Heuristic Strategies, Firm Behavior and Industry Information

C.-Y. Cynthia Lin and Erich J. Muehlegger[1]

May 2009

Abstract

Firms often lack knowledge of the nature of the uncertainty they or their opponents face and use heuristics or approximations to determine their strategy. We define and analyze one type of "heuristic strategy", in which firms choose strategies based on the expectation of their opponents' private information rather than the full information about the distribution of that private information. We find that, in equilibrium, the degree to which the heuristic strategy differs from the Bayesian (or "full-information") strategy depends on convexity and strategic complementarity or substitutability. Under certain conditions, firms' equilibrium profits are greater when all firms use heuristics than when all firms use the full information. Our results provide insight into incentives firms may have to either facilitate or impede access to industry information.

JEL Classification: D21, D80, L10

Keywords: firm behavior, heuristics, strategic complements, strategic substitutes, bounded rationality, asymmetric information, cognitive costs, information provision

Working paper.

1 Introduction

In the classical economic theory of how firms behave when their competitors have private information, each firm has beliefs about the distribution of their competitors' private information and uses the full information about the distribution to choose its optimal strategy. In actuality, firms often use heuristics or approximations to determine their strategy. For example, on the cover of a brochure advertising the Amadeus Business Intelligence Portfolio, a consulting service which acquires for its client firms information about competitors in the airline industry, the following slogan admonishes firms for this very tendency: "Don't just guess! be sure" (Amadeus, 2006).

There are several possible explanations for why a firm might use a heuristic to determine its strategy. First, firms may lack knowledge of the distribution of the competitors' private information and find it difficult or costly to acquire information about the distribution. Simon (1959) cites evidence from surveys that, when making dynamic decisions, firms often acquire information about the mean of the distribution of future events, but not the entire the distribution itself. Firms may find it especially difficult to collect information about the likelihood with which outcomes occur in the tails of the distribution, and may therefore only have information about the mean.[2] Absent complete information about the distribution of each firm's private information, firms may base their strategy upon the moments of the distribution, in particular the first moment, which may be more easy to estimate accurately. A second reason why firms might use a heuristic is that, even if firms know the full distribution, they might choose to base their strategy on a limited set of information. A firm may choose not to use all the available information in situations where the costs or computational difficulties of calculating the Bayesian strategy exceed the incremental profits accrued by the firm. As Baumol and Quandt (1964, p. 23) state: "the more refined the decision-making process, the more expensive it is likely to be, and therefore ... no more than an approximate solution may be justified." Finally,

---

[2]Policy makers face similar uncertainty - for example, Weitzman (2009) discusses the environmental policy implications of "fat tails," uncertainty about the distribution of climate damages in the tails of the damage distribution.

it is possible that firms might use heuristics because they are actually better off when all firms use heuristics than when all firms calculate the Bayesian equilibria using the full distribution of opponent's shock, even when both acquiring the full information and using it to compute the optimal strategy are costless.

In this paper, we study the impacts of heuristic use on firm behavior. In particular, we examine the conditions under which the use of a heuristic may do comparably well to a strategy based on complete knowledge of the distribution of a competitor's private information. We compare one particular heuristic strategy, in which a firm only uses the expectation of the private information a competitor is likely to receive, to the strategy the firm would adopt if it used the full distribution of competitors' private information in its optimization problem. We define the conditions under which the heuristic strategy chosen when a firm approximates is similar to the full-information strategy that is chosen when a firm maximizes expected profits taken with respect to the distribution of opponent's shocks. We characterize conditions under which firms in a market would prefer that all the firms use the full information, and therefore would have incentives to disclose their private information to each other. We also characterize conditions under which a set of firms may be *better off* when all firms approximate, and therefore may rationally attempt to coordinate, just as a set of firms has the incentive to collectively operate as a cartel. Under these circumstances, firms in the industry have the incentive to collectively withhold information from each other, and like a cartel, create mechanisms to facilitate cooperation.

Our paper expands upon several existing strands of literature. First, the notion that firms would want to disclose their private information to each other if they prefer that all firms in the market use the full information builds upon the work of Fried (1984), Gal-Or (1986) and Shapiro (1986), who examine the incentives of firms to share information with their competitors, for example through trade associations. Their models assume Cournot competition with linear demand and constant marginal costs. In contrast, our model makes minimal assumptions about functional form and no assumptions about the nature of competition, and therefore

2

applies more generally.

Our paper also builds upon literature in behavioral economics that finds that, contrary to classical economic theory, individuals often use heuristics and approximations to determine their behavior when faced with costly cognition or information acquisition. Simon (1955) observed that "the concept of 'economic man' (and, I might add, of his brother 'administrative man') is in need of fairly drastic revision", and that "the task is to replace the global rationality of economic man with a kind of rational behavior that is compatible with the access to information and the computational capacities that are actually possessed" by human decision-makers. Gabaix, Laibson, Moloche and Weinberg (2006) find evidence that the search activity of individuals seems to more closely follow a myopic model of cognition when information is costly. Luttmer and Shue (2006) find evidence in the 2003 California recall election that is consistent with misvoting relating to cognition costs. Several new equilibrium concepts that account for bounded rationality have been recently introduced, including "cursed equilibrium" (Eyster and Rabin, 2005), "analogy-based expectation equilibrium" (Jehiel, 2005) and "behavioral equilibrium" (Esponda, 2007).

The idea that, like individuals, firms may also use rules of thumb or heuristics rather than the full information is noted by Simon (1959), who states that when firms form expectations about the future, surveys of businessmen's expectations show that rather than estimate the joint probability distribution of future events, as would be needed in order to make the expected profit-maximizing decision, firms "have contented themselves with asking for point predictions – which, at best, might be interpreted as predictions of the means of the distributions." Ellison and Fudenberg (1993) examine a model in which firms use exogenously specified, simple and naive rules of thumb in deciding which technology to adopt. Weintraub, Benkard and Roy (2008) posit that firms may approximate by using an "oblivious strategy", which are strategies for which a firm considers only its own state and the long run average industry state, but ignores current information about competitors' states. The notion of approximation in our paper is similar to the notion put forward by Simon (1959), but applied to expectations about

opponents in a static game rather than expectations about future events in a single-agent dynamic decision-making problem: in our case, firms approximate by using the mean of an opponent's private information, rather than the full distribution.

We further augment the literature on heuristics and approximations by introducing the possibility that using a heuristic may actually be economically rational for the firm, since it may do better by doing so, and therefore that the classical economic theory of fully rational behavior can still apply.[3] The idea that having or using less information may make firms better off is also explored by Gal-Or (1988), who finds in her two-period model that a firm in a duopolistic market in which there is incomplete information about cost may benefit from having less precise prior information than its competitor, because having imprecise prior information provides a mechanism that enables the firm to commit to expand production relative to its rival. Similarly, Einy, Moreno and Shitovitz (2002) find that in Cournot competition with otherwise symmetric firms, the less informed firm could have greater profits. Mirman, Samuelson and Schlee (1994) also give conditions under which the value of information can be negative.

While we find that there are many cases in which firms would prefer to share their information with their competitors, as is consistent with the previous literature, we also find that, perhaps surprisingly, under certain conditions, industries have the incentive to coordinate on an equilibrium in which all firms calculate strategies based on heuristics rather than on the full information about the distribution of private information. Consequently, our results enable a better understanding of the incentives firms may have to either facilitate or impede access to industry information.

Our results not only have theoretical implications for the behavior of firms, but also speak to econometric applications. When the econometrician either lacks sufficient information or faces computational costs which prevent the estimation of the Bayesian equilibrium, our results present cases in which using an approximation yields an equivalent solution to that from using the full information. Our results also characterize the approximation error the econometrician

---

[3] Etzioni (1987) argues that the rationality of decision rules, or rules of thumb, is dubious and therefore that these rules logically cannot serve as a basis for rational conduct.

faces in cases in which the two solutions differ. While we focus on the implications of our results for the behavior of firms, similar results could identify whether econometric estimation of an approximation-based equilibrium would over- or under-estimate a Bayesian equilibrium.

Our paper proceeds as follows. Section 2 presents two equilibrium concepts corresponding to the equilibrium in which firms calculate the Bayesian strategies using the true distribution of opponents' shocks and the equilibrium in which both firms use heuristics. In section 3, we determine the conditions under which the use of a heuristic does or does not affect the equilibrium strategies chosen by the firms. In section 4, we apply our results from section 3 to investigate the extent to which use of a heuristic affects the equilibrium profits earned by firms. Appendices B and C apply our results from sections 3 and 4 to the case of Cournot competition followed by the case of competition on the Hotelling line, respectively.

2   Basic Model

Suppose there are two firms in the market. Given its private information ε_i, each firm i chooses a strategy s_i so as to maximize its static one-period profit $\pi_i(s_i, s_{-i}, \varepsilon_i)$, where the subscript $-i$ denotes the other firm. We assume that firm i does not see its opponents' private information ε_{-i}, which has no direct effect on firm i's profits.[4] Each shock ε_i has distribution $f_i(\cdot)$, and the shocks for the two firms are independent.

For simplicity, we choose to focus on a static one-period game. In a repeated game, the history of play provides a signal about the opponent's idiosyncratic shocks. While a firm using a heuristic due to substantial cognitive costs would not adjust their behavior in response, a firm which uses a heuristic due to a lack of information about the distribution of the opponent's shock could use the history of play to update their prior distribution. As a firm collects additional information, it would likely revise its decision-making process and use more complicated heuristics than the one we focus on in this paper. In principle, though, more complicated

---

[4]The opponent's private information ε_{-i} may have an indirect effect on firm i's profits, however, through its effect on the opponent's strategy s_{-i}.

heuristics could be evaluated in a similar fashion.

We assume that the profit function $\pi_i(s_i, s_{-i}, \varepsilon_i)$ satisfies the following:

Assumption 1.

1. $\pi_i(s_i, s_{-i}, \varepsilon_i)$ is second-order differentiable in s_i $\forall s_{-i}, \varepsilon_i$.

