Bargaining, Forward-Looking Buyers and Dynamic Competition

Shanglyu Deng
UMD

Shen Hui
CUHK, Shenzhen

Dun Jia
Peking University

Mario Leccese
UMD

Andrew Sweeting*
UMD, CEPR and NBER

Preliminary. Please Do Not Cite. Color Figures.

November 7, 2022

Abstract

In many industries where supply-side scale economies or demand-side network effects may allow a single firm to become dominant, buyers may not be the atomistic price-takers that are typically assumed. We extend the dynamic "learning-by-doing and forgetting" model of Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) to analyze how two types of buyer sophistication, bargaining power and forward-looking behavior, affect market structure, welfare and the effects of stylized policies that are designed to preserve competition. Absent distortionary policies, even modest buyer sophistication of either type leads to unique equilibria. We show how dynamics can lead to the redistribution of bargaining power between buyers and sellers having non-linear and non-monotonic effects on outcomes, welfare and the impact of policies. We also show how buyer bargaining power and forward-looking behavior interact.

Keywords: dynamic competition, learning-by-doing, strategic buyers, forward-looking buyers, bargaining power, buyer power, multiple equilibria.

JEL codes: C73, D21, D43, L13, L41.

*Corresponding author: atsweet@umd.edu. Authors ordered alphabetically. We are very grateful to David Besanko, Uli Doraszelski, Paul Grieco, Steve Kryukov and seminar participants at the FTC, Penn State and the Washington DC IO Conference for helpful feedback on related work. All errors are our own, and comments are very welcome. We may divide the material in this draft into multiple papers.

1 Introduction

This paper seeks to understand what happens when two "hot topics" in Industrial Organization literature are combined. The first topic is the efficiency of dynamic competition when features of demand or supply could facilitate a single firm achieving a position of market dominance in the medium- or the long-run. A sophisticated understanding of dynamic competition should be crucial when applying merger- or conduct-based antitrust policies to digital platforms that benefit from network effects, and when designing industrial and international trade policies to manufacturing industries with large scale and learning economies, such as semiconductors, solar panels, aircraft and shipbuilding. Even if competition is promoted because of an objective other than the pursuit of efficiency or expected consumer surplus (e.g., limiting the economic or political heft of a single firm, ensuring security of supply), an understanding of dynamic competition is required to evaluate whether policies are actually likely to limit firm size and what their costs will be.

The second topic is how buyer-seller bargaining, rather than more traditional assumptions that sellers choose prices or quantities unilaterally, affect outcomes. The recent literature has shown how bargaining results in parties' outside options playing a critical role in model predictions. However, these analyses have assumed static settings where these outside options are not themselves functions of bargaining power. A consequence is that changes in the relative bargaining power of buyers and sellers (for example, reflected in generalized Nash bargaining weights) will typically change prices and surplus monotonically and often quite linearly. In a dynamic model, anticipated changes in future bargained outcomes will necessarily affect continuation values, potentially giving rise to much more complicated effects.

The existing dynamic competition literature has not considered the impact of bargaining. In this paper we extend a well-known model of dynamic competition, the duopoly seller "learning-by-doing (LBD) and forgetting" computational model of Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) (BDKS). Assuming that there is an atomistic buyer every period and that sellers set prices, BDKS show that, for many parameters, there exist equilibria where a single firm is likely to emerge as dominant (i.e., to have a significant cost advantage and to be much more likely to make sales) for at least some length of time. One of BDKS's most important contributions is to show how to use numerical "homotopies" to trace the equilibrium correspondence as the parameters that control the extent of LBD and the probability that know-how is forgetten are changed. BDKS use homotopies to show that multiple equilibria, sometimes involving qualitatively different pricing strategies and outcomes, exist for many parameters.

We extend the BDKS model to allow for both buyer-seller bargaining and buyers who recognize that they may be in the market in the future and so therefore partially internalize how their choices impact future competition. In many of the industries where LBD and forgetting have been documented, at least some of the buyers are large and sophisticated companies making both types of sophistication plausibly important features.[1]

We assume that sellers engage in a form of generalized Nash bargaining over prices every period, with the distribution of bargaining power determined by a single parameter. In the dynamic game, outside options are endogenous. Following Sweeting, Jia, Hui, and Yao (2022) (SJHY), we model buyers' forward-looking behavior using another parameter that measures the share of future buyer surplus that each buyer expects to appropriate. BDKS's model and the social planner's problem are captured as special cases of these additional parameters, and we are able to use homotopies to study precisely how outcomes change as these buyer sophistication parameters are varied separately or together.

We find several novel results. First, at least absent distortionary policies, the multiplicity of equilibria that is common when sellers set prices and buyers are atomistic is eliminated when buyers have even modest amounts of bargaining power or, as in SJHY, are even modestly forward-looking, or both.

Second, and most strikingly, dynamics lead changes in the distribution of bargaining power to have non-linear and often non-monotonic effects on outcomes. To understand why, consider an example. Much of our analysis will focus on illustrative parameters that imply significant LBD and forgetting probabilities that would not prevent both firms co-existing at the bottom of their cost curves in the long-run. For these parameters (and many others),

[1] The empirical literature has found evidence of significant LBD effects in aircraft (Benkard (2000)), chemicals (Lieberman (1984), Lieberman (1987a)), semiconductors (Irwin and Klenow (1994), Dick (1991)), shipbuilding (Thompson (2001), Thornton and Thompson (2001)), nuclear and renewable power plant construction (Zimmerman (1982), Joskow and Rose (1985), Qiu and Anadon (2012)) and hospital procedures (Gaynor, Seider, and Vogt (2005), Dafny (2005)).

when sellers set prices, the scope for one firm to become dominant is limited by the fact that a firm with a cost advantage will choose a larger markup that will, in expectation, make it reasonably likely that the laggard will catch up. This weakens incentives to compete aggressively for an advantage, so that the difference in outside options from agreeing and not agreeing a price early in the game may not be too large. In the opposite polar case where buyers have all of the bargaining power, sellers can never extract any surplus and there is no incentive for sellers to gain an advantage. On the other hand, suppose that sellers have most, but not all of the bargaining power. A leader knows that its future markups will be constrained by buyers' bargaining power, and that this may tend to perpetuate its dominance. This can increase sellers' incentives to gain a lead, leading to lower initial prices, even as sellers' share of the surplus from any individual bargain falls. We show how this logic will often result in measures of concentration and welfare that vary non-monotonically with bargaining power.

Third, we find that while market structure with price-setting sellers can be too unconcentrated (i.e., it would be socially optimal for one firm to get down its cost curve more quickly), market structure when atomistic buyers have most of the bargaining power will often be too concentrated. In fact, across most of the parameters we consider, market concentration will be close to the expected level under a social planner when sellers have 60%-80% of the bargaining power.

This fact raises the possibility that there will be bargaining power parameters for which policies that try to reduce market concentration will enhance efficiency and/or benefit customers, although one needs to account for how such policies may create additional distortions by softening competition. We study the effects of stylized versions of policies considered in the literature, such as a limitation on seller market shares, as in Benkard (2004), and restrictions on the incentives that firms can consider, as in Besanko, Doraszelski, and Kryukov (2014) (BDK). Our fourth set of results involve identifying combinations of parameters where policies can increase discounted and/or future customer welfare. We show how the distribution of bargaining power between buyers and sellers plays an important role in determining when policies are welfare improving.

Finally, we identify welfare- and policy-relevant interactions between bargaining power and forward-looking buyer behavior. When sellers set prices, forward-looking buyers will make purchase choices that lead to sellers tending to remain more symmetric, in order to encourage competition. This tends to soften seller competition, so that buyer welfare, as well as overall efficiency, can be lower than when buyers are atomistic. However, when buyers have bargaining power, both these competition-reducing effects and buyers' incentives to promote symmetry are weakened.

While our results are developed in the context of a specific model where the parameters have not been chosen to reflect any specific industry, we view them as suggesting that, in industries where demand and/or supply features imply that competition will be dynamic, policy evaluation needs to account for whether prices are set through negotiation, and, if so, how bargaining will be affected by dynamics and forward-looking behavior by buyers as well as sellers. As an example, consider the analysis of customer harm and efficiencies after a merger of aircraft manufacturers in An and Zhao (2019). Allowing for LBD, but assuming atomistic buyers and price-setting sellers, they estimate efficiency gains after the Boeing-McDonnell Douglas merger due to faster learning by the merged firm. The results presented here, albeit for a much simpler model, suggest that a 50/50 split of bargaining power could tend to produce too much concentration, potentially changing the welfare impact of a merger.[2]

As mentioned, our focus on the interactions between bargaining and dynamics are in contrast to the fast-growing applied bargaining literature (see Lee, Whinston, and Yurukoglu (2021) for a recent survey).[3] This literature has been motivated by the need to use bargaining to rationalize observed prices (given marginal costs must be non-negative) and its potential role (for example, in hospital or cable TV markets) in determining the types of contract menus or networks that downstream firms might wish to offer. However, with almost no exceptions, the applications are static in the sense that bargaining is assumed to happen only once, or, if bargaining is repeated, there are no other dynamics that would link outcomes across periods.[4] We completely abstract away from the choices that buyers make downstream

[2] In ongoing work with a richer model, we find that the level of synergy required to offset customer harm can change quite dramatically as assumptions concerning bargaining power are varied.

[3] Theoretical work (e.g., Loertscher and Marx (2022)) is also expanding the horizons of applied work beyond the Nash-in-Nash framework of Horn and Wolinsky (1988)).

[4] Lee and Fong (2013) consider a dynamic network formation game with bargaining, and use it to consider

(or how they may compete downstream), focusing instead on how bargaining affects market structure and incentives.

Before going further, we acknowledge several additional assumptions that we make. First, we follow BDKS in assuming that there are always two sellers, whereas, for example, BDK and SJHY consider models with endogenous entry and exit, and no forgetting. This is a deliberate choice. With endogenous entry and exit, shifting bargaining power to buyers would tend to have the additional effect that entry would be less profitable, and so tend to promote monopoly or the complete closure of the market.[5] While the effect of bargaining power on the number of firms is interesting, it is separate from the bargaining-dynamics interactions that we want to focus on. A second reason for preferring the BDKS model is that, as SJHY illustrate, the nature of equilibria in the BDK model is often that, once both firms have made a sale, they are guaranteed to stay in the market forever and competition is driven by the mutual knowledge that they must end up being symmetric at the bottom of their cost curves. The possibility that know-how can be forgotten enlarges the scope for dynamic competition and the potential importance of policies designed to promote competition, by creating the possibility that one firm could eventually become dominant from any point in the state space.

Second, while we allow for bargaining we continue to assume that firms are only able to negotiate over a single unit at one time, and they cannot sign multi-period contracts, possibly with options, contingencies, break fees, etc.. This is a feature that it would certainly be desirable to relax, although in industries where units have to be customized the ability of buyers and sellers to negotiate contracts that cover unspecified future units may be limited.

The rest of the paper is structured as follows. Section 2 describes our model, our definitions of incentives and outcomes, and our counterfactual policies. Section 3 analyzes polar cases for our illustrative parameters. Section 4 examines what happens when bargaining power is varied, assuming buyers are atomistic. Section 5 briefly explains the effects of forward-looking behavior when sellers set prices (see SJHY for more intuition), while Sec-

the effects of hypothetical hospital mergers. In contrast to our paper, the dynamics in Lee and Fong (2013) result from stochastic changes in the network links, whereas, in our paper, dynamics result from supply-side primitives.

[5]Like BDKS, we assume that a unit will always be bought from one seller. This assumption would result in a price-setting monopolist having unlimited market power.

tion 6 considers the interactions of bargaining power and forward-looking behavior. Section 7 concludes. The Appendices contain methodological details, some (limited) theoretical proofs and additional results that support those reported in the text.

2 Model

In this section we present our extended version of the BDKS model. Readers should consult that paper for additional motivation.

2.1 States and Costs.

There is an infinite horizon, discrete time, discrete state game. Throughout, the common discount factor is $\beta = \frac{1}{1.05}$. There are two ex-ante symmetric but differentiated sellers ($i = 1, 2$), whose marginal costs depend on their past sales. Specifically, each seller has a commonly observed state variable, $e_i = 1, ..., M$, that tracks its "know-how".[6] The state of the industry is $\mathbf{e} = (e_1, e_2)$.

A firm's cost of producing a unit of output is $c(e_i) = \kappa\rho^{\log_2(\min(e_i,m))}$. $\rho \in [0, 1]$ is known as the "progress ratio", and a lower number reflects stronger learning economies. When $\rho = 1$, marginal costs are κ for all e_i. We follow BDKS in assuming that $m = 15$ and $M = 30$.

Dynamics arise from the evolution of the know-how states. As described below, in each period one unit will be purchased from one of the sellers. The state of seller firm i evolves, except at the boundaries of the state space, according to

$$e_{i,t+1} = e_{i,t} + q_{i,t} - f_{i,t} \tag{1}$$

where $q_{i,t}$ is equal to one if firm i makes the sale, and $f_{i,t}$ is equal to one (0 otherwise) with probability $\Delta(e_i) = 1 - (1 - \delta)^{e_i}$ with $\delta \in [0, 1)$.[7] The probability of forgetting (Δ) is therefore increasing in δ and e_i. Equation (1) implies that a firm that makes a sale will

[6] Asker, Fershtman, Jeon, and Pakes (2020), Sweeting, Roberts, and Gedge (2020) and Sweeting, Tao, and Yao (2019) consider dynamic models where serially correlated state variables are private information.

[7] At the boundaries of the state space, the evolution is necessarily restricted. For example, when $e_{i,t} = 1$ and $q_{i,t} = 0$, firm i cannot forget ($f_{i,t} = 0$), and when $e_{i,t} = M$ and $q_{i,t} = 1$, firm i has to forget ($f_{i,t} = 1$).

either have the same or one unit more know-how in the next period, whereas a firm that does not make a sale will have either the same or less know-how.

