property for the Stackelberg-down case holds as well. Referring to Corollary 3, we conclude that, for the monetary-policy game, refined equilibria exist, and in any refined equilibrium $q_1^L(R) = q_1^N$ and thus $\Pi_1(R) = \Pi_1(F) = \pi_1^N$.

Returning to the duopoly context, we note that the Stackelberg-down case also can be associated with strategic settings in which best-response functions are decreasing and a higher action generates a positive cross-firm externality, and in which best-response functions are increasing and a higher action generates a negative cross-firm externality. In the Supplementary Appendix, we illustrate the former setting. We consider there a modified quantity-game setting in which each firm enjoys as well a positive externality when the aggregate industry quantity increases, where this externality is captured in a simple additive way and could correspond to the beneficial effects of higher industry output on future demand or costs or the regulatory environment. This setting has decreasing best-response functions and, when the externality associated with industry output is sufficiently great, is also characterized by a positive cross-firm externality. We show that this application fits into the Stackelberg-down case and that our results again apply.

7 Private information and the probability of resolve

In this section, we briefly consider an alternative model in which Firm 1 is privately informed about the *probability* that it will have resolve and thus be constrained to set is final quantity equal to its announced leader quantity, q_1^L. The novel feature captured by this extension is that, when Firm 1 selects its leader quantity, Firm 1 itself may be uncertain about whether it will be constrained by this selection.

The game begins when Nature selects the probability that Firm 1 will have resolve. Firm 1 is privately informed of this probability and selects its leader quantity, q_1^L. After making this selection, Firm 1 observes the realization of a signal, indicating whether Firm 1 is resolved or flexible, where the probabilities of each realization are determined by Firm 1's true type (i.e., the probability of resolve, as selected by Nature). As in our main model, a resolute Firm 1 must set its final output equal to its leader quantity, but a flexible Firm 1 is unconstrained and sets its period-2 equilibrium quantity as a best response to Firm 2's period-2 equilibrium quantity, which is chosen simultaneously. Firm 2 perfectly observes the leader quantity but does not observe the resolve probability or the signal realization.

Two cases arise, depending on whether the set of possible resolve probabilities selected by Nature includes zero. In the first case, zero is not included. To illustrate this case, we can consider a scenario in which Firm 1 is privately informed as to whether it has resolve with low or high probability, where the low probability is still positive. For this case, it is straightforward to see that the Nash pooling equilibrium outcome is no longer feasible: Firm 1 would gain by deviating from $q_1^L = q_1^N$ to at least a slightly higher leader quantity, since then Firm 2 reduces its period-2 equilibrium quantity somewhat, even if it believes that Firm 1 has resolve with low probability.

In the second case, while Firm 1 may have any of potentially many probability types, there is a positive probability that Firm 1 is privately informed that it has resolve with probability zero. For example, Nature may select a resolve probability ρ from a large set of possible values, where the prior probability that Nature selects $\rho = 0$ is positive but could be arbitrarily small. In this case, the Nash pooling equilibrium remains feasible, since Firm 2 can believe that $\rho = 0$, that is, that Firm 1 is flexible, upon observing any deviation. Furthermore, the Nash pooling equilibrium is also refined, since the flexible (i.e., $\rho = 0$) type of Firm 1 is the type that is always able to best respond against Firm 2 and thus that gains from the deviation under the largest set of possible beliefs.

8 Conclusion

This paper considers a model with a leader and follower in which the leader is privately informed about its resolve to follow through on a proposed course of action. The leader's initial or promised action, referred to as the leader action, may then play both commitment and signaling roles. After putting the game in a reduced form where the leader's payoffs are expressed as a function of the leader action and the follower's belief, we show that a Nash pooling equilibrium exists in which, whether the leader is resolute or flexible, the leader action is set equal to the Nash action that the leader would choose in a pure-strategy equilibrium of the associated simultaneous-move game. Motivated by the D1 refinement for signaling games, we define a refinement for the reduced-form game and show that the Nash pooling equilibrium outcome is the unique refined pooling equilibrium outcome. We show further that in any refined equilibrium the final actions are the Nash actions of the simultaneous-move game, with each player thus earning Nash payoffs. These findings hold even when the leader is almost certain to have resolve. The arguments are first developed in a strategic setting motivated by the Cournot quantity game. We then identify sufficient conditions on general reduced-form payoff functions and provide applications.

The findings here suggest that the strategic value of commitment may be difficult to achieve in a one-time interaction when the leader is privately informed about its resolve and there is even a small chance that the leader lacks resolve. At a broad level, this work thus may offer some additional support for the view that the strategic value of commitment is most readily achieved when reputational or other costs are associated with a failure to maintain an initial action or follow through on a promise. Under the reputational view, commitment power accrues over time to a privately informed leader when the leader has the opportunity to develop a reputation for maintaining initial actions and/or carrying through on promises. A related view is that legal or other costs may be associated with misleading initial actions of unkept promises.[47] Such costs could also enhance the leader's ability to realize the strategic benefits of commitment.

The analysis can be extended in a variety of directions. We discuss above an extension in which the leader is privately informed about the probability that it will have resolve. In addition, the analysis could be easily extended to allow the leader the choice of whether or not to reveal the leader action. The final quantities and payoffs associated with the Nash pooling equilibrium could then also be achieved in a "no-reveal pooling equilibrium" in which both types of leader choose not to reveal.[48] Finally, and more speculatively, the results reported here may suggest possible experimental studies.

[47] See Kartik (2009) for an analysis of strategic communication when lying is costly.

[48] Just as in the construction of the Nash pooling equilibrium, any deviation would then be associated with the flexible type, which ensures in turn that no deviation is appealing.

9 Appendix A

We prove here Lemmas 1, 2 and 3 and also Propositions 1, 2, 3, 4, 6 and 9.

Proof of Lemma 1. Suppose that $q_1^L > q_1^N$. Assume to the contrary that $q_2^*(q_1^L, b) \geq q_2^N$. Since $q_1^N = q_1^{br}(q_2^N)$, $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b))$ by (9), and best-response functions are decreasing, it then follows that $q_1^*(q_1^L, b) \leq q_1^N$. Given the stability of the best-response functions, we have that $q_2^{br}(q_1^*(q_1^L, b)) \leq q_2^*(q_1^L, b)$. Further, since best-response functions are decreasing and $q_1^*(q_1^L, b) \leq q_1^N < q_1^L$, we may conclude that $q_2^{br}(q_1^L) < q_2^{br}(q_1^N) \leq q_2^{br}(q_1^*(q_1^L, b)) \leq q_2^*(q_1^L, b)$. Thus, if Firm 2 were instead to select a period-2 quantity slightly lower than $q_2^*(q_1^L, b)$, then it would suffer at most a second-order loss in profit if Firm 1 is flexible and chooses $q_1^*(q_1^L, b)$ and it would enjoy a first-order gain in profit if Firm 1 is resolute and produces q_1^L. Given $b \in (0, 1)$, Firm 2 thus gains from the deviation, and we thus have a contradiction. We conclude that $q_2^*(q_1^L, b) < q_2^N$. Since $q_1^N = q_1^{br}(q_2^N)$, $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b))$ by (9), and best-response functions are decreasing, it then follows that $q_1^*(q_1^L, b) > q_1^N$.

