Supplementary Appendix to Commitment with Private Resolve

Kyle Bagwell*

This draft: January 28, 2018

Abstract

In this Supplementary Appendix, extensions, proofs and applications referred to in Bagwell (2018) are developed.

*Department of Economics at Stanford University; and NBER.

Introduction

In this Supplementary Appendix, extensions, proofs and applications referred to in Bagwell (2018) are developed.

Elimination of never weak best response strategies

Suppose that we have an outcome that is generated by a sequential equilibrium for the game with private resolve as defined in Section 3.2. Consider the set of all sequential equilibria in this game giving rise to this outcome. We allow for mixed-strategy equilibria under the restriction that the distribution over actions is discrete and has positive support.[1] In such equilibria, the flexible Firm 1 randomizes with respect to its leader quantity. The final outputs of the flexible Firm 1 and Firm 2 are not randomized in equilibrium, given concavity as delivered by (1). Indeed, in any sequential equilibrium, for any given $q_1^L \in Q_1$, the flexible Firm 1 selects $q_1^*(q_1^L, b)$ and Firm 2 selects $q_2^*(q_1^L, b)$ for some belief $b = b(q_1^L)$. Allowing for mixed-strategy sequential equilibria under the restriction just described, let us abuse notation somewhat and define $\Pi_1(t)$ as the payoff for type $t \in \{R, F\}$ in the sequential equilibrium outcome under consideration.

We consider two cases. Assume first that there exists a leader quantity $q_1^L \in Q_1$ that is off the equilibrium path (i.e., selected with zero probability) for all sequential equilibria in the set under consideration and that satisfies $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$ and $D^R(q_1^L) \neq \varnothing$. For present purposes, let us call a leader quantity that satisfies these assumptions the "deviant q_1^L."

Let us suppose that the flexible Firm 1 has a strategy in which the deviant q_1^L is played with positive probability and that is a weak best response to some sequential equilibrium in the set under consideration. Fix this strategy and sequential equilibrium. For this sequential equilibrium, Firm 2's strategy must of course specify a response $q_2(q_1^L)$ to the deviant q_1^L; in particular, we know from above that there exists $b' \in [0, 1]$ such that $q_2(q_1^L) = q_2^*(q_1^L, b')$. If as part of this strategy the flexible Firm 1 were to select $q_1(q_1^L) \neq q_1^*(q_1^L, b')$, where $q_1^*(q_1^L, b') = q_1^{br}(q_2^*(q_1^L, b'))$ by (9), then the flexible Firm 1 could do better yet with an alternative strategy that couples the deviant q_1^L with $q_1(q_1^L) = q_1^*(q_1^L, b')$, as then the flexible Firm 1 would be selecting its best response to Firm 2's selection, $q_2(q_1^L) = q_2^*(q_1^L, b')$. Since by hypothesis the fixed strategy for the flexible Firm 1 is a weak best response to the given sequential equilibrium, the alternative strategy would generate a gain from deviation. Thus, the existence of the given sequential equilibrium requires that $q_1(q_1^L) = q_1^*(q_1^L, b')$ in the fixed strategy under consideration. Our supposition is thus that, for the flexible Firm 1, we have fixed a strategy that selects the deviant q_1^L with positive

[1] Mixed-strategy equilibria are discussed further in the next section of this Supplementary Appendix.

probability, that follows the selection of the deviant q_1^L with $q_1(q_1^L) = q_1^*(q_1^L, b')$, and that is a weak best response to the given sequential equilibrium. This supposition in turn ensures that $b' \in D^F(q_1^L) \cup D_0^F(q_1^L)$. Recall now that, for the deviant q_1^L, $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$. It thus follows that $b' \in D^R(q_1^L)$. We can now conclude that the resolute Firm 1 would deviate and select the deviant q_1^L followed by $q_1(q_1^L) = q_1^*(q_1^L, b')$, which contradicts the existence of the given sequential equilibrium. Hence, for the deviant q_1^L, it must be that any strategy for the flexible Firm 1 that includes the deviant q_1^L with positive probability is never a weak best response to any sequential equilibrium in the given set.