2. $\frac{\partial \pi_i(\cdot)}{\partial s_i}\big|_{s_i=0} > 0$ $\forall s_{-i}, \varepsilon_i$.

3. $\lim\limits_{s_i \to \infty} \frac{\partial \pi_i(\cdot)}{\partial s_i} < 0$ $\forall s_{-i}, \varepsilon_i$.

Assumptions 1.2 and 1.3 are boundary conditions. We now define two different equilibria. The two equilibria, which we denote the *full-information equilibrium* and the *heuristic equilibrium*, correspond to cases in which neither or both firms choose to approximate, respectively.

2.1 Full-Information Strategies

Suppose the distributions $f_i(\varepsilon_i)$ and $f_{-i}(\varepsilon_{-i})$ of both firms' private information are common knowledge and suppose each firm chooses its strategy conditional on this distribution. For all possible realizations of its own private information, each firm chooses its strategy to maximize its expected profits taking the expectation with respect to the opponent's private information. We call this strategy the *full-information strategy*, since each firm uses the full information about the distribution of its opponents' private information. For an opponent's strategy $s_{-i}(\cdot)$, the full-information strategy $s_i^*(\cdot)$ for firm i is given by the best response to $s_{-i}(\cdot)$:

$$s_i^*(\varepsilon_i) = \arg\max_{s_i'} \int \pi_i(s_i', s_{-i}(\varepsilon_{-i}), \varepsilon_i) f_{-i}(\varepsilon_{-i}) \, d\varepsilon_{-i} \quad \forall \varepsilon_i \; \forall i. \tag{1}$$

The first-order necessary condition is given by:

$$\int \frac{\partial \pi_i(s_i^*, s_{-i}(\varepsilon_{-i}), \varepsilon_i)}{\partial s_i} f_{-i}(\varepsilon_{-i}) \, d\varepsilon_{-i} = 0 \quad \forall \varepsilon_i \; \forall i. \tag{2}$$

When both firms employ full-information strategies that best respond to each other's full-information strategy, the equilibrium profile $(s_i^*(\cdot), s_{-i}^*(\cdot))$ that arises is the *full-information equilibrium*.

2.2 Heuristic Strategies

Now suppose that instead of using the distribution $f_{-i}(\varepsilon_{-i})$ of its opponents' private information, firm i only uses the mean $E[\varepsilon_{-i}]$ in determining its optimal strategy. Firms may use a heuristic if, for example, they face costs to acquiring or using the full information.[5] In this paper, we consider only the heuristic which uses the first moment of the distribution of its opponents' private information, based on evidence that decision makers often use the mean (Simon, 1959) or have difficulty understanding higher moments (Weitzman, 2009), although the basic idea could be generalized to a model in which firms select the number of moments with which to formulate their strategy. For an opponent's strategy $s_{-i}(\cdot)$, the *heuristic strategy* $\widehat{s}_i(\cdot)$ for firm i is given by the best response to $s_{-i}(E[\varepsilon_{-i}])$:

$$\widehat{s}_i(\varepsilon_i) = \arg\max_{s_i'} \pi_i(s_i', s_{-i}(E[\varepsilon_{-i}]), \varepsilon_i) \quad \forall \varepsilon_i \; \forall i. \tag{3}$$

The first-order necessary condition is given by:

$$\frac{\partial \pi_i(\widehat{s}_i, \widehat{s}_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} = 0 \quad \forall \varepsilon_i \; \forall i. \tag{4}$$

We call the strategy profile $(\widehat{s}_i(\cdot), \widehat{s}_{-i}(\cdot))$ that arises when both firms employ heuristic strategies that best respond to each other's heuristic strategy evaluated at each other's mean the *heuristic equilibrium*.[6]

---

[5] An econometrician may use such an approximation for either reason as well.

[6] When considering the heuristic equilibrium, we assume that the information acquisition costs and cognitive costs are sufficient to ensure firms cannot strictly improve profits by unilaterally using the full distribution of opponent's private information. In Section 4, we characterize the incremental profits associated with unilateral deviation from the heuristic equilibrium. The incremental profits constitute a lower bound on threshold

3   Effect of Heuristic on Strategy Choice

Whether profits for a firm differ in the two equilibria depends on (1) the degree to which the strategy profile chosen when both firms use the heuristic differ from the strategy profile chosen when both firms play full-information strategies, and (2) the degree to which a firm's profits are affected by deviations from the full-information strategy. In this section, we focus on the first of these questions – we identify the conditions under which the strategies chosen when firms use the heuristic differ from those chosen in the full-information equilibrium.

We first examine conditions under which the heuristic equilibrium is equivalent to the full-information equilibrium. Then, in order to understand how equilibrium strategies change, we compare the heuristic strategy and the full-information strategy holding the opponent's choice of strategy constant. Under- or over-estimation of the full-information strategy depends on whether the derivative of a firm's profit function with respect to its own strategy is convex or concave in the opponent's private information. Finally, we characterize the conditions under which the strategies differ in the full-information equilibrium and the heuristic equilibrium. We find that the degree to which equilibrium strategies differ is sensitive to whether firms' strategies are strategic complements or strategic substitutes.

Under what conditions would the full-information equilibrium $(s_i^*(\cdot), s_{-i}^*(\cdot))$ that arises when each firm uses the full information about the distribution of its opponent's private information be equivalent to the heuristic equilibrium $(\widehat{s}_i(\cdot), \widehat{s}_{-i}(\cdot))$ that arises when each firm uses the expectation of its opponent's private information? By comparing the first-order conditions in the heuristic and full-information equilibria, we can define a set of sufficient conditions under which the strategies chosen in the heuristic equilibrium will be equal to those chosen in the full-information equilibrium.

Proposition 1 *Given assumption 1, $(s_i^*(\cdot), s_{-i}^*(\cdot))$ and $(\widehat{s}_i(\cdot), \widehat{s}_{-i}(\cdot))$ are equivalent when, $\forall \varepsilon_i$ $\forall i$,*

(i) $\pi_i(s_i, s_{-i}, \varepsilon_i)$ is strictly concave in s_i and

acquisition and cognitive costs sufficient to ensure existence of the heuristic equilibria.

8

(ii)

$$\frac{\partial \pi_i(s_i^*, s_{-i}^*(\overline{\varepsilon_{-i}}), \varepsilon_i)}{\partial s_i} - \frac{\partial \pi_i(s_i^*, s_{-i}^*(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} - \int \left(\frac{\partial^2 \pi_i(s_i^*, s_{-i}^*(\varepsilon_{-i}), \varepsilon_i)}{\partial s_i \partial \varepsilon_{-i}} \right) F_{-i}(\varepsilon_{-i}) d\varepsilon_{-i} = 0 \; ,$$
(5)

where $\overline{\varepsilon_{-i}}$ is the upper bound in the support of ε_{-i} and where $F_{-i}(\varepsilon_{-i}) = \int f_{-i}(\varepsilon_{-i}) d\varepsilon_{-i}$.

For all proofs, see appendix A.

Corollary 2 *Given assumption 1, $(s_i^*(\cdot), s_{-i}^*(\cdot))$ and $(\widehat{s}_i(\cdot), \widehat{s}_{-i}(\cdot))$ are equivalent when, $\forall \varepsilon_i \; \forall i$,*

(i) $\pi_i(s_i, s_{-i}, \varepsilon_i)$ is strictly concave in s_i and

(ii) $\pi_i(s_i^*, s_{-i}^*(\varepsilon_{-i}), \varepsilon_i)$ has no interaction term between s_i and ε_{-i} (i.e., $\frac{\partial^2 \pi_i(s_i^*, s_{-i}^*(\varepsilon_{-i}), \varepsilon_i)}{\partial s_i \partial \varepsilon_{-i}} = 0$).

When the conditions of Corollary 2 are satisfied, the mean of the distribution is the only relevant parameter needed for the firm to make an optimal decision; the variance and higher moments would be of no additional use even if they were known to the firm. This result is similar to the result that in a single-agent dynamic programming problem under uncertainty, when the criterion function is quadratic, the mean is a certainty equivalent and therefore a sufficient statistic for the entire distribution (Simon, 1956, 1959).

3.1 Comparing Best Response Functions

In this section, we compare the strategies chosen under the heuristic and full information best response functions holding the opponent's choice of strategy constant. In Proposition 3, we study how a firm's strategy for each value of its private information changes when the firm switches from the Bayesian approach to the heuristic approach holding the strategy of the other firm constant. The important criterion governing whether a firm using an heuristic will over- or under-estimate the full-information strategy is whether the derivative of the firm's profit functions with respect to its own strategy is concave or convex in the other firm's private information.

Proposition 3 *Let $s_{-i}(\cdot)$ be an arbitrary strategy for firm $-i$ with the property that $\frac{\partial \pi_i}{\partial s_i}$ is quadratic in ε_{-i}. If $\frac{\partial \pi_i}{\partial s_i}$ is convex (concave) in ε_{-i}, $s_i^*(\varepsilon_i)$ is greater than (less than) $\hat{s}_i(\varepsilon_i)$ for all s_i, ε_i.*

Proposition 3 identifies a sufficient condition under which the strategies chosen to solve the heuristic first-order condition differ from those chosen to solve the full-information first order condition.[7] Conditional on the opponent's choice of strategy, firm i's solution to the heuristic best response function differs from the solution to the full-information best response function when $\frac{\partial \pi_i}{\partial s_i}$ is convex or concave in the firm $-i$'s shock and firm $-i$'s shock is uncertain (and hence the distribution of firm $-i$'s shock has positive variance). Evaluating three special cases of Proposition 3 provides additional intuition about when a strategy solving the heuristic first-order condition will differ from a strategy solving the full-information first-order condition.