Next, suppose that $q_1^L > q_1^N$, and assume to the contrary that $q_2^*(q_1^L, b) \leq q_2^{br}(q_1^L)$. Given $q_1^L > q_1^N$, it follows from the negative slope of the Firm-2 best-response function that $q_2^*(q_1^L, b) \leq q_2^{br}(q_1^L) < q_2^{br}(q_1^N) = q_2^N$. Using (9) and that Firm 1's best-response function is decreasing, we thus have that $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b)) > q_1^{br}(q_2^N) = q_1^N$. By the stability of the best-response functions, we thus know that $q_2^*(q_1^L, b) < q_2^{br}(q_1^*(q_1^L, b))$. Thus, if Firm 2 were instead to select a period-2 quantity slightly higher than $q_2^*(q_1^L, b)$, then it would suffer at most a second-order loss in profit if Firm 1 is resolute and produces q_1^L and and it would enjoy a first-order gain in profit if Firm 1 is flexible and chooses $q_1^*(q_1^L, b)$. Given $b \in (0, 1)$, Firm 2 thus gains from the deviation, and we thus have a contradiction. We conclude that $q_2^*(q_1^L, b) > q_2^{br}(q_1^L)$. Since Firm 2's best-response function is decreasing, we may define $q_1' \in (q_1^N, q_1^L)$ such that $q_2^*(q_1^L, b) = q_2^{br}(q_1')$. Using the stability of the best-response functions and $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b))$, we then have that $q_1^*(q_1^L, b) \in (q_1^N, q_1')$, and so we may conclude that $q_1^*(q_1^L, b) < q_1^L$.

The proof for the case in which $q_1^L < q_1^N$ is analogous.

Proof of Lemma 2. As observed in (10), when $b = 1$, $q_1^*(q_1^L, 1) = q_1^{br}(q_2^{br}(q_1^L))$ and $q_2^*(q_1^L, 1) = q_2^{br}(q_1^L)$. Thus, if $q_1^L > q_1^N$, then $q_2^*(q_1^L, 1) = q_2^{br}(q_1^L) < q_2^{br}(q_1^N) = q_2^N$, where we use that Firm 2's best-response function is decreasing. It follows that $q_1^*(q_1^L, 1) = q_1^{br}(q_2^{br}(q_1^L)) > q_1^{br}(q_2^N) = q_1^N$, where we use that Firm 1's best-response function is decreasing. Finally, $q_1^L > q_1^*(q_1^L, 1)$ follows from the stability of the best-response functions, given that $q_1^*(q_1^L, 1) = q_1^{br}(q_2^{br}(q_1^L))$ and $q_1^L > q_1^N$.

Proof of Lemma 3. Suppose that $q_1^L = q_1^N$. Assume to the contrary that $q_2^*(q_1^N, b) < q_2^N$. Since $q_1^N = q_1^{br}(q_2^N)$, $q_1^*(q_1^N, b) = q_1^{br}(q_2^*(q_1^N, b))$ by (9), and best-response functions are decreasing, it then follows that $q_1^*(q_1^N, b) > q_1^N$. Given the stability of the best-response functions, we have that $q_2^{br}(q_1^*(q_1^N, b)) > q_2^*(q_1^N, b)$. Further, since best-response functions are decreasing and $q_1^*(q_1^N, b) > q_1^N$, we may conclude that $q_2^{br}(q_1^N) > q_2^{br}(q_1^*(q_1^N, b)) > q_2^*(q_1^N, b)$. Thus, if Firm 2 were instead to select a period-2 quantity slightly higher than $q_2^*(q_1^N, b)$, then it would earn greater profit whether Firm 1 is flexible and chooses $q_1^*(q_1^N, b)$ or is resolute and produces q_1^N. We thus have a contradiction. Since the case in which $q_2^*(q_1^N, b) > q_2^N$ may be handled similarly, we may conclude that $q_2^*(q_1^N, b) = q_2^N$. Finally, with $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b))$ following from (9), we may conclude that $q_1^*(q_1^N, b) = q_1^N$.

Proof of Proposition 1. Consider first the possibility of a pooling equilibrium in which $q_1^L(R) = q_1^N$. In this case, we know from Lemma 3 that $\Pi_1(F) = \Pi_1(R) = \pi_1^N$. Consider second the possibility a pooling equilibrium in which $q_1^L(R) \neq q_1^N$. Firm 2 then selects $q_2^*(q_1^L(R), r)$. We argue now that $q_1^L(R) \neq q_1^*(q_1^L(R), r)$. Assume to the contrary that $q_1^L(R) = q_1^*(q_1^L(R), r)$. In this case, $q_2^*(q_1^L(R), r)$ must be a best response to $q_1^*(q_1^L(R), r)$. Since $q_1^*(q_1^L(R), r)$ is a best response to $q_2^*(q_1^L(R), r)$, it follows that $q_1^*(q_1^L(R), r) = q_1^N$, a contradiction. We conclude that $q_1^L(R) \neq q_1^*(q_1^L(R), r)$. Since $q_1^*(q_1^L(R), r)$ is a best response to $q_2^*(q_1^L(R), r)$, and $q_1^L(R)$ is thus not a best response to $q_2^*(q_1^L(R), r)$, we conclude that $\Pi_1(F) > \Pi_1(R)$.

Proof of Proposition 2. Suppose $q_1^L(R) = q_1^L(F) = q_1^N$. We then have that $b(q_1^N) = r$. Let us now specify that $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$. Thus, since $q_2^*(q_1^L, 0) = q_2^N$, Firm 2 selects the period-2 quantity q_2^N whether Firm 1 selects a leader quantity of q_1^N or $q_1^L \neq q_1^N$. A flexible Firm 1 thus enjoys the payoff π_1^N regardless and has no incentive to deviate. A resolute Firm 1 prefers a leader quantity of q_1^N to $q_1^L \neq q_1^N$, since given the beliefs it cannot alter Firm 2's period-2 quantity.

Proof of Proposition 3. We know from Proposition 2 that the Nash pooling equilibrium exists for any value of $r \in (0, 1)$ and that the Nash pooling equilibrium outcome is thus r-robust. Consider now any $q_1^L \neq q_1^N$ and hypothesize a pooling equilibrium such that $q_1^L(R) = q_1^L(F) = q_1^L$. With q_1^L fixed, let us now consider lower values for r. As r approaches 0, Firm 2's belief $b(q_1^L) = r$ likewise approaches 0; hence, the period-2 equilibrium quantity choices, $q_1^*(q_1^L, r)$ and $q_2^*(q_1^L, r)$, approach the Nash output vector, (q_1^N, q_2^N), as r approaches 0. For r sufficiently close to 0, it thus follows that the resolute Firm 1 obtains below-Nash profits by selecting q_1^L, which contradicts Corollary 1.

Proof of Proposition 4. Suppose $q_1^L(R) = q_1^L(F) = q_1^{gs}(r)$. We then have that $b(q_1^{gs}(r)) = r$. Let us now specify for any deviant leader quantity $q_1^L \neq q_1^{gs}(r)$ that $b(q_1^L) = 0$. We show first that the resolute Firm 1 gains more from deviating to q_1^N than by deviating to any $q_1^L < q_1^N$. To see the argument, recall that $q_2^N = q_2^*(q_1^L, 0)$ and observe for any $q_1^L < q_1^N$ that

$$\pi_1(q_1^L, q_2^*(q_1^L, 0)) = \pi_1(q_1^L, q_2^N) < \pi_1(q_1^N, q_2^N) = \pi_1(q_1^N, q_2^*(q_1^N, 0))$$

where the inequality follows since $q_1^L < q_1^N$. Thus, to confirm that the resolute Firm 1 cannot gain from a deviant leader quantity $q_1^L \neq q_1^{gs}(r)$, it is sufficient to consider $q_1^L \geq q_1^N$.

