



This article appeared in a journal published by Elsevier. The attached copy is furnished to the author for internal non-commercial research and education use, including for instruction at the authors institution and sharing with colleagues.

Other uses, including reproduction and distribution, or selling or licensing copies, or posting to personal, institutional or third party websites are prohibited.

In most cases authors are permitted to post their version of the article (e.g. in Word or Tex form) to their personal website or institutional repository. Authors requiring further information regarding Elsevier's archiving and manuscript policies are encouraged to visit:

<http://www.elsevier.com/copyright>



Commitment and observability in a contracting environment [☆]

V. Bhaskar ^{*}

Department of Economics, University College London, Gower Street, London WC1E 6BT, United Kingdom

Received 27 November 2006

Available online 2 October 2008

Abstract

This paper examines leader-follower games where a leader must purchase an essential input from a price-setting supplier in order to take an action. We show that equilibrium outcomes when the followers perfectly observe the leaders' actions cannot be approximated by mixed equilibrium outcomes of the game where followers imperfectly observe the leaders' actions, i.e. they are not *accessible*. Accessibility fails since in a pure strategy equilibrium, a supplier makes positive profits; however in an equilibrium where a leader randomizes, supplier profits must be zero. Our result follows from a generalized indifference principle that mixed strategies must satisfy in economic environments. While supplier profits cannot be approximated, player action profiles are accessible. Our results also apply to games with costly observation.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C73

Keywords: Commitment; Imperfect observation; Mixed strategies; Pricing games

1. Introduction

Thomas Schelling's classic book, *The Strategy of Conflict* (1960), demonstrates that the ability to commit oneself confers a strategic advantage. While Schelling emphasized the value of commitment in military and social situations, his insight has since been formalized and applied in diverse fields of economics including industrial organization, international trade and political economy. The foundations of this literature have been questioned by Bagwell (1995), who argues that the value of commitment is undermined by the slightest imperfection in observation. Bagwell considers a leader-follower model, where the leader's chosen action (or commitment) is observed noisily by the follower. He shows that the pure strategy equilibria of the noisy observation game coincide with the pure strategy equilibria of the simultaneous move game, where the follower does not observe the leader's action. Bagwell interprets this result as saying that the slightest noise completely undermines the leader's ability to commit.

Bagwell's interpretation and his focus on pure strategy equilibria have been questioned. van Damme and Hurkens (1997) analyze games with one leader and one follower and generic payoffs. Such games have a unique backward

[☆] I am grateful to Olivier Compte, Sanjeev Goyal, George Mailath, Klaus Ritzberger, Larry Samuelson, two anonymous referees and various seminar audiences for comments on earlier versions of this paper.

^{*} Fax: +44 207 916 2774.

E-mail address: v.bhaskar@ucl.ac.uk.

induction (or Stackelberg) outcome if observation is perfect. They show that the Stackelberg outcome is always *accessible*—there exists a mixed strategy equilibrium of the game with imperfect observation, whose outcome converges to the Stackelberg outcome, as the noise in observation vanishes.¹ This result has been substantially generalized by Güth et al. (1998), who consider finite leader-follower games with arbitrary numbers of leaders and followers. If payoffs are generic, they show that there always exists a subgame perfect equilibrium outcome of the game with perfect observation that is accessible.² Their proof relies on fundamental properties of generic extensive form games—the existence of a strategically stable set and an essential component. This suggests that accessible outcomes are likely to exist in a large class of games. In particular, if we have a unique subgame perfect equilibrium outcome, this will generically be accessible. This offers an intellectual justification for the fact that applied theory continues to analyze models of commitment without reference to Bagwell's claims.

This paper argues that in games played in a natural economic environment, accessibility fails. We examine leader-follower games that are played in a private contracting environment, where the payoffs to the leaders are influenced by the prices that they must pay for necessary inputs, and where the suppliers of these inputs have some monopoly power.³ Our main finding is that an accessible outcome fails to exist under very general conditions. The failure of accessibility arises since mixed strategy equilibria in these economic environments have to satisfy a generalized indifference principle. The player randomizing between two actions must be indifferent between these actions, as is usual. Furthermore, a supplier who quotes a price to him must also be indifferent between the player's actions, since he can break the player's indifference by a small reduction in price. This implies that such a supplier must be making zero profits in a mixed strategy equilibrium. On the other hand, a supplier who is making a sale at a pure strategy profile will generally make positive profits. This difference in supplier profits between the noisy game and the noiseless game is the reason for the failure of accessibility.

The rest of this paper is organized as follows. Section 2 sets out a simple example that illustrates our basic arguments. Section 3 sets out our general model of a leader-follower game in a contracting environment, when there is perfect observation of the leaders' actions. Section 4 has our main results, when there is noisy observation. Section 5 discusses extensions and the underlying reasons for our results. It shows that our arguments extend to a model of costly but perfect observation. While Sections 3 and 4 set out the formal analysis, the discussion in Sections 2 and 5 is relatively informal. For the sake of smooth exposition, all proofs are in Appendix A.

2. An entry deterrence example

To illustrate Bagwell's argument, consider the entry deterrence game in Fig. 1. The leader is the incumbent firm, and must choose between I and N , i.e. whether or not to invest. The entrant observes this decision and decides whether to stay out or enter. Investment requires the purchase of equipment—assume first that this is available from a competitive market at price p which equals the cost c . Assume that $v - c > 0$, so that the backward induction outcome has the incumbent choosing I while the entrant stays out. Assume also that $v - c < u$, so that the incumbent invests only in order to deter entry.

Now suppose that investment is imperfectly observed. Specifically, the follower observes signals from the set $\{i, n\}$, where the signal is correct (i.e. i after I or n after N) with probability $1 - \varepsilon$, and incorrect with probability ε . For any $\varepsilon > 0$, there cannot be a pure strategy equilibrium where the incumbent invests. If the incumbent invests with probability one, then the follower must believe that I has been chosen when he sees signal n , and will therefore stay out. But then it is optimal for the leader not to invest, since $v - c < u$. The only pure strategy equilibrium is where the incumbent does not invest, and the entrant enters irrespective of his signal.

However, if ε is small, the noisy game has a mixed strategy equilibrium where the incumbent invests with probability θ , where θ is such that the entrant assigns probability 0.5 to investment having taken place when he sees

¹ The noisy observation game will, in general, have multiple equilibria; van Damme and Hurkens use equilibrium selection theory to argue that the mixed equilibrium supporting the Stackelberg outcome is more likely to be played than any other equilibrium. On the other hand, Oechssler and Schlag (2000) use evolutionary dynamics, which typically favor pure strategy equilibria, to select the pure strategy non-Stackelberg equilibrium. Huck and Müller (2000) present experimental evidence showing that the outcome is close to Stackelberg when the noise is small.

² Not all subgame perfect equilibrium outcomes are accessible, even with generic payoffs—see the example in Güth et al. A similar result is obtained in finitely repeated games with imperfect private monitoring—Bhaskar and van Damme (2002) show that efficient equilibrium outcomes under perfect monitoring may not be accessible with imperfect private monitoring.

