Games with Switching Costs and Endogenous References*

Begum Guney† Michael Richter‡

June 30, 2021

Abstract

We introduce a game-theoretic model with switching costs and endogenous references. An agent endogenizes his reference strategy and then, taking switching costs into account, he selects a strategy from which there is no profitable deviation. We axiomatically characterize this selection procedure in one-player games. We then extend this procedure to multi-player simultaneous games by defining a *Switching Cost Nash Equilibrium* (SNE) notion, and prove that (i) an SNE always exists; (ii) there are sets of SNE which can never be a set of Nash Equilibrium for *any* standard game; and (iii) SNE with a specific cost structure exactly characterizes the Nash Equilibrium of nearby games, in contrast to Radner's (1980) ε-equilibrium. Subsequently, we apply our SNE notion to a product differentiation model, and reach the opposite conclusion of Radner (1980): switching costs for firms may benefit consumers. Finally, we compare our model with others, especially Köszegi and Rabin's (2006) *personal equilibrium*.

Keywords: Switching Cost Nash Equilibrium, Choice, Endogenous Reference, Switching Costs, Epsilon Equilibrium

JEL classification: D00, D01, D03, C72

*We thank Levent Koçkesen, Yusufcan Masatlioglu, George Mailath, Ariel Rubinstein, Rani Spiegler, Roee Teper, and three anonymous referees for their helpful comments and suggestions. The first author acknowledges financial support from the Scientific and Technological Research Council of Turkey, Project #113K370.

†Department of Economics, Ozyegin University. E-mail: begum.guney@ozyegin.edu.tr

‡Department of Economics, Royal Holloway, University of London. E-mail: Michael.Richter@rhul.ac.uk

1 Introduction

Switching costs are a feature of daily life. In individual settings, a new job offer in another state involves social and financial relocation costs; firms build in switching costs for consumers by offering lengthy contracts (or loyalty programs) with penalties for switching providers. In multi-agent settings, such as the competition between firms, firms may face switching costs when changing their production technologies or the composition of their workforce. Such costs play an important role in a variety of decisions and can explain experimental deviations from rational behavior (Guney and Richter, 2018). Moreover, these costs can be economically significant. For example, on June 12, 2020, United Airlines made an 8-K filing where the airline valued its mileage program at \$21.9 billion while its stock valuation on that day was \$11.5 billion.[1] That is, the value of the airline itself, e.g. planes, routes, fares, landing spaces, etc. is negative, and this negative value is more than outweighed by the value of its mileage program. To rephrase an old joke, "United Airlines is a mileage program with an airline attached".

We study a model of switching costs with endogenous references in a game-theoretic environment. An agent evaluates each of his own strategies by viewing it as a reference and considering deviations from it by taking into account both the utility difference and switching cost of such a deviation, given the strategies of other players. In an equilibrium, each agent chooses a strategy from which there is no profitable deviation, given the equilibrium play of others. We term this notion a *Switching Cost Nash Equilibrium* (SNE) and study SNE in two game-theoretic settings: one-player and multi-player games.

In the one-player setting, we provide an axiomatic characterization of our model which uses the classical α and γ axioms of Sen (1971), and two convexity-like axioms. As in the previously mentioned real-life examples, the switching cost function in our model depends on both the reference strategy and the strategy to which the agent switches. Switching costs are non-negative and have a special "excess bilinear" formulation: when deviating from a mixed strategy to another mixed strategy, the agent incurs a switching cost that is bilinear over the *excess* parts of these mixed strategies, rather than over the mixed strategies themselves. To fix ideas, imagine a factory owner who allocates his employees among a set of tasks and for whom it is costly to move employees between a pair of tasks. So, an allocation here can be thought of as a mixed strategy

[1] For the 8-K filing, see `https://www.sec.gov/ix?doc=/Archives/edgar/data/100517/000110465920073190/tm2022354d3_8k.htm` and for historical market caps, see `https://ycharts.com/companies/UAL/market_cap`.

where the probability p_i specifies the fraction of workers assigned to task i. In order to move from one allocation to another, the factory owner does not need to reassign every worker. Rather, he only reassigns workers from tasks with excess supply (in the initial allocation) to tasks that need additional workers (for the final allocation), and this reassignment is done proportionally. An allocation is an equilibrium if the factory owner finds it weakly non-profitable to alter the employees' work assignments.

In the multi-player setting, we study SNE in simultaneous games and show that the *set* of possible standard Nash Equilibria (NE) is a strict subset of the *set* of possible SNE. That is, there are equilibria that arise in our switching cost model which cannot arise in *any* standard game. This is different than saying that "adding switching costs creates more equilibria given the same utility function" (which is always true). Rather, it states something that is not immediately obvious, namely that there are *new* equilibria sets which are not realizable for *any* standard game with possibly different utility functions *and* actions. Next, we show that the only overlap between the SNE notion and Radner's (1980) ε-equilibrium notion is the NE. Then, we find that for a particular cost function, the SNE of a game equals the set of NE of all nearby games. This is a strengthening of Mailath et al.'s (2005) one-way result for the connection between ε-equilibria and the NE of nearby games. Consequently, a reader who is not interested in SNE for its own sake, but who is interested in nearby games, may still find SNE to be a technically useful tool.

In our application, we study a model of vertical product differentiation where firms face switching costs. When switching costs are low, nothing goes beyond NE, and when switching costs are high, anything goes. However, the case of intermediate switching costs is surprising. There are new SNE, and in them, consumers benefit, firms suffer, and the consumer effect dominates, so overall efficiency increases.

Finally, our model of switching costs with endogenous references relates to the *personal equilibrium* notion of Köszegi and Rabin (2006, 2007). However, there are significant behavioral differences. For example, their model does not necessarily nest NE and equilibria may not exist in their setting.

The paper proceeds as follows. Section 2 provides an axiomatic study of one-player games with switching costs and endogenous references. Section 3 studies the SNE of multi-player games and relates it to the notion of ε-equilibrium. Section 4 contains an application to vertical product differentiation. Section 5 analyzes alternative models and Section 6 surveys the related literature. Section 7 discusses some implications of our model. All proofs are presented in the Appendix A and two examples regarding personal equilibrium models are provided in Appendix B .

2 One-Player Games

In this section, we take an axiomatic approach to understanding an agent's choice behavior in one-player games (decision problems). Let X be a finite grand set of actions (pure strategies) and $\mathcal{X}$ be the set of all non-empty subsets of X. A *mixed strategy* m is a convex combination of pure strategies a^j with weights m^j, that is, $m = \sum_j m^j a^j$. The term mixed strategy is used to refer to both trivial mixtures (pure strategies) and non-trivial mixtures (strictly mixed strategies). Typically, we use $a, b, ...$ to denote pure strategies and $m, n, ...$ to denote mixed strategies (whether trivial or not). For any set of pure strategies A, we denote the associated set of mixed strategies by $\Delta(A)$. A choice correspondence C is a non-empty valued map $C : \mathcal{X} \rightrightarrows \Delta(X)$ such that $C(A) \subseteq \Delta(A)$ for any $A \in \mathcal{X}$. In words, from each set of pure strategies, an agent can only choose mixtures of those pure strategies.

We first propose and discuss a choice procedure of costly switching from endogenous references, and then provide an axiomatic characterization for that.

Definition 2.1 *A choice correspondence C represents costly switching from endogenous references if for any $A \in \mathcal{X}$,*

$$C(A) = \{m \in \Delta(A) |\ U(m) \geq U(m') - D(m', m)\ \ \forall m' \in \Delta(A)\},$$

for some (1) $U : \Delta(X) \to \mathbb{R}$ such that $U(m) = \sum_j m^j U(a^j)$ for all $m \in \Delta(X)$ and $a^j \in supp(m)$,

(2) $D : \Delta(X) \times \Delta(X) \to \mathbb{R}_+$ such that $D(m, m) = 0$ for all $m \in \Delta(X)$

Here, U is an expected utility function and we interpret it as the reference-free utility of each strategy. We call D a *switching cost* function, and we interpret $D(m', m)$ as the switching cost the agent incurs when he deviates to strategy m' from a reference strategy m. Note that staying with the same strategy m is associated with zero cost while the cost of switching to another strategy m' is non-negative. In this model, the agent chooses a strategy m if he wouldn't want to deviate to any other strategy m' when m is his reference, taking the utility U and switching costs $D(m', m)$ into account. The choice correspondence C selects all strategies from which the agent cannot profitably deviate.

In the procedure of costly switching from endogenous references, we focus on a specific class of switching cost functions, namely excess bilinear ones. To keep track of excesses, we denote the overlap between any two mixed strategies m and n by the vector $o := \min(m, n)$.

Definition 2.2 *A switching cost function* $D : \Delta(X) \times \Delta(X) \to \mathbb{R}_+$ *is called excess bilinear if for any* $(m, n) \in \Delta(X) \times \Delta(X)$ *with* $m \neq n$ *and* $o = \min(m, n)$*,*

$$D(m,n) = \sum_{j=1}^{|A|}\sum_{k=1}^{|A|} \hat{m}^j \hat{n}^k D(a^j, a^k)\left(1 - \sum_{j=1}^{|A|} o^j\right), \text{ where } (\hat{m}, \hat{n}) = (m - o, n - o)\frac{1}{1 - \sum_{j=1}^{|A|} o^j}.$$

While an excess bilinear switching cost function gives the cost of switching between any pair of mixed strategies, let us note that it is uniquely defined by its behavior on pure strategies. That is, when defining an excess bilinear switching cost function, it suffices to specify the switching costs on pures only.

An excess bilinear switching cost function is simply obtained by calculating the expected switching cost for the vector of normalized excesses. Specifically, for a pair of mixed strategies, (i) the common weight between the two mixed strategies is subtracted from each of them, so excesses are obtained for each, (ii) these excesses are normalized to $\Delta(X)$, and (iii) the renormalized expected switching cost between these excesses is calculated. Renormalization ensures a linear structure such that $D(t, pb + (1-p)t) = pD(t, b)$, that is, the switching cost of moving from $pb + (1-p)t$ to t is the fraction p of the whole cost of switching from b to t. This function is bilinear on the excess strategies, which motivates the name "excess bilinear".

To clarify, when switching from one pure strategy b to another pure strategy t, the incurred switching cost is simply $D(t, b)$. When switching between non-overlapping mixed strategies, say from $\frac{1}{2}b + \frac{1}{2}b'$ to $\frac{1}{2}t + \frac{1}{2}t'$, the excess strategies are the same as the original mixed strategies and therefore the switching cost is bilinear: $D(\frac{1}{2}t+\frac{1}{2}t', \frac{1}{2}b+\frac{1}{2}b') = \frac{1}{4}D(t,b)+\frac{1}{4}D(t',b)+\frac{1}{2}D(t',b)+\frac{1}{4}D(t',b')$. When switching between overlapping mixed strategies, the switching costs are bilinear on the excess strategies, and not on the original ones.

To illustrate the calculation of an excess bilinear switching cost in more detail, we now present three examples. In these examples, mixed strategies are represented by vectors of weights over pure strategies.

Example 1: One Mixed Strategy and One Pure Strategy

Consider a mixed strategy $m = \left[\frac{1}{2}\,,\,\frac{1}{4}\,,\,\frac{1}{4}\right]$ and a pure strategy $n = [0\,,\,1\,,\,0]$. Then, $o = \left[0\,,\,\frac{1}{4}\,,\,0\right]$, $m - o = \left[\frac{1}{2}\,,\,0\,,\,\frac{1}{4}\right]$ and $n - o = \left[0\,,\,\frac{3}{4}\,,\,0\right]$. Therefore, the normalized excesses are $\hat{m} = \frac{4}{3} \cdot \left[\frac{1}{2}\,,\,0\,,\,\frac{1}{4}\right] = \left[\frac{2}{3}\,,\,0\,,\frac{1}{3}\right]$ and $\hat{n} = \frac{4}{3} \cdot \left[0\,,\,\frac{3}{4}\,,\,0\right] = [0\,,\,1\,,\,0] \Rightarrow D(m,n) = (\frac{2}{3}D(a,b) + \frac{1}{3}D(c,b)) \cdot \frac{3}{4}$ $\Rightarrow D(m,n) = \frac{1}{2}D(a,b) + \frac{1}{4}D(c,b)$. This linearity holds in general when one strategy is pure and the other is mixed:

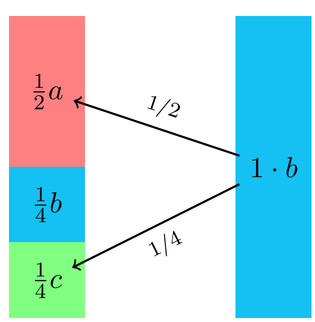

$$D\left(x, \sum_{j=1}^{|A|} n^j a^j\right) = \sum_{j=1}^{|A|} n^j D(x, a^j) \qquad \text{and} \qquad D\left(\sum_{j=1}^{|A|} m^j a^j, y\right) = \sum_{j=1}^{|A|} m^j D(a^j, y).$$

Example 2: Two Mixed Strategies with Non-Overlapping Supports

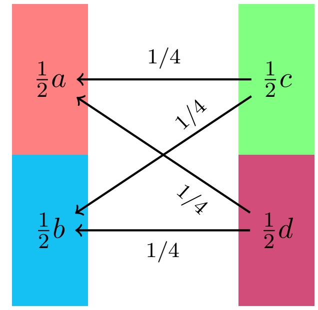

Consider two mixed strategies $m = \left[\frac{1}{2}, \frac{1}{2}, 0, 0\right]$ and $n = \left[0, 0, \frac{1}{2}, \frac{1}{2}\right]$. Then, $o = [0, 0, 0, 0]$ and therefore $\hat{m} = m$ and $\hat{n} = n$. Thus, $D(m,n) = \frac{1}{4}D(a,c) + \frac{1}{4}D(a,d) + \frac{1}{4}D(b,c) + \frac{1}{4}D(b,d)$. Therefore, D is bilinear when the intersection of supports is empty.