Next, suppose that the resolute Firm 1 deviates to $q_1^L = q_1^N$. By Lemma 3, we have that $(q_1^*(q_1^N, 0), q_2^*(q_1^N, 0)) = (q_1^N, q_2^N)$, and so the profit $\pi_1(q_1^N, q_2^N)$ is earned under deviation to $q_1^L = q_1^N$. This profit is also available in the generalized Stackelberg problem; however, as captured in (18), the solution instead is $q_1^{gs}(r) > q_1^N$. We thus conclude that the resolute Firm 1 also cannot gain by deviating to $q_1^L = q_1^N$.

Finally, suppose that the resolute Firm 1 selects $q_1^L > q_1^N$ where $q_1^L \neq q_1^{gs}(r)$. Since $q_2^*(q_1^L, 0) = q_2^N$, the resolute Firm 1 earns $\pi_1(q_1^L, q_2^*(q_1^L, 0))$ after deviating to q_1^L. We can verify that the resolute Firm 1 loses from a deviation to $q_1^L > q_1^N$ where $q_1^L \neq q_1^{gs}(r)$, since

$$\pi_1(q_1^{gs}(r), q_2^*(q_1^{gs}(r), r)) > \pi_1(q_1^L, q_2^*(q_1^L, r)) > \pi_1(q_1^L, q_2^*(q_1^L, 0)),$$

where the first inequality follows from the definition of $q_1^{gs}(r)$ and the associated second-order condition. To confirm the second inequality, we note that Lemma 1 ensures that $q_1^L > q_1^*(q_1^L, b)$ for any $b \in (0, 1)$. It thus follows from (12) that $q_2^*(q_1^L, b)$ rises as b falls from r to lower positive values. Correspondingly, (15) ensures that $\pi_1(q_1^L, q_2^*(q_1^L, b))$ falls as b falls from r to lower positive values. Finally, the inequality is maintained as b reaches 0 by continuity.

The remaining task is to show that the flexible Firm 1 also does not gain from a deviation. Given that $b(q_1^L) = 0$ for $q_1^L \neq q_1^{gs}(r)$, any deviation induces the output of $q_2^*(q_1^L, 0) = q_2^N$ for Firm 2. The flexible Firm 1 then best responds, earning thereby π_1^N from any deviation. The flexible Firm 1's equilibrium payoff exceeds that of the resolute Firm 1, since the flexible Firm 1 can best respond against $q_2^*(q_1^{gs}(r), r)$ and $q_1^{gs}(r) > q_1^N$ by (18). Since as shown above the resolute Firm 1 does not gain from deviating to $q_1^L = q_1^N$ and generating thereby the profit π_1^N, it follows now that the flexible Firm 1 does not gain from deviating to any $q_1^L \neq q_1^{gs}(r)$ under the given belief specification.

Proof of Proposition 6. Assume to the contrary that an equilibrium exists in which $q_1^L(R) < q_1^N$. We establish in Proposition 5 that $q_1^L(R) = q_1^N$ in any separating equilibrium. Suppose then that we have a pooling equilibrium. Using $b(q_1^L(R)) = r \in (0, 1)$ and Lemma 1, it then follows that $q_1^L(R) < q_1^*(q_1^L(R), r) < q_1^N$ and $q_2^N < q_2^*(q_1^L(R), r) < q_2^{br}(q_1^L(R))$.

The profit to the resolute Firm 1 thus satisfies

$$\begin{aligned}
\Pi_1(R) &= \pi_1(q_1^L(R), q_2^*(q_1^L(R), r)) < \pi_1(q_1^{br}(q_2^*(q_1^L(R), r)), q_2^*(q_1^L(R), r)) \\
&< \pi_1(q_1^{br}(q_2^N), q_2^N) = \pi_1(q_1^N, q_2^N) = \pi_1^N
\end{aligned}$$

where the first inequality follows since $q_1^L(R) < q_1^*(q_1^L(R), r) = q_1^{br}(q_2^*(q_1^L(R), r)$ by (9) and the second inequality follows directly from $q_2^N < q_2^*(q_1^L(R), r)$ and (16). We thus have a contradiction to Corollary 1.

Proof that (33) holds. Observe from (28) and (30) that

$$\frac{d\Delta^R(q_1^L, b)}{db} - \frac{d\Delta^F(q_1^L, b)}{db} = [\frac{\partial \pi_1(q_1^L, q_2^*(q_1^L, b))}{\partial q_2} - \frac{\partial \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b))}{\partial q_2}]\frac{\partial q_2^*(q_1^L, b)}{\partial b}$$

We know that $\frac{\partial q_2^*(q_1^L, b)}{\partial b} < 0$ for $b \in (0,1)$ by $q_1^L = q_1^L(R) - \varepsilon$ and (24). Next, we have that

$$\frac{\partial \pi_1(q_1^L, q_2^*(q_1^L, b))}{\partial q_2} - \frac{\partial \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b))}{\partial q_2} = \int_{q_1^*(q_1^L, b)}^{q_1^L} \frac{\partial^2 \pi_1(q_1, q_2^*(q_1^L, b))}{\partial q_1 \partial q_2} dq_1 < 0,$$

where the inequality follows from (2) and $q_1^L = q_1^L(R) - \varepsilon > q_1^*(q_1^L(R), b)$.

Proof of Proposition 9. Consider a pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$, $b(q_1^N) = r$, and, for some $q_1^L \neq q_1^N$, $b(q_1^L) > 0$. For this value of q_1^L, we now examine the gains from deviation for the different types of Firm 1 as different values for b are entertained. As argued in Section 5.2, for the belief $b = 0$, we have that $\Delta^R(q_1^L, 0) < 0 = \Delta^F(q_1^L, 0)$. We may thus conclude once again that (23) holds. Next, suppose that we find $b > 0$ such that $\Delta^R(q_1^L, b) \geq 0$. As shown in the inequality chain from Section 5.2, it then follows that

$$\begin{aligned}
0 &\leq \Delta^R(q_1^L, b) = \pi_1(q_1^L, q_2^*(q_1^L, b)) - \Pi_1(R) \\
&= \pi_1(q_1^L, q_2^*(q_1^L, b)) - \Pi_1(F) \\
&< \pi_1(q_1^*(q_1^L, b), q_2^*(q_1^L, b)) - \Pi_1(F) = \Delta^F(q_1^L, b),
\end{aligned}$$

where the second equality utilizes $\Pi_1(F) = \Pi_1(R)$ at a pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$ and the strict inequality follows since $q_1^*(q_1^L, b) = q_1^{br}(q_2^*(q_1^L, b))$ by (9) and $q_1^L \neq q_1^*(q_1^L, b)$ by Lemmas 1 and 2. Thus, given $q_1^L \neq q_1^N$, we conclude that

$$\text{For } b > 0, \text{ if } b \in D^R(q_1^L) \cup D_0^R(q_1^L), \text{ then } b \in D^F(q_1^L). \tag{36}$$

Given that $0 \notin D^R(q_1^L) \cup D_0^R(q_1^L)$ and (36), we can invoke the refinement once we complete a final step and show that $D^F(q_1^L)$ is non-empty.