Following BDKS, every period there is a buyer that purchases a single unit from one of the firms. Each period the chosen buyer receives flow indirect utility $v - p_i + \sigma\varepsilon_i$ if it buys from seller i, where p_i is the price paid, and σ parameterizes the degree of product differentiation. The ε_is are private information Type I extreme value payoff shocks, which are i.i.d. across buyers, sellers and periods, and do not depend on a buyer's past purchases. We follow BDKS's baseline case by assuming that $\sigma = 1$ throughout this paper. In BDKS buyers are atomistic, and choose based on prices that sellers set simultaneously. However, we depart from BDKS by allowing buyers to be strategic in two senses: (i) a buyer may internalize how its choice affects future buyer surplus (forward-looking behavior), and (ii) they may be able to extract more surplus from a transaction than they would receive when sellers set prices (bargaining). We now explain how we formulate these features of strategic behavior.

Sophisticated Buyers: Forward-Looking Behavior. We model forward-looking behavior by assuming that the buyer in each period makes their purchase choice as if they expect to receive a fraction b^p of future buyer surplus, so that they internalize this share of how much they should expect their choice to change the value of future buyers. BDKS's atomistic buyer assumption corresponds to $b^p = 0$. One way to rationalize this formulation is to imagine that the buyer in each period is chosen, with replacement, from a pool of symmetric potential buyers, and that each buyer expects to be the buyer in any future period with probability $0 \leq b^p \leq 1$. If $b^p = 1$, the model is consistent with a repeat monopsonist buyer, as in Lewis and Yildirim (2002). If $\frac{1}{b^p}$ is an integer, the model is consistent with a pool of this number of symmetric buyers from which a buyer is chosen with replacement each period (e.g., 5 buyers if $b^p = 0.2$). However, we will vary b^p continuously, and one can rationalize values where $\frac{1}{b^p}$ is not an integer using a behavioral interpretation where all buyers over- or under-estimate their future importance.

Sophisticated Buyers: Division of Surplus. We model the formation of prices by assuming that the buyer sends separate agents to each seller, before the buyer knows the realization of its εs. Each agent-seller pair negotiate the price at which a transaction will happen **if** the buyer chooses to purchase from that seller, with no purchase possible if a price is not agreed. We use a Nash-in-Nash approach, so that each agent-seller pair takes the price agreed by the other agent-seller pair as given. The buyer's "bargaining power" is equal to a parameter τ. Once both negotiations are completed, the buyer observes its εs and chooses which seller to purchase from.

There are two benefits of this formulation. First, because information is symmetric during negotiations, agreements will always be reached in equilibrium and outcomes will be consistent with BDKS's assumption of trade in every period, even if transaction prices are different.[8] Second, as described below, the price first-order conditions imply that outcomes equivalent to sellers simultaneously setting prices and the case where the buyer has to pay only marginal production costs are nested as special cases, with $\tau = 0$ and $\tau = 1$ respectively. When $b^p = 1$ and $\tau = 1$, the buyer will make the same choices as a social planner who, in every state, makes the purchase choice that maximizes expected discounted total surplus.

2.2 Equilibrium.

The equilibrium concept is symmetric and stationary Markov Perfect equilibrium (MPE, Maskin and Tirole (2001), Ericson and Pakes (1995), Pakes and McGuire (1994)). Symmetry applies to sellers (i.e., in an equivalent state both buyers will use the same strategies) and to buyers (if multiple buyers exist). An equilibrium can be expressed as the solutions $p^*(\mathbf{e})$ (negotiated prices), $VS^*(\mathbf{e})$ (beginning of period seller values) and $VB^*(\mathbf{e})$ (value of a buyer in the pool before the buyer for the current period is drawn) to the following equations, where symmetry implies that we only need to solve for prices and seller values for firm 1 (i.e., $p_2^*(e_1, e_2) = p_1^*(e_2, e_1)$, $VS_2^*(e_1, e_2) = VS_1^*(e_2, e_1)$) and buyer values in states where

[8] The key feature is that, during bargaining, the buyer's agent and the seller have symmetric information about the realization of the εs, which guarantees that prices will be agreed in equilibrium, allowing trade to happen. With asymmetric information during negotiation, there would be some probability that both negotiations would fail and no purchase would happen, creating a potential non-trivial difference to BDKS's model.

$e_1 \geq e_2$ (i.e., $VB^*(e_1, e_2) = VB^*(e_2, e_1)$).

Beginning of period value for firm 1 (VS):

$$VS_1^*(\mathbf{e}) - D_1^*(\mathbf{e})(p^*(\mathbf{e}) - c(e_1)) - \sum_{k=1,2} D_k^*(\mathbf{e})\mu_{1,k}^S(\mathbf{e}) = 0, \tag{2}$$

where $\mu_{1,k}^S(\mathbf{e})$, is seller 1's continuation value when the buyer chooses to buy from seller k,

$$\mu_{1,k}^S(\mathbf{e}) = \beta \sum_{\forall e'_{1,t+1}|e_{1,t}} \sum_{\forall e'_{2,t+1}|e_{2,t}} VS^*(e'_{1,t+1}, e'_{2,t+1}) \Pr(e'_{1,t+1}|e_{1,t}, k) \Pr(e'_{2,t+1}|e_{2,t}, k),$$

and $\Pr(e'_{i,t+1}|e_{i,t}, k)$ is the probability that i's state transitions from $e_{i,t}$ to $e'_{i,t+1}$ when a purchase is made from k ($q_{k,t} = 1$) given the forgetting probabilities. The probability that the chosen buyer purchases from seller k, $D_k(\mathbf{e})$, given negotiated prices, is

$$D_k(\mathbf{e}) = \frac{\exp\left(\frac{v - p_k + \mu_k^B(\mathbf{e})}{\sigma}\right)}{\exp\left(\frac{v - p_1 + \mu_1^B(\mathbf{e})}{\sigma}\right) + \exp\left(\frac{v - p_2 + \mu_2^B(\mathbf{e})}{\sigma}\right)}, \tag{3}$$

where $\mu_k^B(\mathbf{e})$, the buyer continuation value when it purchases from k, is defined below. $D_k^*(\mathbf{e})$ is this demand function evaluated at equilibrium prices and values.

Beginning of period buyer value (VB):

$$VB^*(\mathbf{e}) - b^p \sigma \log\left(\sum_{k=1,2} \exp\left(\frac{v - p_k^*(\mathbf{e}) + \mu_k^B(\mathbf{e})}{\sigma}\right)\right) - (1 - b^p) \sum_{k=1,2} D_k^*(\mathbf{e})\mu_k^B(\mathbf{e}) = 0 \tag{4}$$

where

$$\mu_k^B(\mathbf{e}) = \beta \sum_{\forall e'_{1,t+1}|e_{1,t}} \sum_{\forall e'_{2,t+1}|e_{2,t}} VB^*(e'_{1,t+1}, e'_{2,t+1}) \Pr(e'_{1,t+1}|e_{1,t}, k) \Pr(e'_{2,t+1}|e_{2,t}, k).$$

Negotiated prices (p): the assumed simultaneous Nash-in-Nash structure of bargaining implies that the equilibrium price negotiated between the buyer and seller 1, $p_1^*(\mathbf{e})$, will be the solution to the following maximization problem where p_2 is treated as fixed.

$$p_1^*(\mathbf{e}) = \arg\max_{p_1} \left(CS(p_1, p_2, \mathbf{e}) - CS(p_2, \mathbf{e})\right)^{\tau} \times ... \tag{5}$$

$$\left(D_1(\mathbf{e})(\mu^S_{1,1}(\mathbf{e}) + p_1 - c(e_1)) + (1 - D_1(\mathbf{e}))\mu^S_{1,2}(\mathbf{e}) - \mu^S_{1,2}(\mathbf{e})\right)^{(1-\tau)} \tag{6}$$

where $CS(p_1, p_2, \mathbf{e}) = \sigma \log\left(\sum_{k=1,2} \exp\left(\frac{v - p_k + \mu^B_k(\mathbf{e})}{\sigma}\right)\right)$ (i.e., the expected future surplus of the buyer when it is able to choose from both firms) and $CS(p_2, \mathbf{e}) = v - p_2 + \mu^B_2$, which is the buyer's expected surplus when the negotiation with seller 1 fails, and it has to buy from seller 2.

$p_1^*(\mathbf{e})$ is therefore the solution to the first-order condition

$$\begin{aligned}
&\tau \frac{\partial CS(p_1^*(\mathbf{e}), p_2, \mathbf{e})}{\partial p_1} \left(D_1^*(\mathbf{e})(\mu^S_{1,1}(\mathbf{e}) + p_1^*(\mathbf{e}) - c(e_1)) + (1 - D_1^*(\mathbf{e}))\mu^S_{1,2}(\mathbf{e}) - \mu^S_{1,2}(\mathbf{e})\right) + ... \\
&(1-\tau)\left(CS(p_1^*(\mathbf{e}), p_2, \mathbf{e}) - CS(p_2, \mathbf{e})\right)\left(D_1^*(\mathbf{e}) + \frac{\partial D_1^*(\mathbf{e})}{\partial p_1}\left(p_1^*(\mathbf{e}) - c(e_1) + \mu^S_{1,1}(\mathbf{e}) - \mu^S_{1,2}(\mathbf{e})\right)\right) = 0,
\end{aligned} \tag{7}$$

where $\frac{\partial CS(p_1^*(\mathbf{e}), p_2, \mathbf{e})}{\partial p_1} = -D_1^*(\mathbf{e})$ and $\frac{\partial D_1^*(\mathbf{e})}{\partial p_1} = -\frac{D_1^*(\mathbf{e})(1 - D_1^*(\mathbf{e}))}{\sigma}$. Algebraic manipulation shows that this can be simplified to

$$-\tau D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - \widehat{c_1}) + (1-\tau)\left[\sigma - (1 - D_1^*(\mathbf{e}))(p_1^*(\mathbf{e}) - \widehat{c_1})\right] \log \frac{1}{1 - D_1^*(\mathbf{e})} = 0. \tag{8}$$

where $\widehat{c_1} = c(e_1) - (\mu^S_{1,1} - \mu^S_{1,2})$ is firm 1's "effective" marginal cost of a sale which accounts for dynamic incentives.

When $\tau = 0$ (seller has all of the bargaining power), the FOC reduces to

$$\sigma - (1 - D_1^*(\mathbf{e}))(p_1^*(\mathbf{e}) - \widehat{c_1}) = 0 \tag{9}$$

which is equivalent to the BDKS first-order condition.

When $\tau = 1$ (buyer has all of the bargaining power), the FOC reduces to

$$D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - \widehat{c_1}) = 0 \tag{10}$$

with a unique solution where $p_1^*(\mathbf{e}) = c(e_1)$ in all states, implying that $VS^* = 0$. Note that in the presence of LBD, prices equal to current production costs will not be efficient when buyers are myopic.

Expressing the effective marginal cost of seller i as $\widehat{c_i}(\mathbf{e}) = c(e_i) - \mu_{i,i}^S(\mathbf{e}) + \mu_{i,j}^S(\mathbf{e})$, i's markup will be the following function of its market share,

$$p_i - \widehat{c_i} = \frac{(1-\tau)\sigma \log \frac{1}{1-D_i}}{\tau D_i + (1-\tau)(1-D_i) \log \frac{1}{1-D_i}}. \tag{11}$$

Note that in our model, the buyer's agent and the seller agree on the price at which trade may happen, not that it will happen, and that the agreed price affects the probability of trade.[9] A price increase reduces an agreement's expected surplus, as it is more likely the other seller will make the sale. This will tend to lower the price for a given τ.

The top panels in Figure 1 show agreed equilibrium prices in a static model (i.e., all buyer and seller continuation values are assumed equal to zero) as a function of τ when the marginal costs of the two sellers are 5 and 5 (so demand is always evenly split), and 3.85 and 7.5 (for the illustrative $\rho = 0.75$ these costs correspond to know-how states 10 and 2 respectively so firm 1 has a large lead), and we assume $\sigma = 1$. Prices equal marginal costs when $\tau = 1$. While equilibrium prices are convex in τ, the curvature of the price-τ relationship is limited.

2.3 Analytical Results.

In our extended model, we can show the following results.

Proposition 1 *1. There will be a unique symmetric MPNE when*

(a) $\delta = 0$ and $b^p = 0$, for all ρ and τ.

(b) $\tau = 1$, and either $b^p = 0$ or $b^p = 1$ for all ρ and δ.

2. If $\tau = 1$, prices will equal marginal production costs in all states, for all b^p, ρ and δ.

[9] In the Nash-in-Nash bargaining applications cited in Lee, Whinston, and Yurukoglu (2021) there is usually a mass of consumers (e.g., a mass of customers for health insurance) so that a known quantity (e.g., insurees using a particular hospital) will be transacted at the agreed prices.

Figure 1: Equilibrium Prices, Consumer Surplus and Total Surplus as a Function of τ (Buyer's Bargaining Power) in a Static Model. Expected consumer surplus equals the expected value of the ε associated with a buyer's purchase choice less the expected price. Expected total surplus equals the expected value of the ε less the expected production cost.