Example 3: Two Mixed Strategies with Overlapping Supports

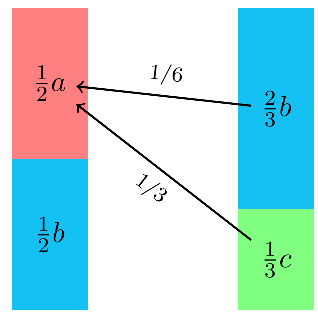

Consider two mixed strategies $m = \left[\frac{1}{2}, \frac{1}{2}, 0\right]$ and $n = \left[0, \frac{2}{3}, \frac{1}{3}\right]$. Then, $o = \left[0, \frac{1}{2}, 0\right] \Rightarrow m - o = \left[\frac{1}{2}, 0, 0\right]$ and $n - o = \left[0, \frac{1}{6}, \frac{1}{3}\right]$. Therefore, the normalized excesses are $\hat{m} = 2 \cdot \left[\frac{1}{2}, 0, 0\right] = [1, 0, 0]$ and $\hat{n} = 2 \cdot \left[0, \frac{1}{6}, \frac{1}{3}\right] = \left[0, \frac{1}{3}, \frac{2}{3}\right]$. So, $D(m,n) = \frac{1}{2}\left(\frac{1}{3}D(a,b) + \frac{2}{3}D(a,c)\right) = \frac{1}{6}D(a,b) + \frac{1}{3}D(a,c)$. In particular, the switching costs are calculated using the probability that needs to be shifted, i.e. $\frac{1}{6}$ to a from b, and $\frac{1}{3}$ to a from c.

Besides the standard understanding, we now present three other interpretations of mixed strategies and excess bilinear switching cost functions.

Interpretation 1: Allocations

To continue the motivating example from Section 1, a mixed strategy refers to the proportion of workers a factory owner assigns to each task, while switching from one mixed strategy to another refers to his reassignment of workers among tasks. The allocation interpretation similarly applies to any problem in which agents divide a continuous amount of resources among different options, such as investors reallocating funds among assets or firms apportioning resources to different product lines. $D(a^j, a^k)$ is the cost that the factory owner incurs when reassigning all workers from task a^k to task a^j. In order to calculate the cost of switching from one mixed strategy to another, the factory owner is assumed to proportionally reassign the excess workers from overstaffed tasks to understaffed tasks. While we presume that switching is proportional, all of our results also carry through for a factory owner who is either talented or untalented at reassigning workers.[2]

Interpretation 2: Population Distributions

Consider a population of identical agents. Agents don't actually randomize, they only play pure strategies. A mixed strategy here refers to the fractions of the population playing each pure strategy. Switching from one mixed strategy to another refers to some (or all) members of the population changing the pure strategies that they play. $D(a^j, a^k)$ is the cost that each individual member of the population incurs when switching from a^k to a^j, and $D(m, n)$ refers to the aggregate cost of switching when the population switches from n to m. In particular, we would like to emphasize that we do not take a representative agent point of view. Instead, this interpretation relies upon *the fundamental lemma* of the paper: a mixed strategy is a best response if and only if all of the pures in its support are best responses. That is, for a mixed strategy to be a best response, it must be that *each* member of the population is individually best responding.

Interpretation 3: State-Dependent Strategies

Consider a world with various payoff-irrelevant states and for each state, the agent chooses a pure strategy to play there. He does not randomize. A mixed strategy refers to the measure of states in which the agent plays that pure strategy. Switching from one mixed strategy to another requires that the agent changes the pure strategy he selects in some (or all) states of the world. He changes his pure strategy in the smallest measure of states possible and the switching costs $D(a^j, a^k)$ he incurs when he changes his pure strategy from a^k to a^j are state-independent. The agent proportionally reassigns his strategy across states such that strategies that are now played less frequently are exchanged for strategies that are now played more frequently. In Figure 1, each state is equally likely and a switch from $\frac{2}{6}a + \frac{4}{6}b$ to $\frac{3}{6}b + \frac{3}{6}c$ is illustrated. In states $1, 2$, and 6, the agent switches his pure strategies, and the switching cost is $\frac{2}{6}D(c, a) + \frac{1}{6}D(c, b)$.

State	Strategy		State	Strategy
1	c		1	a
2	c		2	a
3	b	$\leftarrow$	3	b
4	b		4	b
5	b		5	b
6	c		6	b

Figure 1: Statewise Switching

[2] Suppose that $\tilde{D}_i(m, n) = \sum_j \sum_k w^{jk} D_i(a_i^j, a_i^k)(1 - \|\min(m, n)\|_1)$ where $w^{jk} \geq 0, \sum_j w^{jk} = \hat{n}^k, \sum_k w^{jk} = \hat{m}^j$, and $\hat{m}^j,\ \hat{n}^k$ are as in Definition 2.2. Our proportional weight model has $w^{jk} = \hat{m}^j \hat{n}^k$. A factory owner who has an excellent (or terrible) reassignment ability would choose the weights to reassign workers in the cheapest (or costliest) manner. Importantly, all of our results are robust to any choice of weights.

We now introduce four axioms which will characterize our model. The first two are classical properties introduced by Sen (1971).

Axiom 2.1 *(α) For any $A, B \in \mathcal{X}$ where $B \subseteq A$,*

$$\text{if } m \in C(A) \text{ and } m \in \Delta(B), \text{ then } m \in C(B).$$

Axiom α states that if a mixture is chosen from a set and is choosable from a subset, then it must be chosen from that subset as well.

Axiom 2.2 *(γ) For any $A, B \in \mathcal{X}$,*

$$\text{if } m \in C(A) \cap C(B), \text{ then } m \in C(A \cup B).$$

Axiom γ states that if a strategy is chosen from two sets, then it must also be chosen from their union. Recall that both axioms α and γ are necessary for the classical Weak Axiom of Revealed Preference (WARP), though they are not sufficient (even together).

Axiom 2.3 *(Convexity) For any $A \in \mathcal{X}$ and for any $p \in (0, 1)$,*

$$\text{if } m, m' \in C(A), \text{ then } pm + (1-p)m' \in C(A).$$

Axiom Convexity stipulates that the choice correspondence is convex-valued. This is a standard property of best responses: if two strategies are both best responses, then any mixture of these two strategies is also a best response. Furthermore, it also expresses some risk aversion as the agent weakly prefers to hedge.

Axiom 2.4 *(Support) For any $A \in \mathcal{X}$,*

$$\text{if } m \in C(A) \text{ and } a \in supp(m), \text{ then } a \in C(A).$$

Axiom Support states the following game-theoretic property in a one-player setting: if a mixed strategy is a best response, then so is every pure strategy in its support. It also rules out any strict benefit from randomization, including phenomena like ambiguity aversion, and a preference for randomization (Cerreia-Vioglio et al., 2019).[3]

[3] Recall the standard Independence Axiom: $p \sim q \Leftrightarrow \lambda p + (1-\lambda)r \sim \lambda q + (1-\lambda)r, \ \forall p, q, r$. If an agent's preferences satisfy Independence, then his choices satisfy Axioms Support and Convexity. Hence, Support and Convexity are both weaker than Independence.

Theorem 2.1 *A choice correspondence C satisfies Axioms α, γ, Convexity, and Support if and only if C represents costly switching from endogenous references where costs are excess bilinear.*

If Axioms α and γ are strengthened to WARP, then there is a representation with $D \equiv 0$, that is, the standard model of expected utility. While our model accepts the rational model as a special case, it also allows for agents who do not behave rationally. To see this, suppose that $X = \{x, y, z\}, U(x) = 0, U(y) = 2, U(z) = 4$, and $D(a, a') = 3$ for all $a \neq a'$. Then, $C(X) = \Delta(\{y, z\})$ and $C(\{x, y\}) = \Delta(\{x, y\})$.

Remark 1: The functions U and D in Theorem 2.1 are not necessarily unique, even up to monotonic transformations. Suppose $X = \{t, b\}$ and $C(X) = \Delta(X)$. Then, the external observer can conclude that either $U(t) \geq U(b)$ and $D(t, b) \geq U(t) - U(b)$ holds; or $U(b) > U(t)$ and $D(b, t) \geq U(b) - U(t)$ holds, but cannot tell which. This could be resolved if reference-free choices are also observed, but such decision problems are outside of our model.

Remark 2: A classical argument against mixed strategies is that they are more complicated than pure strategies. Thus, it might be natural that an agent incurs an additional *complexity cost* only when he switches from pure to mixed strategies. As it turns out, the introduction of such complexity costs into our model does not alter our results. The reason is that whenever there is a profitable pure-to-mixed deviation, there is also a profitable pure-to-pure deviation. This argument equally applies to multi-player games and so a surprising byproduct is that adding complexity costs into the standard model does not change the set of NE.

3 Multi-Player Games

3.1 Framework and Results

Consider a simultaneous game setting with $N > 1$ agents. Each agent i has a non-empty finite set of actions (pure strategies) A_i and can play any mixed strategy from $M_i = \Delta(A_i)$. The Cartesian products $A := \prod_{i=1}^{N} A_i$ and $M := \prod_{i=1}^{N} M_i$ denote the spaces of pure and mixed strategy profiles, respectively. The vectors $a = (a_i)_{i=1}^N \in A$ and $m = (m_i)_{i=1}^N \in M$ denote profiles of pure and mixed strategies, respectively. The notations a_{-i} and m_{-i} denote pure and mixed strategy profiles, respectively, excluding agent i's. Each agent i possesses an expected utility function $U_i : M \to \mathbb{R}$. Up to this point, we have described a standard game-theoretic model, which can formally be written as $\Gamma = \langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N \rangle$.

We extend the standard notion of a simultaneous game by supposing that each agent i is endowed with an excess bilinear cost function $D_i : M_i \times M_i \to \mathbb{R}_+$ as in Theorem 2.1. Our main object of study is a *D-game*, which is a tuple $\langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N, (D_i)_{i=1}^N \rangle$. In a D-game, switching costs take place between strategies and this imposes discipline on the model. If switching costs were between outcomes, then the cost incurred by an agent would depend not only on the two strategies he switches between, but also on the entire profile of choices of all other agents. As a result, the agent would bear a switching cost when other agents alter their choices, even though he himself is sticking with the same strategy. Since the number of profiles is much larger than the number of strategies for an individual agent, the switching cost function is more strongly tied down when it operates across strategies rather than across profiles of strategies.

We now define when a strategy is a *switching cost best response* to other agents' strategies. The definition is the same as in the classical theory except that an agent now takes into account the switching cost as well when contemplating deviations from the specified strategy to another one.

Definition 3.3: In a D-game, for agent i, a strategy $m_i \in M_i$ is a *switching cost best response* (SBR) to $m_{-i} \in M_{-i}$ if $U_i(m_i, m_{-i}) \geq U_i(m_i', m_{-i}) - D_i(m_i', m_i)$ holds for all $m_i' \in M_i$. Formally, agent i's SBR correspondence is $SBR_i(m_{-i}) := \{m_i \mid U_i(m_i, m_{-i}) \geq U_i(m_i', m_{-i}) - D_i(m_i', m_i) \ \ \forall m_i' \in M_i\}$.

According to this definition, a strategy m_i is an SBR if, given other agents' strategies m_{-i}, the agent has no incentive to deviate when he plays m_i and views m_i as his reference. The existence of SBRs is trivially guaranteed since any standard best response of the underlying standard game is an SBR. This is because if an agent has no profitable deviation when switching costs are ignored, then there is clearly no profitable deviation when they are taken into account.

The SBR notion satisfies some of the properties that are found in the standard setting. The most important of these can be stated as follows: *A mixed strategy is an SBR if and only if each pure strategy in its support is an SBR* (see the Lemma in Appendix A). According to the population interpretation of mixed strategies, this property states: a mixed strategy is an SBR if and only if each member of the population is playing an SBR. In general, all of the following results hold for any cost function that satisfies this property and that is linear when one strategy is pure and the other is mixed.

We define a *Switching Cost Nash Equilibrium* analogously to the standard Nash Equilibrium (NE).

Definition 3.4: In a D-game, a *Switching Cost Nash Equilibrium* (SNE) is a profile of mixed strategies $m = (m_i)_{i=1}^N \in M$ such that $m_i \in SBR_i(m_{-i})$ for each agent i.

Note that an SNE profile serves two purposes. The standard one is to specify optimal play of all agents, given other agents' behavior. The other is that it serves as the *endogenous reference strategy* from which switching costs are measured. Notice that any NE is an SNE, and therefore the existence of an SNE is trivially guaranteed.