³ Since prices can be chosen from a continuum, the results of Güth et al. do not apply in our context. See the discussion in Section 5.2.1.




Fig. 1. Entry deterrence game.

the signal n . The entrant stays out if he observes the signal i ; if he observes n , he enters with probability γ . The incumbent's payoff from investing is

$$U(I, \gamma) = v(1 - \gamma\varepsilon) - p. \tag{1}$$

His payoff from not investing is

$$U(N, \gamma) = [1 - \gamma(1 - \varepsilon)]u. \tag{2}$$

Equating these payoffs yields $\gamma = \frac{u-(v-p)}{(1-\varepsilon)u-\varepsilon v} \in (0, 1)$, since $u > v - c$ and $p = c$. Note that $\theta = \frac{1-\varepsilon}{1-2\varepsilon}$, so that the probability of investment converges to 1 as $\varepsilon \rightarrow 0$. Since the entrant stays out whenever he observes i , the outcome (I, OUT) occurs with a probability that tends to one as $\varepsilon \rightarrow 0$. Thus the subgame perfect equilibrium outcome is *accessible*.

We now modify the example as follows. Assume that in order to invest, the incumbent must purchase the required equipment from a monopoly supplier, whose cost of production is c . The game with perfect observation, Γ^C , is as follows. The supplier quotes a price, p , to the incumbent (this implies that the payoffs in Fig. 1 for the incumbent are modified, with p not necessarily equal to c). The price is observed by the incumbent alone, who then chooses his action. The entrant observes the incumbent's action and chooses his own action. This game has a unique perfect Bayesian equilibrium, as follows. The supplier chooses $p = v$. The incumbent chooses I if and only if $p \leq v$, and chooses N otherwise. The entrant stays out if he observes I and enters if he observes N . In this equilibrium the supplier earns a profit of $v - c$ and the action profile (I, OUT) is played with probability one—we call this the Stackelberg outcome. The action profile played in this equilibrium is the same as when the investment good is purchased on a competitive market at price c . However, the supplier earns profits that equal $v - c$, his marginal contribution to the incumbent's payoff at this equilibrium minus the cost of production.

Now let us assume that the investment decision is imperfectly observed. We claim that there is no equilibrium with an outcome that is close to the Stackelberg outcome. Assume that there is an equilibrium where the incumbent invests with high probability, and where the supplier chooses a price $p > c$. As in the original Bagwell argument, the incumbent cannot invest with probability one, since the entrant would not vary his behavior with the signal. Let the probability of investment be α , so that the supplier's payoff is $\alpha(p - c)$. We now show that $\alpha < 1$ implies that p cannot be optimal for the supplier. Note that the incumbent must be indifferent between investing and not investing at a price of p . If the supplier reduces his price, then the incumbent will strictly prefer to invest—this follows from the fact that the price reduction is unobserved by the entrant and cannot affect his behavior. The supplier can therefore earn $(p' - c)$ for any $p' < p$, and this exceeds $\alpha(p - c)$ if p' is sufficiently close to p . Hence there cannot exist an equilibrium where the incumbent randomizes and where the supplier chooses a price strictly greater than c .⁴ Since

⁴ There cannot be an equilibrium where the supplier makes positive profits and randomizes across prices either. Suppose that he chooses p, p' with positive probability, where $p > p'$. At p the incumbent must invest with positive probability, since otherwise the supplier's payoff is zero. The argument in the text implies that at any $p'' < p$, the incumbent must invest with probability one. Thus any p'' such that $p > p'' > p'$ must give the supplier strictly higher profits than p' .

supplier profits equal $v - c$ in the Stackelberg outcome with perfect observability, this outcome is not accessible, since it cannot be approximated in the game with imperfect observation.⁵ This argument only requires that the supplier has *some* monopoly power and makes positive profits in the perfect observation game—one can allow for an alternative supplier of equipment, who has a cost greater than c , and the argument would still apply.

The failure of accessibility applies to supplier payoffs, not player actions in the entry deterrence game. To see this, we construct an equilibrium in the noisy game, where the supplier quotes $p = c$, and the incumbent randomizes, choosing I with probability θ . The entrant's probability of entry on observing n , equates the payoffs in (1)–(2), with $p = c$. If the supplier deviates and chooses $p > c$, the incumbent chooses N for sure. In this equilibrium the incumbent's payoff is approximately $v - c$ when the noise is small, which is strictly greater than his payoff under perfect observability, 0. In other words, if the incumbent retains his commitment power under imperfect observation, he also enhances his power vis-a-vis his supplier, and captures all the surplus. Thus the failure of accessibility applies to *outcomes*, where an outcome is the pair consisting of player action profiles and supplier payoffs.

The inability to use mixed strategies implies that we have a failure of accessibility of outcomes in a larger class of games, even in those where the leader has no incentive to deviate from his subgame perfect equilibrium action in the underlying game. Let us return to our entry deterrence example of Fig. 1, but now assume that $v - c > u$. In the base game Γ , where the investment good is provided at cost, the leader has no incentive to deviate, since I is a best response to OUT. This modification does not affect the analysis of the contracting game with perfect observation, Γ^C . If the incumbent does not invest, the entrant will enter and the incumbent's payoff will be zero. The supplier will therefore be able to extract the payoff difference, so that $p = v - 0$. Now consider the game where the leader's action is observed with noise. In any equilibrium where the incumbent invests for sure, the follower will play OUT irrespective of the signal that he observes. That is, if the leader deviates and does not invest, his payoff will be u . Thus the supplier's price must equal $v - u$, rather than $v - 0$. As long as $u \neq 0$, the supplier's payoffs must differ when we compare pure strategy equilibria in the game with perfect observation and the game with imperfect observation. On the other hand, if the leader randomizes then p must equal c , as we have already demonstrated. Thus we have a failure of accessibility, even when the leader has no incentive to deviate—all that is required generically is that the follower's best response varies with the leader's action.⁶

We now show that the insights contained in these two examples are quite general and apply to a large class of games played in a contracting environment, with arbitrarily many leaders and followers.

3. Perfect observation

We first consider a leader-follower game played in a contracting environment, with perfect observation. This follows the set up in Bhaskar (2005), although the exposition here is self-contained. We will use the term player for someone who plays the game in question, and the term supplier to denote someone with whom a player may need to contract with in order to be able to adopt some action in the game. The set of players, I , is partitioned into the set of leaders, $L = \{1, 2, \dots, m\}$ and the set of followers, $F = \{m + 1, \dots, n\}$. Each player i has a finite action set A_i , whose generic element will also be denoted by a_i^j or a_i . Let $A = \times_{i \in I} A_i$ be the set of action profiles, and let $g_i : A \rightarrow R$ be the gross payoff of player i . We extend, in the usual way, the gross payoff function g_i to correlated action profiles: $g_i(\alpha)$ is the payoff to player i when $\alpha \in \Delta(A)$ is the vector of correlated actions played.