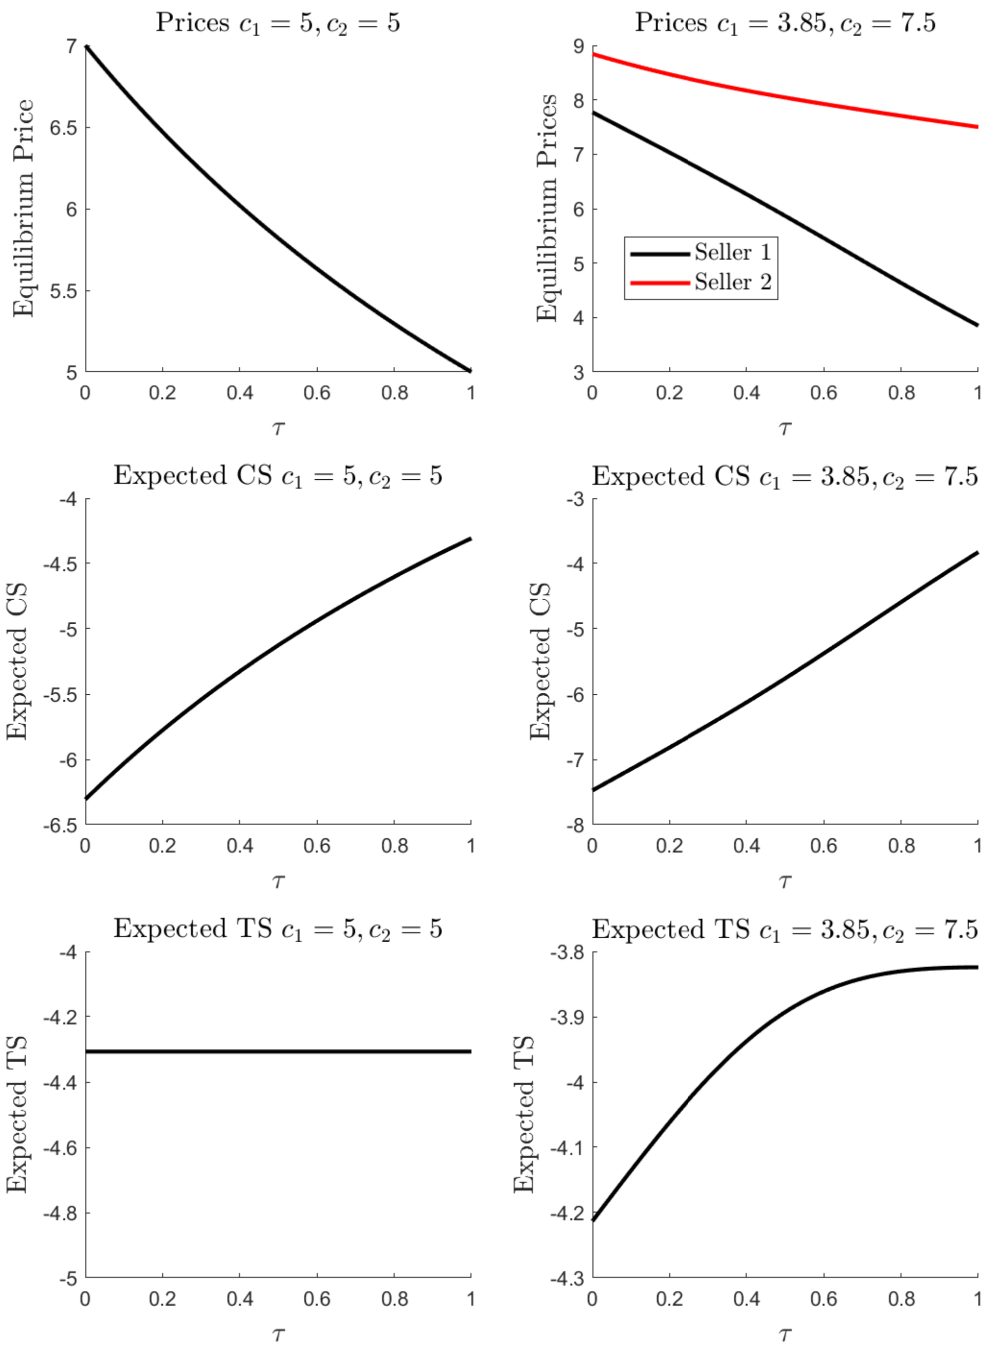

Proof. See Appendix A. ■

The positive results are not surprising. For example, when $\delta = 0$, backwards induction can be applied as the game must end up in state (M, M) and then stay there, and movements through the state space are unidirectional. When buyers are atomistic (completely static), backwards induction can then be applied to prove uniqueness. However, the proposition leaves open the possibility that there could be multiple equilibria when $\tau = 1$ if $0 < b^p < 1$ and $\delta > 0$, because buyers' expectations of their payoffs from particular choices can depend on the buyer behavior they expect in future states, even though prices always equal production costs. That said, we have never been able to construct multiple equilibria when $\tau = 1$ for the state space considered in this paper, although we have done so with very small state spaces.

2.4 Methods for Finding Equilibria and their Classification

When $\delta > 0$, it is always possible for the game to return to states with lower know-how. As BDKS emphasize, this can lead to the existence of many equilibria when $b^p = 0$.

For a given set of parameters we can find *an* equilibrium by solving the system of equations that defines an equilibrium or by using the iterative algorithm of Pakes and McGuire (1994). In order to search systematically for additional equilibria, and to examine what happens to equilibrium strategies and outcomes when we change b^p, τ, ρ or δ we also use BDKS's approach of numerical homotopies, which are, in essence, applications of the implicit function theorem to the system of equilibrium equations. Homotopies trace the equilibrium correspondences through the strategy and value space, as a single parameter is varied, starting from an equilibrium found by the Pakes and McGuire (1994) algorithm or a different homotopy. We will label a homotopy that varies the parameter α as a "α-homotopy".

An advantage of the homotopy approach is that it can identify equilibria that, because of the failure of local stability conditions, cannot be found using the Pakes-McGuire algorithm (unless one started exactly at the equilibrium). However, there is no guarantee that repeated application of homotopies will identify all equilibria, so our results presume that the extensive searches for equilibria, described in Appendix B, are effective.[10]

[10] SJHY show that, in the context of the closely-related BDK model, it is possible to use backwards induction to check whether particular types of equilibria exist. They find that the results from homotopy

2.5 Useful Measures.

We will describe equilibria by looking at prices, welfare, measures of market structure and incentives.

Concentration and Prices. Following BDKS, a summary statistic for expected market structure in period t is HHI^t where

$$HHI^t = \sum_{\forall \mathbf{e}} \mu^t(\mathbf{e}) HHI(\mathbf{e})$$

where

$$HHI(\mathbf{e}) = \sum_{k=1,2} \left(\frac{D_k^*(\mathbf{e})}{D_1^*(\mathbf{e}) + D_2^*(\mathbf{e})} \right)^2$$

and $\mu^t(\mathbf{e})$ is the probability that a game beginning in state (1,1) will be in state $\mathbf{e}$ after t periods. Given two firms, the minimum value of HHI^t is 0.5. All of our calculations assume that the industry begins in state $(1, 1)$ in the first period.

We will focus primarily on HHI^{32} as a measure of medium-run market structure and HHI^{1000} as a measure of long-run market structure (recall our discount factor is $\frac{1}{1.05}$). However, as illustrated in Appendix Figure C.1, expected HHIs can continue to change over hundreds of thousands of periods for some parameters.[11] We can define measures of expected prices after t periods, P^t in a similar way. P^{PDV} is the expected present discounted value of prices paid.

Surplus. As we assume a purchase is always made, the value v does not affect equilibrium strategies and choices, and is, in this sense, arbitrary. We therefore measure welfare ignoring the value of v, so that our measures of expected consumer surpus (CS, CS^t or CS^{PDV}) are the expected value of the buyers' εs associated with their purchases ($E(\varepsilon)$) less the expected price, even though this measure will (almost always) be negative. We will define producer

analysis are consistent with this check, thereby providing new evidence that homotopies are effective in this type of setting.

[11]The intuition for this result is that, for the parameters that we will focus on, once both firms are at the bottom of their cost curves, the expected HHI^t will be 0.5 in all future periods. However, the expected number of periods until this happens can be exceptionally large if a leader is very likely to make almost all sales and the laggard may forget any know-how that it accumulates.

surplus, PS^t or PS^{PDV}, as the expected price less production costs. Total surplus will be defined as $E(\varepsilon)$ for purchased units less production costs (PC^t or PC^{PDV}) (or, equivalently, $CS + PS$).

The lower four panels of Figure 1 show the relationship between τ and these measures in our static example. Expected CS is close to linear in τ in these states, reflecting how τ affects prices in a close to linear way. When the sellers have equal production costs, negotiated prices are equal and the buyer will always purchase from the seller with the highest ε, so expected TS is independent of τ. On the other hand, when marginal costs are different, increases in τ restrict the markup of the lower cost firm so that it becomes more likely to make a sale, which will be efficient.

Incentives. Following BDK, the "advantage building" (AB) and "advantage denying" (AD) incentives of firm 1 are defined as follows:

Definition 1 *The firm 1 AB incentive is* $\mu_{1,1}^S - \mu_{1,0}^S$. *The firm 1 AD incentive is* $\mu_{1,0}^S - \mu_{1,2}^S$.

where $\mu_{1,0}^S$ is seller 1's continuation value if, counterfactually, the buyer was to purchase from neither seller. The AB (AD) incentive measures the gain in firm 1's continuation valuation when firm 1 makes a sale (firm 2 does not make a sale). Firm 2's incentives can be defined similarly. Accounting for dynamics, the effective marginal cost of seller 1 in state $\mathbf{e}$ is therefore $\widehat{c_1}(\mathbf{e}) = c(e_1) - (\mathrm{AB}_1(\mathbf{e}) + \mathrm{AD}_1(\mathbf{e}))$. BDK identify the advantage denying incentive as playing an important role in aggressive pricing, and the existence of multiple equilibria, in their model. We will see that this is also true in the extended BDKS model when $b^p = 0$.

2.6 Policy Counterfactuals

We use our model to understand the welfare effects of stylized versions of two types of policies that might be used to try to maintain competition. Future work will consider the effects of merger and industrial policies using models that are more specifically tailored to those applications.

2.6.1 Restriction on Market Concentration

In his analysis of competition in the market for wide-bodied commercial aircraft, Benkard (2004) considers a counterfactual where limits are imposed on the market share of the largest firm in a given quarter. His estimates imply that market share limits of 60% and 51% slightly reduce consumer surplus and total surplus on average (changes are less than 1%), as prices not only tend to rise in states where the limits are binding, but also in states where firms would otherwise set low prices in order to try to gain a dominant position. While absolute restrictions on market shares are rare, market shares can play an important role in determining potential Sherman Act Section 2 liability for actions that agencies or rivals claim are anticompetitive.[12] The European Union, the UK (e.g., `https://tinyurl.com/hx75cf44`) and legislation in the US Congress have proposed frameworks that impose potentially onerous restrictions or requirements on platforms identified as dominant, which is likely to, at least partially, reflect market shares.

In our duopoly single-buyer-per-period model, we implement the policy as a "soft" market share restriction assuming that a market leader has to pay a compliance cost of $\chi \times \max\{0, D_i - \psi\}^2$ whenever its sale probability is more than $\psi > 0.5$. As χ increases, it becomes more costly for a firm to have a high market share.

Incorporating this penalty, the first-order condition for the negotiated price become the following

$$\begin{aligned} -\tau\left[D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - \widehat{c}_1) - \chi \max\{0, D_1^*(\mathbf{e}) - \psi\}^2\right] + (1-\tau)\log\left(\frac{1}{1 - D_1^*(\mathbf{e})}\right) \times \\ [\sigma - (1 - D_1^*(\mathbf{e}))(p_1^*(\mathbf{e}) - \widehat{c}_1) + 2\chi(1 - D_1^*(\mathbf{e}))\max\{0, D_1^*(\mathbf{e}) - \psi\}] = 0, \end{aligned} \tag{12}$$

and the equation for the seller's value becomes

$$VS_1^*(\mathbf{e}) - D_1^*(\mathbf{e})\left(p_1^*(\mathbf{e}) - c(e_1)\right) - \sum_{k=1,2} D_k^*(\mathbf{e})\mu_{1,k}^S(\mathbf{e}) - \chi \max\{0, D_1^*(\mathbf{e}) - \psi\}^2 = 0. \tag{13}$$

The restriction can affect prices even when $\tau = 1$, because the seller can guarantee that

[12] See the Department of Justice report on single-firm conduct (`https://www.justice.gov/atr/dojs-single-firm-conduct-report-promoting-consumer-welfare-through-clearer-standards-section-2`), although it does not represent official policy.

it will not violate the constraint if it does not agree a price.

2.6.2 Restrictions on Pricing Incentives

BDK and Besanko, Doraszelski, and Kryukov (2019) consider the effects of alternative limitations on pricing in the BDK model, motivated by standards that have been proposed in the antitrust literature for judging that a price is predatory. For example, alternatively requiring that sellers ignore dynamic incentives (i.e., price statically), that they ignore AD incentives and that they ignore incentives arising because their rival may exit.

We consider the effects of similar limitations in the BDKS setting where, instead of exit, there is some possibility that the rival forgets. While the lack of exit may appear to make the BDKS model less attractive for considering rules motivated by the legal literature on predation, it is not the case that the economic literature on investment motivated by anti-competitive intent (e.g., Caves and Porter (1977), Lieberman (1987b)) necessarily assumes that exit has to happen. Further, as shown in SJHY, for many parameters in the BDKS model it is certain that a firm will not exit once it has made a sale, so that firms that have made sales cannot be predated upon, whereas, of course, allegations of anticompetitive pricing often come from rivals that have had some initial successes. In the forgetting model, where a firm's know-how can always decrease if $\delta > 0$, there is potential scope for a leader to act aggressively against a rival over a much wider range of states.

We consider how equilibrium outcomes change when the market leader[13]:

- is assumed to be unable to consider dynamic incentives at all. This is implemented by excluding the seller's continuation values (i.e., the μ^S terms) from the first-order conditions that determine prices, so that the perceived marginal cost is simply the current production cost. We assume that, if $b^p > 0$, buyers still consider future incentives.
- is assumed to be unable to consider AD incentives, but it can consider AB incentives.
- is assumed to be unable to consider dynamic incentives that arise from the possibility that the rival may forget. Specifically we define "AB/Forgetting" (AB/F) and

[13]In the text, we focus on restrictions on a leader where $e_i > e_j$. In the Appendices we provide some results where the restrictions apply to both sellers.

"AD/Forgetting" (AD/F) incentives. If firm 1 makes the sale, the continuation value of firm 1 is

$$(1-\Delta_1)\beta V_1^S(e_1+1,e_2)+(1-\Delta_1)\Delta_2\beta[V_1^S(e_1+1,e_2-1)-V_1^S(e_1+1,e_2)]$$
$$+\Delta_1\beta V_1^S(e_1,e_2)+\Delta_1\Delta_2\beta[V_1^S(e_1,e_2-1)-V_1^S(e_1,e_2)],$$

so the benefit, for seller 1, associated with seller 2 possibly forgetting is

$$(1-\Delta_1)\Delta_2\beta[V_1^S(e_1+1,e_2-1)-V_1^S(e_1+1,e_2)]+\Delta_1\Delta_2\beta[V_1^S(e_1,e_2-1)-V_1^S(e_1,e_2)], \quad (14)$$

where the Δ terms are for state (e_1, e_2). Similarly, if (hypothetically) no sale is made, firm 1's benefit from firm 2 possibly forgetting is

$$(1-\Delta_1)\Delta_2\beta[V_1^S(e_1,e_2-1)-V_1^S(e_1,e_2)]+\Delta_1\Delta_2\beta[V_1^S(e_1-1,e_2-1)-V_1^S(e_1-1,e_2)], \quad (15)$$

If seller 2 makes the sale, firm 1's benefit from firm 2 possibly forgetting is

$$(1-\Delta_1)\Delta_2\beta[V_1^S(e_1,e_2)-V_1^S(e_1,e_2+1)]+\Delta_1\Delta_2\beta[V_1^S(e_1-1,e_2)-V_1^S(e_1-1,e_2+1)]. \quad (16)$$

The AB/Forgetting term is then defined by (14)-(15), and the AD/Forgetting term is defined by (15)-(16).