Proposition 3.1 states that the range of possible equilibrium profiles in D-games is a strict superset of those possible in the standard setting. To formally state that, we define

$$\overline{NE} := \{T \mid T \text{ is the set of NE for some standard game}\}$$
$$\overline{SNE} := \{T \mid T \text{ is the set of SNE for some } D\text{-game}\}$$

Proposition 3.1 $\overline{NE} \subsetneq \overline{SNE}$

The weak inclusion in Proposition 3.1 is trivial, while the strict inclusion requires constructing an example whose SNE are not realizable as a set of NE *for any standard game*. Formally, there is a D-game $\langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N, (D_i)_{i=1}^N \rangle$ whose SNE are different than the NE of any other standard game $\langle N, (A'_i)_{i=1}^N, (U'_i)_{i=1}^N \rangle$, where *both* the action sets and utility functions are allowed to vary between the two games.

Example 1: Consider a Matching Pennies D-game where payoffs are as in Table 1 and both players' switching costs are $D_i(H,T) = D_i(T,H) = 0.1$. Figure 2 depicts both players' SBR correspondences (blue for player 1, red for player 2) and the set of SNE (yellow for the intersection).

	H	T
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Table 1: Matching Pennies

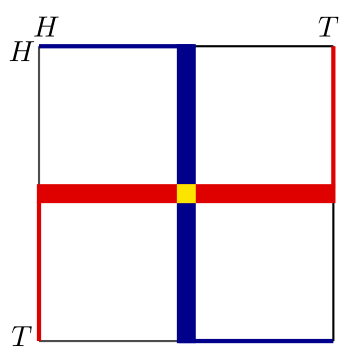

Figure 2: SBR and SNE for Matching Pennies

Example 2: Consider a two-dimensional hide-and-seek game with two players H and S (Rubinstein et al., 1997). Each agent chooses an x-coordinate and a y-coordinate from the set $\{1, 2\} \times \{1, 2\}$. A possible strategy profile is depicted in Figure 3:

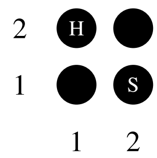

Figure 3: A pure strategy profile where the seeker does not find the hider.

Player H is a hider and receives a payoff 1 if both players choose different locations and 0 otherwise. Player S is a seeker and receives a payoff of 1 if both players choose the same location and 0 otherwise. Formally,

$$U_H([\begin{smallmatrix} x_H \\ y_H \end{smallmatrix}], [\begin{smallmatrix} x_S \\ y_S \end{smallmatrix}]) = \mathbb{1}_{[\begin{smallmatrix} x_H \\ y_H \end{smallmatrix}] \neq [\begin{smallmatrix} x_S \\ y_S \end{smallmatrix}]} \text{ and } U_S([\begin{smallmatrix} x_H \\ y_H \end{smallmatrix}], [\begin{smallmatrix} x_S \\ y_S \end{smallmatrix}])] = \mathbb{1}_{[\begin{smallmatrix} x_H \\ y_H \end{smallmatrix}] = [\begin{smallmatrix} x_S \\ y_S \end{smallmatrix}]}$$

As in the game of Matching Pennies, this game has a unique NE – both players uniformly randomize over all possible locations (with probability ¼), and the seeker finds the hider with probability ¼.

Regarding switching costs, we now suppose that each agent can switch a single coordinate for free, but that it costs at least 1 to switch both coordinates. Without loss of generality, $D_i\left(\begin{bmatrix} x' \\ y' \end{bmatrix}, \begin{bmatrix} x \\ y \end{bmatrix}\right) = \mathbb{1}_{x' \neq x} \mathbb{1}_{y' \neq y}$. This is in the spirit of Arad and Rubinstein (2019) where an agent can only change one dimension. The following proposition characterizes the set of SNE.

Proposition 3.2 *In the above hide-and-seek D-game, besides the standard Nash Equilibrium, there are additional SNE such that:*

1. *The additional SNE consist of any pure profile* $([\begin{smallmatrix} x_H \\ y_H \end{smallmatrix}], [\begin{smallmatrix} x_S \\ y_S \end{smallmatrix}])$ *where* $x_H \neq x_S$ *and* $y_H \neq y_S$.

2. *In any additional SNE, the seeker does not find the hider, so* $U_S = 0$*, and* $U_H = 1$.

As a simple counting exercise, notice that there are 16 pure strategy profiles that the agents may play. In four of them, the seeker finds the hider. In the remaining twelve, the hider hides successfully, and four of those twelve are SNE profiles. In the first panel of Figure 4, we depict the standard Nash Equilibrium. The other three panels depict all of the pure profiles (up to permutation).[4]

[4]An essential feature of this hide-and-seek game and of our model is the heterogeneity in switching costs between different pure

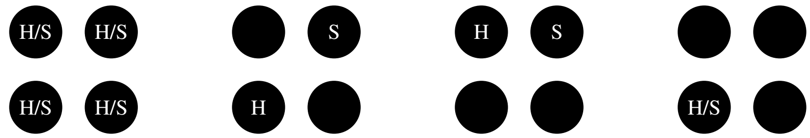

SNE:	✓	✓	✗	✗
NE:	✓	✗	✗	✗
Seeker Finds Hider:	Sometimes	✗	✗	✓

Figure 4: Hide-and-Seek

3.2 A Formal Comparison of SNE and ε-Equilibrium

Notice that the notion of SNE can be equivalently formulated as a profile of mixed strategies $m \in M$ such that for each agent i,

$$U_i(m_i, m_{-i}) \geq U_i(m_i', m_{-i}) - D(m_i', m_i) \quad \forall m_i' \in M_i.$$

This is reminiscent of the ε-equilibrium notion due to Radner (1980). In order to compare that notion with the SNE notion, we need to formally define an ε-game and an ε-best response.

An ε-game is a tuple $\langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N, \varepsilon \rangle$ and an ε-best response is:

$$\varepsilon BR_i(m_{-i}) = \{m_i \mid U_i(m_i, m_{-i}) \geq U_i(m_i', m_{-i}) - \varepsilon, \ \forall m_i' \in M_i\}.$$

An ε-Nash Equilibrium (εNE) is a profile of mixed strategies $m \in M$ such that for each agent i, $m_i \in \varepsilon BR_i(m_{-i})$. This definition of εNE coincides with Radner's (1980) ε-equilibrium definition, which is any profile of mixed strategies $m \in M$ such that for each agent i,

$$U_i(m_i, m_{-i}) \geq U_i(m_i', m_{-i}) - \varepsilon \quad \forall m_i' \in M_i.$$

Thus, εNE can be viewed as an equilibrium with the cost function $D_i(m_i', m_i) = \varepsilon$ for all $m_i',\ m_i \in M_i$. This flat cost function differs fundamentally from ours because *it is not excess bilinear*. To understand differences, first notice that in εNE the switching cost between any pair of distinct pure strategies is a homogenous ε, whereas the cost function in SNE can be heterogenous, that is, costs can differ between any pair of pure strategies. Second, the cost function in SNE scales linearly between a pure strategy and mixtures of itself, i.e. $D(\lambda a + (1-\lambda)a', a') = \lambda D(a, a') = \lambda\varepsilon$, whereas the cost function in εNE remains

strategies. In this particular game, an agent's switching cost depends upon how many dimensions he deviates in. If switching costs were homogenous here, then the set of SNE would be similar to that in Example 1.

ε. That is, the difference between SNE and εNE cost functions is analogous to that of "variable costs" and "fixed costs".

To illustrate the differences between εBR and SBR, let us return to the Matching Pennies game where payoffs are as in Table 1. Figure 5 depicts player 1's εBR correspondences for two different epsilon values[5] and Figure 6 does the same for SBR correspondences.

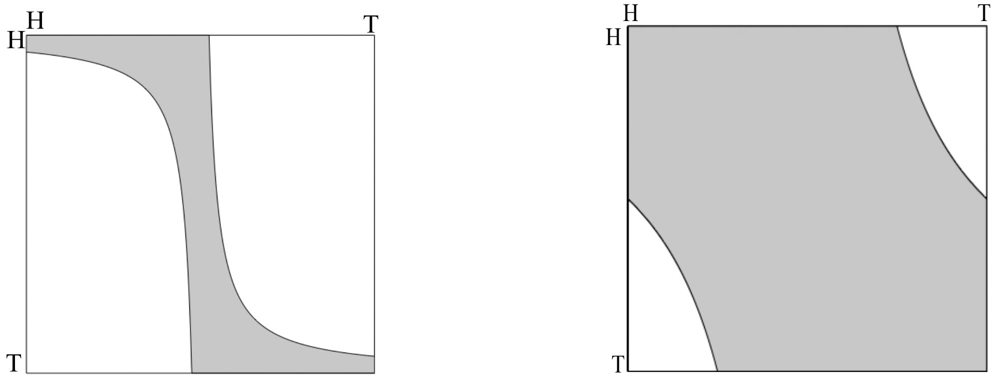

Figure 5: εBR for Player 1 in a Matching Pennies Game with $\varepsilon = 0.1$ (left) and $\varepsilon = 1$ (right)

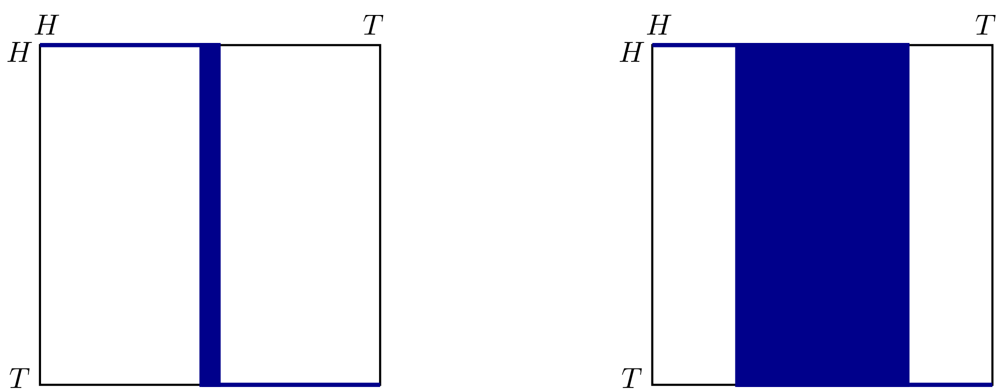

Figure 6: SBR for Player 1 in a Matching Pennies Game with $D = 0.1$ (left) and $D = 1$ (right) between Pures

Clearly, both notions extend the standard best response notion. The next result shows that extensions are completely distinct: *the only correspondences that are both εBR and SBR are standard best responses.*

[5] Algebraic calculations show $U_1(\lambda H + (1-\lambda)T, \gamma H + (1-\gamma)T) = \lambda(4\gamma - 2) - 2\gamma + 1$. Thus, if $\gamma > 1/2$ (so that H is a best response for player 1), then $\lambda H + (1-\lambda)T$ is an ε-best response if and only if $\lambda \geq 1 - \dfrac{\varepsilon}{4\gamma - 2}$. Similarly, if $\gamma < 1/2$, then $\lambda H + (1-\lambda)T$ is an ε-best response if and only if $\lambda \leq \dfrac{\varepsilon}{2 - 4\lambda}$.

To formally state this result, we make the following definitions. In each, the f are best response correspondences and the defined sets consist of all possible best response correspondences for each notion.

$$\overline{BR} := \{f \mid f \text{ is the BR of some player in some standard game}\}$$

$$\overline{SBR} := \{f \mid f \text{ is the SBR of some player in some } D\text{-game}\}$$

$$\overline{\varepsilon BR} = \{f \mid f \text{ is the } \varepsilon\text{BR of some player in some } \varepsilon\text{-game}\}$$

The two best response correspondences depicted in Figure 5 are members of $\overline{\varepsilon BR}$, and not of $\overline{SBR}$ nor $\overline{BR}$. Likewise, the two best response correspondences depicted in Figure 6 are members of $\overline{SBR}$, and not of $\overline{\varepsilon BR}$ nor $\overline{BR}$.

Proposition 3.3 $\overline{BR} = \overline{SBR} \cap \overline{\varepsilon BR}$

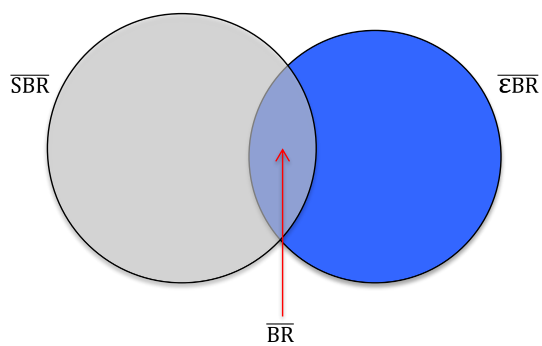

Figure 7: Illustration of Proposition 3.3

Notice that Proposition 3.3 goes beyond simply stating that cost functions cannot be mimicked across the two best response notions. The set $\overline{SBR}$ includes any SBR for *any* D-game and likewise a member of $\overline{\varepsilon BR}$ is any ε-best response for *any* ε-game. Therefore, when a function f is in both $\overline{SBR}$ and $\overline{\varepsilon BR}$, this f may arise from different games with different utilities, different cost functions, and different action sets. However, all of these games must have the same set of other players and other players' strategies because otherwise the common best response correspondence f would have different domains. This proposition shows that even though the D-game and the ε-game may differ in utilities, D-functions, ε-value, and the strategies available to the best responder, any best response correspondence that is both an SBR and an εBR must also be a standard best response correspondence.