Corollary 4 *Let π_i and π_i^H denote two profit functions with $\frac{\partial \pi_i}{\partial s_i}$ quadratic in ε_{-i}. If $\frac{\partial \pi_i^H}{\partial s_i}$ is strictly more convex (concave) in ε_{-i} than $\frac{\partial \pi_i}{\partial s_i}$, $s_i^{H*} - \hat{s}_i^H$ is greater than (less than) $s_i^* - \hat{s}_i$ for all s_{-i}, ε_i.*

Corollary 4 provides additional intuition about when, conditional on the opponent's strategy, employing a strategy solving the heuristic first-order condition would differ substantially from the strategy solving the full-information first-order condition. Two factors affect the degree to which the heuristic based strategy under- or over-estimates the full-information strategy. First, for a given convexity or concavity, the degree of over- or under-estimation is positively related to the variance of the shocks facing the opponent. Second, as $\frac{\partial \pi_i}{\partial s_i}$ becomes more convex (concave) with respect to ε_{-i}, $\hat{s}_i$ under-estimates (over-estimates) s_i^* by a greater amount. As the convexity (concavity) of $\frac{\partial \pi_i}{\partial s_i}$ with respect to ε_{-i} increases or as the variance of the opponent's shock increases, the strategy solving the full-information FOC will differ more from the strategy solving the heuristic FOC. Furthermore, for a special case of Proposition 3, in which

---

[7]Although we focus on a function for which $\frac{\partial \pi_i}{\partial s_i}$ is quadratic in ε_{-i}, it is possible to adapt the previous result to the case of an arbitrary function, by defining quadratic functions which provide upper and lower bounds on the convexity/concavity of the arbitrary function. The analagous result is if $\frac{\partial \pi_i}{\partial s_i}$ is globally convex (concave) in ε_{-i}, s_i^* will be greater (less) than $\hat{s}_i$.

10

$\frac{\partial \pi_i}{\partial s_i}$ is linear in ε_{-i}, heuristic will not affect the best response of firm i for all strategies s_{-i} and distributions of shocks ε_i. In this case, the strategy solving the heuristic equilibrium is identical to that solving the full-information equilibrium.

Corollary 5 *Let $s_{-i}(\varepsilon_{-i})$ be an arbitrarily chosen strategy for firm $-i$. If $\frac{\partial \pi_i}{\partial s_i}$ is linear in ε_{-i}, and a unique solution to the maximization problem exists, then $s_i^* = \hat{s}_i$ for all ε_i.*

Corollary 5 presents an alternative sufficient condition to that put forth in Proposition 1. In combination with Corollary 4, Corollary 5 suggests the cases in which using the heuristic would lead firms to adopt strategies similar to those they would choose if the calculated the full information Bayesian strategy. Specifically, conditional on the opponent's strategy, the strategy solving the approximation FOC will be more similar to the strategy solving the full-information FOC in cases where the variance of the opponent's shock is low and in cases in which it is possible to generalize the preceding result to cases in which $\frac{\partial \pi_i}{\partial s_i}$ is not as convex or concave in ε_{-i}. It is possible to derive an analogous result to that in Corollary 5 for the more general case in in which $\frac{\partial \pi_i}{\partial s_i}$ is linear in a function of the opponent's shock, $g(\varepsilon_{-i})$. In this case, a firm solving the heuristic first-order condition, as if the opponent faced a shock of $E[g(\varepsilon_{-i})]$ would choose the same strategy as if the first had solved the full-information first-order condition.

Corollary 6 *Define $\hat{s}_i = argmax \; \pi_i(s_i, s_{-i}(E[g(\varepsilon_{-i})]), \varepsilon_{-i})$. If $\frac{\partial \pi_i}{\partial s_i}$ is linear in $g(\varepsilon_{-i})$, then $s_i^* = \hat{s}_i$.*

To illustrate these results with a concrete example, we use specifications for cost and demand functions under the assumption of Cournot competition to derive examples in which the derivative of the profit function with respect to the strategy is linear in the opponent's shock, and confirm that for these cases the expression for the derivative could be expressed in the linear form. This linear form was still applicable when the opponent's shock had a direct affect on profits. In particular, some classes of cost-demand function combinations for which the derivative of the profit function with respect to the strategy is linear in the opponent's

shock include: (1) additive demand and cost shocks; (2) shock to slope of cost function; and (3) additive interactive shock to demand. Appendix B analyzes the case of a quadratic shock to the slope of the cost function.

Similarly, we present several specifications of a model of Bertrand competition on the Hotelling line in Appendix C. For cases in which shocks enter the marginal cost of production for the firms or the consumer valuation of the firms' product linearly, we can express the derivative of firm i's profit function with respect to its own strategy as a function which in linear in the opponent's shock. In these cases, the full-information strategy and the heuristic strategy are equivalent. When the shock enters the marginal cost function quadratically, it is no longer possible to express the derivative of firm i's profit function with respect to its own strategy as linear in the opponent's shock. In this case, the full-information strategy and heuristic strategy differ. It is interesting to note that when a firm's shock enters quadratically into a firm's own marginal cost function, the firms' profit functions satisfy the requirements for Corollary 6, where $g(\varepsilon_{-i}) = \varepsilon_{-i}^2$.

3.2 Comparing Equilibrium Strategies

We now extend the results from the previous section to consider how the equilibrium strategies differ in the full-information and heuristic equilibria. In the case of the heuristic equilibria, we assume that fixed cognitive costs are sufficient such that using the heuristic is strictly preferable to calculating the full-information strategy. We find that the equilibria of the three cases is sensitive to whether firms' strategies are strategic complements or strategic substitutes.

For the following analysis, we compare the strategies chosen in the full-information equilibrium, (s_i^*, s_{-i}^*), in which firms maximize their expected profit to those chosen in the heuristic equilibrium, $(\hat{s}_i, \hat{s}_{-i})$, in which firms maximize their profit using only the first moment of the distribution of their competitor's private information.

Let $h_i(\varepsilon_i)$ denote the difference between firm i's heuristic and full-information equilibrium

strategies: $h_i(\varepsilon_i) \equiv \hat{s}_i(\varepsilon_i) - s_i^*(\varepsilon_i)$.

To compare the full-information equilibrium strategies and the heuristic equilibrium strategies, we define $\tilde{s}_i$ as the heuristic best response of firm i to $s_{-i}^*(s_i^*, \varepsilon_{-i})$.[8] We define $\tilde{s}_{-i}$ in an analogous fashion. Note that by our definition of $\tilde{s}_i$ and $\tilde{s}_{-i}$,

$$\frac{\partial \pi_i(\tilde{s}_i, s_{-i}^*(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} = 0 \tag{6}$$

$$\frac{\partial \pi_{-i}(\tilde{s}_{-i}, s_i^*(E[\varepsilon_i]), \varepsilon_{-i})}{\partial s_{-i}} = 0. \tag{7}$$

Then, evaluating the heuristic FOC when firms play $(\tilde{s}_i, \tilde{s}_{-i})$, we have

$$\frac{\partial \pi_i(\tilde{s}_i, \tilde{s}_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} = \int_{s_{-i}^* E[\varepsilon_{-i}]}^{\tilde{s}_{-i} E[\varepsilon_{-i}]} \frac{\partial^2 \pi_i(\tilde{s}_i, s_{-i}, \varepsilon_i)}{\partial s_i \partial s_{-i}} ds_{-i} \tag{8}$$

$$\frac{\partial \pi_{-i}(\tilde{s}_{-i}, \tilde{s}_i(E[\varepsilon_i]), \varepsilon_{-i})}{\partial s_{-i}} = \int_{s_i^* E[\varepsilon_i]}^{\tilde{s}_i E[\varepsilon_i]} \frac{\partial^2 \pi_{-i}(\tilde{s}_{-i}, s_i, \varepsilon_{-i})}{\partial s_{-i} \partial s_i} ds_i. \tag{9}$$

The sign and magnitude of (8) and (9) depend on two factors. First, they depend on whether s_i and s_{-i} are strategic complements or substitutes. In addition, each derivative depends on whether maximizing the heuristic FOC leads firm i and firm $-i$ to adopt higher or lower strategies than does maximizing the full-information FOC. Proposition 3 tells us that if $\frac{\partial \pi_i}{\partial s_i}$ and $\frac{\partial \pi_{-i}}{\partial s_{-i}}$ are convex in ε_{-i} and ε_i, $\tilde{s}_i > s_i^*$ and $\tilde{s}_{-i} > s_{-i}^*$.[9] Thus, if $\tilde{s}_i > s_i^*$ and $\tilde{s}_{-i} > s_{-i}^*$ and if s_i and s_{-i} are strategic complements (substitutes), both derivatives are positive (negative). If $\tilde{s}_i < s_i^*$ and $\tilde{s}_{-i} < s_{-i}^*$ and if s_i and s_{-i} are strategic complements (substitutes), both derivatives are negative (positive).

Now, we relate the evaluated derivatives in (8) and (9) to the heuristic equilibrium by expressing them as Taylor expansions around the heuristic equilibrium. Letting $x_i(\varepsilon_i) =$

---

[8] We assume that use of the heuristic has an effect on the choice of strategy for firm i and firm $-i$. That is, we assume that $\frac{\partial \pi_i}{\partial s_i}$ and $\frac{\partial \pi_{-i}}{\partial s_{-i}}$ are either both convex or concave in ε_{-i} and ε_i respectively.

[9] Furthermore, Corollary 4 shows that the degree to which the heuristic strategy deviates from the Bayesian strategy depends on the degree of concavity or convexity of $\frac{\partial \pi_i}{\partial s_i}$ and $\frac{\partial \pi_{-i}}{\partial s_{-i}}$ are convex in ε_{-i} and ε_i.

13

$\tilde{s}_i(\varepsilon_i) - \hat{s}_i(\varepsilon_i)$ and $x_{-i}(\varepsilon_{-i}) = \tilde{s}_{-i}(\varepsilon_{-i}) - \hat{s}_{-i}(\varepsilon_{-i})$, we can write (8) as

$$
\begin{aligned}
\frac{\partial \pi_i(\tilde{s}_i(\varepsilon_i), \tilde{s}_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} &= x_i(\varepsilon_i) \int_0^1 \frac{\partial^2 \pi_i(\hat{s}_i(\varepsilon_i) + tx_i(\varepsilon_i), \hat{s}_{-i}(E[\varepsilon_{-i}]) + tx_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i^2} dt \\
&+ x_{-i}(E[\varepsilon_{-i}]) \int_0^1 \frac{\partial^2 \pi_i(\hat{s}_i(\varepsilon_i) + tx_i(\varepsilon_i), \hat{s}_{-i}(E[\varepsilon_{-i}]) + tx_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i \partial s_{-i}} dt
\end{aligned}
$$

Similarly, we could write (9) using an analogous Taylor expansion. By definition, $(\hat{s}_i, \hat{s}_{-i})$ jointly satisfy the two Taylor expansions. Using this, we can compare the strategies chosen by the firms in the full-information equilibrium and the strategies chosen by the firms in the heuristic equilibrium.