Gross payoffs will in general differ from the net payoffs of a player since she may have to contract with a supplier in order to be able to play the action. If player i pays a price $p(a_i^j)$ in order to take action a_i^j , while her opponents play a_{-i} , her net payoff is given by

$$u_i(a_i^j, a_{-i}, p(a_i^j)) = g_i(a_i^j, a_{-i}) - p(a_i^j). \tag{3}$$

⁵ An equilibrium outcome (i.e. expected supplier profits and a probability distribution over player action profiles) is an element of Euclidean space. If we let $\varepsilon \rightarrow 0$, there does not exist a sequence of equilibria of the associated games $\Gamma^C(\varepsilon)$, the outcomes of which converge to the Stackelberg outcome in the Euclidean metric.

⁶ This argument requires that the efficient supplier has sufficient monopoly power over the action I , with any alternative supplier having a cost that is greater than $\min\{v, v - u\}$. In the case where $v - c < u$, we only required that the efficient supplier has some monopoly power—an alternative supplier could have any cost strictly greater than c .

A player seeks to maximize her net payoff. In order to simplify exposition, we shall assume that only leaders need contract with a supplier—followers need not do so, and their net payoffs equal gross payoffs. We refer the reader to Bhaskar (2008a), which shows that all the results also apply when followers also need to contract with their suppliers.

We now turn to the market structure on the supplier side. Let $\bar{A}_i \subset A_i$ be the set of actions for which leader i needs a supplier. First, we assume that there are no complementary inputs, so that for any player i and any action a_i^j , no more than one supplier is *required* in order to take this action. For action $a_i^j \in \bar{A}_i$, let $\Lambda(a_i^j)$ denote the set of competing suppliers, which without loss of generality may be assumed to have exactly two members, $\lambda_1(a_i^j)$ and $\lambda_2(a_i^j)$. Let the costs of supply of these suppliers be c_{ij1} and c_{ij2} respectively, where $c_{ij1} < c_{ij2}$. Let B be an upper bound on the gross payoffs in the game, across players and actions. The efficient supplier has full monopoly power when $c_{ij2} = B$. He has limited market power when $c_{ij2} < B$. Note that the case of perfect competition is implicitly covered, this being analytically equivalent to the situation where no supplier is required and the input for the action is supplied at cost. The payoff to a supplier is given by the difference between the price he receives and his cost if he makes a sale, and is zero otherwise.

Let us normalize prices and gross payoffs by measuring them net of the cost of supply of the efficient supplier (equal to c_{ij1} for action a_i^j), so that a zero price corresponds to pricing at minimum cost. Henceforth, the gross payoff $g_i(a_i, a_{-i})$ will denote the payoff when the player pays the minimum cost of action a_i . Let $\Lambda_i = \{\Lambda(a_i^j)\}_{a_i^j \in \bar{A}_i}$ denote the set of suppliers for player i . Let p_{ijh} denote the price which is charged by supplier $\lambda_h(a_i^j)$ for action a_i^j , let $p_i^j = (p_{ijh})_{\lambda_h(a_i^j) \in \Lambda(a_i^j)}$, and let $\mathbf{p}_i = (p_i^j)_{a_i^j \in \bar{A}_i}$.

The leader-follower game with private contracts, Γ^C , is as follows:

1. Each supplier in $\Lambda_L = \bigcup_{i \in L} \Lambda_i$ quotes a price for the action that he supplies.
2. Each leader $i \in L$ observes the price vector \mathbf{p}_i (but not the prices quoted to other players), and chooses an action, and a supplier for the action.
3. Each follower $i \in F$ observes a_L and followers simultaneously choose actions.

We make the following assumptions regarding the game Γ^C .

Assumption A.1. For every player i , there exists an action a_i^0 such that no input is required to play this action.

Assumption A.2. Each supplier supplies a single action of a single player.

Since we have a continuum extensive form game, where agents have infinite strategy sets, we will be explicit in our definition of equilibrium. We focus on Nash equilibria that satisfy sequential rationality. Sequential rationality implies two conditions. First, at any leader action profile a_L , the followers choose actions optimally. Second, at any profile of supplier prices, \mathbf{p}_i ; leader i chooses his action optimally. This requires that we specify the leader's beliefs at out of equilibrium prices. The only beliefs that are directly payoff relevant for leader i are beliefs regarding the action profile played by the other leaders. We assume that a leader's beliefs are *invariant*, i.e. they are the same as his equilibrium beliefs, at any price vector \mathbf{p}_i . This restriction follows naturally from Assumption A.2 (no supplier supplies more than one leader), from the assumption that only leader i observes this deviation by his supplier, and from the fact that different suppliers choose their prices independently. Hence, if i 's supplier deviates, i continues to believe that j 's suppliers have chosen their equilibrium prices, and does not change his beliefs regarding j 's actions. In addition, we rule out equilibria where an inactive supplier (i.e. one who does not make a sale) chooses a price below his cost. Such equilibria are called *cautious*, and can be ruled out by considerations of trembling hand perfection.⁷ We shall call our solution concept perfect Bayesian equilibrium, or equilibrium for short.

⁷ If we discretize the price space, and focus on sequential equilibria, a player's beliefs would be invariant at all price profiles. To justify caution, we may restrict attention to equilibria of the continuum game which are limit points of a sequence of trembling hand equilibria of discrete games, as the grid of prices becomes finer. Equilibria with negative prices will not be a limit of such trembling hand perfect equilibria.

A pure strategy for a supplier $\lambda_h(a_i^j)$ is a price $p_{ijh} \in R_+$. A mixed strategy is a probability measure μ_{ijh} on R_+ . A strategy for leader i is a map $\sigma_i : R^{|A_i|} \rightarrow \Delta(A_i)$.⁸ A strategy for follower j is a map $\sigma_j : A_L \rightarrow \Delta(A_j)$. A strategy profile τ is a collection $((\mu_{ijh})_{\lambda_h(a_i^j) \in A_L}, (\sigma_i)_{i \in L}, (\sigma_j)_{j \in F})$.

In usual terminology, the outcome of a strategy profile τ is the induced distribution over the terminal nodes of the game tree. In price setting games, the set of equilibrium outcomes usually contains considerable redundancy. For example, in the case of Bertrand competition between three firms with differing unit costs, the price set by the highest cost firm is irrelevant, and can therefore be chosen arbitrarily. It will therefore be more useful to focus on outcomes more coarsely defined. The *action outcome* associated with a strategy profile τ is the induced distribution over the set of player-action profiles, A . The supplier payoffs associated with τ is the vector of payoffs to the suppliers under this profile. For the purposes of this paper, the *outcome* of a strategy profile is defined as the pair consisting of the action outcome and the supplier payoffs, and is an element of Euclidean space. Failure of accessibility in terms of our definition of outcome will also entail failure in terms of a more refined definition, such as the distribution over terminal nodes.