We focus on restrictions on the leader as it is leaders who are most likely to be subject to allegations of anticompetitive conduct.

3 The Effects of Sophisticated Buyers in the Extended BDKS Model: Comparison of Polar Cases

Our primary interest will be in what happens when τ and b^p take on values between zero and one, but we begin by considering the four polar cases that can be used as benchmarks. We do so using what we call our "illustrative parameters", $\kappa = 10$, $\sigma = 1$, $\rho = 0.75$ and $\delta = 0.023$. Ghemawat (1985) reports that 79 out of 97 empirical studies of LBD estimated values of ρ between 0.75 and 0.9, so that, our chosen $\rho = 0.75$ implies that learning can reduce costs at a rate towards the higher end of the empirically relevant range, with marginal costs possibly falling as low as 3.25 from a maximum of 10. $\delta = 0.023$ implies that the probabilities of forgetting when e is 15 or 30 are 0.2946 and 0.5024 respectively, implying that, both firms can remain close to the bottom of their cost curves for a sustained period if they split sales equally.[14]

3.1 Equilibrium Outcomes with Polar Values of b^p and τ.

We begin by describing equilibrium outcomes when there are no policies in place.

Equilibria for $b^p = 0$ and $\tau = 0$ (BDKS Baseline). We identify three equilibria, which we will call the "Low-", "Mid-" and "High-HHI" equilibria reflecting the associated value of HHI^{1000}.

The top panel of Figure 2 show equilibrium firm 1 prices. Pricing in the Mid-HHI and High-HHI equilibria is characterized by a "diagonal trench" where firms charge very low prices, which are below production costs, when they are exactly symmetric. Prices also tend to be low when there is a one state difference in know-how. The difference between the High-HHI and Mid-HHI equilibria is that, in the former, the trench extends all the way to (30,30), whereas it stops short in the latter. The Low-HHI is characterized by prices that are more similar across states, apart from a "well" in state (1,1) where firms compete to gain the first unit of know-how. There is a less pronounced "sideways trench" where prices are low when the laggard $e = 2$ for all three equilibria: for example, firm prices are lower in $(10, 2)$ and $(15, 2)$ than in $(10, 3)$ and $(15, 3)$ even though the only difference in costs it that firm 2's costs are lower in the latter states.

Comparing across the equilibria, Low-HHI prices tend to be lower when a firm has a lead of more than two states. For example, in state $(10, 3)$, when marginal costs are (3.85, 6.34), prices are $(5.20, 5.81)$ in the Low-HHI equilibrium and $(5.80, 6.31)$ in the other equilibria. On

[14]BDKS focus much of their discussion on $\rho = 0.85$ and $\delta = 0.0275$. For these parameters, when $b^p = 0$ and $\tau = 0$, only two clearly distinct equilibria are identified (i.e., solutions that are different for a range of numerical tolerances for defining different equilibria). We also only identify two clearly distinct equilibria if $\rho = 0.75$ and $\delta = 0.0275$. As the number of equilibria should generically be odd, we prefer to focus on our illustrative parameters, for which we can identify three equilibria, as a baseline. We discuss what happens as ρ and δ change below.

Figure 2: Equilibrium Prices. The blue curve on the price plot shows firm 1 marginal costs.

(a) $b^p = 0$, $\tau = 0$ High-HHI: $HHI^{1000} = 0.527$

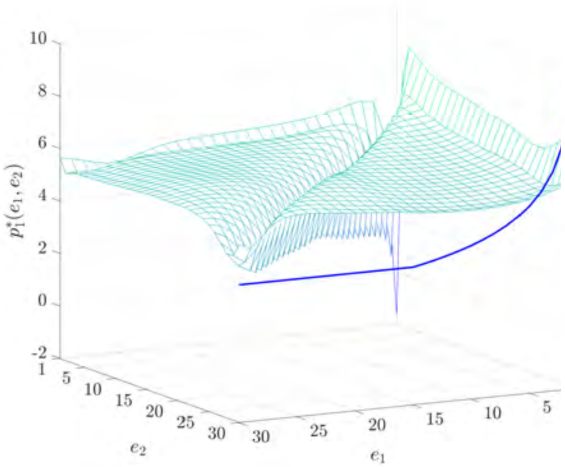

(b) $b^p = 0$, $\tau = 0$ Mid-HHI: $HHI^{1000} = 0.516$

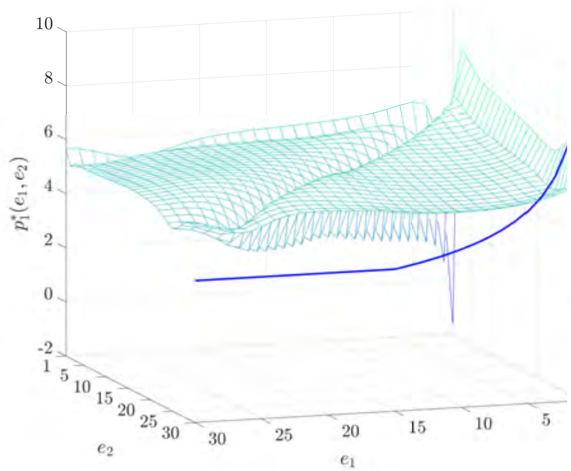

(c) $b^p = 0$, $\tau = 0$ Low-HHI: $HHI^{1000} = 0.500$

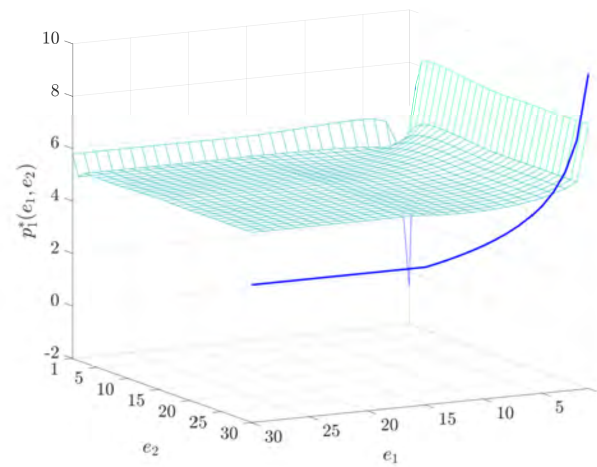

(d) $b^p = 0$, $\tau = 1$: $HHI^{1000} = 0.967$

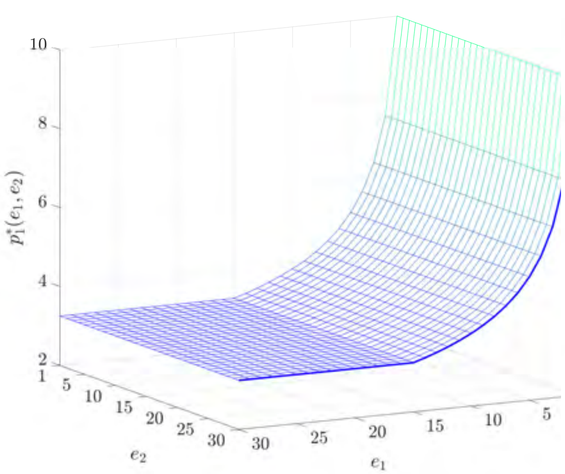

(e) $b^p = 1$, $\tau = 0$: $HHI^{1000} = 0.500$

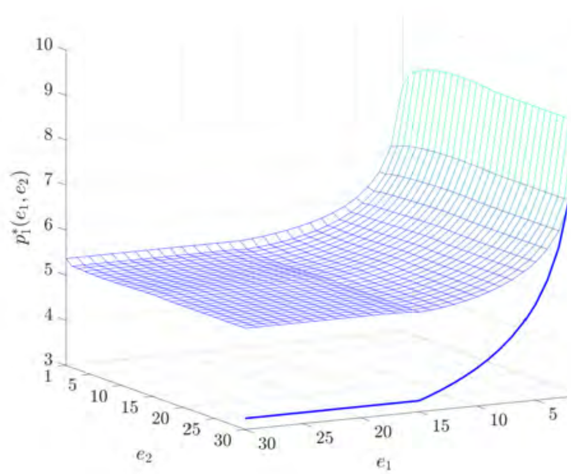

(f) $b^p = 1$, $\tau = 1$: $HHI^{1000} = 0.613$

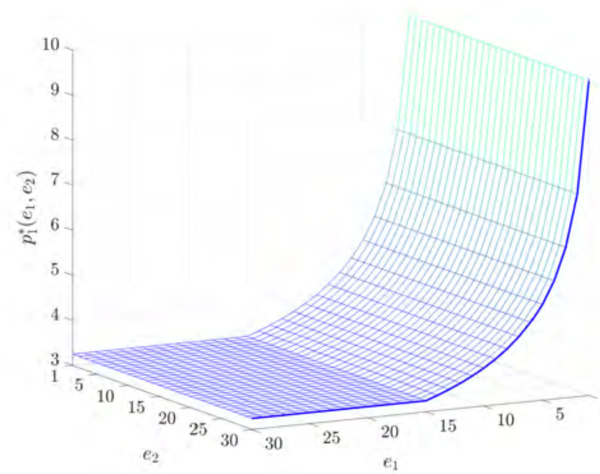

Table 1: Outcome and Welfare Comparisons for Polar b^p and τ Cases for the Illustrative Parameters.

Parameters		$b^p = 0,\ \tau = 0$		$b^p = 1,\ \tau = 1$	$b^p = 0,\ \tau = 1$	$b^p = 1,\ \tau = 0$	$b^p = 0,\ \tau = 0.5$	$b^p = 0.05,\ \tau = 0.5$
Number of Equilibria		3		1	1	1	1	1
	Low-	Mid-	High-					
Expected Concentration (HHI^t) after								
8 periods	0.731	0.669	0.668	0.990	0.972	0.516	0.995	0.992
32 periods	0.537	0.520	0.520	0.976	0.987	0.500	0.993	0.989
100 periods	0.500	0.518	0.519	0.930	0.985	0.500	0.984	0.974
1000 periods	0.500	0.516	0.527	0.613	0.967	0.500	0.877	0.802
Expected Transaction Price (P^t) after								
8 periods	5.818	6.226	6.222	4.397	4.457	6.358	4.322	4.307
32 periods	5.310	5.609	5.616	3.296	3.268	5.368	4.257	4.220
100 periods	5.250	5.419	5.512	3.299	3.262	5.251	4.252	4.215
1000 periods	5.250	5.210	5.376	3.263	3.260	5.250	4.210	4.160
Discounted Total	109.980	115.769	115.902	91.991	92.549	128.401	92.514	92.682
Production Costs								
Avg. across 8 periods	6.566	6.691	6.694	6.123	6.215	7.209	6.115	6.122
Avg. across 32 periods	4.825	4.937	4.936	4.123	4.146	5.039	4.106	4.113
Avg. across 100 periods	3.858	3.868	3.868	3.563	3.547	3.873	3.538	3.547
after 32 periods	3.845	3.779	3.781	3.296	3.268	3.665	3.269	3.276
after 1000 periods	3.250	3.250	3.250	3.263	3.260	3.250	3.268	3.271
Discounted Total	103.840	105.794	105.797	91.991	92.549	108.964	91.627	91.779
Expected ε								
Avg. across 8 periods	0.411	0.451	0.452	0.118	0.194	0.673	0.110	0.117
Avg. across 32 periods	0.533	0.587	0.587	0.056	0.073	0.685	0.041	0.048
Avg. across 100 periods	0.637	0.647	0.646	0.071	0.042	0.690	0.031	0.042
after 32 periods	0.652	0.672	0.672	0.046	0.027	0.693	0.021	0.028
after 1000 periods	0.693	0.677	0.666	0.539	0.052	0.693	0.179	0.283
Discounted Total	11.230	12.110	12.110	1.522	1.895	14.358	1.136	1.288
Expected CS: ε Less Production Costs								
Avg. across 8 periods	-4.552	-4.739	-4.739	-6.005	-6.021	-6.436	-4.607	-4.687
Avg. across 32 periods	-4.817	-5.065	-5.066	-4.067	-4.073	-5.341	-4.312	-4.299
Avg. across 100 periods	-4.647	-4.908	-4.940	-3.492	-3.505	-4.819	-4.254	-4.218
after 32 periods	-4.658	-4.937	-4.944	-3.249	-3.242	-4.675	-4.236	-4.192
after 1000 periods	-4.557	-4.533	-4.710	-2.724	-3.208	-4.557	-4.032	-3.877
Expected TS: ε Less Production Costs								
Avg. across 8 periods	-6.155	-6.240	-6.242	-6.005	-6.021	-6.537	-6.005	-6.005
Avg. across 32 periods	-4.292	-4.349	-4.349	-4.067	-4.073	-4.355	-4.064	-4.065
Avg. across 100 periods	-3.221	-3.221	-3.222	-3.492	-3.505	-3.183	-3.507	-3.505
after 32 periods	-3.193	-3.107	-3.108	-3.249	-3.242	-2.972	-3.248	-3.248
after 1000 periods	-2.557	-2.573	-2.584	-2.724	-3.208	-2.557	-3.089	-2.987
PDV Welfare Measures								
Total Surplus:	-92.610	-93.684	-93.687	-90.469	-90.653	-94.606	-90.491	-90.491
Consumer Surplus	-98.750	-103.659	-103.792	-90.469	-90.653	-114.043	-91.378	-91.394
Producer Surplus	6.140	9.975	10.105	0	0	19.437	0.887	0.903

Figure 3: State Distribution After 32 and 1000 Periods for $b^p = 0$ and $\tau = 0$.