Proposition 7 *If s_i and s_{-i} are strategic complements and $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}| \; \forall s_i, s_{-i}$, then $\hat{s}_i(\varepsilon_i) > \tilde{s}_i(\varepsilon_i) > s_i^*(\varepsilon_i) \; \forall \varepsilon_i$ when use of the heuristic leads firm i and $-i$ to increase their strategy, and $\hat{s}_i(\varepsilon_i) < \tilde{s}_i(\varepsilon_i) < s_i^*(\varepsilon_i) \; \forall \varepsilon_i$ when use of the heuristic leads firm i and $-i$ to lower their strategic variable.*

We find that so long as use of the heuristic affects both firm strategies in the same direction, strategies are complements and $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}|$, the strategies played in the heuristic equilibrium will be further from the full-information equilibrium strategies than the strategies adopted when a single firm approximates and the second firm plays a full-information strategy. In addition, under these conditions, h_i and h_{-i} will have the same signs. Note that the assumption $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}|$ holds for Cournot competition and for basic formulations of price competition on the Hotelling line.[10]

Now consider the case of strategic substitutes. In this case, the result is ambiguous. In some situations, $h_i(\varepsilon_i)$ and $h_{-i}(\varepsilon_{-i})$ are of different signs. The situations where this is most likely to occur would be situations where $\tilde{s}_i$ is relative close to s_i^* and $\tilde{s}_i$ is distant from s_i^*.

Proposition 8 *If s_i and s_{-i} are strategic substitutes and $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}|$ then we can sign $h_i(\varepsilon_i)$ and $h_{-i}(\varepsilon_{-i})$.*

---

[10] In the basic two-firm Cournot model, $\pi_1 = (1 - q_1 - q_2)q_1 \Rightarrow \frac{\partial^2 \pi_1}{\partial q_1^2} = -2, \frac{\partial^2 \pi_1}{\partial q_1 \partial q_2} = -1$. In the Hotelling model (omitting the cost term), $\pi_1 = \frac{1 + p_2 - p_1}{2t} p_1 \Rightarrow \frac{\partial^2 \pi_1}{\partial p_1^2} = \frac{-1}{t}, \frac{\partial^2 \pi_1}{\partial p_1 \partial p_2} = \frac{1}{2t}$.

In the case of strategic substitutes, we can make fewer strong remarks about the signs of $h_i(\varepsilon_i)$ and $h_{-i}(\varepsilon_{-i})$. In fact, the signs of $h_i(\varepsilon_i)$ and $h_{-i}(\varepsilon_{-i})$ depend on the parameters of the problem. What we can learn, though, is that for the case worked through in the proof $\tilde{s}_i > s_i^* \leq \hat{s}_i$ is more likely to happen when the use of the heuristic has a small effect on firm i's optimal strategy but the use of the heuristic has a large effect on firm $-i$'s optimal strategy.

4 Effect of Heuristic on Equilibrium Profits

Now we consider the effect of heuristic use on the profits realized by the firms. As with the previous section, we first examine how switching from a Bayesian approach to an heuristic-based strategy affects the profits of a single firm holding the strategy of the other firm constant. We then turn to investigate how the profits differ for the firms under the full information and heuristic equilibrium. In particular, we define the conditions under which firms would have an incentive to attempt to coordinate on either the heuristic equilibrium or the full information equilibrium. In both sections, we again sign the terms of a Taylor expansion to sign the difference in profits. In particular, note that for a given vector of private information $(\varepsilon_i, \varepsilon_{-i})$, we can express the difference in profits for firm i for two sets of strategies $(s_i(\varepsilon_i), s_{-i}(\varepsilon_{-i}))$ and $(s_i'(\varepsilon_i), s_{-i}'(\varepsilon_{-i}))$ as

$$
\begin{aligned}
\pi_i(s_i, s_{-i}, \varepsilon_i) - \pi_i(s_i', s_{-i}', \varepsilon_i) &= (s_i - s_i') \frac{\partial \pi_i(s_i', s_{-i}', \varepsilon_i)}{\partial s_i} + \\
&\quad (s_{-i} - s_{-i}') \frac{\partial \pi_i(s_i', s_{-i}', \varepsilon_i)}{\partial s_{-i}} + \\
&\quad \int_0^1 (1-t)(s_i - s_i')^2 \frac{\partial^2 \pi_i(\tilde{x})}{\partial s_i^2} dt + \\
&\quad \int_0^1 (1-t)(s_{-i} - s_{-i}')^2 \frac{\partial^2 \pi_i(\tilde{x})}{\partial s_{-i}^2} dt + \\
&\quad 2 \int_0^1 (1-t)(s_i - s_i')(s_{-i} - s_{-i}') \frac{\partial^2 \pi_i(\tilde{x})}{\partial s_i \partial s_{-i}} dt, \quad (10)
\end{aligned}
$$

where $\tilde{x} \equiv t(s_i, s_{-i}, \varepsilon_i) + (1-t)(s_i', s_{-i}', \varepsilon_i)$.

15

4.1 Out-of-Equilibrium Profits

We first consider the effect on profits from using the heuristic holding the strategy of its opponent constant. In particular, we compare the profits firm i would earn in the full information equilibrium with the profits it would earn it chose to unilaterally deviate by playing the heuristic best response.

Proposition 9 *Suppose firm $-i$ plays strategy $s'_{-i}(\varepsilon_{-i})$. Absent cognitive costs, firm i's expected profits are weakly lower when it plays the heuristic best response than when it plays its full-information equilibrium strategy.*

Holding the other firm's strategy constant and absent any cognition costs, a firm would always want to use all of the information about the distribution of the opponent's shock it had rather than use a heuristic. This follows the standard intuition — the Bayesian strategy uses a larger set of information for the optimization than the heuristic strategy. A firm maximizing expected profits could always opt to use the heuristic strategy if that were the strategy which maximized expected profits. Thus, ignoring equilibrium effects and cognition costs, the profits under the full-information Bayesian strategy must be weakly greater than the profits firm i could earn by unilaterally deviating to the heuristic best response. It is important to note that this will not necessarily be true when we compare the profits earned by the firms in the heuristic and full-information equilibria.

The Taylor expansion in equation (10) also provides a lower bound on the cognitive costs necessary to prevent firm i from deviating from the heuristic equilibrium by playing the full-information best response.

Corollary 10 *If the cognitive costs associated with playing the full information best response are greater than $E[\int_0^1 (1-t)(\hat{s}_i - s_i^*)^2 \frac{\partial^2 \pi_i(\bar{x})}{\partial s_i^2} dt | \varepsilon_i]$, firm i will not deviate from the heuristic equilibrium by playing the full information.*

16

The lower bound for cognitive costs varies depending on the realization of ε_i. The lower bound in increasing in both the difference between $\hat{s}_i$ and s_i^* and in the convexity of π_i with respect to s_i. Both of these increase the incremental profit earned by using the full information best response over the heuristic best response.

4.2   Equilibrium Profits

We now compare the profits for firm i in the heuristic equilibrium with those in the full-information equilibrium. In the previous section, we found that holding the opponent's strategy constant, the use of the heuristic always weakly lowers profit relative to the Bayesian strategy with correct beliefs about the distribution of the opponent's private information. In equilibrium, though, this is not always true. In this section we identify the conditions under which firms have the incentive to coordinate on either the heuristic or full-information equilibrium.

Consider

$$E[\pi_i(\widehat{s}_i(\varepsilon_i), \widehat{s}_{-i}(\varepsilon_{-i}), \varepsilon_i) - \pi_i(s_i^*(\varepsilon_i), s_{-i}^*(\varepsilon_{-i}), \varepsilon_i)].$$

which can be rewritten as

$$E[\int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i} h_i dt + \int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_{-i}} h_{-i} dt] \tag{11}$$

The first of the two integrals above can be expanded into first and second order terms as:

$$E[\int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i} h_i dt] = E[\frac{\partial \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)}{\partial s_i} h_i \tag{12}$$
$$+ \int_0^1 (1-t) \frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i^2} h_i^2 dt$$
$$+ \int_0^1 (1-t) \frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i \partial s_{-i}} h_i h_{-i} dt]$$

Noting that the first term of (12) is equal to zero by construction of s_i^* , and rearranging the

terms of (11), we have:

$$E[\pi_i(\hat{s}_i(\varepsilon_i), \hat{s}_{-i}(\varepsilon_{-i}), \varepsilon_i) \quad - \quad \pi_i(s_i^*, s_{-i}^*(\varepsilon_{-i}), \varepsilon_i)] \tag{13}$$

$$= \quad E[\int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_{-i}} h_{-i} dt]$$

$$+ E[\int_0^1 (1-t) \frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i^2} h_i^2 dt]$$

$$+ E[\int_0^1 (1-t) \frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i \partial s_{-i}} h_i h_{-i} dt]$$

Equation (13) defines a sufficient condition for firm i to prefer either the full-information equilibrium or the heuristic equilibrium.[11]

Proposition 11 *If*

$$0 \leq E[\int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_{-i}} h_{-i} dt | \varepsilon_i] + E[\int_0^1 (1-t) \frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i^2} h_i^2 dt | \varepsilon_i]$$

$$+ E[\int_0^1 (1-t) \frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i \partial s_{-i}} h_i h_{-i} dt | \varepsilon_i]$$

for all ε_i, then firm i will (weakly) prefer the heuristic equilibrium to the full-information equilibrium.

Conversely, if the above equation is weakly negative for all ε_i, then firm i will (weakly) prefer the full-information equilibrium to the heuristic equilibrium.

