How much would you pay to change a game before playing it?

by David Wolpert

Santa Fe Institute

1399 Hyde Park Road, Santa Fe, NM 87501

MIT, Arizona State University, `david.h.wolpert@gmail.com`

and

Justin Grana *

Santa Fe Institute

1399 Hyde Park Road, Santa Fe, NM 87501

`justin.grana@santafe.edu`

and

Duncan K. Foley

Department of Economics, New School for Social Research

6 East 16th, New York, NY 10003

`foleyd@newschool.edu`

April 28, 2017

Abstract

We proposes a differential framework for determining how much a decision maker (DM) is willing to pay to alter the parameters of a strategic scenario they will encounter. Our formulation decomposes DM's willingness to pay a given amount to change the parameters as the sum of three factors: 1) the direct effect the parameter change would have on DM's payoffs, holding strategies of all players constant; 2) the effect there would be on DM's strategy as they react to the change, with the strategies of the other players held constant; and 3) the effect there would be on the strategies of all other players as they react to the change, with the strategy of DM held constant. We use this decomposition to analyze how DM's decision to pay for a change to the parameters depends on whether that decision is observed by the other players (e.g., how their decision to pay varies depending on whether the payment would occur as lobbying or bribery). In addition, we show that when players are bounded rational, new factors arise that contribute to DM's decision of whether to pay that are absent under full rationality. We illustrate these results with several examples.

Keywords— Game Theory, Comparative Statics, Envelope Theorem

---

*Corresponding author

1 Introduction

A set of results known collectively as the "envelope theorems" have proven invaluable in advancing neoclassical consumer and producer theory [Milgrom and Segal, 2002]. Envelope theorems describe how a function's extreme value varies with changes in a parameter of the function. Applied to the case of a single decision maker (DM) faced with an exogenous objective function, they describe how the optimal decision by DM varies with changes to that objective function. Classic examples include changes in a consumer's optimal consumption bundle choices due to changes in their income, and changes in a firm's optimal mix of capital and labor due changes in the exogenous real wage.

The envelope theorems can be used to perform comparative statics. In particular, the derivative of the maximal value of DM's objective function as a parameter of that function is changed is often interpreted as the highest *price* that DM would be willing to pay (to a regulator, for example) to enact such a change before they make their decision. Under differentiability assumptions, the envelope theorems can be used to determine this maximal price.

Going beyond basic producer and consumer theory leads us to consider scenarios involving multiple agents that reason strategically, i.e., to non-cooperative games. Unfortunately, the strategic nature of noncooperative games precludes a straightforward extension of the envelope theorem to them. This makes it difficult to know best how to perform comparative static analysis for games. Consequently, without knowing how the equilibrium adjusts in response to a parameter change, it is difficult to determine how much a player would be willing to pay to change the parameter value. While we will discuss some related work that attempts to extend envelope theorems and comparative statics to games, there is no broadly applicable method for determining the price of a parameter change in a strategic scenario.

There are at least three main challenges that arise when determining the price an agent is willing to pay to change the parameter values in strategic scenarios. The first issue has to do with the multiplicity of equilibria. When determining how an equilibrium changes as a result of a parameter change, the modeler must also specify *which* of the equilibria arise after the parameter change. Secondly—and related to the first issue—is that even the number of equilibria may be different for different values of the parameters. This means that after small parameter changes, it is possible that the qualitative characteristics of the *set* of equilibria are different, making it even more obscure how to quantify the effects of the change. The final issue that arises when determining how much a player is willing to pay to change the parameter values is how best to account for whether the player's payment decision is observed by other players in the game.

In this paper we extend the envelope theorems to strategic scenarios. In other words, we extend analysis of how extrema of a parameterized function depend on the values of the parameters of that function to analysis of how the fixed points of a parameterized noncooperative game depend on the parameters of that game. This allows us to investigate how a game's parameter value can arise *endogenously*, as the maximal price a player would agree to pay to change such values. Specifically, we derive a general yet simple expression for how much a decision maker would be willing to pay to change the parameter values of a given game.[1] Our approach allows us to circumvent the issues of multiple equilibria and instead we focus on comparing and contrasting the case in which DM's decision to change the parameter values is observed by other players with the case in which DM's decision is private information.

A major advantage of our expression for how much DM is willing to pay is that it explicitly disentangles the factors that collectively determine a player's willingness to pay to change the parameter values. This permits a simple comparison between how much DM would be willing to pay when its decision to change the parameter

---

[1] Throughout, we assume there is a single player who has the option to change the parameter values, and refer to them as the "decision maker", DM. We refer to all other game participants simply as "players".

values is unobserved with the case of an observed decision. When DM's decision to change the parameter values is unobserved by all other players, DM's willingness to pay to change the parameter values is given by the sum of two factors. The first factor is how the change directly affects DM's payoffs, holding the strategies of *all* players constant (including DM's). The second factor is how DM optimally adjusts their strategy in response to the parameter change. On the other hand, when DM's decision of whether to change the parameter values is observed by all players, a third factor is added to the sum that determines DM's willingness to make that change. That third factor is how *other players'* strategies would adjust in response to the parameter change.

In section 2, we review the literature on envelope theorems, comparative statics and endogenous information structures and elaborate on how our work relates to the current literature. Section 3 presents the main contribution in which we derive the most a decision maker would be willing to pay to make an offered change to the parameter values of a game. Specifically, we derive an expression for the most DM is willing to pay when the payment decision is observed and when the payment decision is unobserved by all other players. This allows us to determine how changing the observability of DM's decision affects its willingness to pay to change the parameter values. We show that when DM's decision to change the parameters is unobserved by all other players, DM's willingness to pay is the same as in non-strategic scenarios and is given by the standard envelope theorem. However, when DM's decision is observed by all other players, a new term arises that accounts for all other players taking into account a change in the parameter values. In section 4 we illustrate the applicability of our results through several examples. In section 5 we show how our framework can be extended to include bounded rational agents, and then provide an example using the quantal response solution concept. Specifically, we show that when players are boundedly rational, there is a term in the expression for DM's willingness to pay that represents DM's anticipation of its own future irrationality. In section 6, we show the difficulty in extending our framework to dynamic games and suggest directions of future work.

2 Envelope Theorems, Comparative Statics and Information Structures

Broadly speaking, our work is related to—and partly synthesizes—three branches of the current literature: comparative statics in games [Milgrom and Shannon, 1994, Acemoglu and Jensen, 2013], value of information in games [Neyman, 1991], and endogenous information structures [Hurkens and Vulkan, 2006].

Comparative static analysis is commonly used to analyze how equilibria of a model adjusts in response to a change in parameters, without an explicit description of the dynamic adjustment process to go from the initial value of the parameters to their final one. The envelope theorem is related to comparative static analysis in that often times, the envelope theorem is used as a *tool* when performing comparative statics.

While the envelope theorems and comparative static analyses of a single agent's decision problem are well-understood, more recent effort has focused on comparative static analysis and associated envelope theorems in strategic scenarios [Milgrom and Shannon, 1994, Villas-Boas, 1997, Sunanda Roy, 2008, Acemoglu and Jensen, 2013, Caputo, 1996, Lippman et al., 1987, Echenique, 2002]. The comparative statics literature focuses on establishing properties of utility functions that allow for a direct comparison among equilibrium strategy profiles before and after a change in the parameter values. Much of the comparative statics in games literature is concerned with partial orders of joint profiles, seeking to determine if a joint profile "increases" when a parameter changes. Our work is related to this branch of research in that in order to determine whether a decision maker is willing to pay to change the parameter values, it is necessary to know how the equilibrium adjusts in response to a change in such values. However, our work is distinct in that we focus on establishing the conditions in which a player is willing to pay to change parameter values instead of analyzing the equilibrium effects of exogenous changes in the parameter values. While there have been efforts into combining envelope theorems and comparative static analy-

sis in games [Caputo, 1996], such results are limited to games with a unique equilibrium and convex constrained utility functions. Our work not only circumvents these limitations but is also the first to provide envelope theorems (representing DM's willingness to pay to change the parameter values) under the case when DM's decision to change the parameter values is observed *and* when it is unobserved.

Secondly, our work has straightforward applications to investigating the value of information *to a particular player* in strategic scenarios. Often times, the amount of "information" available to a player is controlled by a parameter, such as the variance of a normal random variable (see Ui and Yoshizawa [2015] for an example).[2] Our framework allows us to determine a player's willingness to pay to change the parameter that controls the amount of information available to the player. Related fundamental work in the value of information includes Neyman [1991], Bassan et al. [2003], Bertschinger et al. [2014] while an applied example can be found in Morris and Shin [2002]. It is important to keep in mind that our formulation is *not* limited to the value of information and that we do not restrict the parameters to *only* control information. Instead, the value of information is just one application of our main contribution.

Finally, our work is related to that of endogenous information structures in games [Hurkens and Vulkan, 2006, Amir and Lazzati, 2016]. That strand of research focuses on classifying which information structures can arise endogenously when players can pay to acquire information. Furthermore, Hurkens and Vulkan [2006] compare and contrast the case where a player's decision to acquire more information is observed with the case of an unobserved decision. We also make this kind of comparison. However our framework (and so the associated comparison) is more general. It is not only concerned with endogenous information structures and partitions of information sets, but also provides a general formulation to examine whether *any* parameter specification can arise endogenously.