Our results will relate the equilibrium action outcomes in Γ^C (the game in a contracting environment) to those in the leader-follower game Γ , where supplier prices are exogenously fixed at zero and players' net payoffs equal their gross payoffs. A strategy for a leader in Γ is a mixed action $\alpha_i \in \Delta(A_i)$, while a follower's strategy is a map $\beta_j : A_L \rightarrow \Delta(A_j)$. Let E^Γ denote the set of subgame perfect equilibria of Γ . The outcome of a strategy profile (α_L, β) is the induced distribution over the elements of A . Let $\Omega^\Gamma \subset \Delta(A)$ denote the set of subgame perfect equilibrium outcomes of the game Γ . Let $\alpha_L = (\alpha_i)_{i \in L}$, $\beta = (\beta_j)_{j \in F}$, and let $(\alpha_L, \beta) \in E^\Gamma$. Given a vector $a = (a_i)_{i=1}^n$, we use the notation $a \setminus a_i'$ to denote the vector that results when the i th component a_i is replaced by a_i' . For a leader $i \in L$ who plays a pure action \hat{a}_i as part of the equilibrium (α_L, β) , we define his deviation loss $\delta_i(\alpha_L, \beta)$ as:

$$\delta_i(\alpha_L, \beta) = g_i(\alpha_L, \beta(a_L)) - \max_{a_i \neq \hat{a}_i} g_i(a_L \setminus a_i, \beta(a_L \setminus a_i)). \tag{4}$$

If the leader plays a mixed action, $\delta_i(\alpha_L, \beta) = 0$.

The *marginal contribution* of an active supplier $\lambda_1(a_i^j)$ at the profile (a_L, β) equals $\min\{\delta_i(a_L, \beta), c_{ij2}\}$. This is the gross payoff loss that player i suffers from not choosing this supplier and choosing the next best alternative, which could be either purchasing from the inefficient supplier or choosing a different action.

Let $\Omega^{\Gamma^C} \subset \Delta(A)$ denote the set of equilibrium action outcomes of the game Γ^C .

Theorem 1. $\Omega^{\Gamma^C} = \Omega^\Gamma$, i.e. the cautious perfect Bayesian equilibrium action outcomes of Γ^C coincide with the subgame perfect equilibrium outcomes of Γ . In any equilibrium of Γ^C with action outcome $(a_L, \beta(a_L))$, an active supplier for leader i earns his marginal contribution $\min\{\delta_i(a_L, \beta), c_{ij2}\}$.

This establishes that given any subgame perfect equilibrium of the game Γ , there is associated a unique outcome ω of the game Γ^C , since the supplier payoffs are uniquely defined by their marginal contributions. The following corollary is immediate:

Corollary 2. Any one-leader one-follower game Γ^C with generic payoffs has a unique equilibrium outcome. If a supplier is required for the leader to take his equilibrium action, this supplier makes strictly positive profits.

4. The noisy leader-follower game

We now assume that leaders' actions are observed with some noise: given that leader i chooses $a_i^k \in A_i$, nature chooses a signal $a_i^h \in A_L$ with probability π_{hk}^i , where the signals are chosen independently across leaders. This defines a stochastic matrix π on $A_L \times A_L$, where A_L denotes the set of action profiles of leaders. Let Π be the set of possible signal structures that satisfy the independence assumption, and let $Int(\Pi)$ be the subset of signal structures with full support, i.e. the set of stochastic matrices π such that every entry is strictly positive. Let π_0 denote the identity matrix—this corresponds to perfect observation. The game $\Gamma^C(\pi)$ is defined as follows:

⁸ To economize on notation, we do not specify explicitly a player's choice of supplier in her strategy, assuming that she always chooses the lowest price for any action.

1. Each supplier in A_L quotes a price for the action he supplies.
2. Each leader $i \in L$ observes the price vector \mathbf{p}_i , and chooses an action and a supplier for that action.
3. Given the action profile chosen by the leaders, nature chooses a signal in A_L , according to the stochastic matrix π .
4. Each follower $i \in F$ observes the signal chosen by nature and followers simultaneously choose actions.

The gross payoffs to players depend only on action profile realized, and not upon the signal. As before, net payoffs to leaders are equal to gross payoffs minus the prices paid.

Note that all agents have exactly the same strategy sets in the games Γ^C and $\Gamma^C(\pi)$ —only the payoffs associated with strategy profiles differ in the two games. We restrict attention to cautious perfect Bayesian equilibria of the game $\Gamma^C(\pi)$. Fix an equilibrium of the game $\Gamma^C(\pi)$. The action outcome of the equilibrium is the element of $\Delta(A)$ induced by the equilibrium. The outcome of the equilibrium is the pair consisting of the action outcome and the profile of supplier payoffs. Let $\mathcal{E}(\pi)$ denote the set of equilibrium outcomes of $\Gamma^C(\pi)$, and let $\mathcal{E}(0)$ denote the set of equilibrium outcomes of Γ^C . Expected supplier payoffs are real numbers, while the player action outcomes are probability distributions over a finite set. Thus $\mathcal{E}(\pi) \subset R^{|A_L|} \times \Delta(A)$, a subset of Euclidean space, and we may use the usual norm in order to define convergence. We say that a sequence $\omega_n \rightarrow \omega$ if this convergence is in the usual topology.

Definition 3. $\omega \in \mathcal{E}(0)$ is accessible if \exists countable sequences $(\pi, \Gamma^C(\pi), \omega(\pi)), \omega(\pi) \in \mathcal{E}(\pi), \pi \in \text{Int}(\Pi), \pi \rightarrow \pi_0$ such that $\Gamma^C(\pi) \rightarrow \Gamma^C(\pi_0)$ and $\omega(\pi) \rightarrow \omega$.

Note that this is a weak definition of accessibility, since it only requires the existence of *some* sequence π converging to π_0 . Since our main result shows the failure of accessibility, it would also hold under a stronger definition.

4.1. The main result

Consider a subgame perfect equilibrium (a_L^*, β) of the base game Γ , where $a_L^* \in A_L$ is a pure action profile, and β denotes the profile of the followers' equilibrium strategies. Denote the outcome of this equilibrium by (a_L^*, α_F^*) , where $\alpha_F^* = \beta(a_L^*)$, the followers' equilibrium response to a_L^* .

Definition 4. Leader $i \in L$ has an incentive to deviate at (a_L^*, α_F^*) if

$$\max_{a_i} g_i(a_L^* \setminus a_i, \alpha_F^*) > g_i(a_L^*, \alpha_F^*).$$

A leader has an incentive to deviate at a subgame perfect equilibrium if he can increase his payoff, given the choices of other leaders and given that followers do not respond to this deviation. In standard leader-follower games, the question of commitment is only relevant if the leader has an incentive to deviate.