The intuition of the first term in (13) is clear: the partial derivative captures whether the profits of firm i rise with an increase in the strategy chosen by firm $-i$, while h_{-i} captures whether or not the use of the heuristic leads firm $-i$ to choose a higher or lower strategy. If these two have the same sign, then the use of the heuristic by firm $-i$ causes firm $-i$ to choose a strategy more beneficial to firm i. The second term, the product of the the concavity of firm i's profit function with respect to firm i's own strategy and h_i^2, is always negative. The third term is

---

[11] Note that expanding out the second term in (11) as well leaves the familiar second-order Taylor expansion.

18

easily signed knowing whether strategies are complements or substitutes and whether the use of the heuristic leads both firm i and firm $-i$ to adjust strategies in a similar direction.

Being able to easily sign the terms in (13) leads to a sufficient condition for firms to prefer (ex ante) coordinating on the full-information equilibrium.

Corollary 12 *Consider a game in which firm strategies are strategic substitutes and sign(h_i)=sign(h_{-i}). If sign($\frac{\partial \Pi_i}{\partial s_{-i}}$)$\neq$ sign(h_{-i}), firms will earn more in expectation if they coordinate on (s_i^*, s_{-i}^*) rather than $(\hat{s}_i, \hat{s}_{-i})$.*

Using Cournot competition as an example, if using an approximation-based strategy leads firms to adopt higher output strategies, firms will always prefer (ex ante) to coordinate on the full information equilibrium. In this case, if firms could costlessly make public the information about the distribution of their own private information, firms would choose to do so.

While there are many cases in which firms would prefer to share their information with their competitors, as is consistent with the previous literature, there are cases in which the sum of the terms in (13) is positive, and under these conditions, industries have the incentive to coordinate on the equilibrium in which all firms calculate strategies based on the heuristic rather than on the full information. This is a surprising result, because it suggests that firms may prefer to coordinate on an equilibrium in which firms have little information about their competitor. For example, in the model of Bertrand competition on the Hotelling line with nonlinear shocks presented in Appendix C, firms would prefer to approximate rather than use all the information when $\alpha < 1$. In addition, signing the terms in (13) above also provide sufficient conditions for situations in which firms would, ex ante, prefer to coordinate on the heuristic equilibrium rather than the full information equilibrium. As noted above, with the maintained assumption that $\frac{\partial^2 \Pi_i}{\partial s_i^2} < 0$, the second term of equation (13) is always negative. The contrapositive to Corollary 12 provides necessary conditions for firm i to prefer the heuristic equilibrium to the full information equilibrium. Alternatively, in a repeated context, firm may coordinate on the heuristic equilibrium, using the full-information equilibrium to enforce coordination.

Corollary 13 *Consider a game in which sign(h_i)=sign(h_{-i}). If $E[\pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i) - \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)] \geq$*

0. then at least one of the following conditions must hold:

(i) sign($\frac{\partial \Pi_i}{\partial s_{-i}}$)=sign($h_{-i}$)

(ii) firms compete in strategic complements.

The first condition captures whether the use of the heuristic leads firms to choose strategies beneficial to their opponent. For example, if in the heuristic equilibrium, firm i and firm $-i$ adopt strategies which maximize joint profits, the firm i and firm $-i$ would both prefer to coordinate on the heuristic equilibrium rather than the full information equilibrium.

5 Conclusion

In this paper, we study the effects of using heuristics on firm strategic behavior and profits. In particular, we study heuristic strategies in which a firm only uses the expected private information a competitor is likely to receive, rather than the distribution of competitors' private information, in its optimization problem. We define the conditions under which the heuristic strategy chosen when a firm approximates is similar to that which is chosen when a firm uses the entire distribution of its opponent's private information and maximizes expected profits. We characterize conditions under which firms in a market would prefer that all the firms in the market use the full information, and would therefore have incentives to disclose their private information to each other, for example through a trade association. We also examine conditions under which a firm or set of firms may actually be *better off* when they approximate, and therefore may rationally choose to do so. Under these circumstances, a firm has the incentive to ignore the information it has about its opponent, and the firms in the industry have the incentive to collectively withhold information from each other.

Our results identify the conditions under which heuristic strategies differ from full-information strategies. We find that the holding the other firms' strategy constant, heuristic strategies differ from full-information strategies when the derivative of a firm's profit with

respect to its own strategy is nonlinear in the opponent's shock. Moreover, if concavity (convexity) in the opponent's shock is greater, an approximation-based strategy will over-estimate (under-estimate) the expected profit maximizing full-information strategy to a greater degree. The equilibrium strategies chosen when one or both firms approximate depend on the degree of strategic complementarity or substitutability.

Although we focus on a one-period game in this paper, the incentives for firm implied by our results map to a context in which the game is repeated. Repetition would allow a firm in the limit to reach the full information strategy, either by experimentation in the case of a firm facing substantial cognitive costs or by inferring the distribution of opponents' shocks in the case of a firm lacking information about an opponents' distribution. Since unilateral deviation to the full-information strategy weakly dominates the heuristic strategy, this suggests that if the game is repeated, that firms may profitably deviate from both playing heuristic strategies as they either experiment or learn the distribution of their opponents' shocks. Repeating the game, though, does not affect the relationship between the degree of strategic complementarity or substitutability and the distance between the full and heuristic strategies or the incentives to coordinate on either the heuristic or full-information equilibria.

While we find that there are many cases in which firms would prefer to share their information with their competitors, as is consistent with the previous literature, we also find that under certain conditions, industries have the incentive to coordinate on the equilibrium in which all firms calculate strategies based on the heuristic rather than on the full information. This is a surprising result, for it suggests that firms may find it economically rational not to use all the information about their competitor, even if they could acquire and use it costlessly. Consequently, our results enable a better understanding of the incentives firms may have to either facilitate or impede access to industry information.

References

[1] Amadeus IT Group SA. 2006. Amadeus Business Intelligence Portfolio. http://www.amadeus.com/documents/airline/BROCHURE%20BUSINESS%20INTELLIGENCE%20PORTFOLIO%20Dec06.pdf. (accessed Jan. 22, 2007).

[2] Baumol, William J., and Richard E. Quandt. 1964. "Rules of thumb and optimally imperfect decisions." *American Economic Review*, 54: 23-46.

[3] Einy, Ezra, Diego Moreno, and Benyamin Shitovitz. 2002. "Information advantage in Cournot oligopoly." *Journal of Economic Theory*, 106: 151-160.

[4] Ellison, Glenn, and Drew Fudenberg. 1993. "Rules of thumb for social learning." *The Journal of Political Economy*, 101 (4): 612-643.

[5] Esponda, Ignacio. 2007. "Behavioral equilibrium in economies with adverse selection." Working paper, New York University.

[6] Etzioni, Amitai. 1987. "On thoughtless rationality (rules-of-thumb)." *KYKLOS*, 40: 496-514.

[7] Eyster, Erik, and Matthew Rabin. 2005. "Cursed equilibrium." *Econometrica*, 73 (5): 1623-1672.

[8] Fried, Dov. 1986. "Incentives for information production and disclosure in a duopolistic environment." *Quarterly Journal of Economics*, 99: 367-381.

[9] Gabaix, Xavier, David Laibson, Guillermo Moloche, and Stephen Weinberg. 2006. "Costly information acquisition: Experimental analysis of a boundedly rational model." *American Economic Review*, 96 (4): 1043-1068.

[10] Gal-Or, Esther. 1986. "Information Transmission: Cournot and Bertrand Equilibria." *Review of Economic Studies*, 52: 85-92.

[11] Gal-Or, Esther. 1988. "The advantages of imprecise information." *RAND Journal of Economics*, 10 (2): 266-275.

[12] Jehiel, Philippe. 2005. "Analogy-based expectation equilibrium." *Journal of Economic Theory*, 123: 81-104.

[13] Luttmer, Erzo, and Kelly Shue. 2006. "Who misvotes?: The effect of differential cognition costs on election outcomes." KSG Working paper RWP06-048.

[14] Mirman, Leonard J., Larry Samuelson, and Edward E. Schlee. 1994. "Strategic information manipulation in duopolies." *Journal of Economic Theory*, 62: 363-384.

[15] Shapiro, Carl. 1986. "Exchange of cost information in oligopoly." *Review of Economic Studies*, 52: 433-446.

[16] Simon, Herbert A. 1955. "A behavioral model of rational choice." *The Quarterly Journal of Economics*, 69 (1): 99-118.

[17] Simon, Herbert A. 1956. "Dynamic programming under uncertainty with a quadratic criterion function." *Econometrica*, 24 (1): 74-81.

[18] Simon, Herbert A. 1959. "Theories of decision-making in economics and behavioral science." *American Economic Review*, 49 (3): 253-283.

[19] Weintraub, Gabriel, C. Lanier Benkard, & Benjamin Van Roy. 2008. "Markov Perfect industry dynamics with many firms." *Econometrica*, 76 (6): 1375-1411.

[20] Weitzman, Martin. 2009. "On Modeling and Interpreting the Economics of Catastrophic Climate Change." *Review of Economics and Statistics*, 91 (1), 1-19.