We now analyze the following Prisoner's Dilemma game (Table 2). This analysis not only exhibits the difference between the two equilibrium notions (SNE and εNE) more clearly, but also forms the basis for our next proposition, which we view as a fundamental result regarding SNE.

	Coop	Defect
Coop	$5, 5$	$0, 6$
Defect	$6, 0$	$1, 1$

Table 2: Prisoner's Dilemma

First, notice that there is a unique NE, namely $(Defect, Defect)$. Additionally, if $D(a_i, a_i') \leq \varepsilon < 1$ for every pair of actions $a_i, a_i' \in \{Coop, Defect\}$, then $(Defect, Defect)$ is also the unique SNE. However, for any $\varepsilon > 0$, the set of εNE includes not only $(Defect, Defect)$ but also other mixed profiles. To see this, notice that $(\lambda Coop + (1-\lambda)Defect,\ \lambda Coop + (1-\lambda)Defect)$ is an εNE as long as $\lambda \leq \varepsilon$. The difference is because according to SNE, costs incurred to switch between two very similar mixed strategies are very small whereas the cost remains flat according to εNE. Thus, when moving from a standard game to one with ε switching costs, all mixed strategies that were best responses remain so, but so are nearby mixed strategies since the switching cost to the best response cannot be overcome. This need not be the case for the SNE notion because as a mixed strategy gets closer to a standard best response, the switching costs to that standard best response also decrease.

Mailath et al. (2005) develop a notion of ε-equilibrium for extensive form games that extends Radner's (1980) definition for normal form games. Mailath et al. (2005) show that for any extensive form game G, the following are true: (i) all NE of nearby games are ε-equilibria of G, and (ii) all pure ε-equilibria of G are NE of nearby games. Thus, the equivalence between ε-equilibria of a game and NE of nearby games holds for pure strategy profiles, but not for mixed profiles since mixed ε-equilibria may not be NE of nearby games. This is because, for any pure NE and small enough ε, nearby mixed profiles are ε-equilibria, but not NE of nearby games. The above Prisoner's Dilemma game is an example of this: not only does the game have $(Defect, Defect)$ as the unique NE, but also every nearby game has $(Defect, Defect)$ as its unique NE, whereas the set of ε-equilibria is strictly larger than the singleton $\{(Defect, Defect)\}$. On the contrary, since $(Defect, Defect)$ is the unique SNE of the above Prisoner's Dilemma game, there is a chance that the equivalence between SNE of the original game and NE of nearby games holds for both pure and mixed strategy profiles. The next proposition shows that this is indeed the case.

To measure the distance between games, we employ the following notion of Mailath et al. (2005).

Definition 3.6. Given two games G and G' with the same number of players and the same sets of strategies, the distance between them is $\Delta(G, G') = \max_{i,m} |U_i(m) - U'_i(m)|$.

Proposition 3.4 *For any standard game* $G = \langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N \rangle$*, define the D-game* $G_\varepsilon = \langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N, (D_i)_{i=1}^N \rangle$ *where* D *is the excess bilinear cost function such that* $D_i(a_i, a'_i) = \varepsilon$ *for all* i *and all action pairs* $a_i \neq a'_i$*. Then,*

$$SNE(G_\varepsilon) = \bigcup_{\{G' | \Delta(G,G') \leq \varepsilon/2\}} NE(G').$$

This proposition considers a *specific* excess bilinear switching cost function which takes the same ε value across all different pure strategy pairs. Under this specific cost function, an equivalence is obtained between SNE of a game and NE of nearby games for *both pure and mixed strategy profiles*. Hence, Proposition 3.4 is a strengthening of Mailath et al.'s (2005) result. It implies that the SNE notion can serve as a computational tool for calculating the set of NE of all nearby games. Specifically, calculating the NE of all nearby games can be complicated because there are infinitely many nearby games, each of which has a set of NE to be solved for. This proposition provides an alternative approach: one can simply calculate the SNE of the game with ε costs (between pure strategies) which only requires the calculation of equilibria for one game. Therefore, a reader who is not interested in SNE for its own sake may still find SNE to be a useful tool for studying NE of nearby games.

4 An Application to Vertical Differentiation

In a setting with firms and consumers, either party can face switching costs. When consumers face switching costs, such as reward programs, the impact is fairly trivial: consumers are locked in, so firms can charge them higher prices, which benefits firms and harms consumers. What happens when firms face switching costs? In this section, we investigate a standard model of vertical differentiation where firms face switching costs and analyze its welfare/efficiency implications.

Consider two firms, 1 and 2, who choose quality levels $q_1, q_2 \in [\underline{q}, \overline{q}]$ respectively and face a continuum of consumers distributed uniformly on $[\underline{\theta}, \overline{\theta}]$ where $\overline{\theta} > 2\underline{\theta}$.[6] A type θ consumer's utility from consuming a

[6]Without this assumption, the model collapses to a Bertrand competition where firms compete on price and make zero profit.

good of quality q_i is $U(\theta, q_i) = \theta q_i - p_i$, where p_i is the price that firm i charges for the good. We further assume that every consumer buys a good. Both firms face a linear cost of production, c, and their profits are $\Pi_i = (p_i - c)\mu_i$, where μ_i is the measure of agents who shop at firm i. For simplicity, we restrict our attention to pure strategies. In our analysis, without loss of generality, we focus on the case where the second firm produces a good of higher quality, that is, $q_2 \geq q_1$.

Our model is based upon Tirole's (1987) treatment of Shaked and Sutton (1982), who study a two-stage model with firms choosing quality in the first stage and price in the second stage. They demonstrate that there is a unique Subgame Perfect Nash Equilibrium, where the pricing functions are $p_1(q_{1,}q_2) = c + \frac{\overline{\theta} - 2\underline{\theta}}{3}\Delta q$ and $p_2(q_1, q_2) = c + \frac{2\overline{\theta} - \underline{\theta}}{3}\Delta q$ with $\Delta q = q_2 - q_1$. In our model, we focus only on the first stage, namely the competition over quality, and assume that firms use these pricing functions.

In the absence of switching costs, there is a unique NE (up to permutation) and it features maximal differentiation: $q_1 = \underline{q}$ and $q_2 = \overline{q}$. This is because both firms' objectives are aligned and their profits are increasing in the firms' quality difference.

We now introduce linear switching costs into the model so that it is costly for each firm to deviate from its planned quality level, $D_i(\hat{q}_i, q_i) = \lambda|\hat{q}_i - q_i|$, which might be due to expenses that firms face in retooling factories or training. With such switching costs, firms now face a trade-off when they are not maximally differentiated. Firms would like to further differentiate themselves to increase their profits, but doing so would lead them to incur a switching cost.

The following proposition establishes that there are three regions which govern the structure of equilibria, and analyzes their welfare and efficiency implications. We use the phrase "welfare" to refer to overall consumer welfare and "profits" to refer to total firm profits ($\Pi_1 + \Pi_2$). Our efficiency criterion is the standard one: welfare plus profits. In the low switching cost region, the differentiation incentive dominates the switching costs and so the only SNE is the NE. In the intermediate switching cost region, the high-quality firm picks the maximum quality while the low-quality firm picks any quality (not necessarily the lowest one). All additional SNE have lower profits, higher welfare, and higher efficiency than the NE. Finally, in the high switching cost region, switching costs dominate the differentiation incentive, and anything goes: every quality profile is an SNE, including ones which are simultaneously worse for both consumers and firms.

Proposition 4.1 (Vertical Differentiation) *There are thresholds,* $\lambda^l = \dfrac{(\overline{\theta} - 2\underline{\theta})^2}{9(\overline{\theta} - \underline{\theta})}$ *and* $\lambda^h = \dfrac{(2\overline{\theta} - \underline{\theta})^2}{9(\overline{\theta} - \underline{\theta})}$*, such that when the switching cost parameter* λ *is*

Strictly below λ^l*: There is a unique SNE (up to permutation) and it is equal to the NE,* $(\underline{q}, \overline{q})$*.*

Between λ^l *and* λ^h*: The SNE are all profiles* $(q_1, \overline{q})$ *where* $\underline{q} \leq q_1 \leq \overline{q}$*. Furthermore, in comparison to the NE:*
All additional SNE are strictly less profitable, strictly more efficient, and strictly welfare improving.

Strictly above λ^h*: Any quality profile is an SNE. Furthermore, in comparison to the NE:*
All additional SNE are strictly less profitable. Some SNE are strictly more efficient and some are strictly less efficient. Finally, if $\overline{\theta} < (3 + \sqrt{2})\underline{\theta}$*, then all additional SNE are strictly welfare improving, whereas if* $\overline{\theta} > (3 + \sqrt{2})\underline{\theta}$*, then there are also SNE which are strictly welfare worse.*

As discussed earlier, firms' profits are increasing in their quality difference. Without switching costs, this pushes firms to maximally differentiate their qualities. With switching costs, it is possible to support interior choices where the firms are not pushed to the extremes, and Proposition 4.1 deduces which interior outcomes are possible and when. The *key observation* is that while both firms profit from differentiation, they do not profit from it equally as the high-quality firm's profit function has a higher slope than that of the low-quality firm. Thus, there is an intermediate switching cost region where the high-quality firm chooses the highest quality possible, but the low-quality firm need not choose the lowest quality. This implies that all additional SNE in this region have the high-quality firm producing the maximum quality. Relative to the NE benchmark, these additional SNE are more efficient (because the new quality profile Pareto-dominates it), yield less profit for the firms (because the closer the firms' qualities are, the less profits they earn), and therefore, are better for consumers. However, when switching costs are high enough, then even the high-quality firm can be dissuaded from altering its production in favor of higher-quality products, and thus, anything goes. Table 3 summarizes Proposition 4.1.

5 Alternative Models

In this section, we contrast our model to one with fully bilinear switching costs and to Köszegi and Rabin's (2006) personal equilibrium model. We find that for each model class, at least one of the following fails: our Proposition 3.4, existence of an equilibrium, or extension to the NE.

Switching Costs	SNE	Welfare	Profits	Efficiency
Low	Only $(\underline{q}, \overline{q})$	$=$	$=$	$=$
Intermediate	Any $(q_1, \overline{q})$	$\uparrow$	$\downarrow$	$\uparrow$
High	Anything Goes	$\uparrow$ if $\overline{\theta} < (3+\sqrt{2})\underline{\theta}$ $\uparrow,\downarrow$ if $\overline{\theta} > (3+\sqrt{2})\underline{\theta}$	$\downarrow$	$\uparrow$,$\downarrow$

Table 3: Summary of Proposition 4.1

5.1 Fully Bilinear Switching Costs

A fully bilinear switching cost function is $D : M \times M \to \mathbb{R}_+$ such that for any two mixed strategies $m = \sum_j m^j a^j$ and $n = \sum_k n^k a^k$, $D(m, n) := \sum_{j,k} m^j n^k D(a^j, a^k)$. Notice that a fully bilinear cost function differs from our excess bilinear cost function in that the common part of any two mixed strategies is no longer subtracted. Thus, unlike our model, an agent with a fully bilinear cost function may face positive switching costs even when he does not change his strategy. For example, $D\left(\frac{1}{2}t + \frac{1}{2}b, \frac{1}{2}t + \frac{1}{2}b\right) = \frac{1}{4}D(t, b) + \frac{1}{4}D(b, t) > 0$. Now, consider a one-player game setting with $C(A) = \{m \mid U(m) - D(m, m) \geq U(m') - D(m', m), \forall m' \in \Delta(A)\}$ as in our model but suppose that D is fully bilinear. This agent's choice behavior satisfies Axioms α, γ, and Support. However, since he faces a positive switching cost in order to remain at his reference strategy, his choice correspondence C may not satisfy the Convexity Axiom and thus, he may have a dispreference for hedging.[7]

Similarly, in a multi-player game setting with fully bilinear costs, the convex-valuedness of best responses need not hold, which means that this model is quite different from ours. A more important distinction is regarding our fundamental result that demonstrates the equality of SNE of a game to NE of nearby games (Proposition 3.4). This result does not hold when the switching cost function is fully bilinear. To see this, consider the following Battle of the Sexes game shown in Table 4.

Let D be a switching cost function such that $D_i(a_i, a_i') = \varepsilon$, $\forall i, a_i \neq a_i'$ as in Proposition 3.4 and let $\varepsilon = 0.1$. Figure 8 displays the players' best response correspondences and the equilibria when D is excess bilinear (left panel) and when D is fully bilinear (right panel). Furthermore, notice the difference in best

[7] To see this, suppose that $A = \{t, b\}$ and $U(t) = 2, U(b) = 1, D(t, b) = D(b, t) = 2$. Then, $t, b \in C(A)$, but $m = \frac{1}{2}t + \frac{1}{2}b \notin C(A)$. To see why, note that $2 - 1 = U(t) - D(t, m) > U(m) - D(m, m) = \frac{3}{2} - 1$.

	B	O
B	2,1	0,0
O	0,0	1,2

Table 4: Battle of the Sexes

response correspondences: the SBRs (left panel) are thick and the best responses in the fully bilinear case (right panel) have 0 width. Likewise the set of equilibria are quite different: there is a square region of mixed SNE (left panel), but in the fully bilinear case, there is exactly one mixed equilibrium (right panel). Therefore Proposition 3.4 does not hold for fully bilinear switching costs. One can also see that in the case of fully bilinear costs, the best response correspondence (right panel) is not convex-valued as it has a positive slope.

