

On the use of heuristics to approximate competitors' private information

C.-Y. Cynthia Lin and Erich J. Muehlegger¹

December 2012

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 strategy depends on: (1) the convexity of a firm's marginal profits with respect to their opponents' private information, and (2) whether firm strategies are complementary or substitutable. 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

¹Lin: Agricultural and Resource Economics, University of California at Davis, One Shields Avenue, Davis, CA 95616; cclin@primal.ucdavis.edu; (530) 752-0824. Muehlegger: John F. Kennedy School of Government, Harvard University, 79 John F. Kennedy Street, Cambridge, MA 02138; Erich_Muehlegger@harvard.edu; (617) 335-6083. We thank Drew Fudenberg, Daniel Hojman, Ting Liu, Gabriel Weintraub, and participants at the 2007 International Industrial Organization Conference for helpful comments and conversations. All errors are our own.

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.

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 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, 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.

²Policy 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.

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 Bayesian 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 about the distributions of each other's private information, and therefore would have incentives to disclose 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 model is that it is agnostic about the actual mechanism underlying the why firms may use heuristics; our results therefore do not hinge on the verity of any particular story, but rest only on the assumption that some story exists that leads the firms to use heuristics. Thus, the model is general enough to capture the reduced-form implications of any of a number of stories for why firms may use heuristics rather than the full information, including information acquisition costs and limited cognition.

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), Shapiro (1986), Raith (1996), and Jansen (2008), who examine the incentives of firms to share information with their competitors, for example through trade associations. In previous lit-

erature concerning the incentives to share information, results seem to depend sensitively on the specific assumptions of the model: a change from Cournot to Bertrand, from substitutes to complements, from demand to cost uncertainty, or from a common value to private values, may lead to completely different outcomes (Vives, 1990; Raith, 1996). Raith (1996) presents a general model which encompasses many of the existing literature on information sharing as special cases and is able to explain the factors leading to varying results. Similar to Raith (1996), our model allows for both Cournot and Bertrand models as special cases, and we allow for either demand or cost uncertainty. Unlike Raith (1996) and other papers in the literature (see Jansen, 2008), we use a general form for the firm's profit function instead of specifying a particular functional form. Like the previous literature, our results depend on the details of product market interaction, and in particular on the degree of strategic complementarity or substitutability. Our results also depend on the actual distribution of the private information, which determines whether heuristics overestimate or underestimate the expected profit maximizing Bayesian strategy. Although our approach is specific in assuming that the heuristic is to use the mean of the distribution of private information, our approach is general in that we do not assume what the full distribution is. Thus, in contrast to the previous literature, our model makes minimal assumptions about functional form and no assumptions about the nature of competition, and therefore 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 costly 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. Lacetera, Pope and Sydnor (forthcoming) find that heuristic information processing affects the used car market, and in particular that the tendency to focus on the left-most digit of a number affects how used car buyers incorporate odometer values in their purchase decisions. 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). Minor (2011) analyzes the implications of players' using heuristics in a contest environment. He finds that when players react optimally against the average opponent strategies are "more extreme" than when optimally reacting under full consideration of the distribution of types. He also presents experimental evidence in favor of subjects using these kinds of heuristics is presented.

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. Furubotn (2009) develops a non-maximizing approach to the theory of the firm based on the presumption that a systematic search over any significant fraction of the potentially usable operating configurations found in the firm's technological or organizational choice set cannot possibly be undertaken, and that, consequently, some simplified, less costly, decision method must be used. 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.

Consistent with the finding of the previous literature on information exchange that results typically depend on the details of product market interaction, we find that the degree to which the heuristic strategy differs from the Bayesian strategy depends on strategic complementarity or substitutability. Consistent with the finding of the previous literature that results typically

³Etzioni (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.

depend on the nature of the information to be exchanged, we find that even though we allow for both cost and demand shocks, the degree to which the heuristic strategy differs from the Bayesian strategy depends on how the shock enters into the profit function, and in particular on 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.

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 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. Sections 5 and 6 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. Section 7 concludes.

2 Basic Model

Suppose there are two firms in the market. The type ε_i of each firm i embodies the firm's private information.⁴ Given its private information ε_i , each firm i chooses a strategy s_i so as

⁴We denote the type as ε_i since the particular form of private information we focus on is a perturbation to the payoff function, and ε_i is commonly used to denote such shocks.

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.⁵ 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 heuristics could be evaluated in a similar fashion.

Our model is general enough to allow the shocks to profit to be any type of shock to profit. We model the shocks as independently distributed across firms, as independent shocks translate to common empirical models. These shocks can be either cost shocks or firm-specific demand shocks. In principle, the demand shocks could be correlated, so long as they have some firm-specific component.

We assume that the profit function $\pi_i(s_i, s_{-i}, \varepsilon_i)$ satisfies the following assumptions: $\pi_i(s_i, s_{-i}, \varepsilon_i)$ is second-order differentiable in $s_i \forall s_{-i}, \varepsilon_i$, $\frac{\partial \pi_i(\cdot)}{\partial s_i} \Big|_{s_i=0} > 0 \forall s_{-i}, \varepsilon_i$, and $\lim_{s_i \rightarrow \infty} \frac{\partial \pi_i(\cdot)}{\partial s_i} < 0 \forall s_{-i}, \varepsilon_i$.