A   Appendix: Proofs

Proposition 1

Proof. Under (1) and (2), $(s_i^*(\cdot), s_{-i}^*(\cdot))$ and $(\widehat{s}_i(\cdot), \widehat{s}_{-i}(\cdot))$ are equivalent when the left-hand side of equation (2) equals the left-hand side of equation (4):

$$\int \left(\frac{\partial \pi_i(s_i^*, s_{-i}^*(\varepsilon_{-i}), \varepsilon_i)}{\partial s_i} - \frac{\partial \pi_i(\widehat{s}_i, \widehat{s}_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} \right) f_{-i}(\varepsilon_{-i}) \, d\varepsilon_{-i} = 0$$

Integrate by parts and set $\widehat{s}_i(\cdot) = s_i^*(\cdot) \; \forall i$ to yield equation (5). ∎

Corollary 2

Proof. (ii) is a sufficient condition for equation (5). ∎

Proposition 3

Proof. Let $s_{-i}(\cdot)$ be a strategy played by the other firm such that, holding $s_{-i}(\cdot)$ constant, $\frac{\partial \pi_i}{\partial s_i}$ can be written as a function quadratic in ε_{-i}. Given this we can express $\frac{\partial \pi_i}{\partial s_i}$ as

$$\frac{\partial \pi_i}{\partial s_i} = \alpha(s_i, \varepsilon_i) + \beta(s_i, \varepsilon_i)\varepsilon_{-i} + \gamma(s_i, \varepsilon_i)\varepsilon_{-i}^2 \tag{14}$$

By definition of $s_i^*(\varepsilon_i)$ and $\hat{s}_i(\varepsilon_i)$,

$$\frac{\partial E[\pi_i(s_i^*(\varepsilon_i), s_{-i}(\cdot), \epsilon_i)]}{\partial s_i} = \alpha(s_i^*, \varepsilon_i) + \beta(s_i^*, \varepsilon_i)E[\varepsilon_{-i}] + \gamma(s_i^*, \varepsilon_i)E[\varepsilon_{-i}^2] = 0 \tag{15}$$

$$\frac{\partial \pi_i(\hat{s}_i, s_{-i}(E(\varepsilon_{-i})), \varepsilon_i)}{\partial s_i} = \alpha(\hat{s}_i, \varepsilon_i) + \beta(\hat{s}_i, \varepsilon_i)E[\varepsilon_{-i}] + \gamma(\hat{s}_i, \varepsilon_i)E[\varepsilon_{-i}]^2 = 0 \tag{16}$$

Evaluating the first order condition for the heuristic equilibrium at the full information strategy, s_i^*, we have

$$\frac{\partial \pi_i(s_i^*, s_{-i}(E(\varepsilon_{-i})), \varepsilon_i)}{\partial s_i} = -\gamma(s_i^*, \varepsilon_i)Var(\varepsilon_{-i})$$

s_i^* fails to satisfy the first order condition for the heuristic equilibrium for nontrivial ε_{-i} and $\gamma(s_i^*, \varepsilon_i) \neq 0$.

Maintaining the earlier assumption that $\frac{\partial^2 \pi_i}{\partial s_i^2} < 0$, we know that for nontrivial ε_{-i},

$$\gamma(s_i^*, \varepsilon_i) > 0 \leftrightarrow s_i^* > \hat{s}_i. \tag{17}$$

24

■

Corollary 4

Proof. Let π_i be defined as in Proposition 3 and let $s_{-i}(\cdot)$ denote a strategy such that $\frac{\partial \pi_i}{\partial s_i}$ is quadratic in ε_{-i}.

Consider a function, π_i^H which is strictly more convex than π_i given by

$$\frac{\partial \pi_i^H}{\partial s_i} = \alpha(s_i, \varepsilon_i) - \tau(E[\varepsilon_{-i}^2]) + \beta(s_i, \varepsilon_i)\varepsilon_{-i} + (\gamma(s_i, \varepsilon_i) + \tau)\varepsilon_{-i}^2 \tag{18}$$

where $\tau > 0$ is a constant.

Note that

$$\frac{\partial E[\pi_i^H(s_i^*, s_{-i}(\cdot), \varepsilon_i)]}{\partial s_i} = \alpha(s_i^*, \varepsilon_i) - \tau(E[\varepsilon_{-i}^2]) + \beta(s_i^*, \varepsilon_i)E[\varepsilon_{-i}] + (\gamma(s_i^*, \varepsilon_i) + \tau)E[\varepsilon_{-i}^2] = 0 \tag{19}$$

which implies s_i^* maximizes the expected profits of both π_i and π_i^H.

By the fundamental theorem of calculus, we can express $\hat{s}_i$ as the value satisfying

$$\int_{\hat{s}_i}^{s_i^*} \frac{\partial^2 \pi_i(x, s_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i^2} dx = -\gamma(s_i^*, \varepsilon_i)Var(\varepsilon_{-i}) \tag{20}$$

Consider the value $\hat{s}_i^H$, defined as the analogue to $\hat{s}_i$ for π_i^H. From equation (20), we know $\hat{s}_i^H$ satisfies

$$\int_{\hat{s}_i^H}^{s_i^*} \frac{\partial^2 \pi_i^H(x, s_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i^2} dx = -(\tau + \gamma(s_i^*, \varepsilon_i))Var(\varepsilon_{-i}) \tag{21}$$

$$< \int_{\hat{s}_i}^{s_i^*} \frac{\partial^2 \pi_i(x, s_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i^2} dx$$

Since $\frac{\partial^2 \pi_i}{\partial s_i^2} = \frac{\partial^2 \pi_i^H}{\partial s_i^2}$, it follows that $\hat{s}_i^H < \hat{s}_i$. ■

Corollary 5

Proof. Let $s_{-i}(\cdot)$ be an arbitrary strategy for firm $-i$ with the property that $\frac{\partial \pi_i}{\partial s_i}$ is linear in ε_{-i}.

25

$\frac{\partial \pi_i}{\partial s_i}$ is linear in ε_{-i} implies $\gamma(s_i^*, \varepsilon) = 0$ equating equations (15) and (16).

Thus, for all s_i, ε_i

$$\frac{\partial E[\pi_i]}{\partial s_i} = \frac{\partial \pi_i(\hat{s}_i, s_{-i}(E[\varepsilon_{-i}]), \varepsilon_{-i})}{\partial s_i}. \tag{22}$$

implying

$$s_i^*(s_{-i}, \varepsilon_i) = \hat{s}_i(s_{-i}, \varepsilon_i) \; \forall \varepsilon_i. \tag{23}$$

∎

Corollary 6

Proof. This is an application of Corollary 5 using $g(\varepsilon_{-i})$ instead ε_{-i}.

Let $g(\varepsilon_{-i})$ be a function of ε_{-i} with $\varepsilon_{-i} \sim f(\varepsilon_{-i})$.

If $\frac{\partial \pi_i}{\partial s_i}$ is linear in $g(\varepsilon_{-i})$, we can write

$$\frac{\partial \pi_i}{\partial s_i} = \alpha(s_i, \varepsilon_i) + \beta(s_i, \varepsilon_i)g(\varepsilon_{-i}) \tag{24}$$

Note for arbitrary ε_i,

$$\frac{\partial E[\pi_i(s_i, s_{-i}, \varepsilon_i)]}{\partial s_i} = \frac{\partial \pi_i(s_i, s_{-i}(E(g(\varepsilon_{-i}))), \varepsilon_i)}{\partial s_i}.$$

Since s_i^* sets $\frac{\partial E[\pi_i(s_i, s_{-i}, \varepsilon_i)]}{\partial s_i}$ equal to 0 and $\hat{s}_i$ sets $\frac{\partial \pi_i(s_i, s_{-i}(E(g(\varepsilon_{-i}))), \varepsilon_i)}{\partial s_i}$ equal to 0, $s_i^* = \hat{s}_i$ for arbitrary ε_i. ∎

Proposition 7

Proof. Assume that s_i and s_{-i} are strategic complements, and $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}| \; \forall s_i, s_{-i}$.

For notational convenience, let

$$d_i(\varepsilon_i) = \frac{\partial \pi_i(\tilde{s}_i(\varepsilon_i), \tilde{s}_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i}$$

$$a_i(\varepsilon_i) = \int_0^1 \frac{\partial^2 \pi_i(\hat{s}_i(\varepsilon_i) + tx_i(\varepsilon_i), \hat{s}_{-i}(E[\varepsilon_{-i}]) + tx_{-i}(\varepsilon_{-i}), \varepsilon_i)}{\partial s_i^2} dt$$

$$b_i(\varepsilon_i) = \int_0^1 \frac{\partial^2 \pi_i(\hat{s}_i(\varepsilon_i) + tx_i(\varepsilon_i), \hat{s}_{-i}(E[\varepsilon_{-i}]) + tx_{-i}(\varepsilon_{-i}), \varepsilon_i)}{\partial s_i \partial s_{-i}} dt.$$

26

and define $d_{-i}(\varepsilon_{-i}), a_{-i}(\varepsilon_{-i})$ and $b_{-i}(\varepsilon_{-i})$ in n analagous fashion.

Thus, we can write (8) and the analogous expression for firm $-i$ as:

$$d_i(\varepsilon_i) = x_i(\varepsilon_i)a_i(\varepsilon_i) + x_{-i}(E[\varepsilon_{-i}])b_i(\varepsilon_i) \tag{25}$$

$$d_{-i}(\varepsilon_{-i}) = x_{-i}(\varepsilon_{-i})a_{-i}(\varepsilon_{-i}) + x_i(E[\varepsilon_i])b_{-i}(\varepsilon_{-i}). \tag{26}$$

Evaluating (25) at $E[\varepsilon_i]$ and (26) at $E[\varepsilon_{-i}]$, we can express:

$$x_i(E[\varepsilon_i]) = \frac{a_{-i}d_i - b_i d_{-i}}{a_i a_{-i} - b_i b_{-i}} \tag{27}$$

$$x_{-i}(E[\varepsilon_{-i}]) = \frac{a_i d_{-i} - b_{-i}d_i}{a_i a_{-i} - b_i b_{-i}}. \tag{28}$$

Signing the terms in (27) and (28): (1) $a_i, a_{-i} < 0$ by concavity of the profit function with respect to a firms own strategy, (2) $b_i, b_{-i} > 0$ by assumption that s_i, s_{-i} are strategic complements.