For the given set of sequential equilibria, suppose now that we prune strategies for Firm 1 in the game with private resolve that are never a weak best response to any sequential equilibrium in the given set. Since the deviant q_1^L satisfies $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$, we argue above that it must be that any strategy that entails selecting q_1^L with positive probability is never a weak best response for the flexible Firm 1. Given that the deviant q_1^L also satisfies $D^R(q_1^L) \neq \varnothing$, we can easily argue that a strategy that selects the deviant q_1^L is a weak best response for the resolute Firm 1 to some sequential equilibrium in the given set. Hence, in any sequential equilibrium that generates the given outcome in the game after never weak best responses for Firm 1 are removed, we must have that $b(q_1^L) = 1$. In other words, for any leader quantity $q_1^L \in Q_1$ that is off the equilibrium path for all sequential equilibria leading to a given outcome, if $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$ and $D^R(q_1^L) \neq \varnothing$, then the sequential equilibrium outcome exists after never-weak-best-response strategies are pruned for the flexible Firm 1 only if the associated beliefs satisfy $b(q_1^L) = 1$. For pure-strategy sequential equilibria, we may directly transfer this implication to the reduced-form game defined in the paper and impose that, for any leader quantity $q_1^L \in Q_1$ that is off the equilibrium path in all equilibria leading to a given outcome, if $D^F(q_1^L) \cup D_0^F(q_1^L) \sqsubseteq D^R(q_1^L)$ and $D^R(q_1^L) \neq \varnothing$, then $b(q_1^L) = 1$ is required. We observe that this requirement matches exactly the requirement imposed by our refinement in (22) for the case of $t = F$ and $t' = R$. An analogous argument holds in mixed-strategy equilibria wherein the flexible Firm 1 randomizes over leader quantities, given the restriction mentioned above.

We turn now to the second case, which is more straightforward. For this case, we assume that there exists a leader quantity $q_1^L \in Q_1$ that is off the equilibrium path (i.e., selected with zero probability) for all sequential equilibria in the set under consideration and that satisfies $D^R(q_1^L) \cup D_0^R(q_1^L) \sqsubseteq D^F(q_1^L)$ and $D^F(q_1^L) \neq \varnothing$. Focusing now on this second case, let us call a leader quantity that satisfies these assumptions the "deviant q_1^L."

Let us suppose that the resolute Firm 1 has a strategy in which the deviant q_1^L is played with positive probability and that is a weak best response to some sequential equilibrium in the set under consideration. Fix this strategy and sequential equilibrium. As before, for this sequential equilibrium, Firm 2's strategy must specify a response $q_2(q_1^L)$ to the deviant q_1^L; in particular, we know from above that there exists $b' \in [0, 1]$ such that

$q_2(q_1^L) = q_2^*(q_1^L, b')$. For the resolute Firm 1, if it selects the deviant q_1^L as its leader quantity, then the deviant q_1^L is also its final output. Our supposition is thus that, for the resolute Firm 1, we have fixed a strategy in which the deviant q_1^L is selected with positive probability and that is a weak best response to the given sequential equilibrium. This supposition in turn ensures that $b' \in D^R(q_1^L) \cup D_0^R(q_1^L)$. Recall now that, for the deviant q_1^L, $D^R(q_1^L) \cup D_0^R(q_1^L) \sqsubseteq D^F(q_1^L)$. It thus follows that $b' \in D^F(q_1^L)$. We can now conclude that the flexible Firm 1 would deviate and select the deviant q_1^L followed by $q_1(q_1^L) = q_1^*(q_1^L, b')$, which contradicts the existence of the given sequential equilibrium. Hence, for the deviant q_1^L, it must be that any strategy for the resolute Firm 1 that includes the deviant q_1^L with positive probability is never a weak best response to any sequential equilibrium in the given set.