We now present two different types of strategies, *Bayesian strategies* and *heuristic strategies*, which correspond to cases in which neither or both firms choose to approximate, respectively.

⁵The 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} .

2.1 Bayesian 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. This is the standard Bayesian approach which we use as a benchmark to which we compare strategies and profits using the heuristic. For an opponent's strategy $s_{-i}(\cdot)$, the Bayesian strategy $s_i^*(\cdot)$ for firm i is given by the best response to $s_{-i}(\cdot)$:

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

The first-order necessary condition is given by:

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

When both firms employ strategies that best respond to each other's Bayesian strategy, the equilibrium profile $(s_i^*(\cdot), s_{-i}^*(\cdot))$ that arises is the *Bayesian 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.⁶ 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 estimating or understanding higher moments (Weitzman, 2009), although the basic idea could be generalized to a model in which firms select the number of mo-

⁶An econometrician may use such an approximation for either reason as well.

ments with which to formulate their strategy.⁷ 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 $s_{-i}(E[\varepsilon_{-i}])$:

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

The first-order necessary condition is given by:

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

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*.⁸

There are many possible reasons why firms may use heuristics, including information acquisition costs and limited cognition. Although calculating the heuristic strategy may impose cognitive costs, it is still computationally less difficult than computing the Bayesian strategy, which requires computing an integral to calculate an expected value. Thus, firms with limited cognition may be incapable of computing the Bayesian strategy, and therefore use the heuristic strategy instead. Moreover, the heuristic strategy still requires less information than the

⁷It is possible that firms may also choose to exchange the “full” information about their actual realized type. We choose to focus on the exchange of information on the distribution rather than the “full” information for several reasons. First, exchanges of the “full” information are most relevant when considering the incentives of a group of firms to voluntarily share information. In the other examples that we consider (e.g., the incentives of a group of firms to make information more costly to collect, or the incentives of a firm unilaterally collecting costly private information or the decision by an econometrician to model a simplified version of the firm optimization problem), this alternative is less relevant. Second, this issue is more fully explored in the existing literature (see, e.g., Amir, Jin and Troege, 2010; Raith, 1996). Third, models of the exchange of full information in which players have to commit ex ante in the first stage to reveal their private information before the private information is realized in the second stage have a possible ex post commitment problem, since once the shock has been realized some firms might have an incentive to reveal their type even if they committed ex ante not to. In such a model, when an opponent does not reveal his type ex post, this in of itself is information that enables a firm a first to update its prior. Since in this situation, the firm still does not know the type of its opponent, but rather has at best a (posterior) distribution, we are back to the situation we analyze in our paper where firms may base their decision either on a distribution or on an approximation.

⁸When 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 acquisition and cognitive costs sufficient to ensure existence of the heuristic equilibria.

Bayesian strategy, so information acquisition costs are still a possible explanation for the use of heuristics even if limited cognition is not.

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 Bayesian strategies, and (2) the degree to which a firm's profits are affected by deviations from the Bayesian 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 Bayesian equilibrium.

First we compare the heuristic strategy to the optimal Bayesian strategy holding the opponent's strategy constant. We then extend the results about best-response functions to results about the equilibrium strategies when both firms use the heuristic approach or both firms use the optimal Bayesian approach. In addition to the conditions on above, we find that the degree to which equilibrium strategies differ is also sensitive to whether firms' strategies are strategic complements or strategic substitutes.

3.1 Comparing Best Response Functions

We begin by holding the opponent's choice of strategy constant and examine the best response functions using the heuristic. We first find that if the derivative of a firm's profit function with respect to its own strategy is convex (concave) with respect to the opponent's private information, the heuristic strategy will be less (greater) than the Bayesian strategy.

The intuition for this result is relatively straightforward. When the derivative of a firm's profit function with respect to its own strategy is nonlinear in an opponent's private information, deviations in an opponent's private information from the mean shock have an asymmetric effect on a firm's profits. A firm using the heuristic will fail to incorporate this asymmetry into its

decision-making and, consequently, will choose different strategies than a firm using the optimal Bayesian approach would.

On the other hand, when the derivative of a firm's profit function is linear in an opponent's private information, no asymmetry exists, and the heuristic and Bayesian approaches yield identical strategies. In this case, 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 in a single-agent dynamic programming problem under uncertainty that, when the criterion function is quadratic, the mean is a certainty equivalent and therefore a sufficient statistic for the entire distribution (Simon, 1956, 1959).