Now, consider the case in which $\tilde{s}_i > s_i^*$ and $\tilde{s}_{-i} > s_{-i}^*$. By equations (8) and (9), $d_i, d_{-i} > 0$. Furthermore, by the assumption $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}|$ $\forall s_i, s_{-i}$, $a_i a_{-i} - b_i b_{-i} > 0$. Thus, the numerators of expressions (27) and (28) are negative and the denominators are positive, implying $x_i(E[\varepsilon_i]) < 0, x_{-i}(E[\varepsilon_{-i}]) < 0$. Now consider $x_i(\varepsilon_i), x_{-i}(\varepsilon_{-i})$ for arbitrary $\varepsilon_i, \varepsilon_{-i}$. Rewriting (25) and (26), we have:

$$x_i(\varepsilon_i) = \frac{d_i(\varepsilon_i) - b_i(\varepsilon_i)x_{-i}(E[\varepsilon_{-i}])}{a_i(\varepsilon_i)} \tag{29}$$

$$x_{-i}(\varepsilon_{-i}) = \frac{d_{-i}(\varepsilon_{-i}) - b_{-i}(\varepsilon_{-i})x_i(E[\varepsilon_i])}{a_{-i}(\varepsilon_{-i})}. \tag{30}$$

The numerator of expressions (29) and (30) are positive and the denominators are negative. Thus, $x_i(\varepsilon_i) < 0, x_{-i}(\varepsilon_{-i}) < 0$ $\forall \varepsilon_i, \varepsilon_{-i}$, which implies for all $\varepsilon_i, \varepsilon_{-i}$

$$\hat{s}_i(\varepsilon_i) > \tilde{s}_i(\varepsilon_i) > s_i^*(\varepsilon_i), \hat{s}_{-i}(\varepsilon_{-i}) > \tilde{s}_{-i}(\varepsilon_{-i}) > s_{-i}^*(\varepsilon_{-i}). \tag{31}$$

If we consider the case in which $\tilde{s}_i < s_i^*$ and $\tilde{s}_{-i} < s_{-i}^*$, d_i, d_{-i} will be negative. In this case, $x_i(\varepsilon_i) > 0, x_{-i}(\varepsilon_{-i}) > 0 \ \forall \varepsilon_i, \varepsilon_{-i}$, which implies for all $\varepsilon_i, \varepsilon_{-i}$

$$\hat{s}_i(\varepsilon_i) < \tilde{s}_i(\varepsilon_i) < s_i^*(\varepsilon_i), \hat{s}_{-i}(\varepsilon_{-i}) < \tilde{s}_{-i}(\varepsilon_{-i}) < s_{-i}^*(\varepsilon_{-i}). \tag{32}$$

■

Proposition 8

Proof. Using the notation in Proposition 7, we can evaluate equations (27) and (28) at $E(\varepsilon_i)$ and $E(\varepsilon_{-i})$.

$$x_i(E[\varepsilon_i]) = \frac{a_{-i}d_i - b_i d_{-i}}{a_i a_{-i} - b_i b_{-i}}, x_{-i}(E[\varepsilon_{-i}]) = \frac{a_i d_{-i} - b_{-i}d_i}{a_i a_{-i} - b_i b_{-i}}. \tag{33}$$

Signing the terms in (33): (1) $a_i, a_{-i} < 0$ by concavity of the profit function with respect to a firms own strategy, (2) $b_i, b_{-i} < 0$ by assumption that s_i, s_{-i} are strategic substitutes.

Consider the case in which $\tilde{s}_i > s_i^*$ and $\tilde{s}_{-i} > s_{-i}^*$. By equations (8) and (9), $d_i, d_{-i} < 0$. Furthermore, by the assumption $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}| \ \forall s_i, s_{-i}, a_i a_{-i} - b_i b_{-i} > 0$. The denominators of the expressions in (33) are positive, but the sign of the numerators is ambiguous. Working out the cases evaluating firm i's strategy at $E(\varepsilon_i)^{12}$, we have:

$$\hat{s}_i \geq \tilde{s}_i > s_i^* \iff a_{-i}d_i \leq b_i d_{-i}$$

$$\tilde{s}_i > \hat{s}_i > s_i^* \iff (b_i d_{-i} < a_{-i}d_i < b_i d_{-i} + (\tilde{s}_i - s_i^*)[a_i a_{-i} - b_i b_{-i}] \tag{34}$$

$$\tilde{s}_i > s_i^* \geq \hat{s}_i \iff a_{-i}d_i \geq b_i d_{-i} + (\tilde{s}_i - s_i^*)[a_i a_{-i} - b_i b_{-i}].$$

Note that in the first two cases, $h_i > 0$ and in the last case $h_i \leq 0$. Similarly we can define a symmetric set of conditions for s_{-i}. ■

Proposition 9

Proof. Assume firm -i plays strategy profile $s_{-i}(\varepsilon_{-i})$. Let $\tilde{s}_i$ and s_i^* denote the best responses to $s_{-i}(\varepsilon_{-i})$ using the heuristic and full information best response functions, respectively. Writing

---

[12]If $\tilde{s}_i < s_i^*$ and $\tilde{s}_{-i} < s_{-i}^*$, the inequalities on the left side of the if and only ifs will be reversed, but the inequalities on the right side will be identical.

out the taylor expansion in equation (10), we can express, the difference in firm i's profits from using the heuristic best response function rather than using the full information best response function as

$$E\left[\pi_i(\tilde{s}_i, s_{-i}, \varepsilon_i) - \pi_i(s^*, s_{-i}, \varepsilon_i)\right] = E\left[\int_0^1 (1-t)\frac{\partial^2 \pi_i(\tilde{x})}{\partial s_i^2}(\tilde{s}_i - s_i^*)^2 dt|\varepsilon_i\right]$$
$$\leq 0,$$

where $\tilde{x} \equiv t(\tilde{s}_i, s_{-i}, \varepsilon_i) + (1-t)(s^*, s_{-i}, \varepsilon_i)$. ∎

Corollary 10

Proof. Assume firms are currently playing the heuristic equilibrium strategies. For a given realization of ε_i, the the benefits to firm i from deviating to the full information best response function is given by

$$E[\pi_i(s_i^*, \hat{s}_{-i}, \varepsilon_i) - \pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i)|\varepsilon_i] - C \tag{35}$$

where s_i^* is the full information best response of firm i to $\hat{s}_{-i}$ and C is the cognitive cost associated with using the full information best response. The lower bound follows directly from the taylor expansion in (10). ∎

Proposition 11

Proof. Let

$$0 \leq E\left[\int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_{-i}}h_{-i}dt|\varepsilon_i\right] + E\left[\int_0^1 (1-t)\frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i^2}h_i^2 dt|\varepsilon_i\right]$$
$$+ E\left[\int_0^1 (1-t)\frac{\partial^2 \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i \partial s_{-i}}h_i h_{-i} dt|\varepsilon_i\right]$$

for all ε_i. By definition, $E[\pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i) - \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)|\varepsilon_i] \geq 0$ for all ε_i. Taking the expectation across all ε_i, we find that $E[\pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i) - \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)] \geq 0$. The converse follows identically.

∎

Corollary 12

Proof. If sign($\frac{\partial \Pi_i}{\partial s_{-i}}$)$\neq$ sign(h_{-i}) all the terms in Proposition 11 are negative implying

$E[\pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i) - \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)|\varepsilon_i] < 0$ for all ε_i. Taking the expectation across all ε_i, we find that $E[\pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i) - \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)] < 0$. ∎

B Appendix: Cournot Competition

Here we examine a case in which there is a quadratic shock to the slope of the cost function and the goods are imperfect substitutes. We examine the effects of varying the degree of substitutability on the equilibrium quantity choice and also on the difference between the heuristic equilibrium and full-information equilibrium quantity choice.

With strategic substitutes, we expect there to be two countervailing forces acting upon the equilibrium quantity choice. First, in the absence of strategic substitution considerations, i.e., when only one firm is using the heuristic and the other firm is using the full information, then using the heuristic rather than the full information will lead the heuristic equilibrium quantity choice to diverge from the full-information equilibrium quantity choice. Second, when strategic considerations are present, i.e., when all firms are using the heuristic and therefore are reacting to each other's approximation, then we would expect the divergence to be offset by the fact that quantities are strategic substitutes. For example, if approximating leads one firm to decrease its quantity, then this will cause the other firm to raise its quantity in response, and when both firms are reacting to each other in this way, it may be the case that quantities in the heuristic equilibrium are actually closer to the full-information equilibrium than they would be if only one firm were using the approximation.

More formally, let the inverse demand function be given by

$$P_i(q_i, q_j) = a - b\left(q_i + \lambda q_j\right),$$

where $b \geq 0$, $\lambda \in [0, 1]$. Let the cost function be given by:

30

$$C_i(q_i) = \left(c_{i0} + \varepsilon_i^2\right) q_i.$$

Let the profit function be given by:

$$\pi_i = P_i(q_i, q_j)q_i - C_i(q_i) = (a - b(q_i + \lambda q_j))q_i - \left(c_{i0} + \varepsilon_i^2\right)q_i.$$

Thus, the derivative of the profit function with respect to the strategy (which is quantity) is:

$$\frac{\partial \pi_i}{\partial q_i} = P_i'(q_i, q_j)q_i + P_i(q_i, q_j) - C_i'(q_i), \tag{36}$$

which under the above functional form assumptions becomes:

$$\frac{\partial \pi_i(q_i, q_j, \varepsilon_i, \varepsilon_j)}{\partial q_i} = -2bq_i + a - \lambda b q_j - c_{i0} - \varepsilon_i^2. \tag{37}$$

Total differentiating the above FOC evaluated at the equilibrium, one can solve for the degree of strategic substitutability, which is given by:

$$\frac{\partial q_i}{\partial q_j} = \frac{-E\left[\frac{\partial^2 \pi_i}{\partial q_i \partial q_j}\right]}{E\left[\frac{\partial^2 \pi_i}{\partial q_i^2}\right]} = -\frac{\lambda}{2}.$$

So the greater is λ, the greater the degree of strategic substitutability.