To complete this final step, suppose that $q_1^L > q_1^N$. Consider $b = \varepsilon$ for $\varepsilon > 0$ and small. Referring to Lemma 1, we thus know that $q_1^L > q_1^*(q_1^L, \varepsilon) > q_1^N$ and $q_2^N > q_2^*(q_1^L, \varepsilon) > q_2^{br}(q_1^L)$. It follows that

$$\begin{aligned}
\Delta^F(q_1^L, \varepsilon) &= \pi_1(q_1^*(q_1^L, \varepsilon), q_2^*(q_1^L, \varepsilon)) - \Pi_1(F) \\
&= \pi_1(q_1^*(q_1^L, \varepsilon), q_2^*(q_1^L, \varepsilon)) - \pi_1(q_1^*(q_1^L, 0), q_2^*(q_1^L, 0)) \\
&> 0,
\end{aligned}$$

where the inequality follows from (16) since $q_1^*(q_1^L, \varepsilon)$ is a best response to $q_2^*(q_1^L, \varepsilon)$, $q_1^*(q_1^L, 0)$ is a best response to $q_2^*(q_1^L, 0)$, and $q_2^*(q_1^L, \varepsilon) < q_2^*(q_1^L, 0)$ under (12) using continuity and given $q_1^L > q_1^N$. Thus, for $q_1^L > q_1^N$,

$$b = \varepsilon \in D^F(q_1^L). \tag{37}$$

Given (36) and (37), we may now invoke (22) and draw the following conclusion: for any pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$, the equilibrium is refined only if $b(q_1^L) = 0$ for any $q_1^L > q_1^N$.[49]

Finally, we consider $q_1^L < q_1^N$. Such a deviation cannot generate a gain for either type of Firm 1, regardless of the specification of beliefs, so the specification of beliefs for such a deviant leader quantity is immaterial. To complete the proof, we note that, since $\Delta^F(q_1^L, 0) = 0 > \Delta^R(q_1^L, 0)$ and since $\Delta^R(q_1^L, b) < \Delta^F(q_1^L, b) < 0$ for $b > 0$, the sets $D^F(q_1^L)$ and $D^R(q_1^L)$ are empty when $q_1^L < q_1^N$. Referring to (22), we see that the refinement then has no bite. Hence, for any pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^N$, the refinement imposes no conditions on the belief function for $q_1^L < q_1^N$.

10 Appendix B

In this Appendix section, we prove Propositions 11, 12, 13 and 14.

Proof of Proposition 11. We begin with Corollary 1. Suppose to the contrary that $\Pi_1(t) < \pi_1^N$ for $t = R$ or F. Since part 2 of the baseline assumption gives that $\pi^F(q_1^N, b) = \pi^R(q_1^N, b) = \pi_1^N$, it then must be that Firm 1 of type t selects $q_1^L(t) \neq q_1^N$. It follows that Firm 1 of type t could gain by deviating to $q_1^L = q_1^N$ and thereby receiving the payoff π_1^N.

We consider next Proposition 1. For a pooling equilibrium, suppose first that $q_1^L(R) \neq q_1^N$. Using part 3 of the baseline assumptions, we then know that $\Pi_1(F) = \pi^F(q_1^L(R), r) > \pi^R(q_1^L(R), r) = \Pi_1(R)$. Suppose second that $q_1^L(R) = q_1^N$. We then know by part 2 of the baseline assumptions that $\Pi_1(F) = \pi^F(q_1^L(R), r) = \pi_1^N = \pi^R(q_1^L(R), r) = \Pi_1(R)$.

[49] Since $\Delta^R(q_1^L, 0) < 0$ and $\Delta^R(q_1^L, b)$ is continuous, we also have that $\Delta^R(q_1^L, \varepsilon) < 0$ for ε sufficiently small. Thus, for $q_1^L > q_1^N$, we can observe further that $b = \varepsilon \notin D^R(q_1^L) \cup D_0^R(q_1^L)$ and $b = \varepsilon \in D^F(q_1^L)$.

To confirm Proposition 2, we establish the existence of the Nash pooling equilibrium: $q_1^L(R) = q_1^L(F) = q_1^N$, $b(q_1^N) = r$ and $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$. Each type of Firm 1 then earns π_1^N in equilibrium by part 2 of the baseline assumptions. If the flexible Firm 1 were to deviate to $q_1^L \neq q_1^N$, then it would receive $\pi^F(q_1^L, 0) = \pi_1^N$ and thus experience no gain by part 4 of the baseline assumptions. If the resolute Firm 1 were to deviate, then it would receive $\pi^R(q_1^L, 0) < \pi_1^N$ and thus experience a loss since part 1 of the baseline assumptions ensures that $\pi^R(q_1^L, 0) < \pi^R(q_1^L, 0) = \pi_1^N$.

Consider now Proposition 3. We establish the existence of the Nash pooling equilibrium in the preceding paragraph, and our proof holds for any $r \in (0, 1)$. Thus, the Nash pooling equilibrium outcome is r-robust. Consider now any pooling equilibrium outcome such that $q_1^L(R) \neq q_1^N$. As r goes to 0, the payoffs in this pooling equilibrium for a resolute Firm 1 approach $\pi^R(q_1^L(R), 0) < \pi^R(q_1^L, 0) = \pi_1^N$, where the inequality again follows from part 1 of the baseline assumptions. By continuity of the payoff functions, for r sufficiently small, $\pi^R(q_1^L(R), r) < \pi_1^N$. Using part 2 of the baseline assumptions, it follows that the resolute Firm 1 would deviate to $q_1^L = q_1^N$ and thereby receive the payoff $\pi^R(q_1^N, b(q_1^N)) = \pi_1^N$.

Finally, consider Proposition 7. For the Nash pooling equilibrium, we may use part 2 of our baseline assumptions to confirm that the equilibrium payoffs are $\Pi_1(R) = \Pi_1(F) = \pi_1^N$. For any deviation $q_1^L \neq q_1^N$, this equilibrium specifies that $b(q_1^L) = 0$. Using part 1 of the baseline assumptions, we observe that $\pi^R(q_1^L, 0) < \pi^R(q_1^N, 0) = \pi_1^N$; furthermore, part 4 of the baseline assumptions gives $\pi^F(q_1^L, 0) = \pi_1^N$. We conclude that $b = 0 \notin D^R(q_1^L) \cup D_0^R(q_1^L)$ and $b = 0 \in D_0^F(q_1^L)$. We conclude that our specification of $b(q_1^L) = 0$ for all $q_1^L \neq q_1^N$ satisfies the refinement condition (22).

Proof of Proposition 12. To confirm Proposition 4, we specify a pooling equilibrium in which $q_1^L(R) = q_1^L(F) = q_1^{gs}(r)$, $b(q_1^{gs}(r)) = r$ and $b(q_1^L) = 0$ for all $q_1^L \neq q_1^{gs}(r)$. From part 3 of the additional assumptions, we know that $q_1^{gs}(r) > q_1^N$. Consider the resolute Firm 1. A deviation to any $q_1^L \leq q_1^N$ induces the belief $b(q_1^L) = 0$. By part 1 of the baseline assumptions, the best such deviation for the resolute Firm 1 is $q_1^L = q_1^N$, which delivers the payoff $\pi^R(q_1^N, 0) = \pi_1^N$. But using part 3 of the additional assumptions and part 2 of the baseline assumptions, we know $\pi^R(q_1^{gs}(r), r) > \pi^R(q_1^N, r) = \pi_1^N$. Next, a deviation to any $q_1^L > q_1^N$ with $q_1^L \neq q_1^{gs}(r)$ likewise induces the belief $b(q_1^L) = 0$. We then have that $\pi^R(q_1^{gs}(r), r) > \pi^R(q_1^L, r) > \pi^R(q_1^L, 0)$, where the first (second) inequality follows from part 3 (part 2) of the additional assumptions. Thus, the resolute Firm 1 loses from any deviation. Consider the flexible Firm 1. Using $q_1^{gs}(r) > q_1^N$ and part 3 of the baseline assumptions, we know that $\pi^F(q_1^{gs}(r), r) > \pi^R(q_1^{gs}(r), r)$. As just shown, $\pi^R(q_1^{gs}(r), r) > \pi^R(q_1^N, r) = \pi_1^N$. Thus, using part 4 of the baseline assumptions, $\pi^F(q_1^{gs}(r), r) > \pi_1^N = \pi^F(q_1^L, 0)$ for any $q_1^L \neq q_1^{gs}(r)$. Hence, the flexible Firm 1 loses from any deviation.