In Proposition 1, 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 Bayesian 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 1 *Let $s_{-i}(\cdot)$ be an arbitrary strategy for firm $-i$ such that $\frac{\partial \pi_i}{\partial s_i}$ is quadratic in ε_{-i} . If $\frac{\partial \pi_i}{\partial s_i}$ is linear (strictly concave / strictly convex) in ε_{-i} , $s_i^*(\varepsilon_i)$ is equal to (less than / greater than) $\hat{s}_i(\varepsilon_i)$ for all s_i, ε_i .*

Proof: See Appendix

Proposition 1 identifies a sufficient condition under which the strategies chosen to solve the heuristic first-order condition differ from those chosen in the benchmark Bayesian case.⁹ In the most simple case, if $\frac{\partial \pi_i}{\partial s_i}$ is linear in the opponent's cost shock, using the heuristic strategy is equivalent to the Bayesian approach. Although the heuristic strategy doesn't use information

⁹Although 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 analogous 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 .

about the higher-order moments of a competitor’s shock, if the shock is linear potential realizations of a competitor’s shock above the mean offset potential realizations of the competitor’s cost shock below the mean. Consequently, the Bayesian strategy is identical to the heuristic strategy regardless of the distribution of a competitor’s cost shock. As examples, we present a simple Cournot and differentiated-products Bertrand competitions in Sections 5 and 6. When firms face linear cost shocks with linear demand curves, $\frac{\partial \pi_i}{\partial s_i}$ is linear in the opponent’s cost shock.

If, on the other hand, $\frac{\partial \pi_i}{\partial s_i}$ is strictly convex or strictly concave in firm $-i$ ’s shock, firm i ’s best response using the heuristic will differ from the optimal Bayesian best-response. Two factors affect the degree to which the heuristic based strategy under- or over-estimates the Bayesian best-response benchmark. The first is rather straightforward. For a given convexity or concavity, the degree of miscalculation is positively related to the variance of the shocks facing the opponent. If the variance of a firm’s shocks increase, all else equal, the heuristic strategy will differ more from the Bayesian best response benchmark. 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 Bayesian FOC will differ more from the strategy solving the heuristic FOC.

3.2 Comparing Equilibrium Strategies

We now extend the results to examine the equilibrium strategies when both firms use the heuristic to make strategic decisions. Specifically, we compare the strategies chosen in the Bayesian benchmark, (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.¹⁰

¹⁰Implicitly, 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 Bayesian strategy.

Let $h_i(\varepsilon_i)$ denote the difference between firm i 's heuristic and Bayesian equilibrium strategies: $h_i(\varepsilon_i) = \hat{s}_i(\varepsilon_i) - s_i^*(\varepsilon_i)$.

Proposition 2 *Assume s_i and s_{-i} are strategic complements, $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}| \forall s_i, s_{-i}$. If the use of the heuristic leads each firm to increase (decrease) their strategic variables, then the strategies chosen in the heuristic equilibrium are greater (less) than the strategies chosen in the Bayesian equilibrium.*

Proof: See Appendix

The assumption that $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}|$, is relatively innocuous and, as we illustrate, holds for Cournot competition and for basic formulations of price competition on the Hotelling line.¹¹ The intuition for the condition is straightforward - a firm's marginal profit typically changes more with a firm's own strategy since a change directly affects both revenues and costs, whereas a commensurate change in a competitor's strategy affects the former, but typically not the latter. This assumption ensures that a game with strategic complements has a stable equilibrium. It is also a sufficient condition for the equilibrium to be unique (Shapiro, 1989).

When the heuristic causes both firms to adjust their strategy in 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 Bayesian equilibrium strategies than the strategies adopted when a single firm approximates and the second firm plays a Bayesian strategy. Again, the intuition is relatively straightforward. If using the heuristic led each firm to slightly increase (decrease) the value of their strategic variable holding their competitor's strategy fixed, strategic complementarity would further increase the incentive for the firms to increase (decrease) their strategic variable. As a consequence, under these conditions, h_i and h_{-i} will have the same sign.

Following a similar intuition, it is easy to see why the result is ambiguous if strategies are substitutes. Moreover, with strategic substitutes, $h_i(\varepsilon_i)$ and $h_{-i}(\varepsilon_{-i})$ may be of different sign. All else equal, if a competitor's use of a heuristic leads them to increase (decrease) their strategic

¹¹In 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), $\pi_1 = \frac{1+p_2-p_1}{2t} p_1 \Rightarrow \frac{\partial^2 \pi_1}{\partial p_1^2} = \frac{-1}{t}, \frac{\partial^2 \pi_1}{\partial p_1 \partial p_2} = \frac{1}{2t}$.

variable, a firm has the incentive to move their strategic variable in the opposite direction due to strategic substitutability. Whether the direct effect of a firm's use of the heuristic or the indirect strategic effect of the competitor's use of the heuristic dominates depends on the relatively magnitude of the two effects.

We illustrate these results using simple Cournot and differentiated-product Bertrand examples in Sections 5 and 6.

4 Heuristic Use and Profits

Now we consider the effect of heuristic use on the profits realized by the firms. Our objective is to define the conditions under which firms would have an incentive to attempt to coordinate on either the heuristic equilibrium or the Bayesian equilibrium. Holding a competitor's strategy constant, it is important to note that a firm always has the incentive to adopt the Bayesian strategy if there are no fixed costs of cognition or information gathering. 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 Bayesian strategy must be weakly greater than the profits firm i could earn by unilaterally deviating to the heuristic best response.