3 Differential Value of Parameter Changes

In this section we derive the maximum price—defined in terms of utility per unit change of the game parameters—such that it is a Nash equilibrium for DM to pay that amount to enact a given infinitesimal change in the parameters. We proceed in several steps. First we introduce what we call the "inner game". This is a traditional simultaneous move game. Since our framework is not limited to fully rational agents, we then introduce notation for considering arbitrary player "response functions" and associated equilibria. Next we show how to transform the inner game to a two-stage extensive form game that we call the "outer game". In the first stage of the outer game, DM is given the opportunity to pay to enact a change in the parameter values of the inner game. In the second stage, all players play the inner game, where the parameter values that define the inner game are determined by whether or not DM chose to pay to change the parameter values. We restrict attention to scenarios of imperfect information, where all players other than DM know that DM had the opportunity to make the payment to enact the change, but do not have any information about whether they actually did so. The associated "private offer price" to DM is the highest payment such that there exists a Nash equilibrium of the *outer* game in which DM makes that payment in order to enact the associated change in the parameter values.

We then modify this construction to define the public offer price, as the most that DM would be willing to pay to enact an offered change in the parameter values of the inner game *when their decision of whether to pay is observed by all other players.* Formally, the difference between the two prices is that in the private offer price, all players other than DM only have one information set in the outer game since they do not observe whether or not DM paid to change the parameter values. On the other hand, for the public offer price all players have 2

---

[2]We do not propose a particular measures of information. Instead, we just assume that there is a parameter that intuitively sets the amount of information available to a player.

information sets in the second stage of the outer game, corresponding to whether DM did or did not choose to pay to change the parameter values.

3.1 Framework

We consider an arbitrary N-player simultaneous move **inner game**, γ. We write the strategy spaces of the players as $\{X_i : i = 1..N\}$ with associated pure strategies x_i. In general, strategy spaces can be either discrete or continuous. Let $\triangle X_i$ be the set of player i's mixed strategies with element σ_i. With the exception that the mixed strategy is a valid probability distirbution/density, we do not place any restrictions on the strategies. Denote a joint (possibly mixed) strategy profile as σ.

We assume that the inner game is parameterized by a k-dimensional parameter vector θ. We write the game for a specific value of that game parameter as $\gamma(\theta)$. The parameter vector θ can be any compact set of real-valued vectors. For example, θ might specify a tax rate that lies in some closed interval, in particular one that depends on the joint pure strategy of the players. As another example, it might represent real-valued regulatory fines that lie in some closed intervals. As a final example, if $\gamma(\theta)$ is the normal form representation of a Bayesian game [Osborne and Rubinstein, 1994], θ could parameterize the distribution of players' signals conditional on the state of nature. We now complete the definition of the game with a definition of utility, response functions and equilibrium:

Definition 1. *Given any joint profile σ, let σ_i be the strategy of player i and σ_{-i} the strategy profile of all players other than i. We make the following definitions:*

1. *$\tilde{U}_i(\theta, \sigma_i, \sigma_{-i}) \to \mathbb{R}$ is the **expected utility** function of player i at parameter value θ when i plays strategy σ_i and all other players' strategy is specified by σ_{-i}.*

2. *$B_i : (\theta, \sigma_{-i}) \to \sigma_i$ is player i's **response function**.*

3. *$\sigma^*(\theta)$ is a joint strategy profile in the game $\gamma(\theta)$ such that for all i, $\sigma_i^*(\theta) = B_i(\theta, \sigma_{-i}^*(\theta))$. We call $\sigma^*(\theta)$ an **equilibrium** of the game $\gamma(\theta)$.*

The definitions of response functions and equilibrium are standard generalizations of the usual definitions for the case of fully rational players, to allow bounded rationality.

3.2 Unobserved Decision to Change Parameters - The Private Offer Price

In this section we derive how much DM would be willing to pay for a particular change in the parameter values when their decision of whether to pay is not observed by the other players. To start we focus on the Nash equilibrium solution concept, extending the analysis to allow bounded rational players in section 5. Formally, we restrict attention to the case where

$$B_i(\theta, \sigma_{-i}) \to \sigma_i \text{ such that } \tilde{U}(\theta, \sigma_i, \sigma_{-i}) \geqslant \tilde{U}_i(\theta, \sigma_i', \sigma_{-i}) \ \ \forall i \in 1...N, \sigma_i' \in \triangle X_i$$

For notational convenience, we drop subscripts and let $\tilde{U}$ and B be the expected utility and response function for DM unless otherwise noted.

Given any set of parameterizedinner games $\{\gamma(\theta)\}$ and a particular value θ^*, consider an associated **outer** game $\overline{\gamma}(\theta^*)$ that has two stages. In the first stage only DM moves, making a binary decision of whether to:

- Pay an amount c to purchase a change in the parameter values from θ^* to $\theta^* + \delta\vec{v}$, where δ is a positive scalar and $\vec{v}$ is a k-dimensional vector of unit length;

- Pay nothing, leaving the game parameter values as θ^*.

In the second stage the players play $\gamma(\theta)$, where θ equals either θ^* or $\theta^* + \delta\vec{v}$, depending on whether DM purchased the change in parameter values in the first stage. The utility of DM in $\overline{\gamma}(\theta^*)$ is their utility in $\gamma(\theta^* + \delta\vec{v})$, minus c iff they decide to purchase the change to θ^* in the first stage.

To distinguish moves in the inner game and the outer game, we use $\Sigma_{\overline{\gamma}(\theta^*)}$, or just $\Sigma_{\theta*}$ for short, to indicate some joint strategy profile of the outer game $\overline{\gamma}(\theta^*)$. $\sigma(\Sigma_{\theta*})$ indicates those components of $\Sigma_{\theta*}$ that specify the profile of strategies employed by the players *after* DM's decision of whether to pay to change the parameters, i.e. the strategy profile in the second stage of $\overline{\gamma}(\theta^*)$.

We want to relate Nash equilibria in the inner game, γ with Nash equilibria in the outer game $\overline{\gamma}$. To ground our analysis, we first present the almost trivial result that says that a Nash equilibrium profile in the second stage of the outer game is a Nash equilibrium of the inner game:

Lemma 1. *Let $\Sigma_{\theta*}$ be a Nash equilibrium profile of $\overline{\gamma}(\theta^*)$ in which DM pays c to change the parameters from θ^* to $\theta^* + \delta\vec{v}$. Then $\sigma(\Sigma_{\theta*})$ is a Nash equilibrium of $\gamma(\theta^* + \delta\vec{v})$.*

Proof. Suppose the lemma were not true. Then some player would be able to change their action in $\gamma(\theta^* + \delta\vec{v})$ and receive a higher utility. However, that same player would also be able to change their action in the second stage of $\overline{\gamma}(\theta^*)$ and receive a higher utility. This contradicts the hypothesis that $\Sigma_{\theta*}$ is a Nash equilibrium of $\overline{\gamma}(\theta^*)$. □

However, our goal is to determine when a Nash equilibrium of the inner game is also a Nash equilibrium of the outer game in which DM is offered the opportunity to pay to change the parameter values. The following lemma establishes such conditions:

Lemma 2. *Suppose $\sigma^*(\theta^*)$ and $\sigma^*(\theta^* + \delta\vec{v})$ are Nash equilibria of $\gamma(\theta^*)$ and $\gamma(\theta^* + \delta\vec{v})$, respectively. Then it is a Nash equilibrium for DM to pay c in the outer game $\overline{\gamma}(\theta^*)$ to change the parameters from θ^* to $\theta^* + \delta\vec{v}$ iff*

$$
\begin{aligned}
c \;\leqslant\; & \tilde{U}\left(\theta^* + \delta\vec{v}, \; B\left(\theta^* + \delta\vec{v}, \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right), \; \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right) \right) \\
& - \tilde{U}\left(\theta^*, \; B\left(\theta^*, \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right), \; \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right) \right)
\end{aligned}
$$

Proof. By assumption, $\sigma^*(\theta^* + \delta\vec{v})$ is a Nash equilibrium of the game $\gamma(\theta^* + \delta\vec{v})$, so if DM decides to pay the price to change the game parameters to $\theta^* + \delta\vec{v}$ in the first stage of $\overline{\gamma}$, neither DM nor the other players can change their action in the second stage of $\overline{\gamma}$ and thereby increase their payoff. The only remaining possible change to a move by a player in $\overline{\gamma}$ would be for DM to choose not to purchase the change in parameters. However, if DM does that, then since (by definition of Nash equilibrium) none of the other players will change their moves in the second stage of $\overline{\gamma}$, DM's net change in utility is

$$
\tilde{U}\left(\theta^*, \; B\left(\theta^*, \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right), \; \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right) - \tilde{U}\left(\theta^* + \delta\vec{v}, \; B_{DM}\left(\theta^* + \delta\vec{v}, \sigma^*_{-DM}(\theta^* + \delta\vec{v})\right), \; \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right) + c
$$
(1)

This establishes that no player has an incentive to change their move given that the hypothesized inequality holds, i.e., that it is a Nash equilibrium for DM to pay to change the parameter values. The converse direction of the proof follows similarly. □

Lemma 2 gives the maximal amount c such that it is a Nash equilibrium for DM to pay to change the parameter value by $\delta\vec{v}$, given that their choice of whether to pay is not observed by the other players. If that maximal amount is negative, then DM would have to *be paid* to agree to the associated change in the parameter values. Intuitively, lemma 2 says that it is a Nash equilibrium for DM to pay to change the parameter values if and only if DM cannot do better by not paying, saving the cost c and adjusting their strategy so that it maximizes its payoff when parameter values are θ^* and all other players are playing $\sigma_{-i}(\theta^* + \delta\vec{v})$.

We want to quantify how much DM would be willing to pay to change the game parameter in the broadest sense, without specifying c and / or δ. To see how to do this, first recall that the standard way to quantify the willingness of a decision-maker to pay for a change to the amount x of some *good*, in a non-strategic scenario, is as the derivative of the their utility function with respect to x. (Such a quantity is often analyzed via the envelope theorem.) In other words, we consider the limit of an infinitesimal change to x.