B.1   Full-information equilibrium

The full-information FOC $E\left[\frac{\partial \pi_i}{\partial q_i}|\varepsilon_i\right] = 0$ implies that the best response function for i as a function of j's strategy is:

$$q_i = \frac{1}{2b}\left(a - \lambda b E\left[q_j^*(\varepsilon_j)\right] - c_{i0} - \varepsilon_i^2\right). \tag{38}$$

Taking the conditional expectation of (38) with respect to ε_i, we get:

$$E\left[q_i^*(\varepsilon_i)\right] = \frac{1}{2b}\left(a - \lambda b E\left[q_j^*(\varepsilon_j)\right] - c_{i0} - E\left[\varepsilon_i^2\right]\right).$$

Solving the system of equations for $E\left[q_i^*(\varepsilon_i)\right]$ and $E\left[q_j^*(\varepsilon_j)\right]$, we get:

$$E\left[q_i^*(\varepsilon_i)\right] = \frac{1}{(4-\lambda^2)\,b}\left((2-\lambda)\,a + \lambda\left(c_{j0} + E\left[\varepsilon_j^2\right]\right) - 2\left(c_{i0} + E\left[\varepsilon_i^2\right]\right)\right).$$

Substituting back into (38), the full-information equilibrium, where each player is best responding to the other, is:

$$q_i^* = \frac{1}{2b}\left(a - \left(c_{i0} + \varepsilon_i^2\right) - \frac{\lambda}{4-\lambda^2}\left(-\lambda\left(a - \left(c_{i0} + E\left[\varepsilon_i^2\right]\right)\right) + 2\left(a - \left(c_{j0} + E\left[\varepsilon_j^2\right]\right)\right)\right)\right) \quad \forall i.$$

(39)

Equation (39) gives the equilibrium strategy for each player i as a function of its shock ε_i^d. We've written this equation as a weighted sum of the expected maximum markup terms (i.e., expected maximum price minus marginal cost) $a - \left(c_{i0} + \varepsilon_i^2\right)$, $a - \left(c_{i0} + E\left[\varepsilon_i^2\right]\right)$ and $a - \left(c_{j0} + E\left[\varepsilon_j^2\right]\right)$.

B.2 Heuristic equilibrium

The heuristic equilibrium FOCs are $\frac{\partial \pi_i(\widehat{q}_i, \widehat{q}_j(E[\varepsilon_j]), \varepsilon_i)}{\partial q_i} = 0$, yielding as the heuristic equilibrium:

$$\widehat{q}_i = \frac{1}{2b}\left(a - \left(c_{i0} + \varepsilon_i^2\right) - \frac{\lambda}{4-\lambda^2}\left(-\lambda\left(a - \left(c_{i0} + (E\left[\varepsilon_i\right])^2\right)\right) + 2\left(a - \left(c_{j0} + (E\left[\varepsilon_j\right])^2\right)\right)\right)\right) \quad \forall i.$$

(40)

Define h_i as the difference between the heuristic quantity choice $\widehat{q}_i$ that would be chosen if both firms played the heuristic strategy and the full-information quantity choice q_i^* that would be chosen if both firms played the full-information strategy. Then,

$$h_i \equiv \widehat{q}_i - q_i^*$$

$$= \frac{-\lambda}{4 - \lambda^2} \left(\lambda \left((E\left[\varepsilon_i\right])^2 - E\left[\varepsilon_i^2\right] \right) - 2 \left((E\left[\varepsilon_j\right])^2 - E\left[\varepsilon_j^2\right] \right) \right).$$

Suppose the shocks were drawn from the same distribution. Then,

$$h_i = \frac{\lambda}{2 + \lambda} \left((E\left[\varepsilon_i\right])^2 - E\left[\varepsilon_i^2\right] \right),$$

which is increasing in λ over $\lambda \in [0, 1]$. Thus, the greater the strategic substitutability, the greater the divergence between the full-information and heuristic strategies.

C Appendix: Bertrand Competition on the Hotelling Line

Consider a differentiated products Bertrand model in which firms receive private information regarding their marginal cost of production and choose prices to maximize expected profit. Consistent with the assumption of our model, the unobservable "shock" to firm i only affects firm $-i$ through the choice of strategy by firm i.

Let two firms be located at either end of the unit interval and be denoted firm 0 and firm 1 by their respective positions.

Let v_0, p_0 and c_0^α denote the value, price and marginal cost of firm 0's product. We take c_0 to be a cost shock for firm 0 drawn from a distribution and α a parameter which determines whether firm 0's marginal costs are linear, convex or concave in the cost shock c_0. Define v_1, p_1 and c_1^α in a like manner for firm 1. Let the utility of a consumer located at γ and looking to

33

purchase a single unit of the good be given by

$$U(\gamma) = \begin{cases} v_0 - p_0 - t\gamma & \text{if good 0 is purchased.} \\ v_1 - p_1 - t(1 - \gamma) & \text{if good 1 is purchased.} \\ 0 & \text{if neither 0 nor 1 is purchased.} \end{cases} \tag{41}$$

For purposes of this exercise, assume that v_0 and v_1 are sufficiently high to ensure that all consumers on the unit interval are willing to purchase and that the travel cost $t = 1$. Given v, p and c, define $\hat{\gamma}$ as the location of the consumer who is indifferent between purchasing good 1 and good 0,

$$\hat{\gamma} = \frac{v_0 - v_1 - (p_0 - p_1) + 1}{2}. \tag{42}$$

For simplicity, assume $v_0 = v_1$ and $c_0, c_1 \sim U[0, 1]$. In addition, assume that c_0 and c_1 are independent. In this case, we can express the expected profits of firm 0 as

$$E[\pi_0] = E[\frac{1 - (p_0 - p_1)}{2}(p_0 - c_0^\alpha)]. \tag{43}$$

C.1   Example 1: Linear Shock

First consider the case in which the shock to costs enters linearly. That is, consider the case in which $\alpha = 1$. That is, for given vectors v, p, c, let the expected profit of firm 0 be given by:

$$E[\pi_0] = E[\frac{1 - (p_0 - p_1)}{2}(p_0 - c_0)]. \tag{44}$$

Solving for the full-information equilibrium strategies, we have

$$p_0^* = \frac{c_0}{2} + \frac{5}{4} \tag{45}$$

$$p_1^* = \frac{c_1}{2} + \frac{5}{4}. \tag{46}$$

It is easy to check given the strategy $p_{-i}^*(c_{-i})$ chooses, it is optimal for i to play $p_i^*(c_i)$.

34

Now we consider the heuristic based strategy. In this case, we'll consider firm 0 assuming that firm 1 receives a mean realization of cost and plays strategy $p_1^*(E(c_1)) = p_1^*(\frac{1}{2})$. First note that firm 0's profit (if firm 1 plays p_1^*) is given by:

$$\pi_0 = \left(\frac{1-p_0+p_1}{2}\right)(p_0-c_0) = (\frac{5}{8} + \frac{c_1}{4} + \frac{1-p_0}{2})(p_0-c_0). \tag{47}$$

and

$$\frac{\partial \pi_0}{\partial p_0} = \frac{9}{8} + \frac{c_0}{2} - p_0 + \frac{c_1}{4}. \tag{48}$$

Note that the $\frac{\partial \pi_0}{\partial p_0}$ is linear in the shock to firm 1, which from Corollary 5 means that firm 0's full-information strategy and heuristic strategy are the same. Indeed, if firm 0 approximates by assuming that firm 1 receives a mean realization of cost and plays strategy $p_1^*(E(c_1)) = p_1^*(\frac{1}{2})$, we find:

$$\hat{p}_0 = \frac{c_0}{2} + \frac{5}{4}, \tag{49}$$

which is equal to the expected profit maximizing strategy.

C.2 Example 2: Nonlinear Shock

C.2.1 Full-information equilibrium

For given vectors v, p, c, let the expected profit of firm 0 be given by:

$$E[\pi_0] = E\left[\frac{1-(p_0-p_1)}{2}(p_0-c_0^\alpha)\right]. \tag{50}$$

Solving for the full-information equilibrium strategies, we have

$$p_0^* = \frac{c_0^\alpha}{2} + \frac{2\alpha+3}{2\alpha+2} \tag{51}$$

$$p_1^* = \frac{c_1^\alpha}{2} + \frac{2\alpha+3}{2\alpha+2}. \tag{52}$$

35

C.2.2 Heuristic equilibrium

Now, we consider the heuristic equilibrium strategies in which players play $\hat{p}_0 = argmax\pi_0(p_0, \hat{p}_1(E(c_1)), c_0)$ and $\hat{p}_1 = argmax\pi_1(p_1, \hat{p}_0(E(c_0)), c_1)$. Assuming firm 1 plays $\hat{p}_1$, profit of firm 0 (when firm 1 faces an "average" cost shock) is given by:

$$\pi_0 = \frac{1 - p_0 + \hat{p}_1(E(c_1))}{2}(p_0 - c_0^\alpha). \tag{53}$$

Solving for the heuristic equilibrium strategies, we have

$$\hat{p}_0(E[c_0])^* = 1 + (\frac{1}{2})^{\alpha+1} + \frac{c_0^\alpha}{2} \tag{54}$$

$$\hat{p}_1(E[c_1])^* = 1 + (\frac{1}{2})^{\alpha+1} + \frac{c_1^\alpha}{2}. \tag{55}$$

C.3 Comparing the Heuristic and Full-Information Equilibria

Solving for expected profits for firm 0 and firm 1 in the full information equilibrium and the heuristic equilibrium, we have

$$\pi_0^* = (\frac{\frac{c_1^\alpha}{2} - \frac{c_0^\alpha}{2} + 1}{2})(\frac{2\alpha + 3}{2\alpha + 2} - \frac{c_0^\alpha}{2}) \tag{56}$$

$$\pi_1^* = (\frac{\frac{c_0^\alpha}{2} - \frac{c_1^\alpha}{2} + 1}{2})(\frac{2\alpha + 3}{2\alpha + 2} - \frac{c_1^\alpha}{2}). \tag{57}$$

and

$$\widehat{\pi}_0 = (\frac{\frac{c_1^\alpha}{2} - \frac{c_0^\alpha}{2} + 1}{2})(\frac{2^{\alpha+1} + 1}{2^{\alpha+1}} - \frac{c_0^\alpha}{2}) \tag{58}$$

$$\widehat{\pi}_1 = (\frac{\frac{c_0^\alpha}{2} - \frac{c_1^\alpha}{2} + 1}{2})(\frac{2^{\alpha+1} + 1}{2^{\alpha+1}} - \frac{c_1^\alpha}{2}). \tag{59}$$