In equilibrium, though, this is not always true. To compare the two equilibria, we sign the terms of a Taylor expansion of firm i 's profit function. Consider the difference in expected profits earned by firm i under the two equilibrium concepts given by:

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

which can be rewritten as:

$$E\left[\int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_i} h_i dt + \int_0^1 \frac{\partial \pi_i(s_i^* + th_i, s_{-i}^* + th_{-i}, \varepsilon_i)}{\partial s_{-i}} h_{-i} dt\right], \quad (5)$$

where, as before, $h_i(\varepsilon_i) = \hat{s}_i(\varepsilon_i) - s_i^*(\varepsilon_i)$.

The expansion in (5) calculates the difference in profits under the two equilibrium concepts. The first term in the expansion is the product of the average derivative of π_i with respect to s_i along the ray between (s_i^*, s_{-i}^*) and $(\hat{s}_i, \hat{s}_{-i})$ and h_i , while the second term in the expansion is the product of the average derivative of π_i with respect to s_{-i} along the ray between (s_i^*, s_{-i}^*) and $(\hat{s}_i, \hat{s}_{-i})$ and h_{-i} . Consequently, the different in profits depends on the sign and magnitude of the two terms in the expansion.

The signs of the first and second terms are relatively intuitive. For the first term, $s_i^*(s_{-i})$ is the Bayesian best response function to each value of s_{-i} . By definition, the first derivative of i 's expected profit with respect to s_i is 0 if evaluated at s_i^* . Moreover, any deviation (such as the use of a heuristic) from s_i^* will cause profits to fall since the first derivative of i 's profit with respect to s_i is positive if evaluated at $s_i < s_i^*$ and negative if evaluated at $s_i > s_i^*$. The first term is therefore negative. The sign of the second term depends on whether the approximation moves firm $-i$'s strategic variable in a direction that increases firm i 's profits. For example in the Cournot (Bertrand) case, this would correspond to the situation where approximation by a competitor leads them to decrease production (increase prices).

One can therefore think of the difference between (s_i^*, s_{-i}^*) and $(\hat{s}_i, \hat{s}_{-i})$ as having two effects. The first effect, which corresponds to the first term in (5) and is weakly negative, arises when firm i 's approximation causes it to deviate from the Bayesian equilibrium. The second effect, which corresponds to the second term in (5) and can be of either sign, is the effect of a competitor's approximation on firm i 's profits and depends on the signs of $\frac{\partial \pi_i}{\partial s_{-i}}$ and h_{-i} . The overall impact depends on which of these effects dominates. A special case illustrates the intuition quite nicely. Consider, for example, the case where approximation leads only firm i to

change its strategy. In this case, the second term of (5) is zero. Approximation in this context causes firm i to deviate from the Bayesian best response, strictly lowering expected profits. This suggests to a relatively simple necessary condition for the firms to prefer the heuristic to the Bayesian equilibrium.

Proposition 3 *If the firms prefer to coordinate on the heuristic equilibrium (i.e. $E[\pi_i(\hat{s}_i, \hat{s}_{-i}, \varepsilon_i) - \pi_i(s_i^*, s_{-i}^*, \varepsilon_i)] \geq 0$), then $\text{sign}(\frac{\partial \Pi_i}{\partial s_{-i}}) = \text{sign}(h_{-i})$ and $\text{sign}(\frac{\partial \Pi_{-i}}{\partial s_i}) = \text{sign}(h_i)$.*

For firms to prefer the heuristic equilibrium, it must be the case that the use of approximation by $-i$ shifts $-i$'s strategy in a direction which increases i 's expected profits and vice-versa. This also suggests a possible sufficient condition for the firms to collectively prefer the heuristic equilibrium to the Bayesian equilibrium. Assuming the necessary condition above holds for both firms, the second term in (5) measures the increase in profits associated with the firm i 's opponent changing its strategy in a direction that benefits the firm i . The first term captures the degree to which approximation by firm i lowers its own profit because it is no longer playing the Bayesian best-response to its opponent's strategy.

A possible sufficient condition relates to the relative magnitudes of h_i and h_{-i} . Consider, for example, a case in which approximation leads firm i to barely deviate from the Bayesian best response by causes firm $-i$ to change its strategy substantially. In a Bertrand game, this corresponds to an opponent that raises prices substantially when they approximate, but the firm's own strategy remains relatively close to what would be optimal if they used all the information. In the notation of (5), h_{-i} is large and $\hat{s}_i(\hat{s}_{-i})$ is close to $s_i^*(\hat{s}_{-i})$. In this case, firm i is likely to prefer the heuristic equilibrium since the second term likely outweighs the first term in (5) for firm i . In contrast, firm $-i$ is unlikely to prefer the heuristic equilibrium since the second term in (5) for firm $-i$ (driven by h_i) is likely to be close to zero, while the first term is likely to be large and negative. For the approximation equilibrium to be attractive to both firms, the necessary condition above must hold. In addition, the magnitudes of h_i and h_{-i} may need to be close to ensure that both firms prefer the heuristic equilibrium to the Bayesian equilibrium. 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 competitors.

5 Example: Cournot Competition

We now illustrate the intuition behind our results by illustrating them in the context of simple Cournot and Bertrand models. Vives (1999) provides a detailed discussion of (linear) Cournot and Bertrand models with asymmetric information.

Consider a Cournot 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 .

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 Bayesian 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 Bayesian equilibrium than they would be if only one firm were using the approximation.