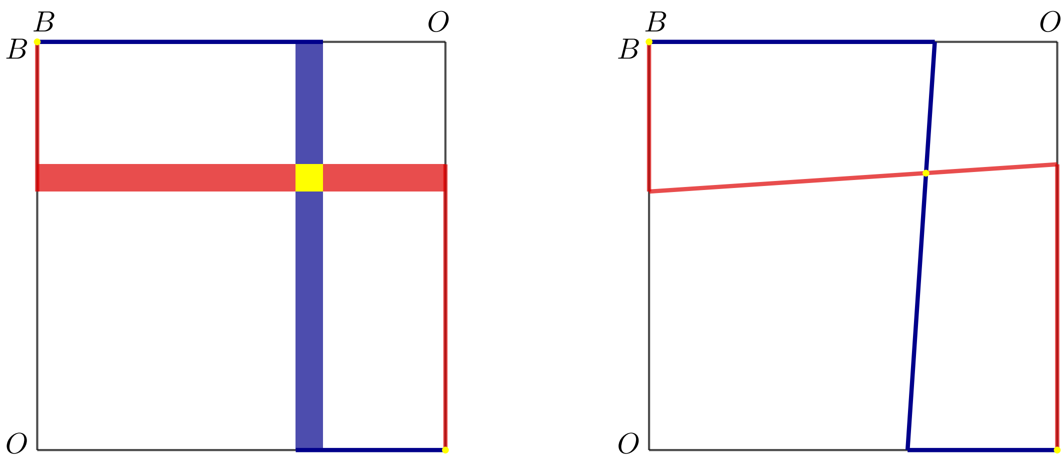

Figure 8: Best Responses and Equilibria under Excess Bilinear Costs (left) and Fully Bilinear Costs (right)

5.2 Personal Equilibria

Based on Köszegi and Rabin (2006, 2007) and Köszegi (2010), Freeman (2017) characterizes a general version of personal equilibrium (and preferred personal equilibrium) in an abstract setting and Freeman (2019) characterizes an expected utility preferred personal equilibrium model of choice under risk. Throughout this section, we focus on the general version of personal equilibrium which we refer to as PE, and the Personal Equilibrium with Linear Loss Aversion (PE-LLA).[8] Table 5 summarizes our findings.

[8]Köszegi and Rabin (2007) also introduce the notion of choice-acclimating personal equilibrium, which is characterized by Masatlioglu and Raymond (2016). The choice-acclimating personal equilibrium model differs from ours in that (i) it reduces

	PE model	PE-LLA model
one-player	Existence Issues Failure of Convexity and Support Axioms	Reduces to Rationality
multi-player	Existence Issues Does not always nest NE	Does not nest NE

Table 5: Summary of Differences between the Switching Cost and Personal Equilibrium Models

Personal Equilibrium (PE)

In the PE model, a decision maker has a reference-dependent utility function $v(m'|m)$ which conveys his utility from choosing the strategy m' when his reference is m. Given v and an available set of mixed strategies S, the set of personal equilibria is:

$$PE_v(S) = \{m \mid v(m|m) \geq v(m'|m),\ \forall m' \in S\}$$

This is a generalization of our model by setting $v(m'|m) = U(m') - D(m', m)$. However, there is a cost to this generalization as there can be existence issues, unlike our model. Freeman (2019) provides an example where a PE does not exist: $S = \{m, m'\}$, $v(m'|m) > v(m, m)$ and $v(m|m') > v(m'|m')$. This example is impossible in our framework. To see why, notice that in our model, the non-negativity of switching costs implies $v(m'|m') \geq v(m'|m)$ because the utility is the same on both sides of the inequality, but on the left side the agent doesn't face switching costs whereas on the right side he does. Likewise, $v(m|m) \geq v(m|m')$. Figure 9 displays these conditions and those of the Freeman example, and shows that they form a cycle. This demonstrates the incompatibility of the Freeman example with our model. Hence, unlike in our paper, there must be a *negative switching cost* (i.e. *boost*) in the Freeman example. That is, for some different $x, y \in \{m, m'\}$, the agent is better off choosing x when his reference is y rather than when his reference is x itself, formally $v(x|y) > v(x|x)$. Just as non-negative switching costs correspond to the status quo bias (Masatlioglu and Ok, 2005, 2014), boosts correspond to a status quo aversion, which in turn can lead to existence issues.

to the rational theory when alternatives are deterministic; and (ii) it satisfies a "Mixture Aversion" property: the mixture of two indifferent non-trivial lotteries is always weakly (and often strictly) less preferred than either of them, whereas agents in our model are "Mixture Indifferent" (choices are convex-valued).

$$\begin{array}{cc|cccc} & & \multicolumn{4}{c}{\text{Reference}} \\ & & m & & m' & \\ \hline \text{Choice} & m & v(m|m) & \geq & v(m|m') & \\ & & \wedge & & \vee & \\ & m' & v(m'|m) & \leq & v(m'|m') & \end{array}$$

Figure 9: The Incompatibility of the Freeman's (2019) Non-Existence Inequalities (red) and the Switching Cost Inequalities (blue)

Köszegi (2010) prove that in the one-player PE model, these existence issues can be resolved by allowing the agent to choose from convex sets. For example, suppose that the agent has two pure strategies t and b, and chooses from the simplex $\Delta(\{t, b\})$. Let $v(t|t) = v(b|b) = 0$ and $v(t|b) = v(b|t) = 1$. That is, the agent has 0 utility from both options, but he likes switching and gets a boost of 1 whenever he does. In this example, there is no personal equilibrium in pure strategies, but by continuity, there is a mixed strategy personal equilibrium. Naturally, this is a violation of our Axiom Support. In our one-player setting with non-negative costs, mixing is not vital for existence, and there is always a pure strategy equilibrium.

This fix does not work for the multi-player PE model. In Appendix B, we provide a 2x2 example where the reference-dependent utility function is continuous and yet, an equilibrium does not exist. This non-existence is due to the fact that the best response correspondences need not be convex-valued.

Personal Equilibrium with Linear Loss Aversion (PE-LLA)

In the PE-LLA model, when choosing an alternative m with reference m', the agent receives (i) expected utility from m, and (ii) gain-loss utility from m, resulting from its comparison to m'. Moreover, choice outcomes can be multi-dimensional, for example, a bundle of consumption goods or a multi-attribute good. Unlike our model, the one-player PE-LLA model with single-dimensional outcomes reduces to rational choice.[9] For multi-player games, the PE-LLA equilibria does not extend NE because it often has different mixed equilibrium. In Appendix B, we show this in a simplified 11-20 money request game. Furthermore, the mixed PE-LLA equilibrium there is unappealing because it converges to the globally welfare-worst outcome as loss aversion increases.

[9] Suppose there is a unique utility-maximizing pure strategy a. It is chosen in the PE-LLA model because deviating to any other strategy (pure or mixed) gives lower utility and negative gain-loss utility. Furthermore, no other b can be chosen because the agent could profitably deviate to a since $v(a|b) > u(a) > u(b) = v(b|b)$. The case with indifferences is similar.

6 Related Literature

Guney and Richter (2018) also model reference effects via costly switching, but from *exogenous* status quos. This is in contrast to the current paper which studies endogenous reference points.[10] An advantage of the endogeneity here is that the current model is portable to standard game-theoretic settings where agents do not naturally have observable references. Furthermore, Guney and Richter (2018) consider pure strategies only, whereas mixed strategies are of fundamental importance for the current paper. To study mixed strategies, the present paper introduces the excess bilinear switching cost function. Applications in both papers also differ in flavor. Guney and Richter (2018) focus on Prisoners' Dilemma games, and derive necessary and sufficient conditions to explain cooperation observed in experiments, using their model as well as other models in the literature. In this paper, we instead perform a general game-theoretic analysis and then explore specific games, including the games of matching pennies, hide-and-seek, 11-20 money request, as well as a vertical product differentiation application.

In the game theory literature, the two closest papers are Shalev (2000) and Arad and Rubinstein (2019). Shalev (2000) considers a simultaneous game setting in which each agent takes his expected utility as a reference. Each agent evaluates outcomes by applying loss aversion to the generated lottery of payoffs relative to his reference utility. Thus, these agents are always rational in deterministic environments, unlike our model. In Arad and Rubinstein (2019), agents play a symmetric simultaneous game with multi-dimensional actions. Each agent only considers uni-dimensional deviations. In the language of switching costs, these agents face zero switching costs between strategies varying in at most one dimension and infinite switching costs otherwise. In contrast, our model allows for costs with intermediate values. A more significant difference is that in their model, each agent's actions are partitioned into *cells* and an equilibrium is a profile of cells such that each agent's cell contains an action from which there is no profitable uni-dimensional deviation, given that the agent believes that his opponents randomize uniformly over their own cells.

There is also a significant literature on repeated games where agents play the same game in each period and pay a cost to switch actions across periods. Prominently, Klemperer (1987,1995), Beggs and Klemperer (1992), Farrell and Shapiro (1989), and Von Weizäcker (1984) study firms' pricing decisions in models where customers face switching costs when moving between firms. Caruana and Einav (2008)

[10] For choice-theoretic works on endogenous references, see Ok et al. (2014) and Guney et al. (2018).

and Libich (2008) consider time-varying switching costs, whereas Lipman and Wang (2000, 2009) establish folk-theorem-like results for certain games. All of the these papers differ from ours in the following ways: (i) they study repeated games and we study simultaneous one-shot games; (ii) the reference point there is the previous period's action, whereas in our model each equilibrium is a reference for itself; (iii) switching costs there are flat across actions as opposed to the varying costs in our model; (iv) they generally consider specific games in contrast to the general class of games considered here; and (v) mixed actions generally do not figure into their analysis whereas they play a fundamental role here.

7 Discussion

We introduced and studied an equilibrium model with switching costs. Even though not framed in this language, the ε-equilibrium models of Radner (1980) and Mailath et al. (2005) can also be viewed as switching cost models where the cost is a flat ε for all switches. Our results are related as well. While Mailath et al. (2005) find the one-way result that all NE of nearby games are ε-equilibria, our model achieves a stronger two-way result: for a particular cost function, the SNE fully characterizes the set of all NE of nearby games. This suggests, in some sense, that our switching cost model may be the ”correct” version of ε-equilibria.

Furthermore, Radner's (1980) classic result for a finitely repeated Cournot game is that ε-equilibria can sustain collusive outcomes which are good for firms and bad for consumers, provided that the end of the game is sufficiently distant. In a standard NE, collusion typically unravels; but if the last period is sufficiently far away, then the discounted benefit of deviating in that far off future is less than ε, and so collusion comprises an ε-equilibrium. Mailath et al. (2005) point out that this use of ε-equilibrium to sustain collusion deserves some healthy skepticism as today's ε is being used to measure the far future's benefit. Their contemporaneous ε-equilibria notion avoids these timing issues by keeping the switching costs in the same time period as when deviations are being made. Moreover, their one-way result on the NE of nearby games implies that substantial collusion is no longer supported in the finitely repeated Cournot game.

Focusing on simultaneous games, we analyzed a switching cost model that is free of the timing issues raised by Mailath et al. (2015) and went beyond skepticism to fully overturning Radner's (1980) classic conclusion: switching costs in our model benefit consumers. Furthermore, our conclusions are robust and unambiguous. In contrast, the ε-equilibrium notion does not actually deliver such a clear message: in a simultaneous Cournot game, ε-equilibrium can support both outcomes where firms produce less (which benefit firms and harm consumers) and outcomes where firms produce more (which harm firms and benefit consumers).

In our view, models where firms are constrained by switching costs fit into the recent literature on Bounded Rationality and Industrial Organization (Spiegler, 2011), except it is now the firms that are bounded. Can firms mitigate their weakness from switching costs? Will they lose because of it or can they even turn it into an advantage? After all, there are now models in both directions. Further research remains to discover the general welfare implications of boundedly rational firms.

References

[1] Ayala Arad and Ariel Rubinstein. The 11-20 Money Request Game: A Level-k Reasoning Study. *American Economic Review*, 102:3561–3573, 2012.

[2] Ayala Arad and Ariel Rubinstein. Multi-Dimensional Reasoning in Games: Framework, Equilibrium, and Applications. *American Economic Journal: Microeconomics*, 11:285–318, 2019.

[3] Alan Beggs and Paul Klemperer. Multi-Period Competition with Switching Costs. *Econometrica*, 60:651–666, 1992.

[4] Guillermo Caruana and Liran Einav. A Theory of Endogenous Commitment. *The Review of Economic Studies*, 75:99–116, 2008.

[5] Simone Cerreia-Vioglio, David Dillenberger, Pietro Ortoleva, and Gil Riella. Deliberately Stochastic. *American Economic Review*, 109:2425–2445, 2019.

[6] Joseph Farrell and Carl Shapiro. Optimal Contracts with Lock-In. *The American Economic Review*, 79:51–68, 1989.

[7] David J Freeman. Preferred Personal Equilibrium and Simple Choices. *Journal of Economic Behavior and Organization*, 143:165–172, 2017.

[8] David J Freeman. Expectations-Based Reference-Dependence and Choice Under Risk. *The Economic Journal*, 129:2424–2458, 2019.

[9] Begum Guney and Michael Richter. Costly Switching from a Status Quo. *Journal of Economic Behavior and Organization*, 156:55–70, 2018.