(a) 32 Periods High-HHI

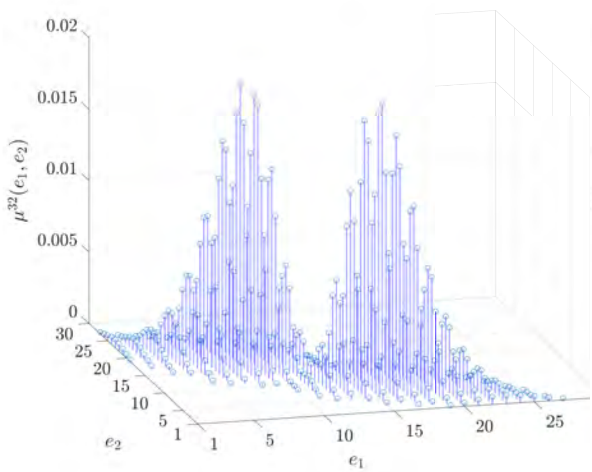

(b) 32 Periods Mid-HHI

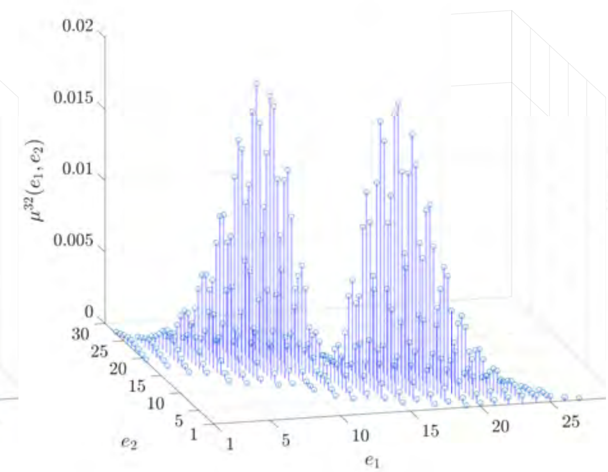

(c) 32 Periods Low-HHI

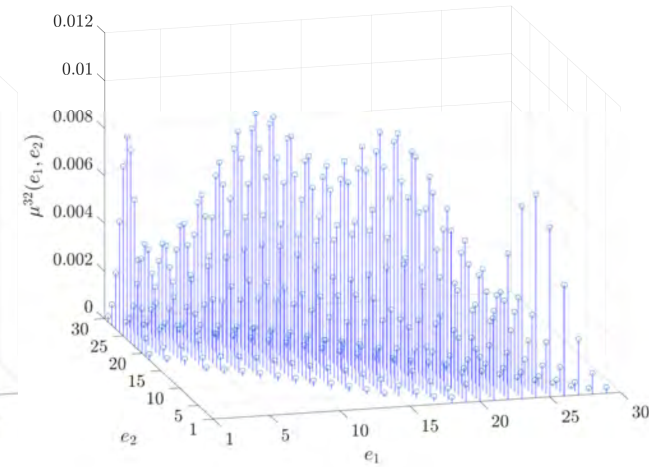

(d) 1,000 Periods High-HHI

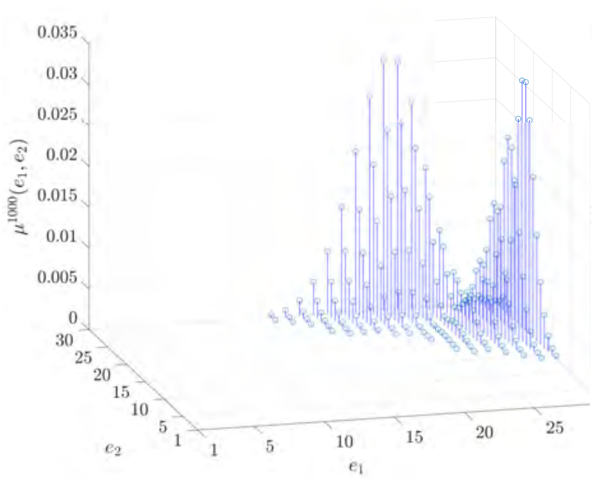

(e) 1,000 Periods Mid-HHI

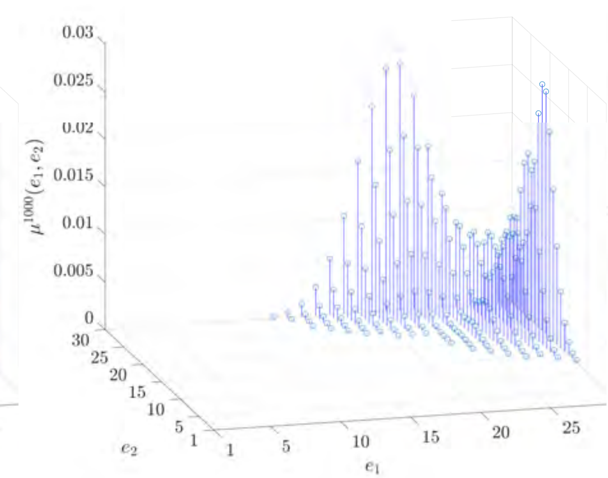

(f) 1,000 Periods Low-HHI

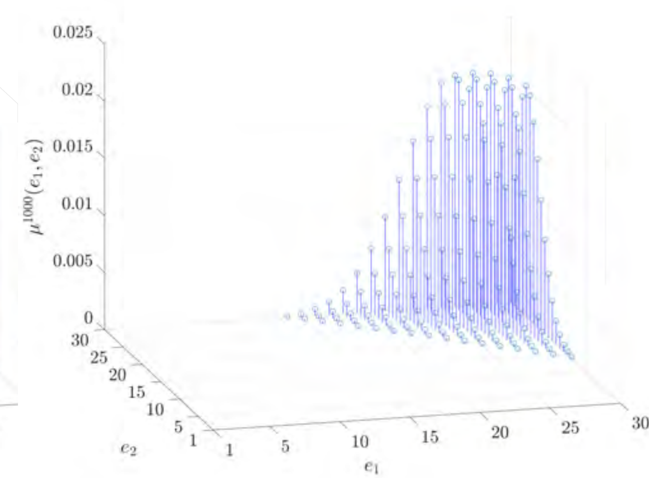

the other hand, as illustrated in Figure 5, which shows the probability firm 1, with $e_1 = 5$, makes a sale, the High- and Mid-HHI are associated with much higher probabilities that the leader makes sales when the lead is small, and smaller probabilities when the lead is large, reflecting the larger markups of a firm with a large lead. All three equilibria imply higher sales for firm 1 when it is the leader ($e_2 \leq 4$) than static pricing (where firms set prices to maximize their current profits in every state), but also when firm 2 is at, or close to, the bottom of its cost curve ($e_2 \geq 15$), as in these states firm 1 sets smaller than static margins in order to try to catch-up.

The left-hand columns of Table 1 report what these equilibria imply for market concentration and the components of welfare over different periods of time. When interpreting the welfare numbers, it is worth remembering that if a firm expects to become a leader, with large margins, later in the game then this will often lead to very low prices, as the sellers compete to become the leader, at the start of the game. Figure 3 shows the probability distribution of states after 32 and 1,000 periods (probabilities of less that 10^{-4} are not marked) .

The probability that one firm makes most of the early sales is actually higher in the (long-run) Low-HHI equilibrium, reflecting how the more aggressive pricing of a leader with a significant advantage will tend to reinforce its lead. However, the leader's aggressive pricing close to the diagonal trench in the Mid- and High-HHI equilibria leads to a lower probability that the firms will be very close to symmetric in the long-run. As the expected value of the εs is maximized when the choice probabilities equal one-half, these tend to be lower early in the game, and higher later in the game, in the Low-HHI equilibrium. The combination of a quicker reduction in production costs and lower expected prices lead to the PDV of CS and TS being higher in the Low-HHI equilibrium.

Equilibrium for $b^p = 1$ and $\tau = 1$ (Social Planner). These parameters support strategies that maximize expected discounted social welfare. Prices, shown in Figure 2(f), equal production costs and one firm tends to move down its cost curve quickly (Figure 4). The social planner sacrifices expected utility from εs at the start of the game in order to lower future production costs. However, once one firm has know-how beyond m and some favorable

Figure 4: State Distribution After 32 and 1000 Periods for Polar Cases.

(a) 32 Periods $b^p = 1$, $\tau = 1$.

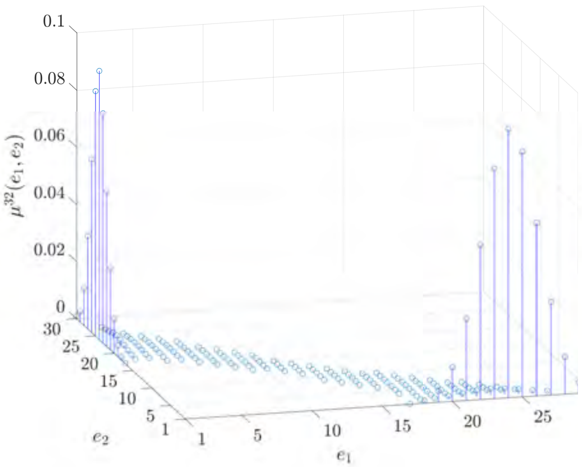

(b) 32 Periods $b^p = 0$, $\tau = 1$.

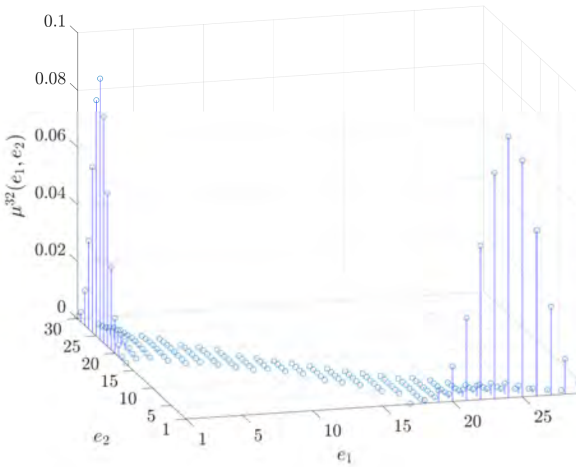

(c) 32 Periods $b^p = 1$, $\tau = 0$.

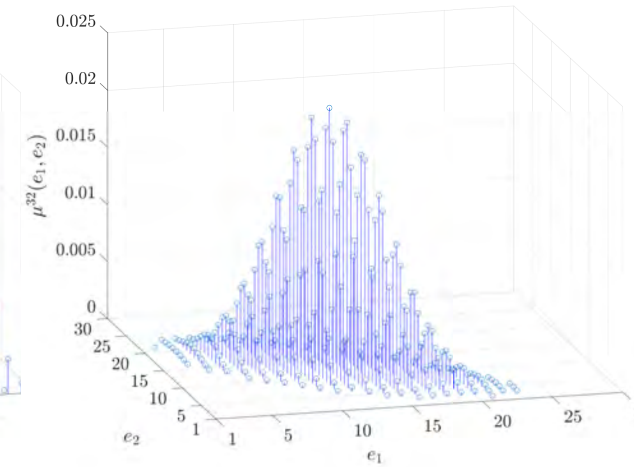

(d) 1,000 Periods $b^p = 1$, $\tau = 1$.

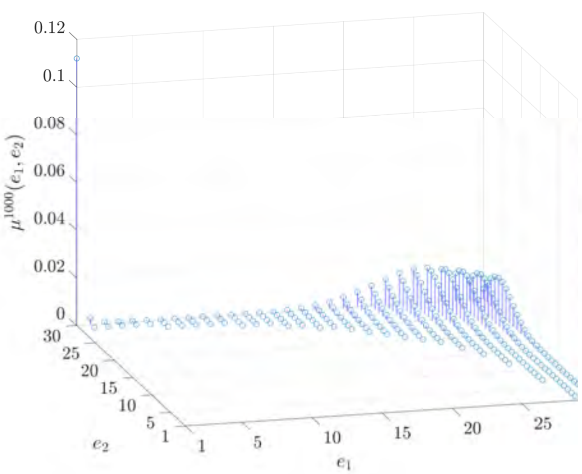

(e) 1,000 Periods $b^p = 0$, $\tau = 1$.

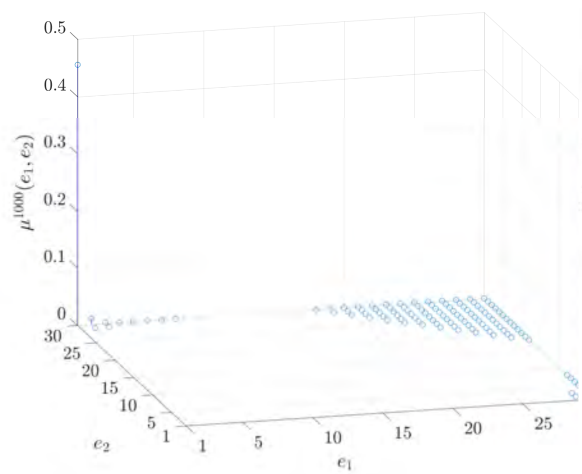

(f) 1,000 Periods $b^p = 1$, $\tau = 0$.