More formally, let the inverse demand function be given by:

$$P_i(q_i, q_j) = a - b(q_i + q_j),$$

where $b \geq 0$. Let the cost function be given by:

$$C_i(q_i) = (c_i^\alpha) q_i.$$

where c_{i0} is private information known by the firm, but unobserved by the rival and α allows for arbitrary concavity or convexity.

For simplicity, we assume that both c_i, c_j and independently drawn from $U[0, 1]$ and that $a > 1$.

5.1 Bayesian equilibrium

Conditional on a firm i 's own cost shock, expected profit of firm i is given by:

$$E[\pi_i | c_i] = [a - b[q_i(c_i) + E[q_j(c_j)]] - c_i^\alpha] q_i(c_i). \quad (6)$$

Consequently, the best response function for a firm playing a Bayesian strategy is:

$$q_i^*(c_i) = \frac{a - bE[q_j(c_j)] - c_i^\alpha}{2b}, \quad (7)$$

which simplifies into the tradition result for $\alpha = 1$ and certain, known c_j . Solving for the Bayesian equilibrium strategy for firm i , we have:

$$q_i^*(c_i) = \frac{2a - E[c_i^\alpha] + 2E[c_j^\alpha] - 3c_i^\alpha}{6b}. \quad (8)$$

The intuition is relatively straightforward. The expectation of c_i^α enters into the strategy, since firm j does not observe c_i .

Under the assumption that c_i, c_j are independently drawn from $U[0, 1]$, simplifies to:

$$q_i^*(c_i) = \frac{2a + \frac{1}{1+\alpha} - 3c_i^\alpha}{6b}. \quad (9)$$

5.2 Heuristic equilibrium

Now, we consider the strategies chosen if both players use the heuristic to guide their decisions. Under the assumption that firm j faces an average shock, the best response function for firm i is given by:

$$\widehat{q}_i(c_i) = \frac{a - bq_j(E[c_j]) - c_i^\alpha}{2b}. \quad (10)$$

Solving for the heuristic equilibrium strategy for firm i , we have:

$$\widehat{q}_i(c_i) = \frac{2a - E[c_i]^\alpha + 2E[c_j]^\alpha - 3c_i^\alpha}{6b}.$$

Under the assumption that c_i, c_j are independently drawn from $U[0, 1]$, simplifies to:

$$\widehat{q}_i(c_i) = \frac{2a + \frac{1}{2^\alpha} - 3c_i^\alpha}{6b}. \quad (11)$$

5.3 Comparing Heuristic and Bayesian Equilibrium Strategies

First, we illustrate the results of Proposition 2 in the context of Cournot competition. In the simple example above, the parameter α reflects the convexity or concavity of a firm's marginal profit with respect to the private information parameters c_i and c_j .

The simple example above, where c_i and c_{-i} are i.i.d. $U[0,1]$, is a case in which the proposition 2 holds, despite the firms competing in strategic complements. When the opponent's private information enters a firm's first-order-condition linearly, ($\alpha = 1$), the heuristic and Bayesian equilibria are identical. In each equilibrium, firm i chooses:

$$q_i(c_i) = \frac{2a + \frac{1}{2} - 3c_i}{6b}. \quad (12)$$

When a firm's first-order-condition is convex in an opponent's private information, the firms choose lower levels of output than in the Bayesian equilibrium. Similarly, if a firm's first-

order-condition is concave in an opponent's private information, the firms choose higher levels of output than in the Bayesian equilibrium.

Proposition 2 does not hold generally for games in which firms strategies are substitutes. An easy illustration of this is to examine the case in which opponent's private information enters in quadratically to a firm's first-order-condition, $\alpha = 2$. In this case, the difference between firm i 's strategy in the heuristic and Bayesian equilibria is given by:

$$h_i(c_i) = \text{var}(c_i) - 2\text{var}(c_{-i}). \quad (13)$$

In this example, so long as the variance of a firm's private information is similar to the variance of the opponent's private information, Proposition 2 holds. But if the variance of a firm's own private information is sufficiently high, the strategic response to a competitor using a heuristic outweighs the direct effect of the heuristic on firm i 's choice of strategy. In the context of Cournot competition in quantities, if the variance of a firm's own information is sufficiently high, the opponent produces much less using the heuristic than it would if it used the optimal Bayesian approach. Even if the use of the heuristic would tend to lower a firm's level of production, the fact that quantities are strategic substitutes may increase firm i 's production in the heuristic equilibrium relative to the Bayesian equilibrium.

5.4 Profits under the Heuristic and Bayesian Equilibria

Finally, we illustrate the results with respect to equilibrium profits. Simplifying the expression for profits under the Bayesian and heuristic equilibria, we can express the sign of difference in profits as:

$$\text{sign}(E[\hat{\Pi}_i] - E[\Pi_i^*]) = \text{sign} \left[\frac{1}{2^\alpha} - \frac{1}{1 + \alpha} \right] \left[-a - \frac{1}{2^\alpha} + \frac{2}{1 + \alpha} \right] \quad (14)$$

First, note that for $\alpha = 1$, the expected profits under the Bayesian and heuristic approaches do not differ. Second, so long as a is greater than the upper bound of the distribution of c_i and

c_j , the sign of the difference in profits is a simple function of α . When the use of the heuristic strategy leads firms to increase production related to the Bayesian strategy (that is, when $\alpha < 1$), profits are lower when firms both use the heuristic strategy. In this case, firms will prefer (ex ante) to coordinate on the Bayesian equilibrium. If the cost of information acquisition is a hurdle for firms to use the Bayesian strategy, firms may choose to costlessly provide information about the distribution of their own private information to facilitate coordination on the Bayesian equilibrium. When the use of the heuristic strategy leads firms to lower production, ($\alpha > 1$), profits increase in the heuristic equilibria. In this case, firms jointly agree to limit information provision as a way to help coordinate on the use of the heuristic strategies.