[10] Paul Klemperer. Markets with Consumer Switching Costs. *The Quarterly Journal of Economics*, 102:375–394, 1987.

[11] Paul Klemperer. Competition When Consumers Have Switching Costs: An Overview with Applications to Industrial Organization, Macroeconomics, and International Trade. *The Review of Economic Studies*, 62:515–539, 1995.

[12] Botond Köszegi. Utility from Anticipation and Personal Equilibrium. *Economic Theory*, 44:415–444, 2010.

[13] Botond Köszegi and M. Rabin. Reference-Dependent Risk Attitudes. *American Economic Review*, 97:1047–1073, 2007.

[14] Botond Köszegi and Matthew Rabin. A Model of Reference-Dependent Preferences. *The Quarterly Journal of Economics*, 121:1133–1165, 2006.

[15] Jan Libich. An Explicit Inflation Target as a Commitment Device. *Journal of Macroeconomics*, 30:43–68, 2008.

[16] Barton L Lipman and Ruqu Wang. Switching Costs in Frequently Repeated Games. *Journal of Economic Theory*, 93:149–190, 2000.

[17] Barton L Lipman and Ruqu Wang. Switching Costs in Infinitely Repeated Games. *Games and Economic Behavior*, 66:292–314, 2009.

[18] George J. Mailath, Andrew Postlewaite, and Larry Samuelson. Contemporaneous Perfect Epsilon-Equilibria. *Games and Economic Behavior*, 53:126–140, 2005.

[19] Yusufcan Masatlioglu and Efe Ok. Rational Choice with Status Quo Bias. *Journal of Economic Theory*, 121:1–29, 2005.

[20] Yusufcan Masatlioglu and Efe Ok. A Canonical Model of Choice with Initial Endowments. *The Review of Economic Studies*, 81:851–883, 2014.

[21] Yusufcan Masatlioglu and Collin Raymond. A Behavioral Analysis of Stochastic Reference Dependence. *The American Economic Review*, 106:2760–2782, 2016.

[22] Efe Ok, Pietro Ortoleva, and Gil Riella. Revealed (P)reference Theory. *The American Economic Review*, 105:299–321, 2014.

[23] Martin J. Osborne. An Introduction to Game Theory. *Oxford University Press*, 2003.

[24] Roy Radner. Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives. *Journal of Economic Theory*, 22:136–154, 1980.

[25] Ariel Rubinstein, Amos Tversky, and Dana Heller. Naive Strategies in Competitive Games. *In: Understanding Strategic Interaction.*, pages 394–402, 1997.

[26] Amartya Sen. Choice Functions and Revealed Preference. *The Review of Economic Studies*, 38:307–317, 1971.

[27] Avner Shaked and John Sutton. Relaxing Price Competition Through Product Differentiation. *The Review of Economic Studies*, 49:3–13, 1982.

[28] Johnathan Shalev. Loss Aversion Equilibrium. *International Journal of Game Theory*, 29:269–287, 2000.

[29] Ran Spiegler. *Bounded Rationality and Industrial Organization*. Oxford University Press, 2011.

[30] Jean Tirole. *The Theory of Industrial Organization*. MIT Press, 1988.

[31] C Christian Von Weizsäcker. The Costs of Substitution. *Econometrica*, 52:1085–1116, 1984.

Appendix A

We first present a fundamental lemma (and its proof) that will be useful for the proof of Theorem 2.1, and throughout. The one-player version states:

"a mixed strategy m is chosen if and only if every pure strategy in its support is chosen".

After proving the Lemma, we present its multi-player version.

Lemma: Given any expected utility function U, excess bilinear function D, and $m \in \Delta(A)$, then
(One-Player)

$$U(m) \geq U(n) - D(n, m) \;\; \forall n \in \Delta(A) \Leftrightarrow$$

$$U(a) \geq U(n) - D(n, a) \; \forall a \in supp(m), \; \forall n \in \Delta(A).$$

Proof of Lemma (One-Player):

For any $m, n \in \Delta(A)$, define $o = \min(m, n)$ and $\|o\|_1 := \sum_{j=1}^{|A|} o^j$. Given any $n \neq m$, define $\hat{n} = \dfrac{n - o}{1 - \|o\|_1}$. Notice that

$$U(n) - U(m) > D(n, m) \Leftrightarrow$$

$$(1 - \|o\|_1)(U(\hat{n}) - U(\hat{m})) > (1 - \|o\|_1)D(\hat{n}, \hat{m}) \Leftrightarrow$$

$$U(\hat{n}) - U(\hat{m}) > D(\hat{n}, \hat{m}) \Leftrightarrow$$

$$\sum_{j=1}^{|A|} \sum_{k=1}^{|A|} \hat{n}^j \hat{m}^k (U(a^j) - U(a^k)) > \sum_{j=1}^{|A|} \sum_{k=1}^{|A|} \hat{n}^j \hat{m}^k D(a^j, a^k)$$

[$\Leftarrow$] If $U(m) \ngeq U(n) - D(n, m)$ for some $n \in \Delta(A)$, then $U(n) - U(m) > D(n, m)$. By the above equivalence, for some $a^j \in supp(n)$ and some $a^k \in supp(m)$, $U(a^j) > U(a^k) - D(a^j, a^k)$. In words, if m is not chosen, then one of it's underlying pure strategies a^k is also not chosen.

[$\Rightarrow$] If $U(a^k) \ngeq U(z) - D(z, a^k)$ for some $z \in \Delta(A)$ and $a^k \in supp(m)$, then $U(z) - U(a^k) > D(z, a^k)$. Since U is linear, and D is linear in the first coordinate when its second coordinate is a pure strategy, $\exists a^j \in supp(z)$ so that $U(a^j) - U(a^k) > D(a^j, a^k)$. Then, $U(\bar{m}) - U(m) > D(\bar{m}, m)$ where $\bar{m} = m - m^k a^k + m^k a^j$, and thus, $U(m) \ngeq U(\bar{m}) - D(\bar{m}, m)$. In words, if some a^k in the support of m is not chosen, then m is also not chosen. $\square$

We now present the multi-player version of the above fundamental lemma. It states:

"a mixed strategy m is a best response if and only if every pure strategy in its support is a best response".

Lemma: (Multi-Player) For any D-game $\langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N, (D_i)_{i=1}^N \rangle$, player i, and mixed strategy m_i,

$$m_i \in BR_i(m_{-i}) \Leftrightarrow a_i \in BR_i(m_{-i}) \;\; \forall a_i \in supp(m_i)$$

Proof of Lemma (Multi-Player): The proof is exactly the same as the one-player version with the following replacements made everywhere: $U(m)$ by $U_i(m_i, m_{-i})$, and $D(n, m)$ by $D_i(n_i, m_i)$. □

Proof of Theorem 2.1

[$\Rightarrow$] Assume that C satisfies Axioms α, γ, Convexity, and Support. First, we define a binary relation R on X: namely, for any a, b, define aRb if $b \notin C(ab)$[11]. Intuitively, "a rules out b". Notice that, by Axiom α, if $a \neq b$ and aRb, then b can never be chosen from any set of actions that also contains a.

We now show that R is anti-symmetric and acyclic. To show that R is anti-symmetric, suppose aRb. If a strictly mixed strategy $m \in \Delta(ab)$ is in $C(ab)$, then by Axiom Support $b \in C(ab)$ which violates aRb. Therefore, $C(ab) = \{a\}$ and so $b \not R\, a$. To show that R is acyclic, suppose $a_1 R a_2 R \ldots R a_n R a_1$. Take $m \in C(a_1 \ldots a_n)$ and $a_i \in supp(m)$. Then, by Axiom Support, $a_i \in C(a_1 \ldots a_n)$ and by Axiom α, $a_i \in C(a_{i-1} a_i)$. Therefore, $a_{i-1} \not R\, a_i$.

Since R is acyclic and anti-symmetric, there exists a $U : X \to [0,1]$ that represents R in the sense that for any pair $a \neq b$, $aRb \Rightarrow U(a) > U(b)$. Next, extend U through expected utility to $U : \Delta(X) \to [0,1]$. Now, for every a, b, define $D(a,b) = 0$ if aRb or $a = b$, and $D(a,b) = 1$ otherwise. Next, extend D to $\Delta(X) \times \Delta(X)$ so that for any m, $D(m,m) = 0$, and for any pair $m \neq n$, let $o = \min(m,n)$ and define:

$$D(m,n) = \sum_{j=1}^{|A|} \sum_{k=1}^{|A|} \hat{m}^j \hat{n}^k D(a^j, a^k) \left(1 - \sum_{j=1}^{|A|} o^j\right), \text{ where } (\hat{m}, \hat{n}) = (m - o, n - o) \frac{1}{1 - \sum_{j=1}^{|A|} o^j}.$$

It finally remains to be shown that for the above U and D:

$$C(A) = \{m \in \Delta(A) |\;\; U(m) \geq U(m') - D(m', m) \;\; \forall m' \in \Delta(A)\}.$$

[11] For simplification, we drop the curly brackets and write $C(ab)$ instead of $C(\{a,b\})$.

[$\subseteq$] Take $z \in C(A)$. By Axiom Support, for every $b \in supp(z)$, $b \in C(A)$. By Axiom α, for all $a \in A$, $b \in C(ab)$ and so $a \not R b$. Therefore, $D(a,b) = 1$ for all $a \in A$. This implies $U(b) \geq U(a) - D(a,b)$ for all $a \in A$ because $1 \geq U(b), U(a) \geq 0$. Then, by linearity, $U(b) \geq U(m') - D(m', b) \ \ \forall m' \in \Delta(A)$ and recall that this holds for any $b \in supp(z)$. By the Lemma, $U(z) \geq U(m') - D(m', z) \ \ \forall m' \in \Delta(A)$.

[$\supseteq$] Take $z \in \Delta(A)$ such that $U(z) \geq U(m') - D(m', z) \ \ \forall m' \in \Delta(A)$. By the Lemma, for any $b \in supp(z)$, $U(b) \geq U(m') - D(m', b) \ \ \forall m' \in \Delta(A)$. Therefore, in particular, for any $a \in A$, $U(b) \geq U(a) - D(a,b)$. If aRb, then by definition $U(a) > U(b)$ and $D(a,b) = 0$, a contradiction. Thus, it must be that $a \not R b$. Therefore, $b \in C(ab)$ for all $a \in A$. By repeated applications of Axiom γ, one obtains $b \in C(A)$ for all $b \in supp(z)$. By Axiom Convexity, $z \in C(A)$.

[$\Leftarrow$] We now assume there exist U and D with the properties outlined in the statement of the theorem and show that axioms are satisfied.

To show that Axiom α holds, take any $m \in C(A)$. Then, $U(m) \geq U(m') - D(m', m)$ for all $m' \in \Delta(A)$ and in particular for any $m' \in \Delta(B)$ such that $B \subseteq A$. Hence, $m \in C(B)$ as well.

To show that Axiom γ holds, take $m \in C(A) \cap C(B)$. Then, by definition, $U(m) \geq U(m') - D(m', m)$ for all $m' \in \Delta(A) \cup \Delta(B)$. Now, suppose $U(m) < U(n) - D(n, m)$ for some $n = \sum_{i,j} \alpha^i a^i + \beta^j b^j$ with $a^i \in A$ and $b^j \in B$. Then, there must exist $a^k \in supp(m)$ such that $U(a^k) < U(n) - D(n, a^k)$. Then, $U(a^k) < U(n) - D(n, a^k) = \sum_i \alpha_i [U(a^i) - D(a^i, a^k)] + \sum_j \beta_j [U(b^j) - D(b^j, a^k)] < U(x) - D(x, a^k)$ for some $x \in \{a^i, b^j\}_{i,j}$, where the equality follows from the linearity of D when one strategy is pure and the other is mixed. Recall that $m \in C(A) \cap C(B)$ implies $a^k \in C(A) \cap C(B)$ and the last inequality above is a contradiction to either $a^k \in C(A)$ or $a^k \in C(B)$.

To show that Axiom Convexity holds, take any $m, m' \in C(A)$. By definition, $U(m) \geq U(n) - D(n, m)$ and $U(m') \geq U(n) - D(n, m')$ for all $n \in \Delta(A)$. Suppose there exists $p \in (0,1)$ such that $U(pm + (1-p)m') < U(n) - D(n, pm + (1 - pm'))$ for some $n \in \Delta(A)$. Then, there must exist $a^k \in supp(m) \cup supp(m')$ such that $U(a^k) < U(n) - D(n, a^k)$. But this is a contradiction to either $m \in C(A)$ or $m' \in C(A)$.

For Axiom Support, it is enough to recall the Lemma we proved earlier. Axiom Support is the same as one direction of the if and only if statement in the Lemma. $\square$

Proof of Proposition 3.1

The weak inclusion $\overline{NE} \subseteq \overline{SNE}$ is trivial. For any standard game G, there is a corresponding D-game G' where for every player i and for every $m_i, m_i' \in M_i$, $D_i(m_i, m_i') = 0$. Then, $BR_i^G = SBR_i^{G'}$ for all i and consequently $NE(G) = SNE(G')$.