For the given set of sequential equilibria, suppose now that we prune strategies for Firm 1 in the game with private resolve that are never a weak best response to any sequential equilibrium in the given set. Since the deviant q_1^L satisfies $D^R(q_1^L) \cup D_0^R(q_1^L) \sqsubseteq D^F(q_1^L)$, we argue above that it must be that any strategy that entails selecting q_1^L with positive probability is never a weak best response for the resolute Firm 1. Given that the deviant q_1^L also satisfies $D^F(q_1^L) \neq \varnothing$, we can easily argue that a strategy that selects the deviant q_1^L is a weak best response for the flexible Firm 1 to some sequential equilibrium in the given set. Hence, in any sequential equilibrium that generates the given outcome in the game after never weak best responses for Firm 1 are removed, we must have that $b(q_1^L) = 0$. In other words, for any leader quantity $q_1^L \in Q_1$ that is off the equilibrium path in all sequential equilibria leading to a given outcome, if $D^R(q_1^L) \cup D_0^R(q_1^L) \sqsubseteq D^F(q_1^L)$ and $D^F(q_1^L) \neq \varnothing$, then the sequential equilibrium outcome exists after never-weak-best-response strategies are pruned for Firm 1 only if the associated beliefs satisfy $b(q_1^L) = 0$. For pure-strategy sequential equilibria, we may directly transfer this implication to the reduced-form game defined in the text and impose that, for any leader quantity $q_1^L \in Q_1$ that is off the equilibrium path for all equilibria leading to a given outcome, if $D^R(q_1^L) \cup D_0^R(q_1^L) \sqsubseteq D^F(q_1^L)$ and $D^F(q_1^L) \neq \varnothing$, then $b(q_1^L) = 0$ is required. We observe that this requirement matches exactly the requirement imposed by our refinement in (22) for the case of $t = R$ and $t' = F$. An analogous argument holds in mixed-strategy equilibria wherein the flexible Firm 1 randomizes over leader quantities, given the restriction mentioned above.

Refined mixed-strategy equilibria

We characterize here the mixed-strategy sequential equilibria for the game with private resolve as defined in Section 3.2. The sequential equilibrium concept is defined for games with finite action spaces, and the D1 refinement is likewise defined for finite signaling games. We analyze here a game with a continuum of actions. The definitions extend in natural ways when players use pure strategies, as shown in Sections 4 and 5. Our approach is to restrict attention to mixed strategies in which the distribution over actions is discrete and has finite support. A mixed strategy for a player then indicates the probability that the player will select a given action in the support. The Bayes' consistency requirement of sequential equilibrium determines beliefs for actions that occur with positive probability in a proposed mixed-strategy equilibrium. We apply the refinement when actions are observed that are zero-probability events under the proposed mixed-strategy equilibrium.

To begin, we note that, 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 selections of period-2 quantities. Any randomization thus must involve the leader quantity choices of the resolute and flexible Firm 1. We can thus embed the period-2 equilibrium quantities into the payoff functions and consider the reduced-form game. The definition of equilibrium is then modified to allow that Firm 1 of type t may randomize over different leader quantities that maximize its payoff, where beliefs are formed via Bayes' rule for leader quantities that arise with positive probability under the equilibrium strategies and where the equilibrium strategies now may be mixed. We will say that an equilibrium so defined is a *mixed-strategy equilibrium* if there exists a type t of Firm 1 such that its equilibrium strategy places positive probability on more than one leader quantity. We now develop our characterizations through a sequence of observations.

A first observation is that there does not exist a mixed-strategy equilibrium in which $q_1^L < q_1^N$ is selected by the resolute Firm 1 with positive probability.