We consider next Proposition 5. Fix a separating equilibrium. We thus have $q_1^L(R) \neq$

$q_1^L(F)$ and $b(q_1^L(R)) = 1 > 0 = b(q_1^L(F))$. It follows from part 4 of the baseline assumptions that $\Pi_1(F) = \pi^F(q_1^L(F), 0) = \pi_1^N$. A separating equilibrium can exist only if the flexible Firm 1 does not gain from deviating to $q_1^L(R)$; thus, it must be that $\pi_1^N = \pi^F(q_1^L(F), 0) \geq \pi^F(q_1^L(R), 1)$. We next observe that $\pi^F(q_1^L(F), 0) = \pi^F(q_1^N, 0) = \pi^F(q_1^N, 1)$, where the first (second) equality follows from part 4 (part 2) of our baseline assumptions. It now follows that a separating equilibrium exists only if $\pi^F(q_1^N, 1) \geq \pi^F(q_1^L(R), 1)$. Using part 1 of the additional assumptions with $b = 1$, we thus have that $q_1^N \geq q_1^L(R)$. Suppose $q_1^N > q_1^L(R)$. Then $\Pi_1(R) = \pi^R(q_1^L(R), 1) < \pi^R(q_1^N, 1) = \pi^R(q_1^N, b(q_1^N)) = \pi_1^N$, where the inequality follows given $q_1^L(R) < q_1^N$ from part 3 of the additional assumptions with $b = 1$ and where the second and third equalities follow from part 2 of the baseline assumptions. It follows that the resolute Firm 1 would deviate to $q_1^L = q_1^N$. Thus, a separating equilibrium can exist only if $q_1^L(R) = q_1^N$. It follows from part 2 of the baseline assumptions that $\Pi_1(R) = \pi^R(q_1^L(R), 1) = \pi_1^N$.

Consider now Proposition 6. Assume to the contrary that an equilibrium exists in which $q_1^L(R) < q_1^N$. By Proposition 5, the equilibrium must be a pooling equilibrium. By part 3 of the baseline assumptions, we know $\pi^F(q_1^L(R), r) > \pi^R(q_1^L(R), r) = \Pi_1(R)$. We also know from part 1 of the additional assumptions that $\pi^F(q_1^L(R), r) < \pi^F(q_1^N, r) = \pi_1^N$, where the equality uses part 2 of the baseline assumptions. We thus have that $\Pi_1(R) < \pi_1^N$, which contradicts Corollary 1.

Proof of Proposition 13. The proof is analogous to the proof of Proposition 12 and is found in the Supplementary Appendix.

Proof of Proposition 14. Consider first the Stackelberg-up case. We know from Proposition 6 that $q_1^L(R) \geq q_1^N$. Thus, let us consider any pooling equilibrium such that $q_1^L(R) = q_1^L(F)$ and $q_1^L(F) > q_1^N$. The equilibrium payoff to Firm 1 of type $t \in \{F, R\}$ is then $\Pi_1(t) = \pi^t(q_1^L(F), r)$. Pick $q_1^L = q_1^L(R) - \varepsilon$ with $\varepsilon > 0$ and $q_1^L = q_1^L(R) - \varepsilon > q_1^N$. Define b' by $\pi^F(q_1^L, b') = \pi^F(q_1^L(F), r)$. For ε small and using parts 1 and 2 of the additional assumptions, we have that $b' \in (r, 1)$. Clearly, $b' \in D_0^F(q_1^L)$. From the single-crossing property, we now have that $\pi^R(q_1^L, b') > \pi^R(q_1^L(F), r)$. Thus, $b' \in D^R(q_1^L)$.

Using part 2 of the additional assumptions, we see that $\Delta^F(q_1^L, b) \equiv \pi^F(q_1^L, b) - \Pi_1(F)$ is increasing in b. Thus, $D^F(q_1^L) \cup D_0^F(q_1^L) = \{b | b \geq b'\}$. Likewise, using part 2 of the additional assumptions, we see that $\Delta^R(q_1^L, b) \equiv \pi^R(q_1^L, b) - \Pi_1(R)$ is increasing in b. Since $\Delta^R(q_1^L, b') > 0$, it follows that $D^R(q_1^L)$ includes $\{b | b \geq b'\}$. We conclude that $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$ and $D^R(q_1^L) \notin \varnothing$, and so the refinement requires that $b(q_1^L) = 1$. This belief in turn induces the resolute Firm 1 to deviate, since $b = 1 \in D^R(q_1^L)$.

The proof for the Stackelberg-down case is analogous and is found in the Supplementary Appendix.

11 Appendix C

We confirm here that the single-crosssing-property assumption holds for our applications.

To confirm that the simple quantity game satisfies the single-crossing property as defined for the Stackelberg-up case, we note that the indifference equation $\pi^F(q_1^L, b) = \pi^F(q_1^L(F), r)$ defines a function $b = \bar{b}(q_1^L)$ such that

$$\frac{d\bar{b}}{dq_1^L}|_{\overline{\pi}^F} = \frac{b(3+b)}{3(q_1^N - q_1^L)} < 0$$

for $b > 0$ and $q_1^N < q_1^L$. As expected, the flexible Firm 1's payoff is held constant exactly when $q_2^*(q_1^L, b)$ is held constant. We may now compute that

$$\frac{d\pi^R(q_1^L, \bar{b}(q_1^L))}{dq_1^L} = \frac{6\beta(q_1^N - q_1^L)}{3+b} < 0$$

for $q_1^N < q_1^L$. Hence, if we start at $(q_1^L(F), r)$ with $q_1^L(F) > q_1^N$ and $\bar{b}(q_1^L(F)) = r$ and then consider (q_1^L, b) with $q_1^L = q_1^L(F) - \varepsilon > q_1^N$ and $\bar{b}(q_1^L) = b' \in (r', 1)$ for $\varepsilon > 0$ sufficiently small, then

$$d\pi^R(q_1^L, \bar{b}(q_1^L)) = \frac{6\beta(q_1^L - q_1^N)}{3+b} \cdot \varepsilon > 0,$$

and so the single-crossing-property assumption holds for the Stackelberg-up case.

To confirm that the simple monetary-policy game satisfies the single-crossing-property assumption as defined for the Stackelberg-down case, we note that the indifference equation $\pi^F(q_1^L, b) = \pi^F(q_1^L(F), r)$ defines a function $b = \bar{b}(q_1^L)$ such that

$$\frac{d\bar{b}}{dq_1^L}|_{\overline{\pi}^F} = \frac{b}{q_1^N - q_1^L} > 0$$

for $b > 0$ and $q_1^N > q_1^L$. We may now compute that

$$\frac{d\pi^R(q_1^L, \bar{b}(q_1^L))}{dq_1^L} = \alpha(q_1^N - q_1^L) > 0$$

for $q_1^N > q_1^L$. Hence, if we start at $(q_1^L(F), r)$ with $q_1^L(F) \in [0, q_1^N)$ and $\bar{b}(q_1^L(F)) = r$ and then consider (q_1^L, b) with $q_1^L = q_1^L(F) + \varepsilon < q_1^N$ and $\bar{b}(q_1^L) = b' \in (r', 1)$ for $\varepsilon > 0$ sufficiently small, then

$$d\pi^R(q_1^L, \bar{b}(q_1^L)) = \alpha(q_1^N - q_1^L) \cdot \varepsilon > 0,$$

and so the single-crossing-property assumption holds for the Stackelberg-down case.