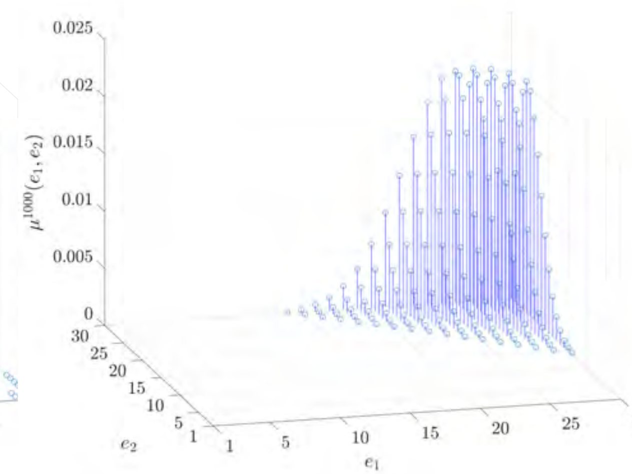

Figure 5: Equilibrium Probability of $e_1 = 5$ Seller 1 Making Sale for the Illustrative Paramet

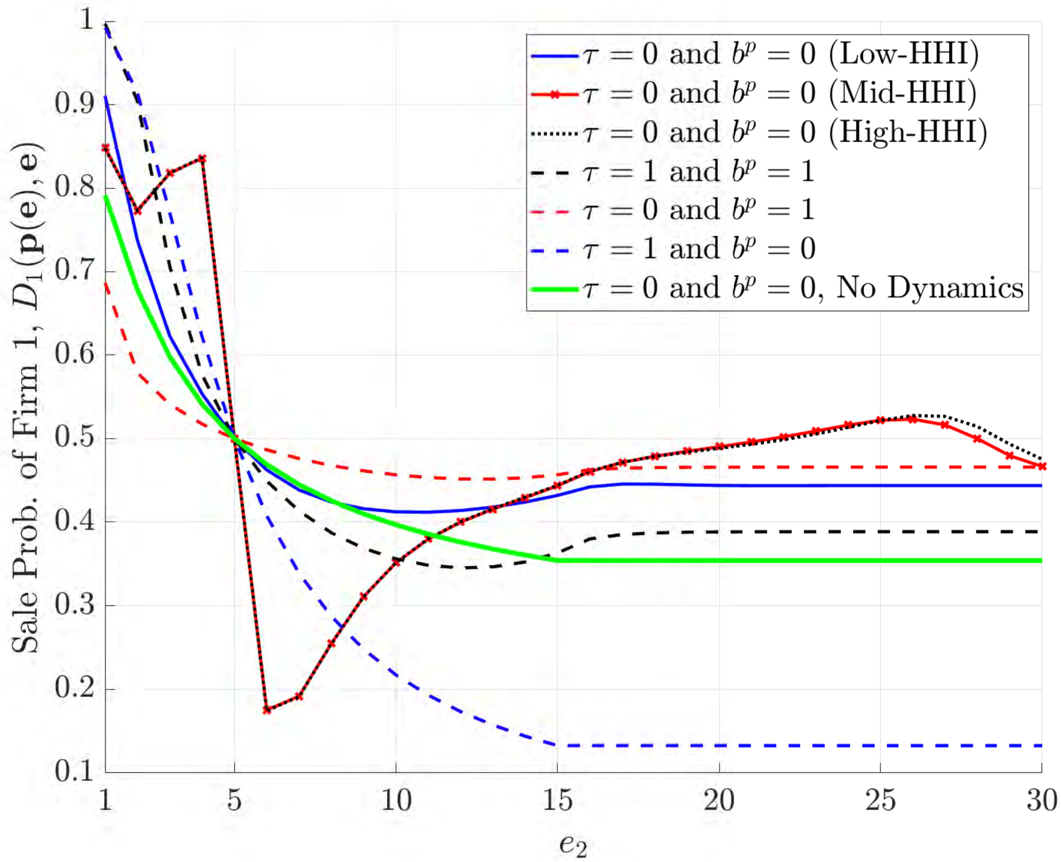

εs and depreciation realizations have given the laggard a few units of know-how, the social planner will exercise some preference to buy from the laggard, in order to create a second low cost supplier. This dynamic is illustrated in Figure 5 by the increase in the probability that an $e_1 = 5$ firm 1 makes a sale when e_2 increases above $m = 15$. On the other hand, the εs that prompt purchases from the laggard, are sufficiently extreme that, in expectation, market structure converges to being symmetric only slowly (Appendix Figure C.1(f)).

Equilibrium for $b^p = 0$ and $\tau = 1$. Transaction prices also equal the marginal production costs in this case, but, unlike in the social planner case, buyers do not internalize how their choices may lower future costs or raise future εs. This results in slower learning/less concentration early in the game, but higher concentration after 1,000 periods as buyers do not try to create a second low cost provider. Reflecting this, the probability that firm 1 makes a sale when it is significantly behind firm 2 (Figure 5) is low, including in comparison to static $\tau = 0$ pricing where the leader's markup raises the probability that the laggard makes a sale. Convergence to symmetric duopoly, which tends to be sustained once it is reached, tends to take hundreds of thousands of periods (Appendix Figure C.1(d)). However, because sellers cannot charge markups, losses in TS and CS, relative to the social optimum, are small.

Equilibrium for $b^p = 1$ and $\tau = 0$. This case reflects the logic of monopsonist procurement policies described by Lewis and Yildirim (2002), Saini (2012) and Anton, Biglaiser, and Vettas (2014) where the buyer seeks to maintain symmetry in order to increase competition. However, this has the effect of softening competition as sellers perceive less benefit to gaining a cost advantage.[15] Our homotopy analysis identifies only a single equilibrium, with probabilities that a laggard makes a sale higher than when static pricing is used (Figure 5), although not necessarily relative to $b^p = 0$ and $\tau = 0$ (because when $b^p = 0$ and $\tau = 0$ a laggard may price aggressively in some states to try to catch-up). Production costs tend to fall slowly and margins tend to be relatively large, so that expected discounted TS and CS are lower than in the other polar cases, and expected seller profits are higher.

[15] Note that these papers assume sellers have some private information about costs which allows them to secure rents even when the buyer uses an optimal mechanism.

3.2 The Effect of Policies to Reduce Dominance.

We now consider the effect of alternative policies for reducing dominance in these polar cases. The left-hand columns of Table 2 report the expected effects on market concentration and welfare of a threshold rule, with $\chi = 50$ and $\psi = 0.75$, and of policies that restrict the types of dynamic incentive that a leader (i.e. a firm with $e_i > e_j$) can consider. Appendix Table C.2 reports the results for a version of the threshold rule with $\chi = 10$ and when the incentives of both sellers, not just the leader, are restricted. We assume that the incentive-type policies can be imposed costlessly, but we report the expected "compliance costs" levied on a firm subject to the threshold rule, and deduct them from PS and TS.

The first panel in Table 2 summarizes outcomes in the polar cases with no policies. The remaining panels show changes in outcomes under the alternative policies. For $b^p = 0$ and $\tau = 0$, three equilibria remain under the threshold rule, and we label the equilibria based on the paths that are traced from the baseline equilibria as χ is increased. For the illustrative parameters, we only identify a single equilibrium for the incentive policies, and the reported changes are those from the listed baseline equilibrium (we will see that for other parameters considerable multiplicity can exist). When $\tau = 1$, sellers' incentives do not affect pricing so that none of the incentive policies affect outcomes, even though the industry tends to be very concentrated without any policy (reflecting the possibly inefficient choices of atomistic buyers). On the other hand, as mentioned before, threshold policies do affect outcomes as a seller with a large lead has to be compensated for expected compliance costs.

Policy Effects when $b^p = 0$ and $\tau = 0$. Figure 6 shows changes in equilibrium prices and Figure 7 shows changes in the distribution of states after 32 periods (changes in probabilities that are less than 10^{-4} are not marked). In the case of the incentive policies, the figures show changes from the Low-HHI equilibrium. All of the policies have the effect of increasing the probability that the sellers will be relatively close to symmetric after 32 periods, consistent with the nominal policy goal, at the cost of slower learning and increased production costs. When long-run concentration with no policy is low (for example, when $b^p = 0$ and $\tau = 0$), the threshold policy has little or no long-run effects. However, this does not prevent significant effects earlier in the life of the industry.

Table 2: Effects of Alternative Policies to Restrict Dominance in the Baseline Cases for the Illustrative Parameters. Δ indicates changes from the values in the upper part of the table. Low, Mid and High denote the relevant baseline equilibria.

Parameters		$b^p = 0,\ \tau = 0$		$b^p = 1,\ \tau = 1$	$b^p = 0,\ \tau = 1$	$b^p = 1,\ \tau = 0$	$b^p = 0,\ \tau = 0.5$	$b^p = 0.05,\ \tau = 0.5$
No Policy Baseline	Low	Mid	High					
HHI^{32}	0.537	0.520	0.520	0.976	0.987	0.500	0.993	0.989
HHI^{1000}	0.500	0.516	0.527	0.613	0.967	0.500	0.877	0.802
PDV CS	-98.750	-103.659	-103.793	-90.469	-90.653	-114.043	-91.378	-91.394
PDV PS	6.149	9.975	10.105	0	0	19.437	0.887	0.903
CS^{1000}	-4.557	-4.533	-4.710	-2.723	-3.208	-4.557	-4.032	-3.877
PS ^{1000}	2.000	1.960	2.2126	0	0	2.000	0.943	0.890
Threshold Rule, $\chi = 50$**,** $\psi = 0.75$								
# of equilibria		3		1	1	1	1	1
ΔHHI^{32}	-0.025	0.001	0.001	-0.449	-0.117	0.000	-0.192	-0.267
Δ HHI^{1000}	0.000	-0.000	-0.000	-0.113	-0.465	0.000	-0.377	-0.302
Δ PDV CS	-6.424	-3.240	-3.242	-4.110	-33.728	0.000	-33.053	-22.447
Δ PDV PS (net compliance)	3.222	1.984	1.978	0	0	0.000	0.758	1.271
PDV Compliance Cost	2.492	0.966	0.980	1.123	32.795	0.000	31.365	19.647
Δ CS^{1000}	0.000	0.000	0.001	0.167	0.640	0.000	0.658	0.504
Δ PS^{1000} (net compliance)	0.000	-0.000	-0.001	0	0	0.000	-0.126	-0.074
Compliance $Cost^{1000}$	0.000	0.000	0.000	0.000	0.008	0.000	0.000	0.000
No Leading Seller Dynamics								
# of equilibria		1		1	1	1	1	1
Δ HHI^{32}	-0.036	-0.019	-0.019	0	0	0.000	-0.351	-0.416
Δ HHI^{1000}	0.000	-0.016	-0.027	0	0	0.000	-0.377	-0.302
Δ PDV CS	-18.384	-13.475	-13.342	0	0	-13.232	-4.375	-5.975
Δ PDV PS	15.398	11.563	11.433	0	0	12.441	1.654	2.448
Δ CS^{1000}	0.002	-0.022	0.155	0	0	0.000	0.663	0.509
Δ PS^{1000}	-0.002	0.039	-0.128	0	0	0.000	-0.131	-0.079
No Leading Seller AD Incentive								
# of equilibria		1		1	1	1	1	1
ΔHHI^{32}	-0.034	-0.017	-0.017	0	0	-0.000	-0.169	-0.252
Δ HHI^{1000}	0.000	-0.016	-0.027	0	0	-0.000	-0.377	-0.302
Δ PDV CS	-8.450	-3.541	-3.408	0	0	-9.767	-1.102	-1.837
Δ PDV PS	7.127	3.292	3.162	0	0	9.360	0.526	0.849
Δ CS^{1000}	0.002	-0.022	0.155	0	0	-0.001	0.664	0.510
Δ PS^{1000}	-0.002	0.039	-0.127	0	0	0.001	-0.132	-0.082
No Leading Seller AB/Forgetting or AD/Forgetting Incentive								
# of equilibria		1		1	1	1	1	1
ΔHHI^{32}	-0.003	0.014	0.013	0	0	0.000	-0.000	-0.000
Δ HHI^{1000}	-0.000	-0.016	-0.026	0	0	0.000	-0.005	0.006
Δ PDV CS	-0.176	4.733	4.866	0	0	0.205	0.001	0.001
Δ PDV PS	0.131	-3.704	-3.834	0	0	-0.201	-0.000	-0.001
Δ CS^{1000}	-0.000	-0.024	0.153	0	0	-0.000	0.005	0.007
Δ PS^{1000}	0.000	0.040	-0.126	0	0	0.000	0.002	0.001

Figure 6: Effects of Policies on Equilibrium Pricing for the Illustrative Parameters When $b^p = 0$ and $\tau = 0$. For the incentive policies the figures show changes from the Low-HHI baseline equilibrium.

(a) Share Thresholds ($\chi = 50, \psi = 0.75$): Price Changes for Low-HHI Eqm.

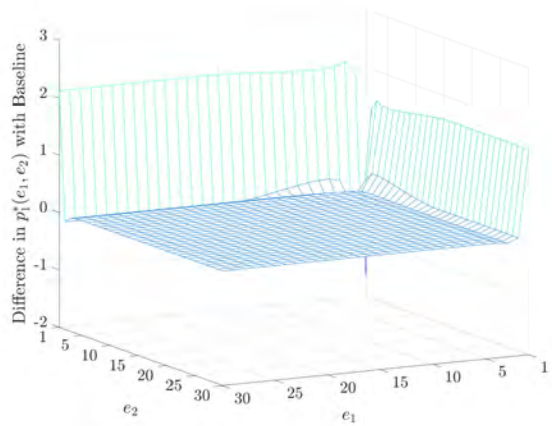

(b) Share Thresholds ($\chi = 50, \psi = 0.75$): Price Changes for Mid-HHI Eqm.

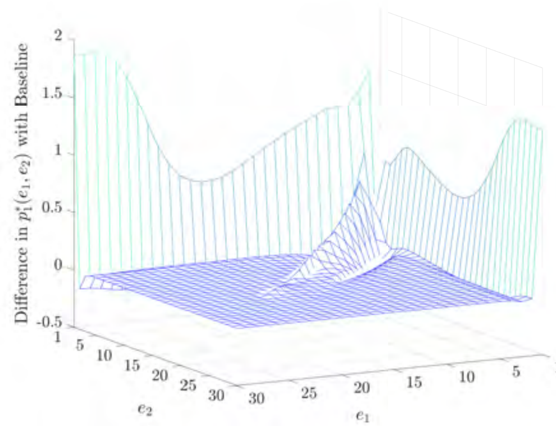

(c) Share Thresholds ($\chi = 50, \psi = 0.75$): Price Changes for High-HHI Eqm.