Definition 5. Leader i 's action is relevant at (a_L^*, β) if

$$\max_{a_i \neq a_i^*} g_i(a_L^* \setminus a_i, \alpha_F^*) \neq \max_{a_i \neq a_i^*} [g_i(a_L^* \setminus a_i, \beta(a_L^* \setminus a_i))].$$

Let $\hat{a}_i \in \arg \max_{a_i \neq a_i^*} g_i(a_L^* \setminus a_i, \beta(a_L^* \setminus a_i))$. If $\beta(a_L^* \setminus \hat{a}_i)$ differs from α_F^* , then leader i 's action will be relevant (for the followers) at (a_L^*, β) , provided that payoffs are generic. In other words, any game where some follower's best response depends upon whether leader i chooses his commitment action a_i^* or deviates (optimally) from this, leader i 's action will be relevant. Clearly, in any game where a leader has an incentive to deviate, his action will be relevant, but the converse is not true. In the entry deterrence example, the leader has an incentive to deviate if $v - c > u$; however, generically (i.e. as long as $u \neq 0$), the leader's action is always relevant, since the follower's best responses to I and N differ.

Theorem 6. Let (a_L^*, β) be a subgame perfect equilibrium of the base leader-follower game Γ where the subgame following a_L^* has a unique equilibrium α_F^* . Let Γ^C be the associated game in a contracting environment, and let ω^* denote the equilibrium outcome of Γ^C with action outcome (a_L^*, α_F^*) . Suppose that there exists a leader $i \in L$

such that either (a) i has an incentive to deviate at (a_L^*, α_F^*) and needs to contract with a supplier to take action a_i^* , or (b) i 's action is relevant at (a_L^*, β) and there is a monopoly supplier for action a_i^* . Then the outcome ω^* is not accessible.

4.2. Approximating action profiles

We now show that it is possible to approximate equilibrium action profiles taken by the players in the game, even if one cannot approximate suppliers' payoffs. Let us consider games Γ with one leader, player 1, and one follower, player 2. We assume that the gross payoffs in the game satisfy the following genericity assumption:

Assumption A.3. For any $a, a' \in A$, $a \neq a'$, $g_1(a) \neq g_1(a')$ and $g_2(a) \neq g_2(a')$.

A.3 implies that the game Γ has a unique subgame perfect equilibrium, the outcome of which we denote by $a^* = (a_1^*, a_2^*)$. By Theorem 1, the associated game Γ^C played in a contracting environment has a unique equilibrium action outcome $\{a^*\}$. The following theorem shows that this equilibrium action outcome is accessible in a strong sense, since the result holds for any sequence π converging to π_0 .

Theorem 7. Let Γ be a one-leader one-follower game that satisfies A.3, with subgame perfect equilibrium outcome a^* . For any countable sequence $\pi \rightarrow \pi_0$, $\pi \in \text{Int}(\Pi)$ and associated sequence of noisy games $\Gamma^C(\pi)$, there exists a sequence $\tilde{\alpha}(\pi) \in \Omega(\Gamma^C(\pi))$ such that $\tilde{\alpha}(\pi) \rightarrow a^*$.

It might be argued that in games played in a contracting environment, imperfect observation has distributional consequences, but has no implications for the actions that are taken. In our view this is not an appropriate interpretation: the payoffs to suppliers will have incentive effects and will therefore affect outcomes in a broader sense. For example, if supplier entry decisions were included in the analysis, these would be influenced by the profits they make. More generally, if one studies economic environments which combine pricing and commitment, outcomes may be very sensitive to the monitoring structure. Bhaskar (2008b) studies a model of dynamic price competition and finds that economic outcomes are very different when one compares perfect monitoring with monitoring that is slightly imperfect.

5. Extensions and robustness

We now turn to some extensions and examine the robustness of our results. In order to simplify exposition, we shall focus throughout this section on extensions and variations on our entry deterrence example.

5.1. Costly observation⁹

Our results on the failure of accessibility also extend to the case where observation is costly, as in Várdy (2004). Várdy considers a model with one leader and one follower, where the follower must choose whether or not to observe the leader's action. Observation entails a small cost, but is perfect. Várdy shows that the Stackelberg action profile of the game where the follower automatically and costlessly observes the leader's action can be approximated by an equilibrium distribution over action profiles of the game with costly observation when the observation cost is small.

Let us now consider the implications of costly observation in the context of our entry deterrence game, where we assume $v - c < u$, so that the leader has an incentive to deviate. There cannot be an equilibrium where the leader invests with probability one, since then it is not optimal for the follower to pay the observation cost. So the commitment equilibrium must be in mixed strategies, where the leader invests with probability less than one, and where the follower randomizes between observing and not observing. Indeed, one can construct an equilibrium, where the probability of

⁹ I am indebted to a referee for suggesting this extension.

the outcome (I , OUT) (the leader investing and the follower staying out) converges to one as the observation cost tends to zero.¹⁰

Now let us embed this game in a contracting environment, where the leader needs to purchase equipment from a monopoly supplier in order to invest. Since the follower's observation decision cannot depend upon the price, at the equilibrium price (p) chosen by the supplier, the leader is indifferent between I and N . Suppose now that the supplier reduces his price slightly. This breaks the leader's indifference, and the leader will invest for sure. Therefore, as long as $p > c$, the supplier has an incentive to reduce the price. The only equilibrium price consistent with leader randomization is $p = c$, and supplier payoffs cannot be positive when observation is costly.

This result can be generalized to arbitrary one-leader one-follower games with generic payoffs played in a contracting environment. Specifically, the equilibrium outcome when the follower costlessly observes the leader's action cannot be approximated by equilibrium outcomes of games with costly observation, as the cost tends to zero. This result obtains under precisely the same conditions as our main Theorem 6. Accessibility fails if the leader has an incentive to deviate at the equilibrium action profile and he needs to contract with a supplier to take this action. It also fails if his action is relevant for the follower (as in Definition 5) and there is a monopoly supplier for this action. We refer the reader to Bhaskar (2008a) for a precise statement and proof of this result.

This result is related to Morgan and Várdy (2007), who analyze a leader-follower game with continuum action sets and strictly concave payoff functions. They show that the subgame perfect equilibrium outcome cannot be approximated in the game with costly observation. This arises since the leader's payoff function is strictly concave and has a unique maximizer, so that randomization becomes impossible. Thus their conclusions are similar to ours, but they follow from considerations that are quite distinct.

5.2. Robustness

We now consider the role of our assumptions in explaining the failure of accessibility. These include the fact that prices can be chosen from an infinite set, that supplier prices are private, and that there is complete information about payoffs.

A key assumption is that contracting is private. In the context of the entry deterrence example, this implies that the price quoted by the supplier to the incumbent is not observed by the entrant. If contracting is public so that the price p is also observed by the entrant, the probability of entry can be conditioned upon p , so as to make the incumbent indifferent for any value of p . This removes the supplier's incentive to undercut, and one may construct equilibria that restore accessibility in this example. Public contracting throws up other interesting possibilities when there is competition between rival suppliers in a mixed strategy equilibrium, and we refer the interested reader to Bhaskar (2005) for an analysis of games with public contracting. In our view, private contracting is a more reasonable assumption than public contracting, since the latter is only viable if the supplier is unable to offer secret price discounts.