Proceeding analogously, we quantify DM's willingness to pay for a change of the parameter values in the inner game by taking the limit of infinitesimal δ of the expression in Lemma 2:

$$\mathscr{V}(\theta^*, \vec{v}) \equiv \lim_{\delta \to 0} \frac{\tilde{U}\left(\theta^* + \delta\vec{v}, B\left(\theta^* + \delta\vec{v}, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right) - \tilde{U}\left(\theta^*, B\left(\theta^*, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right)}{\delta}$$

(2)

If we assume differentiability of response functions and expected utility, we obtain the following result:

Lemma 3. *Assume $B(\theta, \sigma_{DM})$ is differentiable with respect to θ and $\tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})$ is differentiable with respect to both θ and σ_{DM}. Then*

$$\mathscr{V}(\theta^*, \vec{v}) = \sum_j v_j \left(\frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial \theta_j} + \sum_k \frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial [\sigma_{DM}]_k} \frac{\partial [B(\theta, \sigma_{-DM})]_k}{\partial \theta_j} \right)$$

(3)

where for any vector or x, x_j represents the j'th element of x, and the derivatives are evaluated at $\theta = \theta^*$ and $\sigma = \sigma^*(\theta^*)$.

Proof. Adding and subtracting $\frac{\tilde{U}\left(\theta^*, B\left(\theta^* + \delta\vec{v}, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right)}{\delta}$ from equation 2 gives:

$$\mathscr{V}(\theta^*, \vec{v}) =$$
$$\lim_{\delta \to 0} \frac{\tilde{U}\left(\theta^* + \delta\vec{v}, B\left(\theta^* + \delta\vec{v}, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right) - \tilde{U}\left(\theta^*, B\left(\theta^* + \delta\vec{v}, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right)}{\delta} +$$
$$\lim_{\delta \to 0} \frac{\tilde{U}\left(\theta^*, B\left(\theta^* + \delta\vec{v}, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right) - \tilde{U}\left(\theta^*, B\left(\theta^*, \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right), \sigma^*_{-\text{DM}}(\theta^* + \delta\vec{v})\right)}{\delta}$$

Applying the chain rule and the definition of partial derivative gives the claimed result. $\square$

We call $\mathscr{V}(\theta^*, \vec{v})$ the **private offer price in direction $\vec{v}$** and sometimes just refer to it as the private offer price when the context is clear. The first term inside the first sum in equation 3 captures the dependence of the private offer price on how DM's utility changes as a result of changing an element of θ while holding strategies of all players constant. The term inside the second sum captures the composite dependence of private offer price on how each element of DM's best response changes in response to changes in θ when strategies of all other players are constant, and how these changes in DM's best response changes DM's utility. Note that the private offer price does not depend on changes in the strategies of players other than DM, since other players do not observe whether

or not DM chose to change the parameters. We also stress that equation 2 is the definition of the private offer price and that the differential representation of equation 3 is a result of simplifying assumptions. In section 4, we show how it is possible for the private offer price to be defined even when the best response functions are not differentiable with respect to the parameters.

When considering the Nash equilibrium solution concept, equation 3 is further simplified and is given by the following lemma:

Corollary 1. *Assume $B(\theta, \sigma_{DM})$ is differentiable with respect to θ and $\tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})$ is differentiable with respect to both θ and σ_{DM}. Then*

$$\mathscr{V}(\theta^*, \vec{v}) = \sum_j v_j \frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial \theta} \tag{4}$$

Proof. Since B is the best response function, a necessary condition for a Nash equilibrium is that

$$\frac{\partial \tilde{U}}{\partial [\sigma^*_{DM}]_k} = 0 \quad \forall k$$

Therefore, equation 3 becomes

$$
\begin{aligned}
\mathscr{V}(\theta^*, \vec{v}) &= & \sum_j v_j \left(\frac{\partial \tilde{U}(\theta, \sigma_{\text{DM}}, \sigma_{-\text{DM}})}{\partial \theta_j} + \sum_k 0 \frac{\partial [B(\theta, \sigma_{-\text{DM}})]_k}{\partial \theta_j} \right) \\
&= & \sum_j v_j \frac{\partial \tilde{U}(\theta, \sigma_{\text{DM}}, \sigma_{-\text{DM}})}{\partial \theta}
\end{aligned}
\tag{5}
$$

$\square$

Corollary 1 shows that when DM's decision to change the parameters is unobserved by all other players and DM is best responding, the private offer price is the same as in the non-strategic setting. That is to say, the most DM is willing to pay for an infinitesimal unobserved change in the parameters does not depend on how DM would adjust its strategy in response to the parameter change nor how other players would adjust their strategy. This is the same as in the case of the non-strategic scenario. In other words, the private offer price in games is the same as the envelope theorem in non-strategic scenarios.

3.3   Observed Decision to Change Parameters - The Public Offer Price

This section parallels that of section 3.2 but instead considers the case where DM's decision of whether to change parameter values is observed by all other players. Once again, begin with a game $\gamma(\theta^*)$ and consider the outer game $\hat{\gamma}(\theta^*)$ in which DM has the opportunity to change θ from θ^* to $\theta^* + \delta \vec{v}$. In the case of the public offer price, *all players observe DM's decision to change the parameters*. So the difference between $\overline{\gamma}$ and $\hat{\gamma}$ is that in $\hat{\gamma}$ the strategy of *all* players have two singleton information sets in the second stage of the game. That is, each player's strategy specifies a distribution over actions for the case in which DM paid to change the parameters and the case in which DM did not pay. Let $\hat{\Sigma}_{\theta*}$ be a strategy profile of all players in $\hat{\gamma}$. Additionally, let $\sigma_1(\hat{\Sigma}_{\theta*})$ be the joint strategy profile in the second stage of $\hat{\gamma}$ at the information set in which DM does purchase the change in

the parameter values and let $\sigma_0(\hat{\Sigma}_{\theta*})$ be the joint strategy profile in the second stage of $\hat{\gamma}$ at the information set in which DM does *not* purchase the change in the parameter values.

As in the discussion of private offer price, we explicitly relate the equilibrium profiles of the inner and outer games. However, unlike the private offer price where we were not concerned with equilibrium refinements, in this section we will be concerned with subgame perfect equilibria since all players now have 2 singleton information sets since they observe DM's decision to change the parameter values.

Lemma 4. *Suppose* $\hat{\Sigma}_{\theta*}^*$ *is a* subgame perfect *equilibrium of* $\hat{\gamma}(\theta^*)$ *in which DM pays to change the parameters from* θ^* *to* $\theta^* + \delta\vec{v}$. *Then* $\sigma_1(\hat{\Sigma}_{\theta*}^*)$ *is a subgame perfect equilibrium profile of the game* $\gamma(\theta^* + \delta\vec{v})$.

Proof. Since $\hat{\Sigma}_{\theta*}^*$ is a subgame perfect equilibrium, no player can increase their utility by deviating at the information set at which players learn of the parameter change. Since at this information set players play the game $\gamma(\theta^*+\delta\vec{v})$, no player can deviate from $\sigma_1(\hat{\Sigma}_{\theta*}^*)$ in the game $\gamma(\theta^*+\delta\vec{v})$ and earn a higher expected utility. Therefore, $\sigma_1(\hat{\Sigma}_{\theta*}^*)$ is a subgame perfect equilibrium of the game $\gamma(\theta^* + \delta\vec{v})$. $\qquad\square$

Parallel to lemma 1, lemma 4 says that a subgame perfect strategy profile in the game in which DM pays to change the parameters is also a subgame perfect profile in which the parameters are specified exogenously. The associated analog of Lemma 2 is the following:

Lemma 5. *Suppose* $\sigma^*(\theta^*)$ *and* $\sigma^*(\theta^* + \delta\vec{v})$ *are Nash equilibria of* $\gamma(\theta^*)$ *and* $\gamma(\theta^* + \delta\vec{v})$, *respectively. Then, there exists a subgame perfect equilibrium for DM to pay c in the game* $\hat{\gamma}(\theta^*)$ *to change the parameters from* θ^* *to* $\theta^* + \delta\vec{v}$ *iff*

$$c < \tilde{U}\left(\theta^* + \delta v, B(\theta^* + \delta v, \sigma_{-i}^*(\theta^* + \delta v)), \sigma_{-i}^*(\theta^* + \delta v)\right) - \tilde{U}\left(\theta^*, B(\theta^*, \hat{\sigma}_{-DM}^*(\theta^*)), \sigma_{-DM}^*(\theta^*)\right)$$

Proof. Define the strategy profile $\hat{\Sigma}_{\theta*}^*$ of $\hat{\gamma}(\theta^*)$ such $\sigma_1(\hat{\Sigma}_{\theta*}^*) = \sigma^*(\theta^* + \delta\vec{v})$ and $\sigma_0(\hat{\Sigma}_{\theta*}^*) = \sigma^*(\theta^*)$ and DM chooses to pay c to change the parameter values. Since $\sigma^*(\theta^* + \delta\vec{v})$ and $\sigma^*(\theta^*)$ are Nash equilibria of $\gamma(\theta^* + \delta\vec{v})$ and $\gamma(\theta^*)$ respectively, no player can change their strategy at any proper subgame of $\hat{\gamma}(\theta^*)$ after DM's initial decision and increase their utility, conditional on reaching that subgame. Therefore, the only deviation that needs to be considered is the one in which DM does not pay to change the parameters. However, if DM chooses not to pay to change the parameter values DM's net change in utility is given by

$$\tilde{U}(\theta^*, B(\theta^*, \hat{\sigma}_{-DM}^*(\theta^*)), \sigma_{-DM}^*(\theta^*)) - \tilde{U}(\theta^* + \delta v, B(\theta^* + \delta v, \sigma_{-i}^*(\theta^* + \delta v)), \sigma_{-i}^*(\theta^* + \delta v)) + c \qquad (6)$$

This quantity is less than or equal to 0 if the inequality in 6 holds and therefore DM would not have an incentive to deviate at the subgame where they choose to pay to change the parameters. This establishes that no player has any incentive to deviate at any proper subgame, and therefore $\hat{\Sigma}_{\theta*}^*$ is a subgame perfect equilibrium of $\hat{\gamma}(\theta^*)$ $\qquad\square$

Note how the arguments of the $\tilde{U}$ functions and B functions differ between lemma 5 and lemma 2. This difference directly reflects the fact that those lemmas concern observed and unobserved moves by DM in stage 1, respectively. More specifically, since players observe whether or not DM paid to change the parameter values, when DM considers the profitability of a deviation, DM needs to consider how all other players will change their strategy if DM does not pay.