References

Abreu, D. and F. Gul (2000), "Bargaining and Reputation," *Econometrica*, 68.1, 85-117.

Backus, D. and J. Driffill (1985), "Inflation and Reputation," *The American Economic Review*, 75.3, 530-538.

Bagwell, K. (1995), "Commitment and Observability in Games," *Games and Economic Behavior*, 8, 271-280.

Bagwell, K. (2007), "Signalling and Entry Deterrence: A Multidimensional Analysis," *The Rand Journal of Economics*, 38(3), 670-697.

Bagwell, K. and R. W. Staiger (1994), "The Sensitivity of Strategic and Corrective R&D Policy in Oligopolistic Industries," *Journal of International Economics*, 36, 133-150.

Banks, J. S. (1990), "A Model of Electoral Competition with Incomplete Information," *Journal of Economic Theory*, 50, 309-325.

Banks, J. S. and J. Sobel (1987), "Equilibrium Selection in Signaling Games," *Econometrica*, 55(3), 647-661.

Barro, R. and D. Gordon (1983), "Rules, Discretion, and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, July, 12, 101-21.

Bayus, B. L., S. Jain and A. G. Rao (2001), "Truth of Consequences: An Analysis of Vaporware and New Product Announcements," *Journal of Marketing Research,* 38, 3-31.

Bernheim, B. D. (1994), "A Theory of Conformity," *Journal of Political Economy*, 102(5), 841-877.

Bernheim, B. D. and S. Severinov (2003), "Bequests as Signals: An Explanation for the Equal Division Puzzle," *Journal of Political Economy*, 111(4), 733-764.

Besley, T. and K. Suzumura (1992), "Taxation and Welfare in an Oligopoly with Strategic Commitment," *International Economic Review*, 33, 413-431.

Bhaskar, V. (2009), "Commitment and Observability in a Contracting Environment," *Games and Economic Behavior*, 66, 708-720.

Callander, S. (2008), "Political Motivations," *The Review of Economic Studies*, 75, 671-697.

Callander, S. and S. Wilkie (2007), "Lies, Damned Lies, and Political Campaigns," *Games and Economic Behavior*, 60, 262-286.

Chatterjee, K. and L. Samuelson (1987), "Bargaining with Two-Sided Incomplete Information: An Infinite-Horizon Model with Alternating Offers," *The Review of Economic Studies*, 54(2), 175-192.

Chatterjee, K. and L. Samuelson (1988), "Bargaining with Two-Sided Incomplete Information: The Unrestricted Offers Case," *Operations Research*, 36(4), 605-638.

Cho, I.-K. and D. M. Kreps (1987), "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics*, 102(2), 179-221.

Cho, I.-K. and J. Sobel (1990), "Strategic Stability and Uniqueness in Signaling Games," *Journal of Economic Theory*, 50, 380-413.

Christensen, L. R. and R. E. Caves (1997), "Cheap Talk and Investment Rivalry in the Pulp and Paper Industry," *Journal of Industrial Economics*, 45, 47-73.

Compte, O. and P. Jehiel (2002), "On the Role of Outside Options in Bargaining with Obstinate Parties," *Econometrica*, 70, 1477-1517.

Corona, C. and L. Nan (2013), "Preannouncing Competitive Decisions in Oligopoly Markets," *Journal of Accounting and Economics*, 56, 73-90.

Cournot, A. (1963), *Researches into the Mathematical Principles of the Theory of Wealth*, Homewood, Il.: Richard D. Irwin, Inc.

Crawford, V. (1982), "A Theory of Disagreement in Bargaining," *Econometrica*, 50, 607-637.

Crawford, V. and J. Sobel (1982), "Strategic Information Transmission," *Econometrica* 50, 1431-1451.

Dai, Y. (2017), "Consumer Search and Optimal Pricing under Limited Commitment," Chapter 2 in *Essays on Information, Search and Pricing,* University of Iowa thesis, August.

Doyle, M. P. and C. M. Snyder (1999), "Information Sharing and Competition in the Motor Vehicle Industry," *Journal of Political Economy*, 107, 1326-1364.

Fudenberg, D. and D. K. Levine (1989), "Reputation and Equilibrium Selection in Games with a Patient Player," *Econometrica*, 57(4), 759-778.

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

Gilbert, R. J. and M. Lieberman (1987), "Investment and Coordination in Oligopolistic Industries," *The Rand Journal of Economics*, 18, 17-33.

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.

Inderst, R. (2005), "Bargaining with a Possibly Committed Seller," *Review of Economic Dynamics*, 8, 927-944.

Kamada, Y. and M. Kandori (2011), "Revision Games," December 31, working paper.

Kamada, Y. and S. Moroni (2017), "Games with Private Timing," August 31, working paper.

Kambe, S. (1999), "Bargaining with Imperfect Commitment," *Games and Economic Behavior*, 28, 217-237.

Kartik, N. (2009), "Strategic Communication with Lying Costs," *The Review of Economic Studies*, 76, 1359-1395.

Kartik, N. and A. Frankel (2017), "Muddled Information," March 26, working paper.

Kartik, N. and R. P. McAfee (2007), "Signaling Character in Electoral Competition," *The American Economic Review*, 97(3), 852-870.

Kim, K. (2009), "The Coase Conjecture with Incomplete Information on the Monopolist's Commitment," *Theoretical Economics*, 4, 17-44.

Kohlberg, E. and J.-F. Mertens (1986), "On the Strategic Stability of Equilibria," *Econometrica*, 54(1), 1003-1037.

Kreps, D. M. and G. Ramey (1987), "Structural Consistency, Consistency, and Sequential Rationality," *Econometrica*, 55(6), 1331-1348.

Kreps, D. M. and J. A. Scheinkman (1983), "Quantity Precommitment and Bertrand Competition Yield Cournot Outcomes," *The Bell Journal of Economics*, 14(2), 326-337.

Kreps, D. M. and R. Wilson (1982a), "Sequential Equilibria," *Econometrica*, 50, 863-894.

Kreps, D. M. and R. Wilson (1982b), "Reputation and Imperfect Information," *Journal of Economic Theory*, 27, 253-279.

Maggi, G. (1996), "Strategic Trade Policies with Endogenous Mode of Competition," *The American Economic Review*, 86(1), 237-258.

Maggi, G. (1999), "The Value of Commitment with Imperfect Observability and Private Information," *The Rand Journal of Economics*, 30, 555-574.

Mailath, G. J. (1993), "Endogenous Sequencing of Firm Decisions," *Journal of Economic Theory*, 59(1), 169-82.

Milgrom, P. R. and D. J. Roberts (1982), "Predation, Reputation and Entry Deterrence," *Journal of Economic Theory*, 27, 280-312.

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, 1344-53.

Myerson, R. (1991), *Game Theory: Analysis of Conflict*, Cambridge, MA: Harvard University Press.

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.