5.2.1. Discrete prices

Games played in a contracting environment are clearly not generic extensive form games, in the sense of Güth et al. (1998). First, the game is not finite, since prices are chosen from a continuum. Second, payoffs are assumed to be quasi-linear, and so even in a discretized version of the game, the set of feasible payoffs does not have full dimension. Let us discuss these issues in the context of the entry deterrence example, where we assume $v - c < u$. Quasi-linearity is not really critical, since one can relax this assumption. Our argument only requires that the leader's payoff is strictly decreasing in the price that he pays, so that a small reduction in the price suffices to break his indifference. To explore the role of the continuum assumption, let us now consider a discrete price grid. The supplier must choose a price p_i from the set $\{\frac{k}{m}: k \in \mathbb{N} \cup \{0\}\}$, where $m \in \mathbb{N}$ indexes the fineness of the price grid. The noisy game Γ^C can now be parameterized by the pair (ε, m) where ε is the noise in observation. It will be convenient to assume that the payoff v in Fig. 1 is an irrational number. Let $p^*(m)$ be the largest price that is less than v , and let the grid be sufficiently fine so that $p^*(m) > c$. With perfect observation, the unique perfect Bayesian equilibrium outcome has the supplier choosing $p^*(m)$, the incumbent investing at this price, and the entrant staying out, so that supplier profits equal $p^*(m) - c$.

Now let us consider the situation where commitment is imperfectly observed. Fix an equilibrium of the game $\Gamma^C(\varepsilon, m)$, where I is played by the incumbent with positive probability, and let $\hat{p}(m)$ be the largest price that is

¹⁰ For experimental evidence on this type of game, see Morgan and Várdy (2004).

chosen by the supplier in this equilibrium. Clearly, at $\hat{p}(m)$ the incumbent must buy with positive probability, since otherwise \hat{p} will not be chosen by the supplier. Since it is optimal for the incumbent to buy at \hat{p} , it must be strictly optimal to buy at any price strictly below \hat{p} . Hence the supplier can ensure the payoff of $\hat{p}(m) - \frac{1}{m}$ by choosing the price $\hat{p}(m) - \frac{1}{m}$. If α is the probability that the incumbent invests at price \hat{p} , we must have that

$$\alpha \hat{p}(m) \geq \hat{p}(m) - \frac{1}{m}. \tag{5}$$

On the other hand, the probability of investment must equal $\theta = \frac{1-2\varepsilon}{1-\varepsilon}$ so as to provide incentives for the entrant, implying that

$$\frac{1-2\varepsilon}{1-\varepsilon} \geq \alpha \geq 1 - \frac{1}{\hat{p}(m)m}. \tag{6}$$

This gives us the condition:

$$\hat{p}(m) \leq \frac{1-\varepsilon}{\varepsilon m}. \tag{7}$$

The right-hand side of the inequality (7) provides an upper bound on the price that can be charged by the supplier. This bound is independent of v and tends to zero for fixed ε as m tends to infinity. The bound is relevant provided that it is smaller than $p^*(m)$. It is easy to verify that:

1. For any $\varepsilon > 0$, there exists $m^*(\varepsilon)$ such that if we consider the sequence $(\varepsilon, m(\varepsilon))$ where $\varepsilon \rightarrow 0$ and $m(\varepsilon) \geq m^*(\varepsilon)$, the Stackelberg outcome is not accessible.
2. If we fix m , and consider the sequence $\Gamma^C(\varepsilon, m)$ where $\varepsilon \rightarrow 0$, the Stackelberg outcome is accessible.¹¹

As an example, let $c = 0$, and $v = \$100 + \xi$, where ξ is a tiny irrational number. The supplier makes a profit of \$100 when commitment is perfectly observed. In the game with noisy observation, suppose that the probability that the incumbent does not invest must equal ε in order to make the entrant indifferent between entering and staying out when she observes n . If the price grid is in cents, then the supplier will have no incentive to reduce price below \$100 only if $(1 - \varepsilon)100 \geq 99.99$, i.e. ε must be smaller than 0.0001. In other words, if the profits that the supplier makes are large relative to the minimum unit of account, the noise must be very small indeed.

5.2.2. Incomplete information and purification

One comment, which has often been made, is that a small amount of incomplete information regarding payoffs will smooth away the discontinuities associated with Bertrand pricing, and restore accessibility. This intuition is not valid; if the uncertainty about payoffs is small relative to the noise in observation, the failure of accessibility persists. To consider this question most simply, let us modify the payoffs of the incumbent in the entry deterrence game of Fig. 1 by augmenting the payoff to action I by ξz , where ξ is a small positive number, and z is the realization of a random variable which is distributed with density f on support $[0, 1]$. Assume $f(\cdot) > m > 0$ on its support,¹² and that only the incumbent observes the realization of z , so that we have a private payoff shock as in Harsanyi (1973).

Consider first the game where the entrant perfectly observes the incumbent's choice of action. Given a price p , the incumbent's payoff from investing is given by $v + \xi z - p$. The marginal type of incumbent, who is indifferent between investing and not, is given by $z^* = \frac{p-v}{\xi}$. Thus the payoff of the supplier from a price p equals

$$\pi(p) = (p - c) \cdot \left[1 - F\left(\frac{p-v}{\xi}\right) \right]. \tag{8}$$

The derivative of profits with respect to p equals

$$\pi'(p) = \left[1 - F\left(\frac{p-v}{\xi}\right) \right] - \frac{(p-c)f\left(\frac{p-v}{\xi}\right)}{\xi}, \tag{9}$$

¹¹ This result is not implied by the theorem of Güth et al. (1998). With quasi-linear payoffs, one cannot appeal to existence results that apply for generic extensive form games.

¹² The lower bound on $f(\cdot)$ is not essential for our argument, but simplifies its exposition.

which is strictly negative if ξ is sufficiently small, as long as the incumbent invests with probability less than one. Therefore the supplier must choose price equal to v , so that the incumbent invests with probability one.

Let us now suppose that commitment is imperfectly observed. Consider an equilibrium where the supplier chooses a price p^* , and the incumbent chooses to invest for some realizations of his private information, and chooses not to invest for other realizations. The entrant stays out when he observes the signal i , and enters with probability γ on observing the signal n . Since the entrant's payoffs are not perturbed, the probability that the incumbent invests must equal θ . Letting $z(p)$ denote the payoff realization of the type of incumbent who is indifferent between the two actions at an arbitrary price p , the indifference condition for $z(p)$ is

$$(1 - \gamma\varepsilon)v - p + \xi z(p) = [(1 - \gamma)(1 - \varepsilon) + \varepsilon]v. \tag{10}$$

Supplier profits, as a function of p , are given by

$$\pi(p) = (p - c)[1 - F(z(p))]. \tag{11}$$

The derivative of profits is given by

$$\pi'(p) = [1 - F(z(p))] - \frac{(p - c)f(\cdot)}{\xi}, \tag{12}$$

which is strictly negative for ξ sufficiently small, at any interior value of z . Since z must be interior if the incumbent does not invest with positive probability, the equilibrium price p^* cannot be strictly greater than c if ξ is sufficiently small. That is, for a given value of ε , supplier payoffs must converge to zero as $\xi \rightarrow 0$, and small payoff uncertainty, in the manner of Harsanyi, does not restore accessibility. Note however that if payoff uncertainty is large relative to ε (as in Maggi, 1999), then the leader will take his non-commitment action with sufficiently high probability, so that the follower's beliefs will be sensitive to the signal.¹³

5.3. Concluding comments

We have demonstrated that in leader-follower games played in a contracting environment, subgame perfect equilibrium outcomes are not accessible under fairly general conditions, when suppliers have some monopoly power. The basic intuition is as follows. Under perfect observation, the monopoly power of suppliers ensures that they make positive profits at the commitment equilibrium. When there is imperfect observation, a leader may have an incentive to deviate from his commitment action. To provide incentives for the followers to respond appropriately to this deviation, the leader needs to randomize, and must choose his commitment action with probability less than one. However, if a leader randomizes, then his supplier must also be indifferent between his actions, since otherwise the supplier would have an incentive to reduce his price slightly to ensure that the leader takes the commitment action for sure. Thus the supplier cannot make positive profits in any such mixed equilibrium, ensuring a failure of accessibility.