Like in the case of the private offer price, Lemma 5 establishes the maximal amount that DM would be willing to pay for a *discrete* change of the parameter values from θ^* to $\theta^* + \delta\vec{v}$. Further paralleling the analysis that led to the definition of private offer price, we are led to the following definition of **public offer price**, $\mathcal{W}(\theta^*, \vec{v})$:

9

$$\mathscr{W}\left(\theta^*, \vec{v}\right) = \lim_{\delta \to 0} \frac{\tilde{U}\left(\theta^* + \delta\vec{v}, B\left(\theta^* + \delta\vec{v}, \sigma^*_{-\text{DM}}\left(\theta^* + \delta\vec{v}\right)\right), \sigma^*_{-\text{DM}}\left(\theta^* + \delta\vec{v}\right)\right) - \tilde{U}\left(\theta^*, B\left(\theta^*, \sigma^*_{-\text{DM}}\left(\theta^*\right)\right), \sigma^*_{-\text{DM}}\left(\theta^*\right)\right)}{\delta}$$

(7)

which we sometimes abbreviate as $\mathscr{W}_v$.

In equation 7, $\sigma^*_{-\text{DM}}$ does not depend on σ_{DM}. Instead, $\sigma^*_{-\text{DM}}$ is the *reduced form* equilibrium strategy in terms of the parameter vector, θ. For example, in Cournot competition $\sigma_{-\text{DM}}$ gives players' *equilibrium* quantity in terms of cost and demand parameters and not in terms of DM's quantity. This approach is similar to that of [Caputo, 1996].

If we assume differentiability we can re-express the public offer price in differential form as formalized by the following lemma:

Lemma 6. *Assume differentiability of $\tilde{U}$, B and σ_{-DM} with respect to their arguments, then equation 7 is given by*

$$\mathscr{W}_v = \sum_j v_j \left[\frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial \theta_j} \right.$$
$$+ \sum_k \frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial [B(\theta, \sigma_{-DM})]_k} \left(\frac{\partial [B_{DM}(\theta, \sigma_{-DM})]_k}{\partial \theta_j} + \sum_l \frac{\partial [B_{DM}(\theta, \sigma_{-DM})]_k}{\partial [\sigma_{-DM}]_l} \frac{\partial [\sigma_{-DM}]_l}{\theta_j} \right)$$
$$\left. + \sum_l \frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial [\sigma_{-DM}]_l} \frac{\partial [\sigma_{-DM}]_l}{\partial \theta_j} \right]$$

(8)

where the derivatives are evaluated at $\theta = \theta^*$ and $\sigma = \sigma^*$.

Proof. The result follows by applying the chain rule and the definition of directional derivative. □

Once again, since we assume B is a best response function, we can simplify equation 8, which is formalized in the following corollary:

Corollary 2. *Assume B, $\tilde{U}$ and σ_{-DM} are differentiable with respect to all of their arguments. Then equation 8 reduced to*

$$\mathscr{W}_v = \sum_j v_j \left[\frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial \theta_j} + \sum_l \frac{\partial \tilde{U}(\theta, \sigma_{DM}, \sigma_{-DM})}{\partial [\sigma_{-DM}]_l} \frac{\partial [\sigma_{-DM})]_l}{\partial \theta_j} \right]$$

(9)

Proof. All $\frac{\partial \tilde{U}}{\partial [\sigma_{\mathbf{DM}}]_k}$ terms are 0 by the same argument in corollary 1. The claimed result follows. □

The first term in equation 9 represents the change in DM's utility as a result of a change in the parameter values, holding all strategies constant. This term is the same as the first term in equation 4. The second term represents how DM's utility changes as a result of all players other than DM changing their strategy in response to a change in the parameters. It is this term that differentiates the public offer price from the private offer price. When DM is deciding whether or not to pay to change the parameter values, they must anticipate how *all* players will change their strategy in response to the change in parameters since all players observe DM's decision. This effect is not present in the private offer price because DM's decision to change the parameters is unobserved. This can be formalized by the following equation:

$$\mathscr{W}_v - \mathscr{V}_v \;\; = \;\; \sum_l \frac{\partial \tilde{U}(\theta, \sigma_{\text{DM}}, \sigma_{-\text{DM}})}{\partial [\sigma_{-\text{DM}}]_l} \frac{\partial [\sigma_{-\text{DM}}]_l}{\partial \theta_j}. \tag{10}$$

This establishes that the difference between the private offer price and the public offer price is whether or not DM anticipates all other players changing their strategy in response to DM's decision to change the parameter values.

In summary, in this section we derived how much DM would be willing to pay for an infinitesimal change in the parameter values. Specifically, we derived these prices for the case where DM's decision to change the parameters is unobserved and the case where DM's decision is unobserved. When DM's decision is unobserved, under the Nash equilibrium solution concept, the private offer price is the same as that of the envelope theorem in non-strategic settings. In the case when DM's decision is observed by all players, DM's willingness to pay for a parameter change is given by the directional derivative of the utility function. Finally, we showed that the only difference between the two scenarios is that when DM's decision to change the parameter values is observed, the price depends on how other players will adjust their strategies in response to the parameter change. In the next section we show how to apply such results to a variety of scenarios.

4 Examples

We now present several examples to illustrate the concept and applicability of a public and private offer price for a variety of games. We will examine the private and public offer price on variants of Cournot competition, matching pennies and a tragedy of the commons game. We also illustrate how the public and the private offer price may differ. Finally, we will demonstrate how our differential framework can be applied when a decision maker is faced with the decision whether or not to change multiple parameters simultaneously.

4.1 Cournot Competition

We begin with a simple example to clearly illustrate a straightforward application of the private and public offer prices. Consider a two-player game of Cournot competition. Firm i, $i = 1, 2$ chooses $q_i \in [0, \infty]$. Inverse demand is linear and is given by $P = \alpha - (q_1 + q_2)$. Firm i's constant per-unit marginal cost is given by c_i. Firm i's profits are then given by the standard profit function $\tilde{U}_i(q_i, q_{-i}) = (\alpha - (q_i + q_{-i})) q_i - c_i q_i$. Firm i's best response function is given by $B_i(\alpha, q_{-i}) = \frac{\alpha - q_{-i} - c_i}{2}$. Assuming $\alpha + c_i - 2c_{-i} > 0$, the Nash equilibrium of the game is $q_i^* = \frac{\alpha + c_{-i} - 2c_i}{3}$. We now ask what is the private and public offer price to firm i for an infinitesimal change in α. Since the derivatives of the utility function and best response function are differentiable in α and the utility function is differentiable in q_i we can apply equation 3 directly:

$$\begin{aligned}
\frac{\partial \tilde{U}_i}{\partial \alpha} &= q_i \\
\frac{\partial B_i}{\partial \alpha} &= \frac{1}{2} \\
\frac{\partial \tilde{U}_i}{\partial q_1} &= \alpha - 2q_i - q_{-i} - c_i
\end{aligned} \tag{11}$$

Evaluating the derivatives of equation 11 at the Nash equilibrium values gives

$$\begin{aligned}
\frac{\partial \tilde{U}_i}{\partial \alpha} &= \frac{\alpha + c_{-i} - 2c_i}{3} \\
\frac{\partial B_i}{\partial \alpha} &= \frac{1}{2} \\
\frac{\partial \tilde{U}_i}{\partial q_i} &= \alpha - 2\left(\frac{\alpha + c_{-i} - 2c_i}{3}\right) - \frac{\alpha + c_i - 2c_{-i}}{3} - c_i = 0
\end{aligned} \tag{12}$$

Therefore, the private offer price for firm i to pay and change α is

$$\mathscr{V}(\alpha, 1) = \frac{\alpha + c_{-i} - 2c_i}{3} > 0 \tag{13}$$

This means that when two firms are competing in Cournot competition, there is a Nash equilibrium in which firm i pays to increase demand as long as the cost to change the parameter is less than $\frac{\alpha + c_{-i} - 2c_i}{3}$.

Now consider the public offer price. Again, since the requisite derivatives exist and we are considering the Nash equilibrium case, it is possible to apply equation 9. The first term is the same as the public offer price. Therefore, the public offer price is given by:

$$\begin{aligned}
\mathscr{W}(\alpha, 1) &= \frac{\alpha + c_{-i} - 2c_i}{3} + \frac{\partial \tilde{U}_i}{q_{-i}} \frac{\partial q_{-i}}{\partial \alpha} \\
&= \frac{\alpha + c_{-i} - 2c_i}{3} - q_i \frac{1}{3} \\
&= \frac{\alpha + c_{-i} - 2c_i}{3} - \frac{\alpha + c_{-i} - 2c_i}{3} \frac{1}{3} \qquad \text{(Plug in Nash equilibrium value for } q_i) \\
&= \frac{2}{9}(\alpha - 2c_i + c_{-i})
\end{aligned} \tag{14}$$

For comparison, note that player i's equilibrium payoffs are given by

$$\begin{aligned}
\tilde{U}_i^* &= \left(\alpha - (q_i^* + q_{-i}^*)\right) q_i - q_i^* c_i \\
&= \frac{1}{9}(\alpha - 2c_i + c_{-i})^2
\end{aligned}$$

Taking the derivative of equilibrium profits with respect to α gives

$$\frac{\partial \tilde{U}_i^*}{\alpha} = \frac{2}{9}(\alpha - 2c_i + c_{-i}) = \mathscr{W}(\alpha, 1) < \mathscr{V}(\alpha, 1). \tag{15}$$

We see that the public offer price is the same as the derivative of equilibrium profits with respect to α and that the public offer price is less than the private offer price. The difference between the public offer price and private offer price is due to the $\frac{\partial \tilde{U}_i}{q_{-i}} \frac{\partial q_{-i}}{\partial \alpha}$ term.