Poddar, S. and D. Sasaki (2002), "The Strategic Benefit from Advance Production," *European Journal of Political Economy*, 18(3), 579-595.

Reinganum, J. (1983), "Technology Adoption under Imperfect Information," *The Bell Journal of Economics*, 14, 57-69.

Rogoff, K. (2017), "Dealing with Monetary Paralysis at the Zero Bound," *Journal of Economic Perspectives*, 31(3), 47-66.

Ross, S. A. (1977), "The Determination of Financial Structure: The Incentive-Signalling Approach," *The Bell Journal of Economics*, 8(1), 23-40.

Sanktjohanser, A. (2017), "Optimally Stubborn," September 8, working paper, preliminary draft.

Schelling, T. (1960), *The Strategy of Conflict*, Cambridge, MA: Harvard University Press.

Sonnenschein, H. (1968), "The Dual of Duopoly Is Complementary Monopoly: or, Two of Cournot's Theories Are One," *The Journal of Political Economy*, 76(2), 316-318.

van Damme, E. and S. Hurkens (1997), "Games with Imperfectly Observed Commitment," *Games and Economic Behavior*, 21, 282-308.

Vardy, F. (1994), "Stackelberg Leader Games with Observation Costs," *Games and Economic Behavior*, 49, 374-400.

Vardy, F. (2004), "The Value of Commitment in Stackelberg Games with Observation Costs," *Games and Economic Behavior*, 49, 374-400.

von Stackelberg H. (1934), *Marktform und Gleichgewicht*, Springer-Verlag, Vienna/Berlin.

Wolitsky, A. (2012), "Reputational Bargaining with Minimal Knowledge of Rationality," *Econometrica*, 80, 2047-2087.