The present paper has set out a very specific class of games—leader-follower games played in a contracting environment—in order to make this argument precise. However, it is clear that the essential intuition is likely to extend, to contexts where the payoffs to agents who undertake commitment actions are influenced by prices set by other agents. The essential assumptions are that the prices are privately quoted, and are chosen from a continuum. Thus in many economic contexts, one may have a failure of accessibility. Bhaskar (2008b) provides an analysis of dynamic interaction between a strategic buyer and competing sellers, where the buyer may influence future competition by her current purchase decisions. This provides a natural context where the buyer's commitment decisions (i.e. current purchases) are influenced by prices. The finding here is that the set of equilibrium outcomes when the buyer's purchase decisions are perfectly observed is disjoint from the set of outcomes when purchase decisions are imperfectly observed, reinforcing the point of the present paper. The underlying reason is that mixed equilibria require double indifference, for the agent who randomizes and also for the agent who sets a price that influences the former's decision. These results suggest that Bagwell's point, that one should be cautious in focusing upon commitment effects under perfect observation, appears to be valid when we consider games played in economic environments. Many economic applications fall into this category, since the payoffs to agents who choose actions are affected by the prices set by other agents.

¹³ Similar arguments apply when there is private information about the costs of the supplier, c . Large uncertainty about supplier costs will allow accessibility, while small uncertainty implies a failure of accessibility of supplier payoffs.

Appendix A

Proof of Theorem 1. Let μ be a probability distribution on A_L , the set of leader action profiles. Let $G(\mu)$ be the normal form game between the followers in stage 2 where it is common knowledge that the leaders' actions are chosen according to μ , and let $G^C(\mu)$ denote the second stage of the (noisy or noiseless) game Γ^C , when it is common knowledge among the followers that the distribution over leader action profiles is μ . It is immediate that there is an equivalence between the equilibria of $G(\mu)$ and $G^C(\mu)$ —although the prices chosen by the suppliers in stage 1 are not observed by the followers, these prices are payoff irrelevant. With a slight abuse of notation, let $G(a_L)$ denote the game where μ assigns probability one to a_L and let $E(a_L)$ denote the set of Nash equilibria of $G(a_L)$.

We now proceed to the first stage of the game Γ^C , the leader-follower game in a contracting environment. Given any action profile a_L chosen by the leaders, gross payoffs to any leader i are given by $g_i(a_L, \beta(a_L))$ where $\beta(a_L) \in E(a_L)$. Thus any equilibrium strategy profile of the followers defines a strategic form game for the leaders played in a contracting environment. Bhaskar (2005) provides an analysis of this class of game. Specifically, if G^C denotes a normal form game G between a set of players in a contracting environment, Theorem 1 in Bhaskar (2005) establishes that: (a) the set of distributions over player action profiles in G^C coincides with the set of equilibria of G , and (b) if an equilibrium action profile α is played in G^C , each supplier gets his marginal contribution to the player's payoff at α . Applying this theorem to the present context, there is an equivalence between equilibrium distributions over leaders' actions in Γ^C and Nash equilibria of this strategic form game. Finally, the payoff loss to any leader from choosing his best deviant action or an alternative supplier is $\min\{\delta_i(a_L, \beta), c_{ij2}\}$, which must equal the payoff to each active supplier to a leader. \square

Proof of Theorem 6. By Theorem 1, the game Γ^C has an equilibrium with action outcome (a_L^*, α_F^*) , where the price charged by supplier $\lambda_1(a_i^*)$, $p_1(a_i^*) = \delta_i(a_L^*, \beta)$ if this supplier is a monopolist, and strictly positive otherwise. Let ω^* denote the outcome of this equilibrium.

Since $G(a_L^*)$ has a unique equilibrium, α_F^* , $G(\mu)$ also has a unique equilibrium that is close to α_F^* provided that μ is close to a_L^* (i.e. μ assigns probability close to one to a_L^*). This follows from the fact that the Nash equilibrium correspondence has a closed graph, and the uniqueness of equilibrium in $G(a_L^*)$. If the weak inequalities required for Nash equilibrium are violated for any other profile α' in $G(a_L^*)$, they will also be violated at α' in $G(\mu)$ if μ is sufficiently close to a_L^* . Therefore, in stage 2 of $\Gamma^C(\pi)$, the actions of the followers must be close to α_F^* for beliefs μ that are close to a_L^* .

Fix (η, δ) such that $\pi \in \text{Int}(\Pi)$ is η -close to π_0 and the outcome $\omega(\pi)$ of the equilibrium $\tau(\pi)$ is δ -close to ω^* . Let $\mu(a_L^*)$ denote the common beliefs in stage 2, given that the signal profile a_L^* is observed. For any $\varepsilon > 0$, we can ensure that $\tilde{\alpha}_F(a_L^*)$ is ε -close to α_F^* by choosing (η, δ) that are sufficiently small, since then $\mu(a_L^*)$ will be sufficiently close to a_L^* .