The proof follows arguments made in the proof of Proposition 6. Suppose to the contrary that a mixed-strategy equilibrium exists in which the resolute Firm 1 selects $q_1^L < q_1^N$ with positive probability, so that $b(q_1^L) > 0$. If $b(q_1^L) = 1$, then we can argue similarly to the discussion leading up to Proposition 5 that the resolute Firm 1 would gain by reallocating the probability it plays q_1^L from q_1^L to q_1^N. Suppose then that $b(q_1^L) \in (0, 1)$. Lemma 1 then gives $q_1^L < q_1^*(q_1^L, b(q_1^L)) < q_1^N$ and $q_2^N < q_2^*(q_1^L, b(q_1^L)) < q_2^{br}(q_1^L)$. Since the resolute Firm 1 must be indifferent over all leader quantities that are selected with positive probability, its equilibrium profit level must be given by $\pi_1(q_1^L, q_2^*(q_1^L, b(q_1^L)))$. From here, we can follow the proof of Proposition 6 to conclude that $\pi_1(q_1^L, q_2^*(q_1^L, b(q_1^L))) < \pi_1^N$. Since the resolute Firm 1 can always deliver the profit π_1^N by selecting q_1^N, we thus have a contradiction, and so the first observation is established.

Second, we suppose that a mixed-strategy equilibrium exists in which the flexible Firm 1 places positive probability on more than one leader quantity. Let q_1^{Lm} and q_1^{Ln} denote any two distinct leader quantities that are selected with positive probability by the flexible Firm 1 in a mixed-strategy equilibrium. Since the flexible Firm 1 must be indifferent between any leader quantities that it selects with positive probability, and since the flexible Firm 1 ultimately chooses its period-2 equilibrium quantity as a best response to Firm 2's period-2 equilibrium quantity, we conclude from (9) and (16) that q_1^{Lm} and q_1^{Ln} must generate the same period-2 equilibrium quantity from Firm 2: $q_2^*(q_1^{Lm}, b(q_1^{Lm})) = q_2^*(q_1^{Ln}, b(q_1^{Ln}))$.

Third, we claim that a refined mixed-strategy equilibrium does not exist in which the flexible Firm 1 employs a mixed strategy and earns a payoff greater than its Nash payoff, π_1^N. To establish this claim, suppose to the contrary that a mixed-strategy equilibrium exists in which the flexible Firm 1 randomizes and earns a payoff greater than π_1^N. Since the flexible Firm 1 must be indifferent with respect to all leader quantities that it selects with positive probability, and since the flexible Firm 1 is sure to earn π_1^N if it selects the Nash leader quantity, q_1^N, the flexible Firm 1 must not select q_1^N with positive probability. Hence, the flexible Firm 1 must randomize over at least two leader quantities, with each such leader quantity differing from q_1^N. We also know that, for any leader quantity that the flexible Firm 1 selects with positive probability, the resolute Firm 1 must also choose this leader quantity with positive probability; otherwise, the flexible Firm 1 would be revealed when choosing this leader quantity and thus earn π_1^N, which contradicts our initial assumption.

With these observations in hand, we may now conclude that the flexible Firm 1 selects at most two leader quantities with positive probability. This follows since the resolute Firm 1 must select with positive probability any leader quantity that the flexible Firm 1 selects with positive probability, Firm 2's period-2 equilibrium quantity must be constant in response to such quantities as we know from our second observation, and the resolute Firm 1's profit function is concave by (1) and thus has at most two leader quantities that deliver a fixed profit level when Firm 2's period-2 equilibrium quantity is held constant. Our assumption is that these two leader quantities exist. Let us denote them as q_1^{Lm} and q_1^{Ln}, where $q_1^{Lm} > q_1^{Ln}$. Since these leader quantities must differ from q_1^N, we may conclude from our first observation that $q_1^{Lm} > q_1^{Ln} > q_1^N$. We also know that the resolute leader cannot choose any leader quantity just below q_1^{Lm} with positive probability, since otherwise the flexible Firm 1 would deviate to the slightly lower leader quantity and generate a lower period-2 equilibrium quantity from Firm 2 by inducing the belief that Firm 1 is the resolute type.