To understand the intuition, suppose firm i is considering if there is a profitable deviation from paying to change the parameters. If firm i's decision to pay is unobserved and firm i instead does not pay to change the demand parameter, their profits would decrease because demand goes down. At the same time, firm $-i$'s quantity would be "too high" since the decision to change parameters is unobserved and firm i would be better off if firm $-i$ lowered their quantity. On the other hand, if firm i's decision to change parameters is observed, then if firm

	H	**T**
H	(a,-1)	(-1,1)
T	(-1,1)	(1,-1)

Table 1: Payoff matrix for matching pennies example.

i deviates by not paying to change parameters, firm $-i$ will decrease their quantity in response to the change in parameters, which is beneficial for firm i. Since firm i benefits by firm $-i$ changing their price when firm i deviates by not paying, there is *more* of an incentive for firm i to *not* pay to change the parameters when their decision is observed. For this reason, the private offer price is higher than the public offer price.

This may seem counter-intuitive since if firm i had the opportunity to increase demand without firm $-i$ knowing, it would be able to increase its payoff due to a higher demand and also capture more of the market since firm $-i$ did not adjust its quantity to reflect the higher demand. However, this is not the case due to the definition of the Nash equilibrium. It can only be a Nash equilibrium for firm i to pay to change the parameter values if *both* firms best respond after firm i's decision to pay to change the parameter values. Therefore, if firm i chooses to pay to change the parameters, firm $-i$ will best respond to the new parameter values even though it does not observe firm i's decision. In other words, as long as firm $-i$ knows that DM has the *option* to pay to change the parameters, the Nash equilibrium solution concept ensures that firm $-i$'s action is the best response to the resultant parameter values.

4.2 Matching Pennies

Consider a version of the matching pennies game depicted in table 1 where we once again assume best responses. Without loss of generality assume that DM is the row player and a is a parameter that controls DM's preference for matching heads (H). Since players only have two strategies, let σ_{DM} and $\sigma_{\text{-DM}}$ represent the probability that DM plays H and -DM plays H, respectively. When $a = 1$, the unique Nash equilibrium of this game is $\sigma^*_{\text{DM}} = \sigma^*_{-\text{DM}} = \frac{1}{2}$ and equilibrium expected payoffs are 0.

More generally, when $a = a^* > -1$, the equilibrium profile specifies $\sigma^*_{\text{DM}} = \frac{1}{2}$ and $\sigma^*_{\text{-DM}} = \frac{2}{a^*+3}$ and the expected utility for DM is given by

$$
\begin{aligned}
\tilde{U}(a^*, \sigma^*_{\text{DM}}, \sigma^*_{\text{-DM}}) &= \sigma^*_{\text{DM}}\sigma^*_{\text{-DM}}a^* - \sigma^*_{\text{DM}}(1 - \sigma^*_{\text{-DM}}) - (1 - \sigma^*_{\text{-DM}})\sigma^*_{\text{DM}} + (1 - \sigma^*_{\text{DM}})(1 - \sigma^*_{\text{-DM}}) \\
&= \sigma^*_{\text{DM}}\sigma^*_{\text{-DM}}a^* + 3\sigma^*_{\text{DM}}\sigma^*_{\text{-DM}} - 2\sigma^*_{\text{DM}} - 2\sigma^*_{\text{-DM}} + 1 \\
&= \frac{a^* - 1}{a^* + 3}
\end{aligned}
\tag{16}
$$

where the argument to all σ^* terms is implicity a^*. Since at the mixed strategy equilibrium players are indifferent among all actions that have support under the equilibrium profile, B_{DM} is not unique when player $-DM$ plays the mixed strategy equilibrium profile. Therefore in this example we need to restrict B_{DM} to determine the public and private offer price. We do so in the following way: If $B_{\text{DM}}(a, \sigma_{-\text{DM}})$ is not unique, then $B_{\text{DM}}(a, \sigma)$ returns $\sigma_{-\text{DM}}$ such that -DM is indifferent between its actions. In other words, if DM is indifferent between its two possible actions, it randomizes such that -DM is also indifferent between its actions and therefore $B_{\text{DM}}(a, \sigma^*_{-\text{DM}})$ returns the mixed strategy equilibrium profile for DM when -DM randomizes with probability $\sigma^*_{-\text{DM}}$.

Furthermore, B_{DM} is not differentiable with respect to a when -DM chooses the mixed strategy Nash equilibrium profile. Specifically, for any small increase in a, DM's best response would be to place probability 1 on

H, holding -DM's strategy constant. Therefore, we *cannot* apply equation 3 to determine the private offer price. Instead, we use the definition given in equation 2 to derive the private offer price to DM at $a = a^*$:

$$
\begin{aligned}
\mathcal{V}(a^*, 1) &= \lim_{\delta \to 0} \frac{\widetilde{U}(a^* + \delta, B(a^* + \delta, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta)) - \widetilde{U}(a^*, B(a^*, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta))}{\delta} \\
&= \lim_{\delta \to 0} \frac{\widetilde{U}(a^* + \delta, B(a^* + \delta, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta)) - \widetilde{U}(a^*, B(a^* + \delta, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta))}{\delta} \\
&\quad + \lim_{\delta \to 0} \frac{\widetilde{U}(a^*, B(a^* + \delta, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta)) - \widetilde{U}(a^*, B(a^*, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta))}{\delta} \\
&= \sigma^*_{\text{DM}}(a^*)\sigma^*_{-\text{DM}}(a^*) + \lim_{\delta \to 0} \frac{\left(\frac{1}{2} \frac{2}{a^* + \delta + 3} a^* + \frac{3}{2} \frac{2}{a^* + \delta + 3} - \frac{2}{2} - 2 \frac{-2}{a^* + \delta + 3} - 2 \frac{2}{a^* + \delta + 3} \right)}{\delta} \qquad (17) \\
&= \sigma^*_{\text{DM}}(a^*)\sigma^*_{-\text{DM}}(a^*) + \lim_{\delta \to 0} \frac{\frac{-\delta}{a^* + \delta + 3}}{\delta} \\
&= \sigma^*_{\text{DM}}(a^*)\sigma^*_{-\text{DM}}(a^*) - \frac{1}{a^* + 3} \\
&= \frac{1}{2} \frac{2}{a^* + 3} - \frac{1}{a^* + 3} = 0 \qquad (18)
\end{aligned}
$$

Line 17 arises by first taking the derivative of $\widetilde{U}$ with respect to a. The second term results from DM setting $\sigma_{-\text{DM}} = 0$ since that is its best response when -DM is best responding as if $a = a^* + \delta$ but a actually equals a^*.

There are several points worth emphasizing in this example. First, even though DM's best response with respect to a change in a is not continuous, the private offer price is still defined. That is, even though DM's best response function specifies that DM would change their strategy from $\sigma_{\text{DM}} = .5$ to $\sigma_{\text{DM}} = 1$ for an infinitesimally small increase in a, the private offer price is still defined in terms of limits of the utility function. The reason is that the change in expected utility with respect to any change in DM's strategy is 0 *holding -DM's strategy and parameters constant*. Consequently any non-differentiable change in DM's strategy does not affect expected utility.

Second, the private offer price is exactly 0. The reason is that the Nash equilibrium expected utility to DM when $a = a^* + \delta$ is exactly the same as DM's expected utility when the column player best responds as if a is $a^* + \delta$ but a is actually a^* and DM best responds as if $a = a^*$. In other words, the utility gained by DM by exploiting a non-equilibrium strategy of -DM is exactly the same as the utility gained by DM for an increase in a.

Now, we derive the public offer price to DM for the matching pennies game for an increase in a. Again, since B is not differentiable with respect to a, we must apply equation 7 directly:

$$
\begin{aligned}
\mathcal{W} &= \lim_{\delta \to 0} \frac{\widetilde{U}(a^* + \delta, B(a^* + \delta, \sigma^*_{-\text{DM}}(a^* + \delta)), \sigma^*_{-\text{DM}}(a^* + \delta)) - \widetilde{U}(a^*, B(a^*, \sigma^*_{-\text{DM}}(a^*)), \sigma^*_{-\text{DM}}(a^*))}{\delta} \\
&= \lim_{\delta \to 0} \frac{\frac{1}{2} \frac{2}{3 + \delta + a^*} + 3\frac{1}{2} \frac{2}{3 + \delta + a^*} - \frac{2}{2} - 2\frac{2}{3 + \delta + a^*} + 1 - \left(\frac{1}{2} \frac{2}{3 + a^*} + 3\frac{1}{2} \frac{2}{3 + a^*} - \frac{2}{2} - 2\frac{2}{3 + a^*} + 1 \right)}{\delta} \\
&= \lim_{\delta \to 0} \frac{\frac{4}{3 + a} - \frac{4}{3 + \delta + a}}{\delta} = \frac{4}{(3 + a)^2} > 0 \qquad (19)
\end{aligned}
$$

Equation 19 shows that the public offer price to the row player is positive, even though equation 18 says the private offer price is zero. This means that DM would be willing to pay to increase *a only if column player was aware that DM paid for the parameter change*.