6 Example: 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. 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 purchase a single unit of the good be given by:

$$U(\gamma) = \begin{cases} v_0 - p_0 - t\gamma & \text{if good 0 is purchased.} \\ v_1 - p_1 - t(1 - \gamma) & \text{if good 1 is purchased.} \\ 0 & \text{if neither 0 nor 1 is purchased.} \end{cases} \quad (15)$$

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,

$$\hat{\gamma} = \frac{v_0 - v_1 - (p_0 - p_1) + 1}{2}. \quad (16)$$

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

$$E[\pi_0] = E\left[\frac{1 - (p_0 - p_1)}{2}(p_0 - c_0^\alpha)\right]. \quad (17)$$

6.1 Bayesian equilibrium

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

$$E[\pi_0] = E\left[\frac{1 - (p_0 - p_1)}{2}(p_0 - c_0^\alpha)\right]. \quad (18)$$

Solving for the Bayesian equilibrium strategy for firm 0, we have

$$p_0^*(c_0) = \frac{c_0^\alpha}{2} + \frac{2\alpha + 3}{2\alpha + 2}. \quad (19)$$

6.2 Heuristic equilibrium

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

$$\pi_0 = \frac{1 - p_0 + \hat{p}_1(E(c_1))}{2}(p_0 - c_0^\alpha). \quad (20)$$

Solving for the heuristic equilibrium strategy for firm 0, we have

$$\hat{p}_0(c_0) = 1 + \left(\frac{1}{2}\right)^{\alpha+1} + \frac{c_0^\alpha}{2} \quad (21)$$

6.3 Comparing Heuristic and Bayesian Equilibrium Strategies

As in the Cournot example, if a firm's first-order-condition is linear with respect to the opponent's private information ($\alpha = 1$), the strategies chosen in the heuristic and Bayesian equilibria are equivalent. When a firm's first-order-condition is convex in an opponent's private information, the firms choose lower levels of output than in the Bayesian equilibrium. Similarly, if a firm's first-order-condition is concave in an opponent's private information, the firms choose higher levels of output than in the Bayesian equilibrium.

Unlike the Cournot case, though, Proposition 2 holds generally. As an illustrative example, consider generalizing the Bertrand example to allow for arbitrary distributions of c_1 and c_0 . As in the Cournot case, with $\alpha = 2$, the difference between firm i 's strategy in the heuristic and Bayesian equilibria can be represented as a function on the variances of c_1 and c_0 :

$$h_0(c_0) = -var(c_0) - 2var(c_1). \quad (22)$$

In this case, both the direct effect of using the heuristic as well as the strategic response to a competitors use of the heuristic shift the firm's strategy in a similar direction.

6.4 Profits under the Heuristic and Bayesian Equilibria

Solving for expected profits for firm 0 in the Bayesian equilibrium and the heuristic equilibrium, we have:

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

and

$$\hat{\pi}_0 = \left(\frac{\frac{c_1^\alpha}{2} - \frac{c_0^\alpha}{2} + 1}{2} \right) \left(\frac{2^{\alpha+1} + 1}{2^{\alpha+1}} - \frac{c_0^\alpha}{2} \right)$$

Note that if the cost shock enters in linearly, that is $\alpha = 1$, the heuristic strategy and the Bayesian strategy are equivalent, as are the profits earned by the firms.¹²

¹²It's interesting to note that $\hat{\pi}_0 = \pi_0^*$ at $\alpha = 0$ in addition to $\alpha = 1$. In this case, the cost shock does not enter into a firms decision.

Comparing firm 0's profits under the Bayesian and heuristic equilibria, we find that

$$E[\widehat{\pi}_0] > E[\pi_0^*] \iff \alpha < 1. \quad (23)$$

When the use of the heuristic leads firms to set higher prices than in the Bayesian equilibrium, firm profits are higher.¹³ Consequently, this is a case in which firms might have an incentive to jointly withhold information from each other to coordinate on the heuristic equilibrium rather than the Bayesian equilibrium.

7 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 use the full information about the distributions of each other's private information, and therefore would have incentives to disclose 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, 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 Bayesian

¹³Note that in defining the threshold value of γ , we implicitly assumed that the value of the goods was high enough to ensure that all consumers chose to purchase. If good values are insufficiently high to induce all consumers to purchase the good under both the heuristic equilibrium and the full information equilibrium, the above results may not hold.

strategies. We find that holding the other firms' strategy constant, heuristic strategies differ from Bayesian 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 Bayesian 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 Bayesian 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 Bayesian 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 Bayesian 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. Consequently, our results enable a better understanding of the incentives firms may have to either facilitate or impede access to industry information.