Let us thus suppose that there exists a mixed-strategy equilibrium in which there are two distinct leader quantities, q_1^{Lm} and q_1^{Ln} with $q_1^{Lm} > q_1^{Ln} > q_1^N$, such that each is selected with positive probability both by the flexible Firm 1 and the resolute Firm 1 where the probability that the flexible Firm 1 selects q_1^{Lm} or q_1^{Ln} is one. Given $q_2^*(q_1^{Lm}, b(q_1^{Lm})) = q_2^*(q_1^{Ln}, b(q_1^{Ln}))$, the resolute Firm 1 can be indifferent only if $q_1^{Lm} > q_1^{br}(q_2^*(q_1^{Lm}, b(q_1^{Lm}))) > q_1^{Ln}$. Since the resolute leader cannot choose any leader quantity just below q_1^{Lm} with positive probability, for any $\varepsilon > 0$ and sufficiently small, we can find a deviant leader quantity, $q_1^L = q_1^{Lm} - \varepsilon > q_1^{br}(q_2^*(q_1^{Lm}, b(q_1^{Lm})))$, such that q_1^L is played with probability zero in the given mixed-strategy equilibrium. Since $b(q_1^{Lm}) \in (0,1)$, we can also pick $\varepsilon > 0$ sufficiently small so that $q_2^*(q_1^{Lm}, b(q_1^{Lm})) = q_2^*(q_1^L, b')$ for $b' \in (b(q_1^{Lm}), 1)$. The flexible Firm 1 is then indifferent between its equilibrium payoff and a deviation to q_1^L that is associated with the belief b'. The equilibrium payoff for the resolute Firm 1, however, is lower than that which it would receive from a deviation to q_1^L that is associated with the belief b', since the resolute Firm 1 enjoys the benefit of moving its final quantity closer to its best-response value. For $b > b'$, both types of Firm 1 would gain from the deviation. Our refinement thus requires that $b(q_1^L) = 1$, which induces the resolute Firm 1 to deviate.

Fourth, we claim that a refined mixed-strategy equilibrium does not exist in which the flexible Firm 1 employs a pure strategy and earns a payoff greater than its Nash payoff, π_1^N. To establish this claim, we assume to the contrary that a mixed-strategy equilibrium exists in which the resolute Firm 1 randomizes over at least two leader quantities, the flexible Firm 1 selects a single leader quantity $q_1^L(F)$ with probability one, and the flexible Firm 1 earns a payoff that exceeds π_1^N. This case is possible only if the resolute Firm 1 employs a mixed strategy that selects $q_1^L(F)$ with positive probability, so that $b(q_1^L(F)) \in (0,1)$ where $q_1^L(F) > q_1^N$. The resolute Firm 1 then must also select at least one other leader quantity with positive probability, where in this case any leader quantity q_1^{Lm} with $q_1^{Lm} \neq q_1^L(F)$ that the resolute Firm 1 selects with positive probability must satisfy $q_1^{Lm} \geq q_1^N$ and $b(q_1^{Lm}) = 1$. The flexible Firm 1 would deviate and mimic any such q_1^{Lm} if $q_1^{Lm} > q_1^L(F)$, since it would thereby lower the period-2 equilibrium quantity of Firm 2. Hence, we also require for this case that $q_1^L(F) > q_1^{Lm}$ for any q_1^{Lm} that the resolute Firm 1 selects with positive probability that differs from $q_1^L(F)$. Given our assumption that the second-order condition for the standard Stackelberg solution holds with strict inequality, and the fact that the resolute Firm 1 must be indifferent with respect to all leader quantities q_1^{Lm} that induce the belief $b(q_1^{Lm}) = 1$, we know that the mixed strategy of the resolute Firm 1 can put positive probability in this case on at most two leader quantities that differ from $q_1^L(F)$. For any mixed-strategy equilibrium of this kind, we can thus always find q_1^L arbitrarily close to and below $q_1^L(F)$ that is played with zero probability.