4.3 Tragedy of the Commons

We now present an example for a tragedy of the commons game with two players who simultaneously fish the ocean off some coastline. Player i chooses a continuous value $x_i \in [0, \infty]$, which represents how many fish player i catches. There is a non-strategic regulator that patrols the pond that has a fuel budget $\theta \in [0, 1]$. Player i's expected utility is given by:

$$\tilde{U}_i(x_i, x_{-i}) = \begin{cases} \frac{x_i}{x_{-i}^2} & \text{if } x_i \leqslant .25 \\ (1 - \theta)\frac{x_i}{x_{-i}^2} - \theta x_i^2 & \text{if } x_i > .25 \end{cases} \tag{20}$$

Intuitively, the regulator sets a limit that says that players can only fish .25 units.[3] However, the regulator's fuel budget prohibits the regulator from patrolling all of the pond so therefore a player that fishes more than .25 units evades the regulator with probability θ. However, if player i fishes over the limit and is detected, they incur a cost of $-x_i^2$. Depending on the value of θ, there are 4 pure strategy Nash equilibria in this game. One equilibrium is when both players fish .25, two equilibria consist of 1 player fishing over the limit and the other player fishing the limit and one equilibrium where both players fish over the limit.

Suppose θ is low enough so that both players choose to fish over the limit and risk being detected. Then, player i's best response is given by $BR_i(x_{-i}, \theta) = \frac{1-\theta}{2\theta x_{-i}^2}$ and the symmetric Nash equilibrium is for player i to fish $\left(\frac{1-\theta}{2\theta}\right)^{\frac{1}{3}}$. Then, the private offer price to player i is given by:

$$\mathscr{V}(\theta, 1) = \frac{\partial \tilde{U}_i}{\partial \theta} = \frac{-x_i}{x_{-i}^2} - x_i^2 \tag{21}$$

Although we can plug in the equilibrium values of x_i and x_{-i}, since x_i and x_{-i} are strictly positive, it is easy to see that equation 21 is negative. That means that player i would *have to be paid* in order to agree to an increase in θ. To put this result in context, we now consider the public offer price, which is given by:

$$\begin{aligned} \mathscr{W}(\theta, 1) &= \frac{-x_i}{x_{-i}^2} - x_i^2 + 2(1 - \theta)\frac{x_i}{x_{-i}^3}\frac{1}{3\theta^2}\left(\frac{1}{2(\frac{1}{\theta} - 1)^2}\right)^{\frac{1}{3}} \\ &= \frac{1 - 3\theta}{3(4(1 - \theta)\theta^2)^{\frac{1}{3}}} \end{aligned} \tag{22}$$

where equation 22 arises by plugging in the Nash equilibrium values of x_i and x_{-i}. For low enough θ, equation 22 is *positive*. In other words, a player would be willing to increase θ only if their decision to increase the budget was observed by other players.

To see why the sign of the public and private offer price are different, note that this is a tragedy of the commons game. For example, suppose both players fish the same amount. Then the players' utility is decreasing in the amount the other player fishes, but each player has an incentive to increase the amount they fish. As a result, both players fish above the social optimum. However, when θ increases, both players fish less since they are more likely to be detected and the penalty they face for being detected is increasing in the amount they fish. In other words, an observed increase in the fuel budget decreases the amount both players fish. However, the gain in firm i's expected utility due to firm $-i$ fishing outweighs the loss in firm i's expected utility from fishing less and thus the public offer price is positive.

---

[3]Of course, the legal limit can also be a parameter but we set it to .25 for ease of exposition.

Now consider the case when player i's decision to increase the regulator's budget is not observed. For it to be a Nash equilibrium for player i to pay to increase the budget, there must not be an incentive for i to not pay and adjust their strategy accordingly. Suppose i does pay some positive amount to change θ to $\theta + \delta$ and both players best respond accordingly. In this case, the amount that both players fish is less at $\theta + \delta$ than it is at θ. However, player i can benefit by *not paying* to increase the regulator's budget, increase the amount they fish, and decreases the probability that they are detected. Crucially, player $-i$ does *not* increase the amount that they fish because player i's decision to change the parameters are unobserved. Therefore, player i is a victim of their own greed. That is, player i knows that it would have a higher utility when the parameters were $\theta + \delta$ but i would never pay to change the parameters because it would always be profitable for i to avoid the cost of paying and cheat.

4.4 Cournot Competition - Revisited

In the first example of this section, we examined the public and private offer price for a change in the demand parameter in Cournot competition. The two other previous examples demonstrated the public and private offer price for one parameter change. This example will now examine the effect of changing two parameters simultaneously. Reconsider the Cournot game above. Instead of firm i facing the decision of whether or not to increase just α, firm i faces the decision to increase (α, c_i) in direction (x, y). Roughly speaking, this might represent firm i choosing the increase advertising (which increases α) but is also costly, and therefore increases per-unit marginal costs. The private offer price can be written:

$$\mathcal{V}((\alpha, c_i), (x, y)) = (x, y) \cdot (q_1, -q_1) \tag{23}$$

Without plugging in for equilibrium values of q_i, equation 23 shows that as long as $x > y$, player i would be willing to pay a positive amount to change α and c_1. That is, as long as the rate at which the intercept increases is greater than the rate at which costs increase, player i would be willing to change the parameters *if the decision to change parameters was unobserved*.

On the other hand, we can look at the public offer price of the same change:

$$
\begin{aligned}
\mathcal{V}((\alpha, c_i), (x, y)) &= (x, y) \cdot (q_i, -q_i) + (x, y) \cdot \left(-\frac{q_i}{3}, -\frac{q_i}{3}\right) \\
&= q_i \left(\frac{2}{3}x - \frac{4}{3}y\right).
\end{aligned}
\tag{24}
$$

In this case, the rate at which α increases must be at least *twice* the rate at which c_i increases. The difference in the public and private offer price is once again due to the fact that firm $-i$ adjusts their equilibrium action when they observe whether or not $-i$ paid to change the parameters and such an adjustment hurts firm i. While the intuition of this example is relatively straightforward, the important aspect is to see how the the public and private offer prices can also be used to derive the marginal rate of substitutions in games.

5 Private and Public Offer Prices Under Bounded Rationality

In section 3, we defined the private and public offer price when DM is perfectly rational. That is to say, when DM considers whether or not to pay for a change in the parameter values, DM anticipates that all players (including DM) will best respond in the second stage of the game. However, the definition of public and private offer price

are the same for *any* response function specification and corresponding equilibrium concept.[4] However, there are several mathematical as well as intuitive differences that arise. In this section, we illustrate those differences using the (logit) quantal response solution concept. To do so, let $B_i(\theta, \sigma_{-i}) \to \sigma_i$ be a vector valued function taking (θ, σ_{-i}) as an input and outputting a probability distribution on X_i such that for all $x_i \in X_i$,

$$Pr(x_i) = \frac{e^{\beta_i(\tilde{U}(\theta, x_i, \sigma_{-i}))}}{\sum_{j=1}^{|X_i|} e^{\beta_i(\tilde{U}(\theta, x_j, \sigma_{-i}))}} \tag{25}$$

where β_i is a scalar parameter. Note that under this assumption the equilibrium profile $\sigma^*(\theta)$ is a quantal response equilibrium. We can then motivate the private and public offer prices in the same way as in section 3.1. More specifically, DM would only be willing to pay to change the parameters if they could not benefit by saving the cost of changing the parameters and adjusting their strategy accordingly.

Mathematically, the definition of the private and public offer price under bounded rationality is identical to the definitions given in equations 3 and 8. The distinction is that the best response function is replaced by the quantal response function, and the Nash equilibrium strategy profile is replaced by the quantal response equilibrium profile when evaluating derivatives. The other distinction is that the derivative of DM's utility with respect to their own strategy is *not* 0 under the quantal response formulation so the simplification of offer prices given in equations 4 and 9 does not apply under quantal response functions.

The main conceptual difference between the offer prices under bounded rationality and the Nash case has to do with the notion of equilibrium. Recall that in the Nash case, the offer prices were the most DM would be willing to pay to change the parameter values such that it is a Nash equilibrium in the outer game for DM to make such a payment. However, under the quantal response solution concept (and any general solution concept), it is not generally true that the offer price are the most DM would be willing to pay such that it is a quantal response equilibrium for DM to pay in the **outer** game without adding additional structure. The reason is because in the first stage of the outer game, DM only pays to change the parameter values if their utility in the resulting equilibrium is greater than if it did not pay. In other words, DM in the first stage is perfectly rational while it is boundedly rational in the second stage. This discrepancy muddles the notion of equilibrium.

A formal way to clearly define an equilibrium when this discrepancy is present is to adopt the agent-representation of players in extensive form games [McKelvey and Palfrey, 1998]. In this case, the N player game is expanded to an $N + 1$ player game in which DM in the first stage is considered to be a different player than DM in the second stage, but both versions of DM have the same utility function. In our formulation, the version of DM in the first stage of the game is fully rational but anticipates that all players (including DM) would be boundedly rational in the second stage of the game. Formally, this can be done by setting the β parameter for DM in the first stage to ∞. Then, the offer prices under the quantal response equilibrium concept represent the most that DM would be willing to pay such that it is a quantal response equilibrium in the $N + 1$ agent representation of the outer game. This notion of equilibrium implies that DM in the first stage anticipates their own irrationality in the second stage of the game and then chooses the best action in the first stage, knowing that they might not act optimally in the second stage.