We consider first the case where leader i plays a_i^* for sure in $\tau(\pi)$. Let $a_L^* \setminus a_i$ denote the signal profile which differs from a_L^* only in the i th component, with a_i replacing a_i^* . $\mu(a_L^* \setminus a_i) = \mu(a_L^*)$ since nature chooses the components of the signal profile independently, and since the followers believe that i has chosen a_i^* irrespective of the signal that is observed. Since equilibrium is unique at belief $\mu(a_L^*)$, the followers must play $\tilde{\alpha}_F(a_L^*)$ also at the signal $a_L^* \setminus a_i$. If condition (a) holds, leader i has an incentive to deviate at the profile (a_L^*, α_F^*) , and will also have an incentive to deviate from a_i^* in $\Gamma(\pi)$ when ε is sufficiently small, since the follower's behavior does not depend upon his action a_i when the signal profile $a_L^* \setminus a_i$ is realized. If condition (b) holds, then in the noiseless game, supplier $\lambda_1(a_i^*)$ must earn $\delta_i(a_L^*, \beta)$. Let $\gamma_i(a_L^*, \alpha_F^*) = g_i(a_L^*, \alpha_F^*) - \max_{a_i \neq a_i^*} g_i(a_L^* \setminus a_i, \alpha_F^*)$ denote the deviation loss of the leader in the simultaneous move game, where the follower has no observation of the leader's action. If leader i 's action is relevant, $\gamma_i(a_L^*, \alpha_F^*) \neq \delta_i(a_L^*, \beta)$. So if $\gamma_i(a_L^*, \alpha_F^*) < \delta_i(a_L^*, \beta)$ and $p_1(a_i^*) > \gamma_i(a_L^*, \alpha_F^*)$ and ε is sufficiently close to zero, it is optimal for the leader to deviate and choose the action $a_1 \in \arg \max_{a_1 \neq a_1^*} g_1(a_1, a_2^*)$ since the followers must play $\tilde{\alpha}_F(a_L^*)$ after every signal $(a_L^* \setminus a_i)$. Hence the payoff of supplier $\lambda_1(a_i^*)$ cannot be close to $\delta_i(a_L^*, \beta)$. Conversely, if $\gamma_i(a_L^*, \alpha_F^*) > \delta_i(a_L^*, \beta)$, $\lambda_1(a_i^*)$ can increase his price and it will still be optimal for the leader to buy. Thus we cannot have an equilibrium where a_i^* is played with probability one and where supplier's price $p_1(a_i^*)$ is close to $\delta_i(a_L^*, \beta)$.

Finally, we consider the case where either leader i or supplier $\lambda_1(a_i^*)$ randomize. If leader i randomizes, $p_1(a_i^*) = 0$, since otherwise the supplier can reduce his price slightly and ensure that the leader plays a_i^* for sure.

If the supplier randomizes across two prices and makes positive profits, then the lower price cannot be optimal, since the leader must play a_i^* for sure at an intermediate price. Since supplier profits must equal zero at any equilibrium with randomization, the outcome cannot be close to ω^* . \square

Proof of Theorem 7. We consider two separate cases, depending upon whether the leader has an incentive to deviate or not. Suppose that the leader has no incentive to deviate, so that $a_1^* \in \arg \max_{a_1} g_1(a_1, a_2^*)$. In the noisy game, let supplier $\lambda_1(a_1^*)$ price at $\min\{\gamma_1(a_1^*, a_2^*), c_{12}^*\}$ (if no supplier is needed for a_1^* , set this price equal to zero), and let suppliers for other actions choose a price of zero. Let the leader choose a_1^* , and let the follower choose the continuation strategy in Γ^C (the noiseless game) that follows the play of a_1^* , regardless of the signal that is observed. This is clearly optimal given that a_1^* is played with probability one by the leader, since $\pi \in \text{Int}(\Pi)$. Given the follower's behavior, it is optimal for the leader to play a_1^* , since $\max_{a_1 \neq a_1^*} g_1(a_1, a_2^*) = g_1(a_1^*, a_2^*) - \gamma_1(a_1^*, a_2^*) \leq g_1(a_1^*, a_2^*) - p_1(a_1^*)$. Thus a^* is an equilibrium action outcome of the noisy game.

Suppose now that the leader has an incentive to deviate at a^* . From van Damme and Hurkens (1997), we know that if A.3 is satisfied, there exists countable sequences $(\pi, \Gamma(\pi))$, where $\pi \rightarrow \pi_0$, $\pi \in \text{Int}(\Pi)$, such that in each $\Gamma(\pi)$, there exists an equilibrium $(\alpha_1(\pi), \beta_2(\pi))$, where the outcomes of this sequence of equilibria converge to a^* . For any π in this sequence, we shall construct an equilibrium $\tau(\pi)$ in $\Gamma^C(\pi)$, the noisy game played in a contracting environment, with the property that $\tau(\pi)$ induces the same behavior by the players as $(\alpha_1(\pi), \beta_2(\pi))$. Consider $\Gamma^C(\pi)$ and for any signal a_1 , let the follower's strategy $\sigma_2(a_1, \pi)$ be such that $\sigma_2(a_1, \pi) = \beta_2(a_1, \pi)$. Since $\alpha_1(\pi)$ is optimal for the leader in $\Gamma(\pi)$, it is also optimal for the leader to play $\alpha_1(\pi)$ in $\Gamma^C(\pi)$ since the payoffs are the same in the two games. Furthermore, given that the leader has an incentive to deviate at a^* , $\alpha_1(\pi)$ does not assign probability one to a_1^* , i.e. the leader is randomizing between two or more actions. Thus it is optimal for every seller to choose a price of zero, since any active seller who increases his price will fail to sell with probability one.

Finally, if $\beta_2(a_1, \pi)$ is optimal for the follower in the game $\Gamma(\pi)$, then $\sigma_2(a_1, \pi)$ is optimal in $\Gamma^C(\pi)$ since the payoff function of the follower is identical in the two games. \square

References

- Bagwell, K., 1995. Commitment and observability in games. *Games Econ. Behav.* 8, 271–280.
- Bhaskar, V., 2005. Games played in a contracting environment. <http://else.econ.ucl.ac.uk/papers/uploaded/195.pdf>.
- Bhaskar, V., 2008a. Commitment and observability in a contracting environment. <http://else.econ.ucl.ac.uk/papers/uploaded/193.pdf>.
- Bhaskar, V., 2008b. Dynamic countervailing power with public and private monitoring. *Mimeo*.
- Bhaskar, V., van Damme, E., 2002. Moral hazard and private monitoring. *J. Econ. Theory* 102, 16–39.
- van Damme, E., Hurkens, S., 1997. Games with imperfectly observable commitment. *Games Econ. Behav.* 21, 282–308.
- Güth, W., Kirchsteiger, G., Ritzberger, K., 1998. Imperfectly observable commitments in n -player games. *Games Econ. Behav.* 23, 54–74.
- Harsanyi, J., 1973. Games with randomly disturbed payoffs: A new rationale for mixed strategy equilibrium points. *Int. J. Game Theory* 2, 1–23.
- Huck, S., Müller, W., 2000. Perfect versus imperfect observability—An experimental test of Bagwell's results. *Games Econ. Behav.* 31, 174–190.
- Maggi, G., 1999. The value of commitment with imperfect observability and private information. *RAND J. Econ.* 30, 555–574.
- Morgan, J., Várdy, F., 2004. An experimental study of commitment and observability in Stackelberg games. *Games Econ. Behav.* 49, 401–423.
- Morgan, J., Várdy, F., 2007. The value of commitment in contests and tournaments when observation is costly. *Games Econ. Behav.* 60, 326–338.
- Oechssler, J., Schlag, K., 2000. Loss of commitment: An evolutionary analysis of Bagwell's example. *Int. Game Theory Rev.* 2, 83–96.
- Schelling, T., 1960. *The Strategy of Conflict*. Harvard University Press, Cambridge, MA.
- Várdy, F., 2004. The value of commitment in Stackelberg games with observation costs. *Games Econ. Behav.* 49, 374–400.