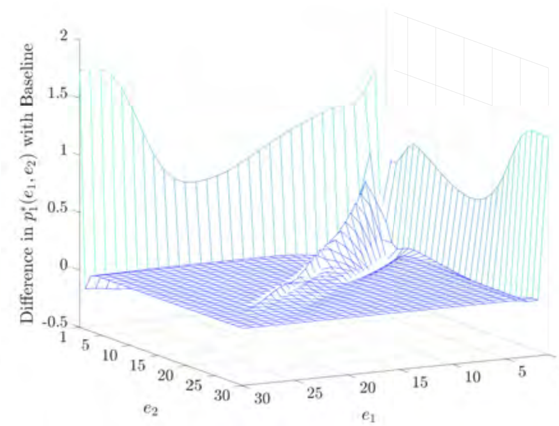

(d) No Leader Dynamic Incentives: Price Changes

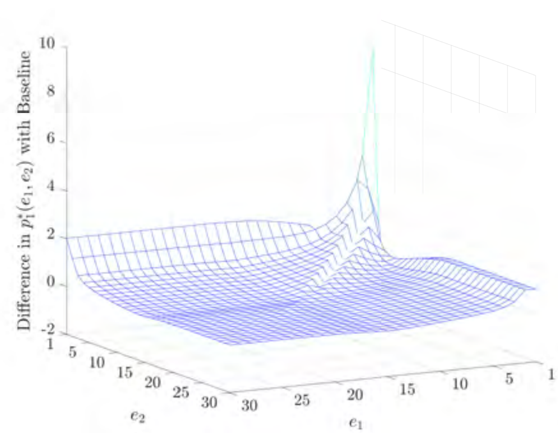

(e) No Leader AD Incentives: Price Changes

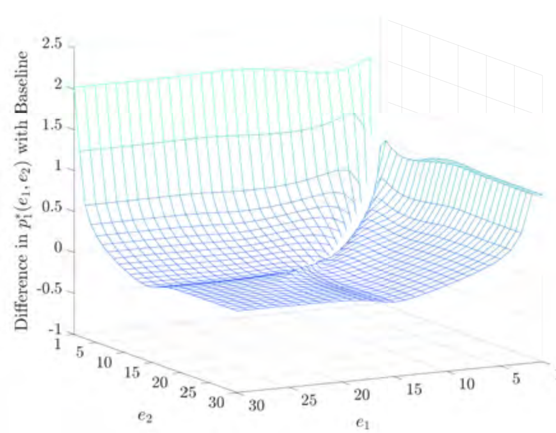

(f) No Leader AB/F and AD/F Incentives: Price Changes

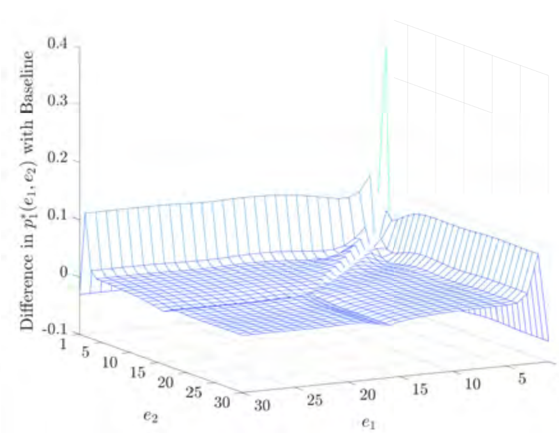

Figure 7: Effects of Policies on the Distribution of States After 32 Periods for the Illustrative Parameters When $b^p = 0$ and $\tau = 0$. For the incentive policies the figures show changes from the Low-HHI baseline equilibrium.

(a) Share Thresholds ($\chi = 50, \psi = 0.75$): State Changes for Low-HHI Eqm.

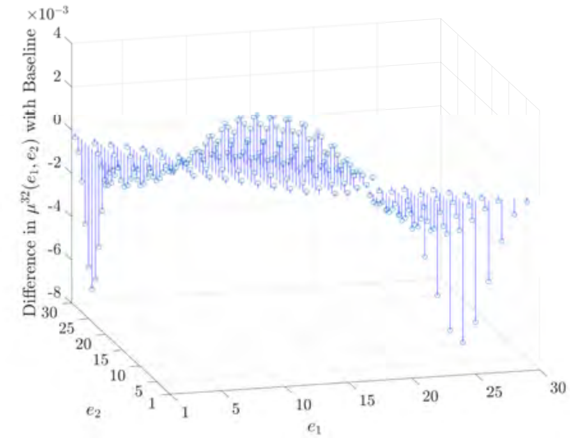

(b) Share Thresholds ($\chi = 50, \psi = 0.75$): State Changes for Mid-HHI Eqm.

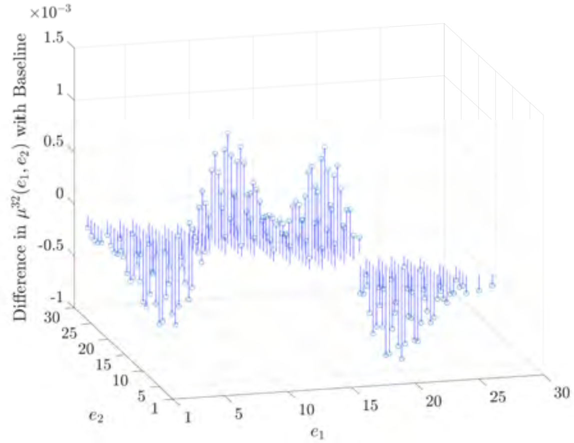

(c) Share Thresholds ($\chi = 50, \psi = 0.75$): State Changes for High-HHI Eqm.

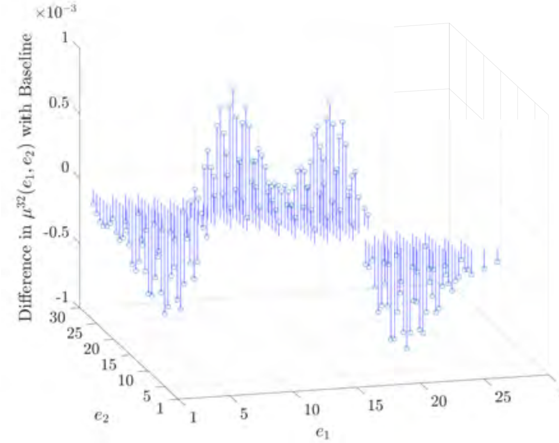

(d) No Leader Dynamics: State Changes

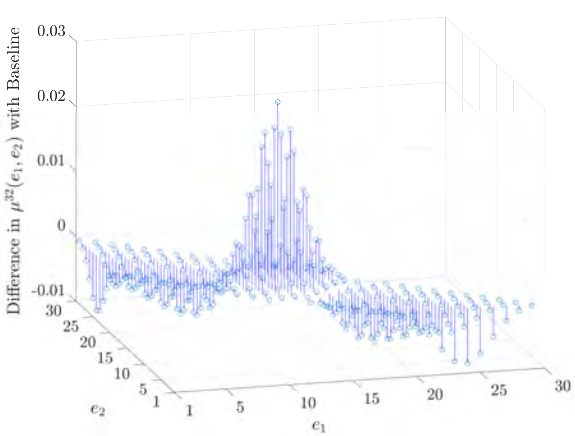

(e) No Leader AD Incentives: State Changes

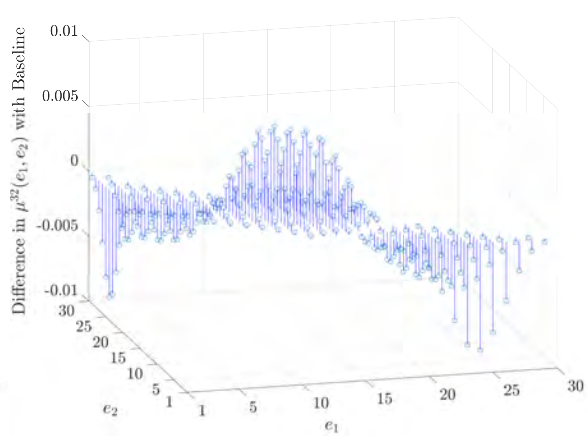

(f) No Leader AB/F and AD/F Incentives: State Changes

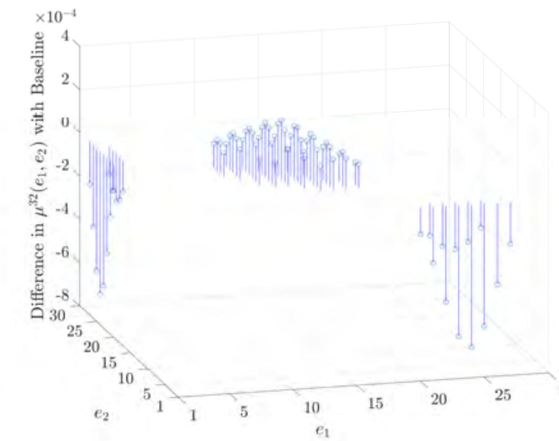

Price changes depend on the specific policy/equilibrium being considered, and the effect on prices in (1,1) varies depending on how the policy affects the ability of a leader to sustain its position and how much it softens competition. For our illustrative parameters, the threshold rules raise prices of both firms in the states where the firms are asymmetric, while partially removing the diagonal trench from the Mid- and High-HHI equilibria. Prices in (1,1) are lower, however, as the softening of competition in states that will be reached later in the game can increase the value to becoming the leader. Removing all of the leader's dynamic incentives raises the (1,1) price, as the leader will be more likely to lose its advantage, whereas eliminating only its AD incentive causes the (1,1) price to fall. Removing the AB/F and AD/F incentives reduces the depth of the sideways trench. This tends to raise the probability that a firm with no know-how will be able to catch a leader that has made several sales (albeit by a fairly small amount), which reduces the value of initial leadership, raising the (1,1) price.

Some intuition for why different flavors of incentive policies have such different, and non-uniform, effects on prices comes from Figure 8 which shows how the level of different dynamic incentives for firm 1 varies across states in the baseline equilibria. A positive number indicates an incentive that lowers firm 1's price in a particular state (of course, this depends on the equilibrium level of prices in other states), whereas a negative incentive contributes to a higher prices. The incentives vary in both sign and size (non-monotonically) across states and with the level of the buyer sophistication parameters. While not shown in these figures, the dynamic incentives can also be large when both firms are close to the bottom of their cost curves, explaining why the incentive policies can effect long-run surplus more than the threshold policies when $b^p = 0$ and $\tau = 0$.[16]

Figure 9 shows how the PDV welfare effects vary with the level of χ in this case: compliance costs are maximized in the Low-HHI equilibrium for a χ of around 30. As χ increases, the policy tends to hurt consumers much more than it lowers total surplus.

[16] For example, suppose that $e_1 = 20$ and $e_2 = 19$. In the High- and Mid-HHI equilibria the sum of firm 1's AB and AD incentives (which are roughly equal in size) equals 4, so that dynamic incentives imply an effective 4 unit reduction in marginal costs.

Figure 8: AB, AB/F, AD and AD/F Incentives for Firm 1 when $e_1 = 5$.

(a) AB Incentive

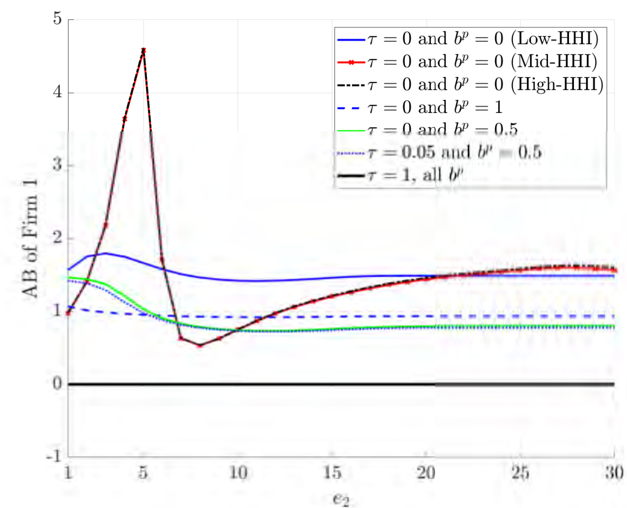

(b) AD Incentive

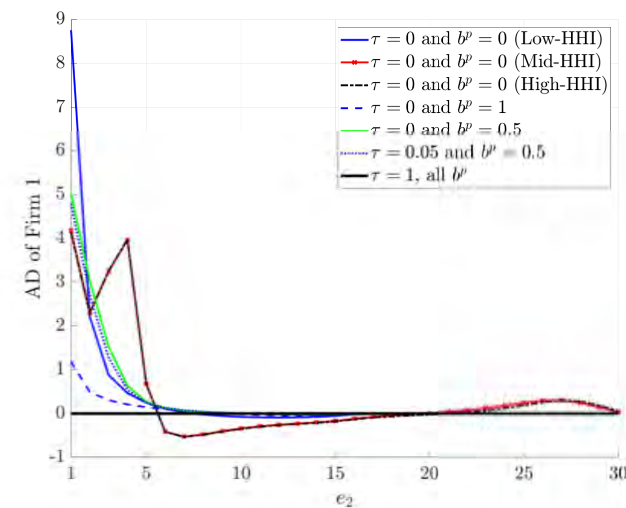

(c) AB/Forgetting Incentive

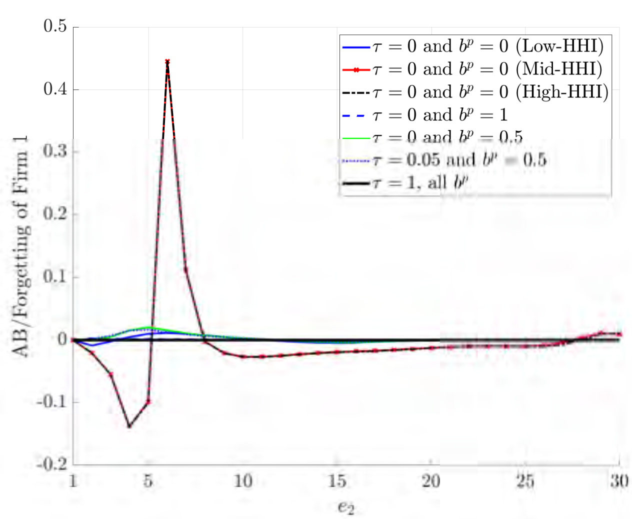

(d) AD/Forgetting Incenitve

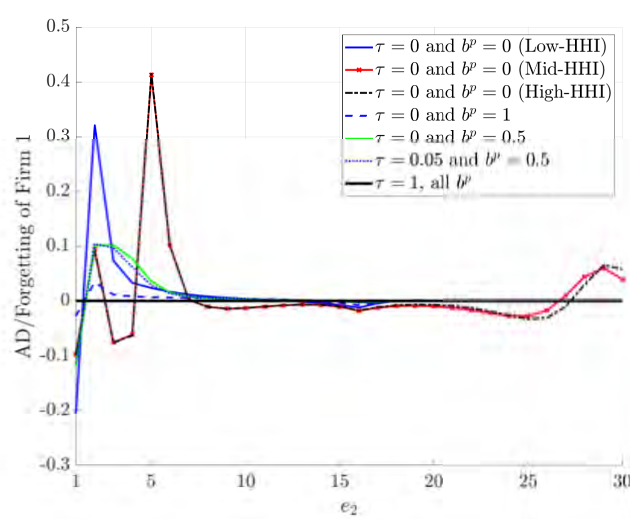

Figure 9: Effects of Changing The Compliance Cost Parameter (χ) For the Illustrative Parameters for $b^p = 0$ and $\tau = 0$. L=Low-HHI, M=Mid-HHI and H=High-HHI equilibrium, but the paths for M and H cannot be distinguished in the Figure.