References

- [1] Amir, R., Jim, J.Y., Troege, M. 2010. "Robust results on the sharing of firm-specific information: Incentives and welfare effects." *Journal of Mathematical Economics*, 46: 855-866.
- [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] Furubotn, Eirik G. 2009. "Heuristics, the non-maximizing firm and efficient allocation." *Metroeconomica*, 60 (1): 1-23.
- [10] 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.

- [11] Gal-Or, Esther. 1986. "Information Transmission: Cournot and Bertrand Equilibria." *Review of Economic Studies*, 52: 85-92.
- [12] Gal-Or, Esther. 1988. "The advantages of imprecise information." *RAND Journal of Economics*, 10 (2): 266-275.
- [13] Jansen, J. 2008. Information acquisition and strategic disclosure in oligopoly. *Journal of Economics and Management Strategy*, 17 (1): 113-148.
- [14] Jehiel, Philippe. 2005. "Analogy-based expectation equilibrium." *Journal of Economic Theory*, 123: 81-104.
- [15] Lacetera, N., D. Pope, and J. Sydnor (forthcoming). "Heuristic Thinking and Limited Attention in the Car Market." *American Economic Review*.
- [16] Luttmer, Erzo, and Kelly Shue. 2006. "Who misvotes?: The effect of differential cognition costs on election outcomes." KSG Working paper RWP06-048.
- [17] Minor, D. (2011). "Coarse thinking and competition." Working paper, Northwestern University.
- [18] Mirman, Leonard J., Larry Samuelson, and Edward E. Schlee. 1994. "Strategic information manipulation in duopolies." *Journal of Economic Theory*, 62: 363-384.
- [19] Raith, M. 1996. "A general model of information sharing in oligopoly." *Journal of Economic Theory*, 71: 260-288.
- [20] Shapiro, Carl. 1986. "Exchange of cost information in oligopoly." *Review of Economic Studies*, 52: 433-446.
- [21] Shapiro, Carl. 1989. "Theories of oligopoly behavior." In R. Schmalensee and R.D. Willig (Eds), *Handbook of Industrial Organization*, Vol. I, pp. 329-414.

- [22] Simon, Herbert A. 1955. "A behavioral model of rational choice." *The Quarterly Journal of Economics*, 69 (1): 99-118.
- [23] Simon, Herbert A. 1956. "Dynamic programming under uncertainty with a quadratic criterion function." *Econometrica*, 24 (1): 74-81.
- [24] Simon, Herbert A. 1959. "Theories of decision-making in economics and behavioral science." *American Economic Review*, 49 (3): 253-283.
- [25] Vives, X. 1990. "Trade association disclosure rules, incentives to share information and welfare." *Rand Journal of Economics*, 21: 409-430.
- [26] Vives, X. 1999. *Oligopoly Pricing: Old Ideas and New Tools*. Cambridge, MA: MIT Press.
- [27] Weintraub, Gabriel, C. Lanier Benkard, & Benjamin Van Roy. 2008. "Markov Perfect industry dynamics with many firms." *Econometrica*, 76 (6): 1375-1411.
- [28] 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. 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

$$\frac{\partial \pi_i}{\partial s_i} = \alpha(s_i, \varepsilon_i) + \beta(s_i, \varepsilon_i)\varepsilon_{-i} + \gamma(s_i, \varepsilon_i)\varepsilon_{-i}^2 \quad (24)$$

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

$$\frac{\partial E[\pi_i(s_i^*(\varepsilon_i), s_{-i}(\cdot), \varepsilon_i)]}{\partial s_i} = \alpha(s_i^*, \varepsilon_i) + \beta(s_i^*, \varepsilon_i)E[\varepsilon_{-i}] + \gamma(s_i^*, \varepsilon_i)E[\varepsilon_{-i}^2] = 0 \quad (25)$$

$$\frac{\partial \pi_i(\hat{s}_i, s_{-i}(E(\varepsilon_{-i})), \varepsilon_i)}{\partial s_i} = \alpha(\hat{s}_i, \varepsilon_i) + \beta(\hat{s}_i, \varepsilon_i)E[\varepsilon_{-i}] + \gamma(\hat{s}_i, \varepsilon_i)E[\varepsilon_{-i}]^2 = 0 \quad (26)$$

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} ,

$$\gamma(s_i^*, \varepsilon_i) > 0 \leftrightarrow s_i^* > \hat{s}_i. \quad (27)$$

■

Proposition 2

Proof. To compare the equilibrium strategies under the heuristic and Bayesian approaches, we define \tilde{s}_i as the heuristic best response of firm i to $s_{-i}^*(s_i^*, \varepsilon_{-i})$.¹⁴ We define \tilde{s}_{-i} in an analogous fashion. By our definition of \tilde{s}_i and \tilde{s}_{-i} ,

$$\frac{\partial \pi_i(\tilde{s}_i, s_{-i}^*(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} = 0 \quad (28)$$

$$\frac{\partial \pi_{-i}(\tilde{s}_{-i}, s_i^*(E[\varepsilon_i]), \varepsilon_{-i})}{\partial s_{-i}} = 0. \quad (29)$$

¹⁴We 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.