Let us thus suppose that there exists a mixed-strategy equilibrium in which the flexible Firm 1 selects $q_1^L(F)$ with probability one and the resolute Firm 1 selects $q_1^L(F)$ with positive probability while also selecting at least one but no more than two other leader quantities with positive probability. Given our preceding discussion, let us suppose, too, that $b(q_1^L(F)) \in (0,1)$ and $q_1^L(F) > q_1^N$. Using Lemma 1 and (9), we may further conclude that $q_1^L(F) > q_1^{br}(q_2^*(q_1^L(F), b(q_1^L(F)))) > q_1^N$. As noted above, we can now find a deviant leader quantity, $q_1^L = q_1^L(F) - \varepsilon > q_1^{br}(q_2^*(q_1^L(F), b(q_1^L(F))))$, such that q_1^L is played with probability zero in the given mixed-strategy equilibrium. Since $b(q_1^L(F)) \in (0,1)$, we can also pick $\varepsilon > 0$ sufficiently small so that $q_2^*(q_1^L(F), b(q_1^L(F))) = q_2^*(q_1^L, b')$ for $b' \in (b(q_1^L(F)), 1)$. As before, the flexible Firm 1 is then indifferent between its equilibrium payoff and a deviation to q_1^L that is associated with the belief b'; however, the equilibrium payoff for the resolute Firm 1 is lower than that which it would receive from a deviation to q_1^L that is associated with the belief b', since the resolute Firm 1 enjoys the benefit of moving its final quantity closer to its best-response value. For $b > b'$, both types of Firm 1 would gain from the deviation. Our refinement thus requires that $b(q_1^L) = 1$, which in turn induces the resolute Firm 1 to deviate.

Fifth, regardless of its type, Firm 1 does not earn less than its Nash payoff, π_1^N, in any equilibrium. This follows since Firm 1 could always achieve its Nash payoff by setting its leader quantity equal to its Nash output, q_1^N. Combining this point with our third and fourth observations, we conclude that any refined mixed-strategy equilibrium the flexible Firm 1 earns a payoff equal to its Nash payoff, π_1^N.

Sixth, we now claim that, in any refined mixed-strategy equilibrium, the resolute Firm 1 must play a pure strategy whereby it selects q_1^N with probability one. To establish this claim, we refer to our first observation to rule out leader quantities for the resolute Firm 1 that are below q_1^N. Furthermore, if the resolute Firm 1 were to select a leader quantity above q_1^N with positive probability in a mixed-strategy equilibrium, then the flexible Firm 1 could mimic this choice and earn a payoff that exceeds π_1^N, which contradicts the conclusion based on our third and fourth observations that the flexible Firm 1 earns a payoff equal to its Nash payoff, π_1^N, in any refined mixed-strategy equilibrium. This means that the resolute Firm 1 also earns its Nash payoff, π_1^N, in any refined mixed-strategy equilibrium.

Last, we observe that Firm 2 earns its Nash payoff, π_2^N, in any refined mixed-strategy equilibrium. This follows since the resolute Firm 1 must select q_1^N with probability one, while the flexible Firm 1 either separates with probability one with randomly determined leader quantities that differ from q_1^N, or pools 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 the final output of Firm 1 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.

Omitted proofs for the Stackelberg-down case

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, $q_1^L(R) \leq q_1^N$.*

Proof. The proof is analogous to the proof of Proposition 12 and is provided here for completeness.

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 $q_1^L \geq q_1^N$ 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 $q_1^N \leq q_1^L(R)$. 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 $\Pi_1(R) = \pi^R(q_1^L(R), 1) = \pi_1^N$.