5.1 QRE Offer Prices and Matching Pennies

Recall the matching pennies formulation in section 4.2 in which the private offer price is 0 for all values of a and the public offer price is given by $\frac{4}{(3+a)^2}$. In this section, we give a numerical representation of the public and

---

[4]Recall, we limit the inner game to a simultaneous move game. We discuss why this is not true of extensive form games in the next section.
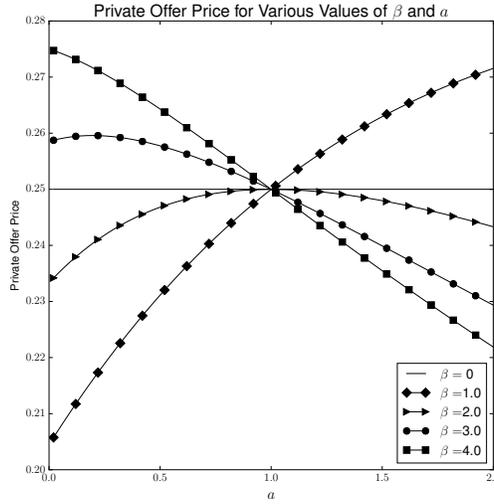
Figure 1: Private offer price for QRE solution concept.

private offer price for the quantal response solution concept and compare it to the Nash case. Figure 1 plots the private offer price for different values of β and a to DM in the matching pennies game under the quantal response equilibrium concept. Recall that a is DM's payoff by matching "heads" and β is the rationality parameter. For simplicity, we assume β is common to both players.

As figure 1 illustrates, the quantal response private offer price contains more richness than in the Nash case. First, when $\beta = 0$, the private offer price is .25 for all values of a. The reason for this is that for all values of a, both players randomize uniformly and the outcome (H, H) occurs with probability .25. Therefore, increasing a by amount δ increases a's expected utility by $.25\delta$ for any value of δ. Since player's strategies don't change when $\beta = 0$, the second term in equation 3 is 0 and the private offer price is .25.

However, when $\beta > 0$, the private offer price depends on both β and a. Figure 1 shows that when β is relatively low, the private offer price is higher for higher values of a but when β is relatively high, the opposite is true.

Figure 2 once again plots the private offer price when $\beta = 1$ as well as the first term ($\frac{\partial U}{\partial a}$) of the private offer price. Therefore, the difference between the solid line and the lined marked with diamonds gives the impact on DM's utility as a result of DM adjusting their strategy in response to a parameter change (i.e. the terms after the second summation in equation 3). As the plot shows, $\frac{\partial U}{\partial a}$ is always positive. However, when a is low, the impact on DM's utility due to DM changing their strategy in response to an increase in a is *negative*. The reason is that when $a < 1$ and $-DM$ is playing the quantal response equilibrium strategy, DM's *best* response would be to always play "T". However, under the quantal response equilibrium concept DM doesn't play "T" but instead randomizes according to the logit probabilities given in equation 25. Under such a solution concept, increasing a induces DM to place a *higher* probability on "H", even though it is not DM's best strategy. On the other hand, when $a > 1$, DM's best strategy *is* to always play "H" and thus increasing a makes DM place more weight on "H" and DM increases their utility by adjusting their strategy in response to a change in a. Put another way, DM would be willing to pay for an increase in a. However, when $a < 1$, DM's willingness to pay is dampened by the fact that increasing a makes a sub-optimal action more desirable and thus increases the probability that a will not best respond.

Figure 3 shows the same phenomenon when $\beta = 4$. That is, $\frac{\partial U}{\partial a} > 0$ and the private offer price is less than $\frac{\partial U}{\partial a}$ when $a < 1$ and greater than $\frac{\partial U}{\partial a}$ when $a > 1$. It is the slope of the $\frac{\partial U}{\partial a}$ curve that accounts for the difference in the
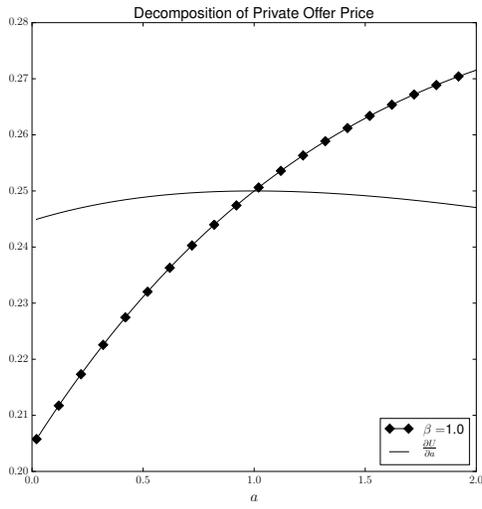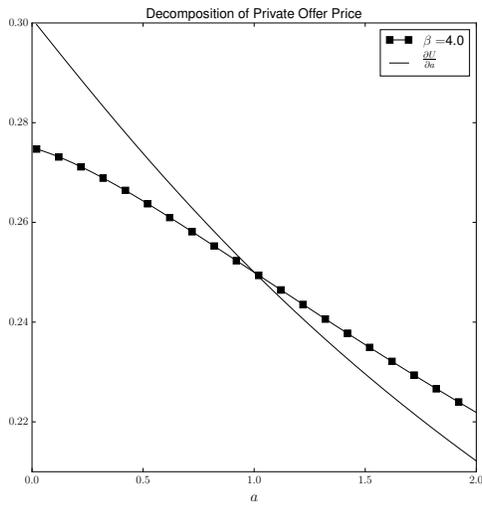
Figure 2: Decomposition of private offer price

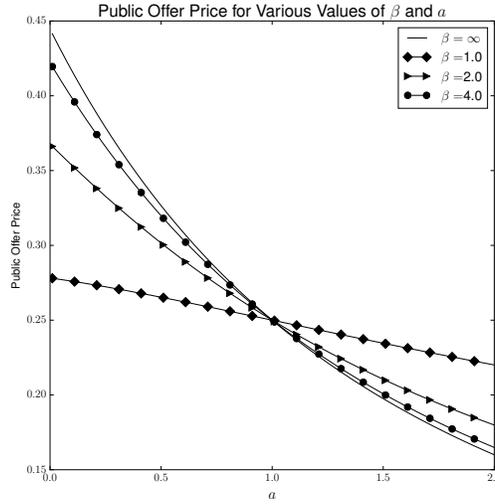Figure 3: Decomposition of private offer price when $\beta = 4$

Figure 4: Public offer price for QRE solution concept.

private offer price when $\beta = 1$ and $\beta = 4$. As figure 3 shows, when $\beta = 1$, $\frac{\partial U}{\partial a}$ is relatively flat but when $\beta = 4$, the curve slopes down. To see why, consider the perfectly rational case in which players play the Nash equilibrium. In that case, as a gets larger, player -DM puts lower probability on 'H' in order to keep DM indifferent between heads and tails while DM randomizes with probability $\frac{1}{2}$ for all values of a. Therefore, the impact of an increase in a on DM's utility (holding strategies constant) is less when a is high. The same logic applies under the QRE solution concept except the player's mixing probabilities do not respond as strongly to an increase in a as they do in the Nash case. In other words, as β gets larger, the players move closer to the Nash equilibrium solution and under the Nash solution, $\frac{\partial U}{\partial a}$ is deceasing in a.

In the most general sense, it is possible to have DM's payoffs for *all* possible outcomes increase but the private offer price actually be *negative*. For example, let θ be a vector such that each element of θ represents DM's payoff in a specific outcome. Furthermore, suppose that $\frac{\partial U}{\partial \theta_j}$ is positive for all components of θ. Even though DM's utility would increase under all possible outcomes, it is possible that the change in utilities for each action makes it more likely that DM chooses a sub-optimal action. If the negative effect of DM adjusting their strategy is greater in magnitude than the increase in expected utility due to the parameter change, the private offer price would be negative. Characterizing the games and solution concepts in which this phenomenon occurs is a direction of future work.

Figure 4 plots the public offer price under the quantal response solution concept for various values of a and β. In this case, the public offer price behaves qualitatively similar under the quantal response solution concept as under the Nash solution concept. Indeed, as $\beta \to \infty$, the public offer price under quantal response equilibrium approaches the public offer price under the Nash equilibrium.

6   Extensions and Future Work

In this work, we introduced a formalism to determine how much a player in a game would be willing to pay to change the parameters of a simultaneous move game. An obvious extension to this work is to expand the inner game to extensive form games and more refined solution concepts (such as the sequential equilibrium). However,

	l	r
t	4,3	6,1
m	-4,0	0,1
b	3,0	5,5

	l	r
t	-4,0	0,1
m	4,3	6,1
b	3,0	5,5

1,4   out / in — 2   2,0   out / in — 1

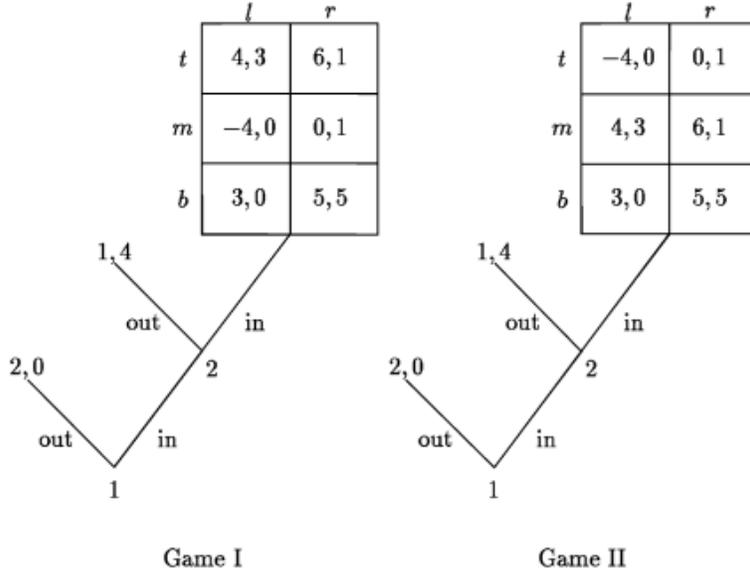1,4   out / in — 2   2,0   out / in — 1

Game I   Game II

Figure 5: Example adopted from Hurkens and Vulkan [2006]

such an extension is non-trivial. This is most evident in the counter-example given in [Hurkens and Vulkan, 2006] and reproduced in figure 5.