To demonstrate the strict inclusion, we now provide an example. Define G' to be $A_1 = \{t, b\}$, $A_2 = \{l, r\}$, $D_1(t, b) = 2$, $D_1(b, t) = 0$, and $D_2(l, r) = D_2(r, l) = 0$. Payoffs are as in Table 6.

	l	r
t	$1, 0$	$0, 0$
b	$0, 0$	$1, 0$

Table 6: Payoffs for Strict Inclusion Proof

Notice that $D_1(\lambda t + (1-\lambda)b, \omega t + (1-\omega)b) = \begin{cases} 2(\omega - \lambda) & \text{if } \omega > \lambda \\ 0 & \text{otherwise} \end{cases}$

In other words, a strategy that deviates by placing more probability on b is never profitable because the potential benefit is capped at $\omega - \lambda$ and the switching cost always outweighs this benefit. Thus, for any given mixed strategy, the only deviations that may be profitable increase the probability placed on t. Thus, a mixed strategy $\lambda t + (1-\lambda)b$ is an SBR to $\gamma l + (1-\gamma)r$ if and only if

$$U_1(\lambda t + (1-\lambda)b, \gamma l + (1-\gamma)r) \geq U_1(t, \gamma l + (1-\gamma)r)$$

$$\lambda\gamma + (1-\lambda)(1-\gamma) \geq \gamma$$

$$(1-\lambda)(1-\gamma) \geq \gamma(1-\lambda)$$

$$\gamma \leq 1/2 \text{ or } \lambda = 1$$

Figure 10 depicts the SBR function of player 1. Notice also that every mixed strategy of player 2 is a best response to any mixed strategy of player 1. Therefore, Figure 10 also depicts $SNE(G')$ of this D-game.

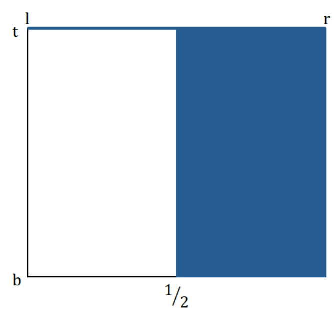

Figure 10: SBR Correspondence depicted in blue.

If there is a game G such that $NE(G) = SNE(G')$, then since $(t, r), (b, r) \in SNE(G')$, it must be that $U_1(t, r) = U_1(b, r)$. Since $\left(t, \frac{1}{2}l + \frac{1}{2}r\right), \left(b, \frac{1}{2}l + \frac{1}{2}r\right) \in SNE(G')$, it must be that $U_1\left(t, \frac{1}{2}l + \frac{1}{2}r\right) = U_1\left(b, \frac{1}{2}l + \frac{1}{2}r\right)$. By expected utility, $U_1(t, l) = U_1(b, l)$. Similarly, because $\left(t, \frac{1}{2}l + \frac{1}{2}r\right), (t, r) \in SNE(G')$, it must be that $U_2(t, l) = U_2(t, r)$ and because $\left(b, \frac{1}{2}l + \frac{1}{2}r\right), (b, r) \in SNE(G')$, it follows $U_2(b, l) = U_2(b, r)$. Notice that if each player has no further actions, then $NE(G) = M \times M \neq SNE(G')$. However, it is possible that there are further actions and then the payoff matrix must look like Table 7.

	l	r	m	...
t	w, x	y, x		
b	w, z	y, z		
n				
$\vdots$				

Table 7: Payoff Matrix with Potential Deviations

Since $(b, l) \notin SNE(G') = NE(G)$ it must be the case that either player 1 or 2 must have a strictly profitable deviation from (b, l). If it is player 1, then $U_1(n, l) > w$, but then $(t, l) \notin NE(G)$, contradicting $NE(G) = SNE(G')$ because $(t, l) \in SNE(G')$. If it is player 2, then $U_2(b, m) > z$, but then $(b, r) \notin NE(G)$, again contradicting $NE(G) = SNE(G')$ because $(b, r) \in SNE(G')$. Since we reach a contradiction in either case, it cannot be that $NE(G) = SNE(G')$ for any standard game G. $\square$

Proof of Proposition 3.2:

Trivially, the standard mixed NE is an SNE. We now show that there is no other SNE where the seeker finds the hider. Let (m_H, m_S) be an SNE where both agents play a common location, say $\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]$, with non-zero probability, that is, $m_H(\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]), m_S(\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]) > 0$. Then, by our fundamental Lemma, $\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]$ is an SBR for the hider, and so it must be that $m_S(\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]) \leq m_S(\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right])$. Thus, $\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right], \left[\begin{smallmatrix}1\\2\end{smallmatrix}\right]$ are SBRs for the seeker, and since the seeker can costlessly switch between them, it be that $m_H(\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]) = m_H(\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right])$. But, then $\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right]$ is an SBR for the hider and so $m_S(\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right]) = m_S(\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right])$. Iterating this argument across all strategies gives $m_S(\left[\begin{smallmatrix}1\\1\end{smallmatrix}\right]) = m_S(\left[\begin{smallmatrix}1\\2\end{smallmatrix}\right]) = m_S(\left[\begin{smallmatrix}2\\1\end{smallmatrix}\right])$. So for this to be an SNE, it is necessary that the seeker uniformly randomizes over all locations and likewise for the hider, which is simply the standard NE.

Any other SNE must have the seeker not finding the hider. This is the first-best for the hider, so it only needs to be checked when the seeker does not have a profitable deviation. The seeker has a profitable deviation if and only if the seeker and hider play adjacent strategies with non-zero probability because then the seeker could costlessly deviate and find the hider with non-zero probability.[12] Since there are only two choices in each dimension, the seeker has no profitable deviations if and only if the seeker and hider play pure strategies which differ in both dimensions. □

Proof of Proposition 3.3

[$\subseteq$] Trivially, $\overline{BR} \subsetneq \overline{SBR}$ and $\overline{BR} \subseteq \overline{\varepsilon BR}$ hold.

[$\supseteq$] Take $f \in \overline{SBR} \cap \overline{\varepsilon BR}$. Then, there is a corresponding game $\langle N, (A_i)_{i=1}^N, (U_i)_{i=1}^N \rangle$ and $\varepsilon \geq 0$ such that f is player i's ε-best response. If $\varepsilon = 0$ in the corresponding game, then f is a standard best response. Therefore, suppose that $\varepsilon > 0$.

By expected utility, for any profile of other agents' strategies m_{-i}, there is a pure strategy a_i such that $U_i(a_i, m_{-i}) \geq U_i(m_i', m_{-i})\ \forall m_i' \in M_i$. Now, take a sequence of strategies $\{m_i^n\}$ which assign positive probability to every pure strategy in A_i such that $m_i^n \to a_i$. By the continuity of U_i, there is a mixed strategy in this sequence so that $U_i(m_i^n, m_{-i}) \geq U_i(a_i, m_{-i}) - \varepsilon$ and thus $m_i^n \in f(m_{-i})$. Therefore, for any profile m_{-i}, $f(m_{-i})$ contains at least one mixed strategy that assigns positive probability to every pure strategy in A_i. Since f is also an SBR for some D-game, the Lemma applies and stipulates that all actions in the support of that strategy are best responses. Thus, all actions in A_i (the pure strategies from the ε-game) are best responses. Therefore, $\forall i, a_i \in A_i, m_{-i}, a_i \in f(m_{-i})$. But, by again applying the Lemma, one obtains that $\forall i$, $f(m_{-i}) = M_i$, a standard best response.[13] □

Proof of Proposition 3.4

[$\subseteq$] Take $m \in SNE(G_\varepsilon)$. Then, by the Lemma, $\forall a_i \in supp(m_i)$ and $\forall a_i' \neq a_i$, $U_i(a_i, m_{-i}) \geq U_i(a_i', m_{-i}) - D(a_i', a_i) = U_i(a_i', m_{-i}) - \varepsilon$. This implies that the utilities of all pure strategies in the support of m are within ε of each other. Let $v_i = \min_{a_i \in supp(m_i)} U_i(a_i, m_{-i}) + \varepsilon/2$. For any action profile a, define

[12] Two locations are adjacent if they coincide in one dimension and differ in the other.

[13] Notice that in the D-game for which f is an SBR, player i may have access to more pure strategies than the ε-game that we analyze, but this is irrelevant for the above analysis.

$$U_i'(a) = \begin{cases} U_i(a) + v_i - U_i(a_i, m_{-i}) & \text{if } a \in supp(m) \\ U_i(a) - \dfrac{\varepsilon}{2} & \text{otherwise} \end{cases}.$$

To show that $m \in NE(G')$, it remains to show for any action a_i, $U_i'(m_i, m_{-i}) \geq U_i'(a_i, m_{-i})$. First, notice that for any $a_i \in supp(m_i)$, $U_i'(a_i, m_{-i}) = U_i(a_i, m_{-i}) + v_i - U_i(a_i, m_{-i}) = v_i$. Thus, by the linear expected utility of the standard setting, $U_i'(m_i, m_{-i}) = v_i$. Therefore, there is no profitable deviation to any action $a_i \in supp(m_i)$. Secondly, for any $\hat{a}_i \notin supp(m_i)$, $U_i'(\hat{a}_i, m_{-i}) = U_i(\hat{a}_i, m_{-i}) - \frac{\varepsilon}{2}$. Now, recall that $m_i \in SNE(G_\varepsilon)$ implies (by the Lemma) that $\min\limits_{a_i \in supp(m_i)} U_i(a_i, m_{-i}) \geq U_i(\hat{a}_i, m_{-i}) - \varepsilon$. Thus, $U_i'(m_i, m_{-i}) = v_i = \min\limits_{a_i \in supp(m_i)} U_i(a_i, m_{-i}) + \frac{\varepsilon}{2} \geq U_i(\hat{a}_i, m_{-i}) - \varepsilon + \frac{\varepsilon}{2} = U_i(\hat{a}_i, m_{-i}) - \frac{\varepsilon}{2} = U'(\hat{a}_i, m_{-i})$. Therefore, there is no profitable deviation to any action $\hat{a}_i \notin supp(m_i)$. Finally, recall a well-known game theory result that if there is no profitable deviation to any pure strategy, then there are no profitable deviations to mixed strategies as well (Proposition 116.2, Osborne (2003)).

[$\supseteq$] Take a game G' such that $\Delta(G, G') \leq \varepsilon/2$ and $m' \in NE(G')$ and take an action $a_i' \in supp(m_i')$. Then in G', $a_i' \in BR'(m_{-i}')$. For any other pure strategy a_i, it must be that $U_i'(a_i, m_{-i}') \leq U_i'(a_i', m_{-i}') \Rightarrow U_i(a_i, m_{-i}') - \varepsilon \leq U_i(a_i', m_{-i}')$. Therefore, $a_i' \in SBR(m_{-i}')$. By the Lemma, $m_i' \in SBR(m_{-i}')$ and applying this argument for each agent i implies $m' \in SNE(G)$. $\square$

Proof of Proposition 4.1:

Given the equilibrium pricing functions, each firm's profit can be computed as

$$\Pi_1(q_1, q_2) = \left(\frac{\overline{\theta} - 2\underline{\theta}}{3}\right)^2 \frac{\Delta q}{\overline{\theta} - \underline{\theta}} \text{ and } \Pi_2(q_1, q_2) = \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{\Delta q}{\overline{\theta} - \underline{\theta}}$$

While there are a continuum of deviations that each firm could make, there are only a few that could be most profitable. If the most profitable deviation is not desirable, then there are no profitable deviations at all. In particular, since switching costs are linear and the profit functions above are linear in the quality difference, the agent's benefit of deviating is piecewise linear with kinks at $\underline{q}, q_1, q_2$ and $\overline{q}$. Furthermore, deviating from q_1 to q_2 (or vice-versa) is never worthwhile, because firms make 0 profit if they produce the same quality. Thus, $\underline{q}$ and $\overline{q}$ are the only points that need to be checked for profitable deviations.

The four best possible deviations are depicted in Figure 11. As it turns out, deviation $1r$ is less profitable than deviation $2r$ while deviation $2l$ is less profitable than deviation $1l$. We now show this for the $1r \leftrightarrow 2r$ case.

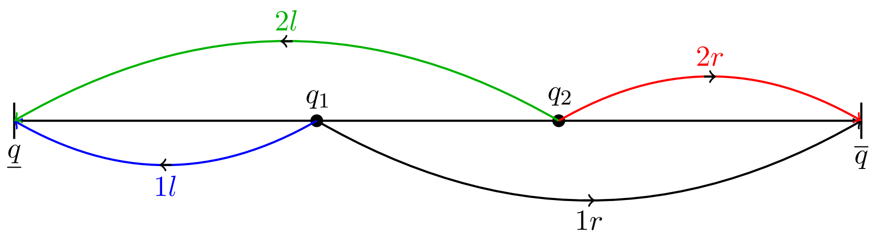

Figure 11: Best Possible Deviations

$$
\begin{aligned}
(1r) \qquad \Pi_1(\overline{q}, q_2) - \Pi_1(q_1, q_2) &= \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{\overline{q} - q_2}{\overline{\theta} - \underline{\theta}} - \left(\frac{\overline{\theta} - 2\underline{\theta}}{3}\right)^2 \frac{q_2 - q_1}{\overline{\theta} - \underline{\theta}} \\
&< \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{\overline{q} - q_2}{\overline{\theta} - \underline{\theta}} \\
&= \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{\overline{q}}{\overline{\theta} - \underline{\theta}} - \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{q_2}{\overline{\theta} - \underline{\theta}} \\
&= \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{\overline{q} - q_1}{\overline{\theta} - \underline{\theta}} - \left(\frac{2\overline{\theta} - \underline{\theta}}{3}\right)^2 \frac{q_2 - q_1}{\overline{\theta} - \underline{\theta}} \\
&= \Pi_2(q_1, \overline{q}) - \Pi_2(q_1, q_2) \qquad (2r)
\end{aligned}
$$

Thus, a deviation by firm 1 to $\overline{q}$ has a lower benefit than firm 2 switching to $\overline{q}$. Furthermore, to deviate to $\overline{q}$, firm 1 has to move further and thus faces a higher switching cost than firm 2 would. Therefore, the deviation $1r$ is always strictly less profitable than the deviation $2r$. Similarly, one can show that the deviation $2l$ is always strictly less profitable than the deviation $1l$. Therefore, it only needs to be checked when the most profitable deviations $1l$ and $2r$ are not profitable.