We now show that Proposition 6 holds with a reversed inequality: in any equilibrium, $q_1^L(R) \leq q_1^N$. 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. ■

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.*

Proof. The Stackelberg-up case is proved in the paper. The proof for the Stackelberg-down case is analogous and is included here for completeness.

For the Stackelberg-down case, we know that Proposition 6 holds with reverse inequality: in any equilibrium, $q_1^L(R) \leq q_1^N$. 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 $b = 1 \in D^R(q_1^L)$. ■

Stackelberg-down: A simple quantity-game setting with a positive aggregate-quantity externality

Consider the quantity-game setting with quadratic payoffs but suppose that each firm enjoys an additional and separable gain when the industry output increases. To simplify the discussion, suppose also that the unit costs are symmetric across the two firms: $c_1 = c_2 \equiv c$. The profit functions thus now take the following form:

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

We can imagine that higher industry output has beneficial effects on future demand or costs or the regulatory environment, which we capture here in a simplistic way with the parameter $\lambda > 0$. For $i, j = 1, 2$ with $i \neq j$, we assume that $\alpha > 0$, $\beta > 0$ and $c \geq 0$ where $\alpha/\beta > c$, $\alpha > c$ and $\lambda > \alpha - c$. The quantity space for Firm i is $Q_i = [0, \overline{q})$ where $\overline{q} = (\alpha + \lambda - c)/\beta$.

The corresponding Nash output levels and Nash profit levels are now

$$\begin{aligned} q_i^N &= \frac{\alpha + \lambda - c}{3\beta} \\ \pi_i^N &= (\frac{\alpha + \lambda - c}{3})^2 \frac{1}{\beta} + \lambda q_j^N, \end{aligned}$$

where $q_i^N \in (0, \overline{q})$ and $i, j = 1, 2$ with $i \neq j$.

For a given leader quantity $q_1^L \in (0, \overline{q})$ 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:

$$\begin{aligned} q_1^*(q_1^L, b) &= \frac{\alpha + \lambda - c + b\beta q_1^L}{\beta(3 + b)} \\ q_2^*(q_1^L, b) &= \frac{(\alpha + \lambda - c)(1 + b) - 2\beta b q_1^L}{\beta(3 + b)}, \end{aligned}$$

where $q_i^*(q_1^L, b) \in (0, \overline{q})$. 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 payoff values take the following form:

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

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

$$q_1^{gs}(b) = \frac{\alpha + \lambda(1 - b) - c}{\beta(3 - b)} \in (0, q_1^N),$$

where $q_1^{gs}(b) > 0$ follows from $\alpha > c$ and where $q_1^{gs}(b) < q_1^N$ for $b > 0$ follows since $\lambda > (\alpha - c)/2$ is implied by our parameter restrictions. Thus, the model belongs to the Stackelberg-down case. It is direct to verify that the model satisfies the additional assumptions for the Stackelberg-down case. We note that the first additional assumption uses $\lambda > \alpha - c$ in order to establish monotonicity for the case where $b = 1$ with q_1^L near $\overline{q}_1$, whereas $\lambda > (\alpha - c)/2$ suffices for the second additional assumption and, as noted, the third additional assumption utilizes $\alpha > c$ and $\lambda > (\alpha - c)/2$.

To confirm the single-crossing property 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 $b = \overline{b}(q_1^L)$ such that

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

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 $q_1^L(F) \in [0, q_1^N)$ 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

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

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

Referring to Corollary 3, we conclude that, when this standard linear model of duopolistic quantity competition is augmented to include a positive aggregate-quantity externality, 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-down case can be associated with strategic settings in which best-response functions are decreasing and a higher action generates a positive cross-firm externality, as in the setting just analyzed, and also in which best-response functions are increasing and a higher action generates a negative cross-firm externality. For an appropriately specified model, we could illustrate the application of our results to Stackelberg-down settings of the latter kind as well.

References

Bagwell, K. (2018), "Commitment with Private Resolve," January, working paper.