Figure 1: Period-2 equilibrium quantities for 0 < b < 1

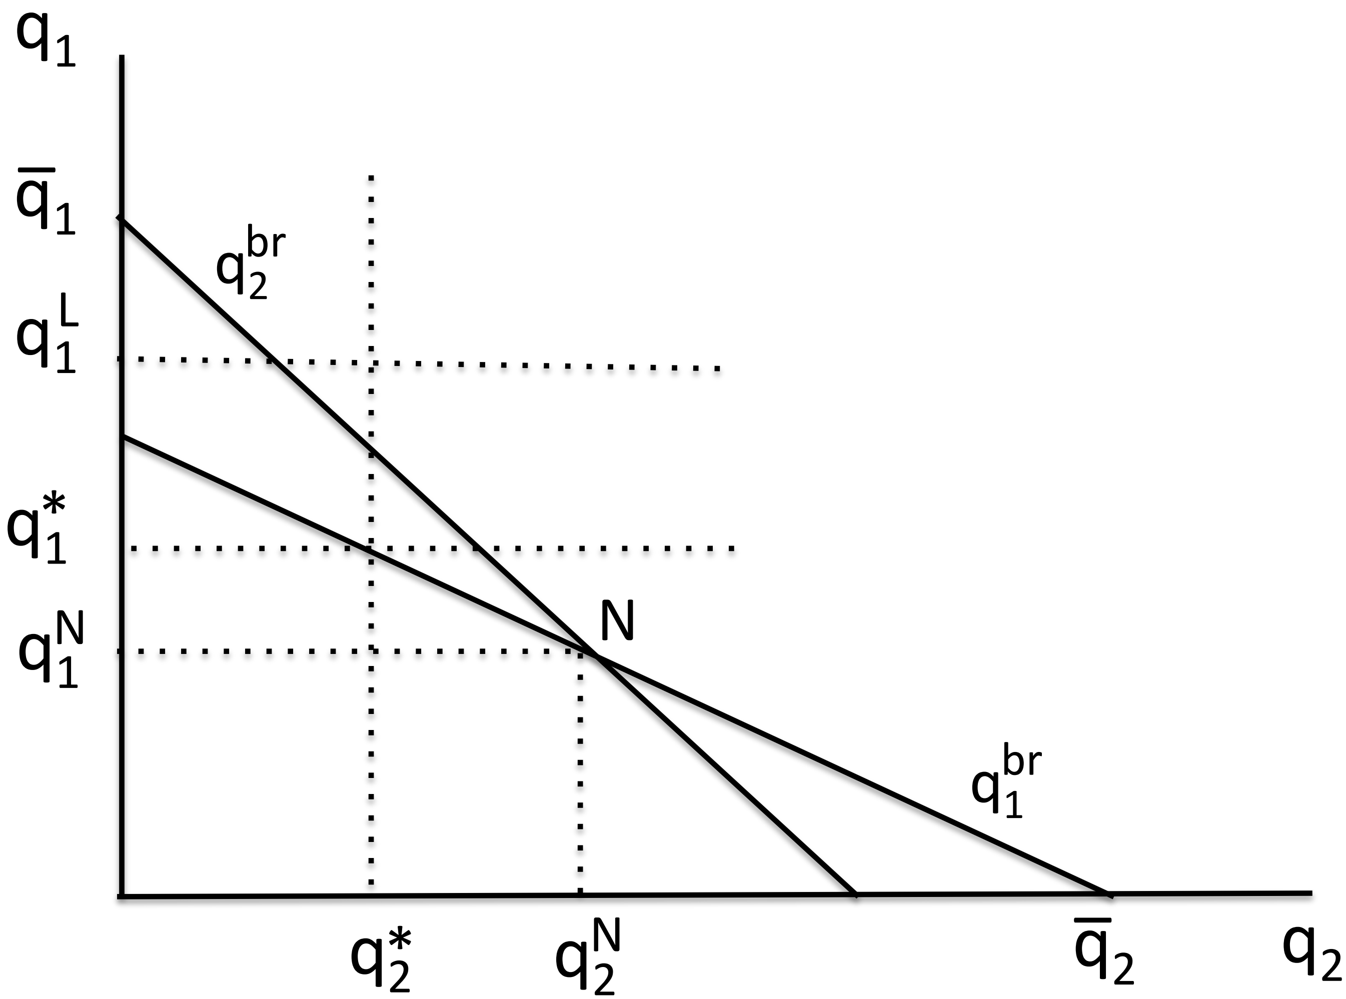

Figure 2: Gains from Deviation

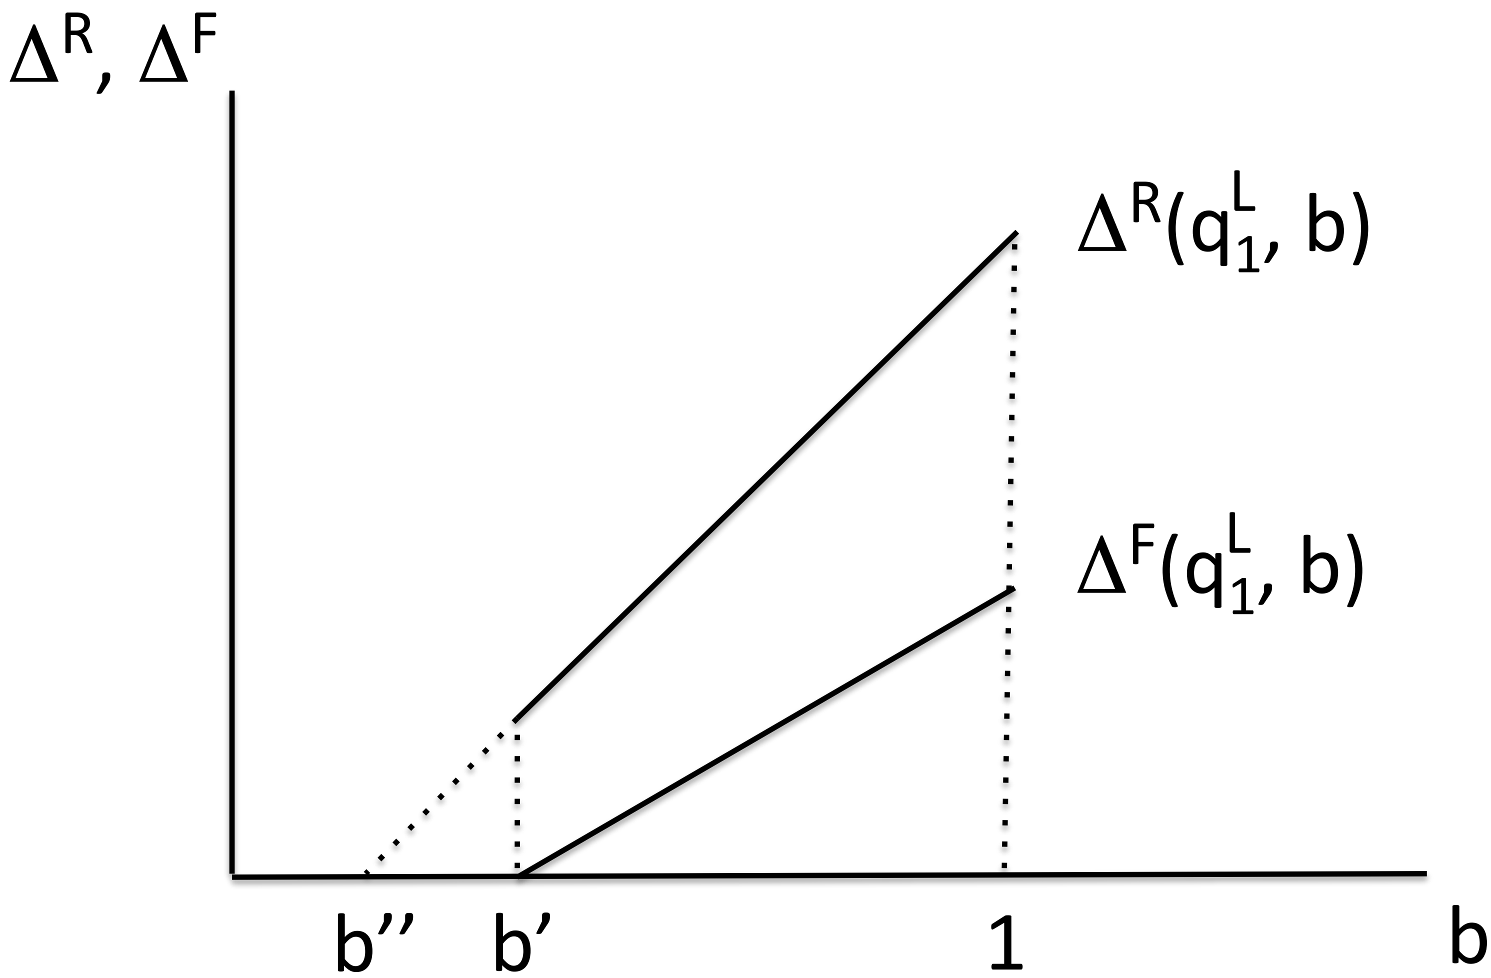

Figure 3: Pooling equilibrium with $q_1^N < q_1^L(R) = q_1^L(F) < q_1^{gs}(r)$

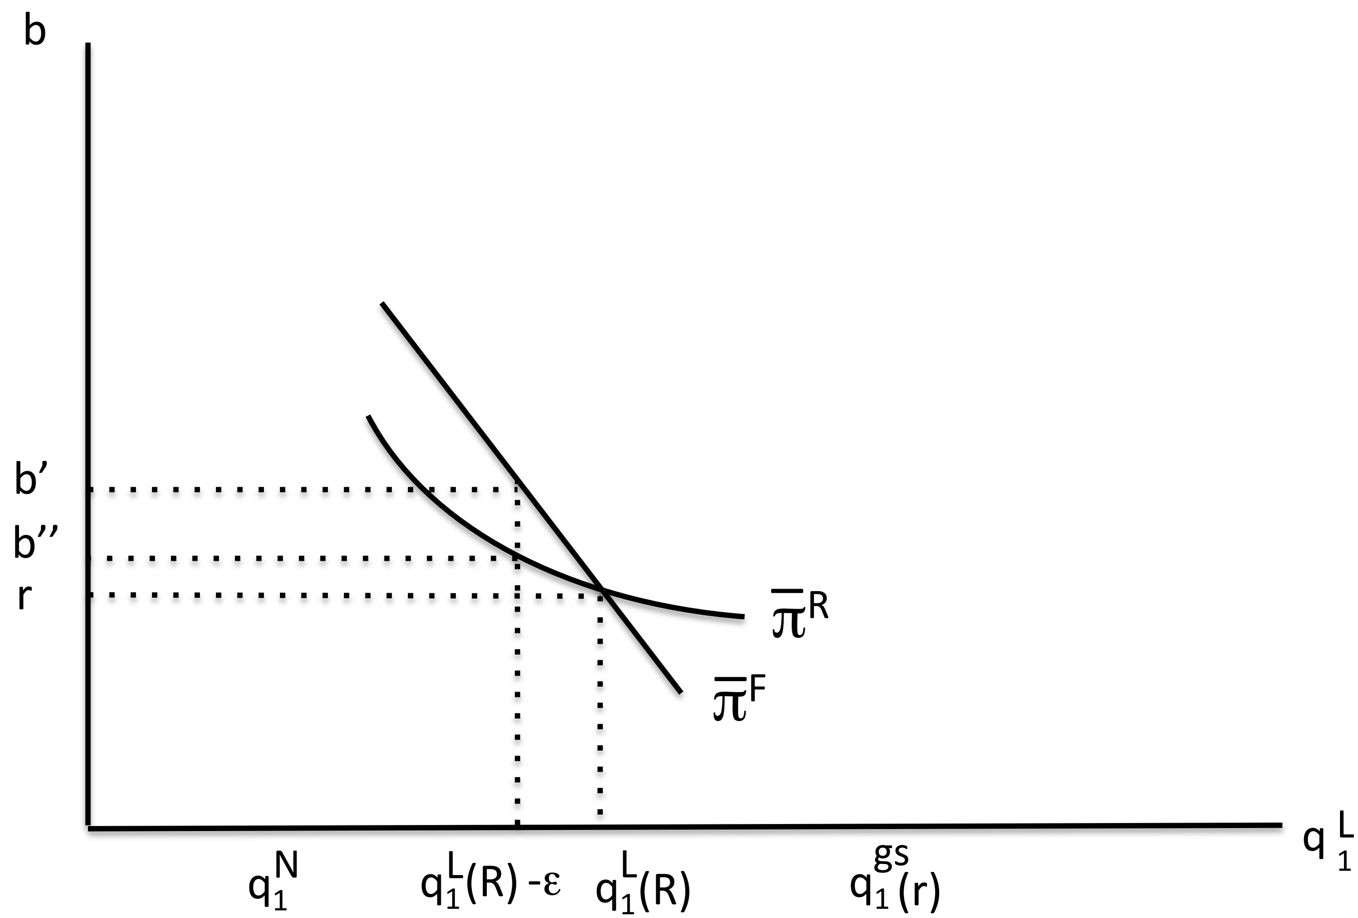

Figure 4: Indifference Curves