The deviation $1l$ is unprofitable when

$$
\begin{aligned}
\Pi_1(q_1, q_2) &\geq \Pi_1(\underline{q}, q_2) - c(\underline{q}, q_1) \qquad (1l) \\
c(\underline{q}, q_1) &\geq \Pi_1(\underline{q}, q_2) - \Pi_1(q_1, q_2) \\
\lambda(q_1 - \underline{q}) &\geq \frac{\left(\overline{\theta} - 2\underline{\theta}\right)^2 (q_1 - \underline{q})}{9(\overline{\theta} - \underline{\theta})}
\end{aligned}
$$

Thus, in equilibrium, we have

$$\lambda \geq \frac{\left(\overline{\theta} - 2\underline{\theta}\right)^2}{9(\overline{\theta} - \underline{\theta})} = \lambda^l \text{ or } q_1 = \underline{q} \tag{1l}$$

Similarly, for the deviation $2r$, we have that in equilibrium

$$\lambda \geq \frac{\left(2\overline{\theta} - \underline{\theta}\right)^2}{9(\overline{\theta} - \underline{\theta})} = \lambda^h \text{ or } q_2 = \overline{q} \tag{2r}$$

The temptingness of deviations and consequently, the set of SNE is surmised in Figure 12.

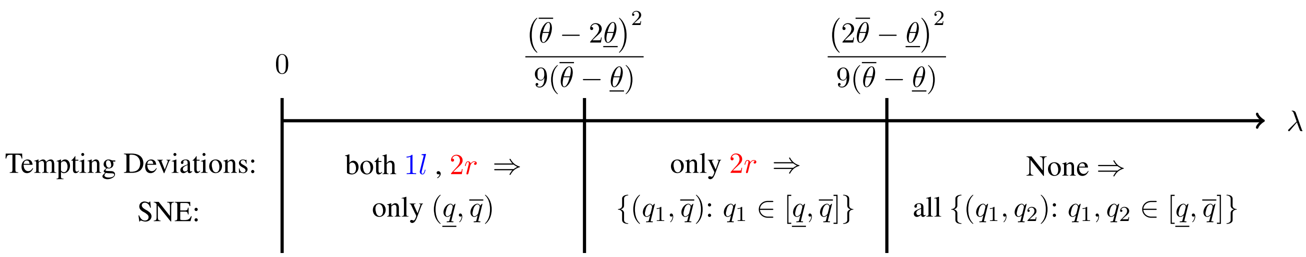

Figure 12: SNE Equilibrium Regions for Vertical Competition

To understand the furthermore statements, note that in any SNE (including the standard NE), $1/3$ of the consumers buy the lower quality good and $2/3$ of the consumers buy the higher quality good. Therefore,

$$\begin{aligned}
\text{Efficiency} = E(q_1, q_2) &= \int_{\underline{\theta}}^{\frac{\overline{\theta}+2\underline{\theta}}{3}} \theta q_1 \frac{d\theta}{\overline{\theta} - \underline{\theta}} + \int_{\frac{\overline{\theta}+2\underline{\theta}}{3}}^{\overline{\theta}} \theta q_2 \frac{d\theta}{\overline{\theta} - \underline{\theta}} - c \\
&= \frac{\overline{\theta} + \underline{\theta}}{2} q_1 + \frac{4\overline{\theta}^2 - 2\overline{\theta}\underline{\theta} - 2\underline{\theta}^2}{9(\overline{\theta} - \underline{\theta})} \Delta q - c \\
\text{Total Profit} = \Pi(q_1, q_2) &= \left(\frac{5\overline{\theta}^2 - 8\overline{\theta}\underline{\theta} + 5\underline{\theta}^2}{9} \right) \frac{\Delta q}{\overline{\theta} - \underline{\theta}}
\end{aligned}$$

Welfare = Efficiency - Total Profit

Intermediate switching cost region ($\lambda^l \leq \lambda < \lambda^h$): As $q_1 \in [\underline{q}, \overline{q}]$ and $q_2 = \overline{q}$, if follows that efficiency is strictly higher in the additional SNE than in the NE. Since Δq is smaller, overall profits are strictly lower. Finally, since welfare is equal to efficiency minus profits, it must be that consumer welfare is strictly higher.

High switching cost region ($\lambda^h \leq \lambda$): Note that efficiency, welfare, and total profit are linear in q_1, q_2. Therefore, only the extreme SNE need to be checked: $(\overline{q}, \overline{q})$ (checked above) and $(\underline{q}, \underline{q})$ (checked now). When $q_1 = q_2 = \underline{q}$, efficiency is minimized. Since $\Delta q = 0$, overall profits are minimized.

Finally, welfare is improved iff

$$0 \leq \frac{\partial W}{\partial q_2} = \frac{\partial E}{\partial q_2} - \frac{\partial \Pi}{\partial q_2} = \frac{4\overline{\theta}^2 - 2\overline{\theta}\underline{\theta} - 2\underline{\theta}^2}{9(\overline{\theta} - \underline{\theta})} - \left(\frac{5\overline{\theta}^2 - 8\overline{\theta}\underline{\theta} + 5\underline{\theta}^2}{9(\overline{\theta} - \underline{\theta})}\right) = \frac{-\overline{\theta}^2 + 6\overline{\theta}\underline{\theta} - 7\underline{\theta}^2}{9(\overline{\theta} - \underline{\theta})}$$

Focusing on the numerator, dividing through by $\underline{\theta}^2$ and letting $r = \overline{\theta}/\underline{\theta}$, we get $0 \leq -r^2 + 6r - 7$. Since we assumed that $r > 2$, the welfare of the extreme SNE strictly increases if $r < 3+\sqrt{2}$, that is, if $\overline{\theta} < (3+\sqrt{2})\underline{\theta}$, and the welfare of the extreme SNE strictly decreases if $r > 3 + \sqrt{2}$, that is, if $\overline{\theta} > (3 + \sqrt{2})\underline{\theta}$. □

Appendix B

Example 1: Non-Existence of an Equilibrium in the Multi-Player PE Model

In the multi-player setting, the most natural extension of the PE model is as follows:

$$BR_i^{PE}(m_{-i}) = \{m_i \mid V_i(m_i, m_{-i}|m_i) \geq V_i(m_i', m_{-i}|m_i), \forall m_i' \in M_i\}$$

An equilibrium is a profile where each agent is best responding. Consider a matching pennies game where player 2 is rational and payoffs are as in Table 1. Denote player 2's strategy as $m_2 = q \cdot H + (1-q) \cdot T$. Define V_1 as below, which implies the best response correspondence BR_1^{PE}.

$$V_1(m_1', m_2|m_1) = \begin{cases} 1 & \text{if } m_1' = H \text{ and } q \geq 1/2 \\ 1 & \text{if } m_1' = T \text{ and } q \leq 1/2 \\ < 1 & \text{otherwise} \end{cases} \qquad BR_1^{PE}(m_2) = \begin{cases} H & \text{if } q > 1/2 \\ \{H, T\} & \text{if } q = 1/2 \\ T & \text{if } q < 1/2 \end{cases}$$

Figure 13 depicts both players' best response correspondences. Since they do not intersect, there is no personal equilibrium. This is in contrast to our model where the existence of an equilibrium is guaranteed by the convexity of best response correspondences (which does not hold in the PE model).

Example 2: PE-LLA Equilibrium in a Simplified 11-20 Money Request Game

The extension of the PE-LLA model to multi-player games could potentially be done in various ways, but one natural method would be to treat other players' strategies as different dimensions. Specifically, an agent takes the other agents' choices as given, and receives (i) expected utility from (m_1, m_{-1}), and (ii) gain-loss utility from the new strategy profile (m_1, m_{-1}) when it is compared to the reference strategy profile

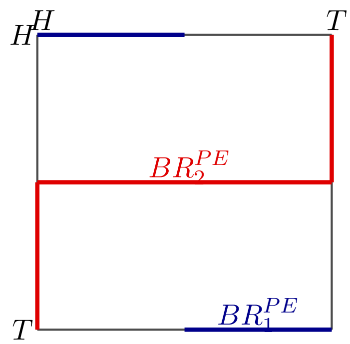

Figure 13: Non-Existence of an Equilibrium in the PE Model.

(m'_1, m_{-1}). In this model, references are utility-based and the gain-loss utility calculation uses not just the agent's own strategy and reference strategy, but also the strategy of his opponents. This is a *key* difference of that model from ours, because switching costs in our model are between strategies and depend only on m_1 and m'_1. The two models are formally written in Table 8.

Switching Costs	$U_1((m'_1, m_{-1})\|m_1) = \mathbb{E}[U_1(m'_1, m_{-1})] - D(m_1, m'_1)$
Personal Equilibrium with	$U_1((m'_1, m_{-1})\|(m_1, m_{-1})) = \mathbb{E}[U_1(m'_1, m_{-1})] - D_{KR}(m_1, m'_1, m_{-1})$
Linear Loss Aversion	where $D_{KR}(m_1, m'_1, m_{-1}) = \sum_{j=1}^{\|A_1\|}\sum_{k=1}^{\|A_1\|}\sum_{l=1}^{\|A_{-1}\|} m_1^j m_1'^k m_{-1}^l \, \mu(u_1(a^j, a^l) - u_1(a^k, a^l))$
(PE-LLA)	and μ is a gain-loss utility function, that is, $\mu(z) = \begin{cases} z & \text{if } z \geq 0 \\ \lambda z & \text{if } z < 0 \end{cases}$ and $\lambda > 1$

Table 8: The Switching Cost and PE-LLA Models

In the multi-player setting, the set of personal equilibrium *does not* extend the set of NE. That is, while an agent behaves rationally in the one-player setting, in a multi-player setting the equilibria are no longer the same as those of the rational model. This is generically true, but to illustrate, we now present a simplified version of the 11-20 money request game (Arad and Rubinstein, 2012). A player can bid either $h = 2$ (high) or $l = 1$ (low). Every player is paid their bid, and in addition, if a player bids low and the other player bids high, then the low bidder receives a bonus of 3. The payoff matrix is presented in Table 9.

This game has two pure NE, (h, l) and (l, h), and one mixed NE, namely $\left(\frac{1}{3}h + \frac{2}{3}l, \frac{1}{3}h + \frac{2}{3}l\right)$. Note that, in the standard model, player 1's mixed strategy must leave player 2 indifferent between pure strategies (and vice versa). However, in the PE-LLA model, when player 1 plays this mixture, player 2 is no longer indifferent, rather he strictly prefers to bid low. To see why, suppose that player 1 plays $m = \frac{1}{3}h + \frac{2}{3}l$. Then,

	h	l
h	2,2	2,4
l	4,2	1,1

Table 9: A Simplified 11-20 Money Request Game

player 2's payoff of deviating from m to l is:

$$U_2\left((m,l)|(m,m)\right) = 2 + \frac{1}{3}\cdot 1\cdot\frac{1}{3}\cdot 2 + \frac{1}{3}\cdot 1\cdot\frac{2}{3}\cdot(-\lambda) = \frac{20}{9} - \frac{2\lambda}{9}$$

and his payoff from remaining at m is:

$$U_2\left((m,m)|(m,m)\right) = 2 + \frac{1}{3}\left(\frac{1}{3}\cdot\frac{2}{3}\cdot(-3\lambda) + \frac{2}{3}\cdot\frac{1}{3}\cdot 3\right) + \frac{2}{3}\left(\frac{1}{3}\cdot\frac{2}{3}\cdot 1 + \frac{2}{3}\cdot\frac{1}{3}\cdot(-\lambda)\right) = \frac{64}{27} - \frac{10\lambda}{27}$$

Since $\lambda > 1$, algebraic calculations show that $U_2((m,l)|(m,m)) > U_2((m,m)|(m,m))$. However, a mixed personal equilibrium exists, it is symmetric, and it differs from the NE. In the mixed personal equilibrium, each player plays $p\cdot h + (1-p)\cdot l$ where $p = \frac{-\sqrt{\lambda^2+2\lambda+33}+\lambda+5}{2(\lambda-1)}$. Thus, the PE-LLA model does not extend the NE (nor does the PE model).

Let us call this mixed equilibrium, $KR(\lambda)$. Notice that, $\lim\limits_{\lambda\downarrow 1} KR(\lambda) = \frac{1}{3}h + \frac{2}{3}l$, that is, it approaches the mixed NE as loss aversion vanishes. Furthermore, $\lim\limits_{\lambda\uparrow\infty} KR(\lambda) = l$. This is counter-intuitive because the mixed equilibrium profile converges to (l,l) as loss aversion increases, but this profile is Pareto-dominated and the welfare-worst outcome.