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

$$\frac{\partial \pi_i(\tilde{s}_i, \tilde{s}_{-i}(E[\varepsilon_{-i}]), \varepsilon_i)}{\partial s_i} = \int_{s_{-i}^*(E[\varepsilon_{-i}])}^{\tilde{s}_{-i}(E[\varepsilon_{-i}])} \frac{\partial^2 \pi_i(\tilde{s}_i, s_{-i}, \varepsilon_i)}{\partial s_i \partial s_{-i}} ds_{-i} \quad (30)$$

$$\frac{\partial \pi_{-i}(\tilde{s}_{-i}, \tilde{s}_i(E[\varepsilon_i]), \varepsilon_{-i})}{\partial s_{-i}} = \int_{s_i^*(E[\varepsilon_i])}^{\tilde{s}_i(E[\varepsilon_i])} \frac{\partial^2 \pi_{-i}(\tilde{s}_{-i}, s_i, \varepsilon_{-i})}{\partial s_{-i} \partial s_i} ds_i. \quad (31)$$

Proposition 1 tells us that if $\frac{\partial \pi_i}{\partial s_i}$ and $\frac{\partial \pi_{-i}}{\partial s_{-i}}$ are concave in ε_{-i} and ε_i , $\tilde{s}_i > s_i^*$ and $\tilde{s}_{-i} > s_{-i}^*$.¹⁵ 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 (30) and (31) to the heuristic equilibrium by expressing them as Taylor expansions around the heuristic equilibrium. Letting $x_i(\varepsilon_i) = \tilde{s}_i(\varepsilon_i) - \hat{s}_i(\varepsilon_i)$ and $x_{-i}(\varepsilon_{-i}) = \tilde{s}_{-i}(\varepsilon_{-i}) - \hat{s}_{-i}(\varepsilon_{-i})$, we can write (30) 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 (31) 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 Bayesian equilibrium and the strategies chosen by the firms in the heuristic equilibrium.

To prove the proposition, assume that s_i and s_{-i} are strategic complements, and $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}| \forall s_i, s_{-i}$.

¹⁵Furthermore, it is possible to show 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 . We omit the formal presentation of the proposition for brevity.

For notational convenience, let

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

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

Thus, we can write (30) and the analogous expression for firm $-i$ as:

$$d_i(\varepsilon_i) = x_i(\varepsilon_i)a_i(\varepsilon_i) + x_{-i}(E[\varepsilon_{-i}])b_i(\varepsilon_i) \quad (32)$$

$$d_{-i}(\varepsilon_{-i}) = x_{-i}(\varepsilon_{-i})a_{-i}(\varepsilon_{-i}) + x_i(E[\varepsilon_i])b_{-i}(\varepsilon_{-i}). \quad (33)$$

Evaluating (32) at $E[\varepsilon_i]$ and (33) at $E[\varepsilon_{-i}]$, we can express:

$$x_i(E[\varepsilon_i]) = \frac{a_{-i}d_i - b_id_{-i}}{a_ia_{-i} - b_ib_{-i}} \quad (34)$$

$$x_{-i}(E[\varepsilon_{-i}]) = \frac{a_id_{-i} - b_{-i}d_i}{a_ia_{-i} - b_ib_{-i}}. \quad (35)$$

Signing the terms in (34) and (35): (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 (30) and (31), $d_i, d_{-i} > 0$. Furthermore, by the assumption $|\frac{\partial^2 \pi_i}{\partial s_i^2}| > |\frac{\partial^2 \pi_i}{\partial s_i \partial s_{-i}}| \forall s_i, s_{-i}$, $a_ia_{-i} - b_ib_{-i} > 0$. Thus, the numerators of expressions (34) and (35) 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 (32) and (33), we have:

$$x_i(\varepsilon_i) = \frac{d_i(\varepsilon_i) - b_i(\varepsilon_i)x_{-i}(E[\varepsilon_{-i}])}{a_i(\varepsilon_i)} \quad (36)$$

$$x_{-i}(\varepsilon_{-i}) = \frac{d_{-i}(\varepsilon_{-i}) - b_{-i}(\varepsilon_{-i})x_i(E[\varepsilon_i])}{a_{-i}(\varepsilon_{-i})}. \quad (37)$$

The numerator of expressions (36) and (37) 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 $\varepsilon_i, \varepsilon_{-i}$

$$\hat{s}_i(\varepsilon_i) > \tilde{s}_i(\varepsilon_i) > s_i^*(\varepsilon_i), \hat{s}_{-i}(\varepsilon_{-i}) > \tilde{s}_{-i}(\varepsilon_{-i}) > s_{-i}^*(\varepsilon_{-i}). \quad (38)$$

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 $\varepsilon_i, \varepsilon_{-i}$

$$\hat{s}_i(\varepsilon_i) < \tilde{s}_i(\varepsilon_i) < s_i^*(\varepsilon_i), \hat{s}_{-i}(\varepsilon_{-i}) < \tilde{s}_{-i}(\varepsilon_{-i}) < s_{-i}^*(\varepsilon_{-i}). \quad (39)$$

■