In this example, the inner game is as follows. Nature chooses either Game *I* or Game *II* with equal probability, then player 1 chooses in or out. Then player 2 chooses in or out. Then, if neither player has chosen out, the players play the simultaneous move game. In all cases, player 2 does not observe the move of nature. Let $\theta = 0$ indicate that player 1 observes nature's move perfectly with no noise and let $\theta = 1$ indicate the case in which player 1's signal provides no information about the game nature chose. The question then is "under what conditions is it a *sequential equilibrium* for player 1 to pay to change θ from 0 to 1 when player 2 does not observe whether or not player 1 paid for the change?"

Intuitively, it appears that it would be possible to establish the following conjecture (parallel to lemma 2) when $\gamma(\theta^*)$ is a general extensive form game:

Conjecture 1. *Suppose $\sigma^*(\theta^*)$ and $\sigma^*(\theta^* + \delta\vec{v})$ are* sequential *equilibria of $\gamma(\theta^*)$ and $\gamma(\theta^* + \delta\vec{v})$, respectively. Then it is a* sequential *equilibrium for DM to pay c in the outer game $\overline{\gamma}(\theta^*)$ to change the parameters from θ^* to $\theta^* + \delta\vec{v}$ iff*

$$
c \;\leqslant\; \tilde{U}\left[\theta^* + \delta\vec{v},\; B\left(\theta^* + \delta\vec{v}, \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right),\; \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right]
$$
$$
- \tilde{U}\left[\theta^*,\; B\left(\theta^*, \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right),\; \sigma^*_{-DM}\left(\theta^* + \delta\vec{v}\right)\right] \tag{26}
$$

Proof. We prove sufficiency while necessity follows by the same argument. By assumption, $\sigma^*(\theta^* + \delta\vec{v})$ is a sequential equilibrium of the game $\gamma(\theta^* + \delta\vec{v})$ so no player can change their action in the second stage of $\overline{\gamma}$ and increase their payoff. The only deviation that might make a player better off is if DM chooses not to purchase the

change in parameters. However, if DM deviates by not purchasing the change in parameters, DM's net change in utility is upper bounded by:

$$\tilde{U}\left(\theta^*,\; B\left(\theta^*, \sigma^*_{-\text{DM}}\left(\theta^* + \delta\vec{v}\right)\right),\; \sigma^*_{-\text{DM}}\left(\theta^* + \delta\vec{v}\right)\right) - \tilde{U}\left(\theta^* + \delta\vec{v},\; \sigma^*_{\text{DM}}\left(\theta^* + \delta\vec{v}\right),\; \sigma^*_{-\text{DM}}\left(\theta^* + \delta\vec{v}\right)\right) + c \quad (27)$$

If the inequality in Lemma 2 holds, then the net change in DM's utility is less than or equal to 0, which implies DM would *not* have an incentive to deviate. This establishes that no player has an incentive to deviate and therefore it is a Nash equilibrium for DM to pay to change the parameter values. □

Unfortunately, conjecture 1 is false. We first illustrate a counter-example and then dissect the proof to see exactly why. In the example given in 5 the unique sequential equilibrium profile when θ is exogenously specified to be 1 in the example described above is $((in, r), (in, b))$ and player 1 earns an expected utility of 5. On the other hand, if player 1 had the opportunity to pay to change the parameter to 1 (at no cost) and play the resulting sequential equilibrium, we see that player 1 would not do so because player 1 would benefit by not paying, learning the move of nature and earning an expected utility of 6 *holding player 2's strategy constant at the sequential equilibrium strategy when θ is exogenously specified to be* 1. This difference in expected utilities for player 1 implies the player 1 would have to be paid 1 unit of utility to choose to not learn the outcome of nature. However, when $c = -.25$ (i.e. player 1 gets paid , 25) the unique sequential equilibrium of the outer game is for player 1 to choose to not learn the outcome of nature and choose out! So, even though c is greater than the quantity determined by the right hand side of equation 26, it is still a sequential equilibrium for player 1 to pay to change θ to 1.

In a broader sense, the reason that 1 is false is because the sequential equilibrium strategy profile of the *outer game* when DM pays to change the parameters is not necessarily a sequential equilibrium of the game in which the parameters are specified exogenously. In the exact wording of the proof, the misleading logic is found in the sentence "By assumption, $\sigma^*(\theta^* + \delta\vec{v})$ is a sequential equilibrium of the game $\gamma(\theta^* + \delta\vec{v})$ so no player can change their action in the second stage of $\bar{\gamma}$ and increase their payoff." While this is true, the crucial element that is hidden is that a sequential equilibrium is a set of actions *and* beliefs. More specifically, $\sigma^*(\theta^* + \delta\vec{v})$ is part of a sequential equilibrium *only because it is justified by correct beliefs*.[5] However, the beliefs that support a sequential equilibrium when the parameters are exogenously specified do not necessarily support a sequential equilibrium in which a player has the opportunity to change the value of the parameters. This implies that the strategy profile $\sigma^*_{-i}(\theta^* + \delta\vec{v})$ is not optimal in the outer game since it is only sequentially rational under the set of beliefs derived from the game with exogenously specified parameters, and such beliefs do not form part of a sequential equilibrium when a player has the opportunity to change the parameter values. Relating this to the example above, when player 1 has the opportunity to set the parameters, the belief supported under a sequential equilibrium for player 2 *if they reaches their first information set and has the opportunity to act* is to believe that player 1 did not pay and thus knows the move of nature, in which case player 2 would choose out. This is contradictory to the case when the parameters are exogenously specified such that player 1 is uninformed. In that case, player 2 (correctly) believes that player 1 is uninformed and those beliefs are part of a sequential equilibrium.

As this example shows, the logic of the public and private offer price does not immediately extend to extensive form games. Therefore one direction of future work is to determine if there are conditions that can be placed on general extensive form games such that the definition of public and private offer prices apply. Such assumptions might include full support assumptions or perfect information assumptions. Another direction of future research is to establish general conditions on a game's primitives that give information about the public and private offer prices. One example question might be what are the properties of utility functions and parameters such that the public offer price and the private offer price always have different signs?

---

[5]By "correct beliefs" we mean beliefs that satisfy the conditions of sequential equilibrium.

7   Acknowledgements

References

Daron Acemoglu and Martin Kaae Jensen. Aggregate comparative statics. *Games and Economic Behavior*, 81: 27–49, 2013.

Rabah Amir and Natalia Lazzati. Endogenous information acquisition in bayesian games with strategic complementarities. *Journal of Economic Theory*, 163:684 – 698, 2016. ISSN 0022-0531. doi: http://dx.doi.org/10.1016/j.jet.2016.03.005. URL `http://www.sciencedirect.com/science/article/pii/S0022053116000399`.

Bruno Bassan, Olivier Gossner, Marco Scarsini, and Shmuel Zamir. Positive value of information in games. *International Journal of Game Theory*, 32(1):17–31, 2003.

Nils Bertschinger, David H. Wolpert, Eckehard Olbrich, and Jürgen Jost. Information geometry of influence diagrams and noncooperative games. *CoRR*, abs/1401.0001, 2014. URL `http://arxiv.org/abs/1401.0001`.

Michael R. Caputo. The envelope theorem and comparative statics of nash equilibria. *Games and Economic Behavior*, 13(2):201 – 224, 1996. ISSN 0899-8256. doi: http://dx.doi.org/10.1006/game.1996.0034. URL `http://www.sciencedirect.com/science/article/pii/S0899825696900342`.

Federico Echenique. Comparative statics by adaptive dynamics and the correspondence principle. *Econometrica*, 70(2):833–844, 2002.

Sjaak Hurkens and Nir Vulkan. Endogenous private information structures. *European Economic Review*, 50(1): 35–54, 2006.

Steven A Lippman, John W Mamer, and Kevin F McCardle. Comparative statics in non-cooperative games via transfinitely iterated play. *Journal of Economic Theory*, 41(2):288 – 303, 1987. ISSN 0022-0531. doi: http://dx.doi.org/10.1016/0022-0531(87)90021-4. URL `http://www.sciencedirect.com/science/article/pii/0022053187900214`.

Richard D McKelvey and Thomas R Palfrey. Quantal response equilibria for extensive form games. *Experimental economics*, 1(1):9–41, 1998.

Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.

Paul Milgrom and Chris Shannon. Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, pages 157–180, 1994.

Stephen Morris and Hyun Song Shin. Social value of public information. *The American Economic Review*, 92 (5):1521–1534, 2002.

Abraham Neyman. The positive value of information. *Games and Economic Behavior*, 3(3):350–355, 1991.

Martin J Osborne and Ariel Rubinstein. *A course in game theory*. MIT press, 1994.

Tarun Sabarwal Sunanda Roy. On the (non-)lattice structure of the equilibrium set in games with strategic substitutes. *Economic Theory*, 37(1):161–169, 2008. ISSN 09382259, 14320479. URL `http://www.jstor.org/stable/40282910`.

Takashi Ui and Yasunori Yoshizawa. Characterizing social value of information. *Journal of Economic Theory*, 158, Part B:507 – 535, 2015. ISSN 0022-0531. doi: http://dx.doi.org/10.1016/j.jet.2014.12.007. URL `http://www.sciencedirect.com/science/article/pii/S0022053114001847`. Symposium on Information, Coordination, and Market Frictions.

J.Miguel Villas-Boas. Comparative statics of fixed points. *Journal of Economic Theory*, 73 (1):183 – 198, 1997. ISSN 0022-0531. doi: http://dx.doi.org/10.1006/jeth.1996.2224. URL `http://www.sciencedirect.com/science/article/pii/S0022053196922243`.