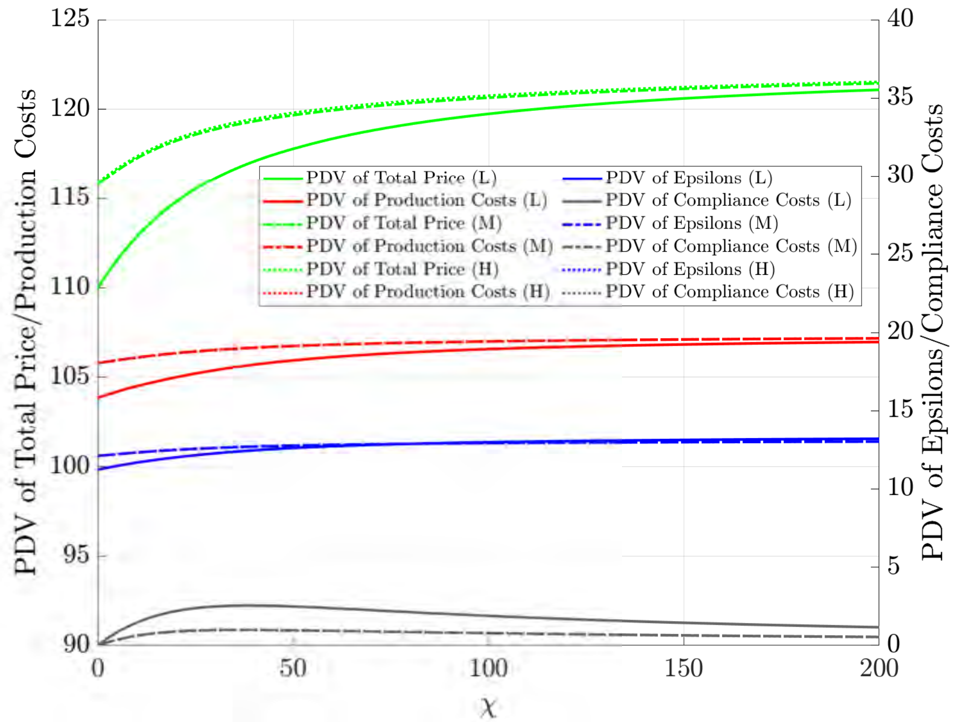

Policy Effects for the Remaining Polar Cases. When $\tau = 1$, the threshold rule increases long-run consumer surplus, as it becomes almost certain that both firms will be at the bottom of their cost curves in the long-run, but losses in CS at the start of the game are large. This is particularly true when $b^p = 0$ as atomistic buyers will tend to buy from the leader because they do not internalize how these choices tend to increase expected future compliance costs. Under the policy, atomistic buyers are actually better off when sellers have all of the bargaining power because forward-looking sellers will choose prices so that future compliance costs are limited.

When $b^p = 1$ and $\tau = 0$, there can also be a long-run gain from eliminating AB/F and AD/F incentives as, for these parameters, the net role of these incentives is to raise prices.[17] In contrast, the complete removal of the leader's dynamic incentives, or its AD incentive, can lead to a significant additional softening of competition, even though firms rarely have a significant lead for these parameters.

4 Bargaining and Equilibrium Market Structure and Welfare.

In this section, we focus on how equilibrium outcomes and the effects of policies change with τ, the measure of buyers' bargaining power. We fix $b^p = 0$ so all buyers act atomistically. We will discuss how τ and b^p interact in Section 6.

Our primary questions are: (i) are the relationships between bargaining power and market outcomes monotonic and linear? For example, are outcomes when $\tau = 0.5$ the approximate average of outcomes when $\tau = 0$ and $\tau = 1$? (ii) if buyer bargaining power can lead to excessive concentration, how does this affect the welfare effects, and incidence, of policies to restrict dominance? (iii) how does bargaining power affect the existence of multiple equilibria, which are a well-known feature of the seller price-setting model?

We begin by looking at the effects of bargaining on outcomes for the illustrative parameters, with and without policies, before considering how these effects vary as the values of ρ

[17] Recall for these parameters, buyers try to keep the sellers symmetric by favoring a laggard. Therefore the possibility that a laggard will forget can lower the expected future profits that a leader will make, so that eliminating the forgetting-related incentives can actually lead to slightly more aggressive pricing.

and δ change.

4.1 Effects of Bargaining Power on Outcomes for the Illustrative Parameters with No Policies.

Figure 10 shows the values of HHI^{32} and HHI^{1000} (panel a), the components of PDV welfare (b), and prices (c), sale probabilities (d) and AB/AD incentives (e) in state (3,2) along the equilibrium paths traced by τ-homotopies from the three baseline equilibria. Panel (f) shows the expected PDV of AB and AD incentives of sellers along the homotopy path that runs from the Low-HHI baseline equilibrium right across the diagrams, allowing the weights on each state to change to reflect the probabilities that the states are visited given equilibrium play, or keeping the weights fixed at the probabilities implied by the Low-HHI baseline ($\tau = 0$) equilibrium.

The homotopy paths from the Mid- and High-HHI baseline equilibria form a loop which does not extend past $\tau = 0.07$. For the illustrative parameters, we cannot identify more than a single equilibrium beyond this value of τ, which represents only a modest departure from the assumption that sellers set prices.

Consistent with the differences between outcomes when $\tau = 0$ and $\tau = 1$, HHI^{32}, HHI^{1000}, PDV CS and TS increase as τ increases from zero, as the leading firm tends to make more sales and its markups are constrained. However, the relationships between τ and outcomes are strikingly non-linear. Table 1 also shows that welfare outcomes when $\tau = 0.5$ are much closer to outcomes when $\tau = 1$, and buyers possess all of the bargaining power, than when $\tau = 0$. In some cases, outcomes are also non-monotonic in τ. For example, PDV TS is maximized for $\tau \approx 0.2$, and prices and sale probabilities in particular states, including (3,2), exhibit U-shaped or inverted U-shaped patterns.

The nonlinearities reflect the different and interacting static and dynamic effects of bargaining power on margins, incentives and market structure. To understand the forces at work, consider state (3,2) and the four panels of Figure 11. Holding continuation values fixed, the static effect of increasing τ is to lower sellers' markups towards the opportunity costs of sale implied by production costs and dynamic incentives (1.43 for firm 1 and 4.02 for firm 2 in the Low-HHI equilibrium). This effect is shown in panel (a).

Figure 10: The Effects of Changes in τ with $b^p = 0$ For the Illustrative Parameters.

(a) HHI^{1000}: τ homotopy

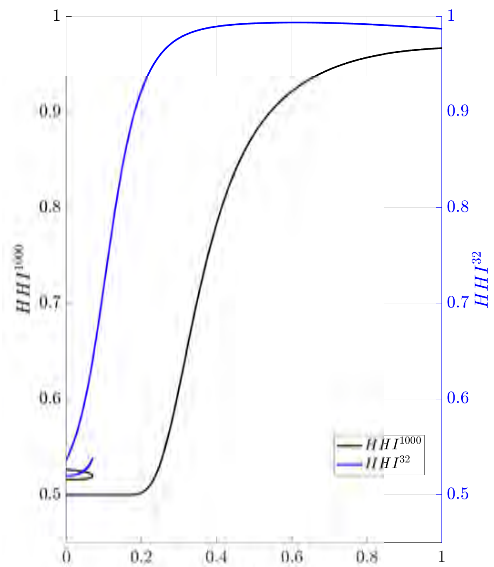

(b) PDV Welfare: τ homotopy

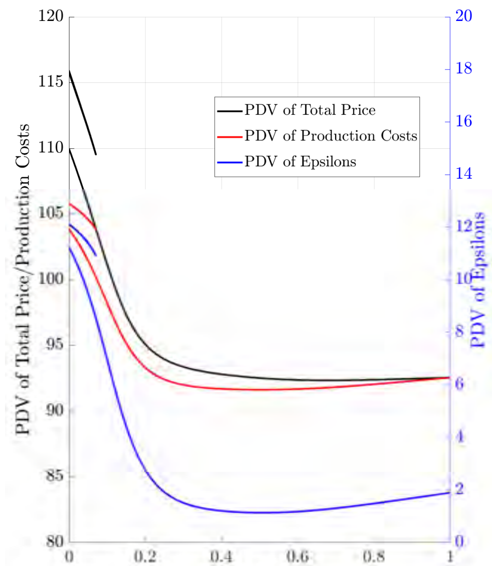

(c) (3,2) Prices: τ homotopy

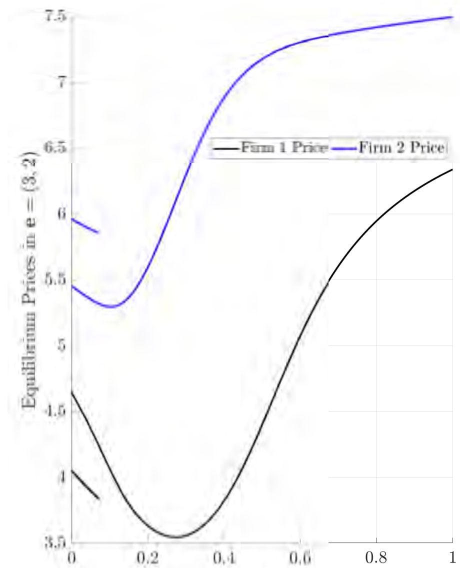

(d) Sale Probabilities: τ homotopy

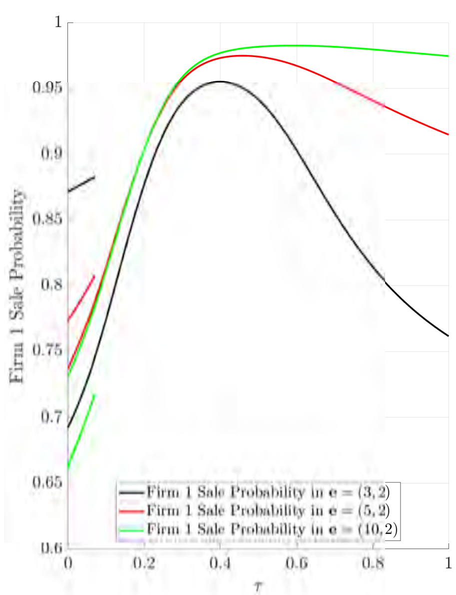

(e) AB/AD Incentives in (3,2): τ homotopy

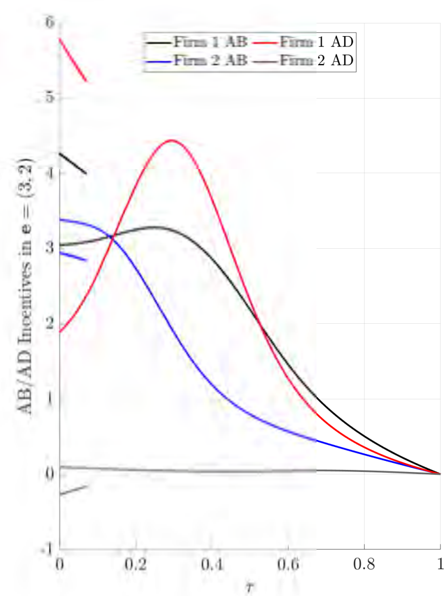

(f) Average PDV AB/AD Incentives

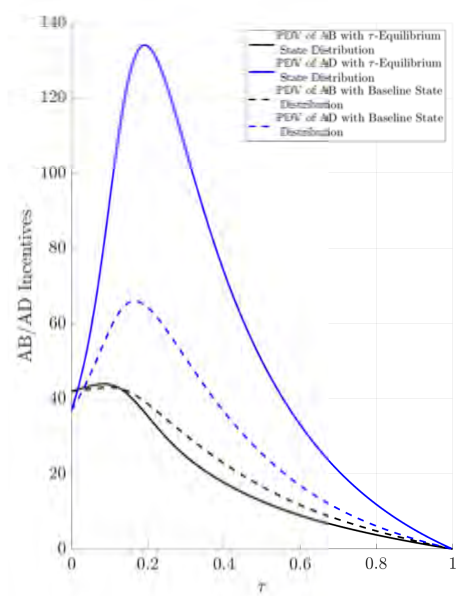

Figure 11: Various Effects on Prices in State (3,2) As a Function of the Bargaining Power Parameter (τ). L=Low-HHI, M=Mid-HHI, H=High-HHI, but M and H cannot be distinguished.

(a) Equilibrium τ Prices in (3,2) with Fixed (Baseline) Continuation Values

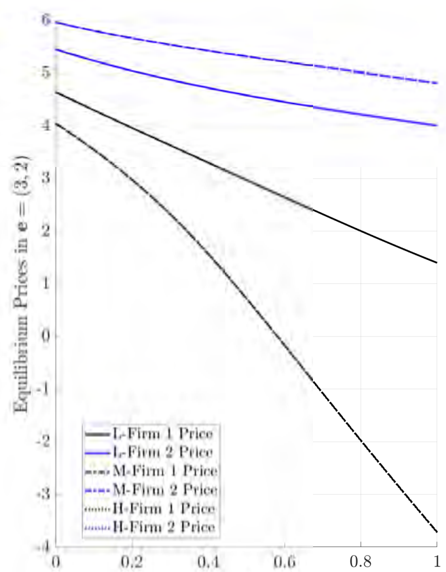

(b) $\tau = 0$ Prices in (3,2) with $(1-\tau)\times$ Baseline Continuation Values

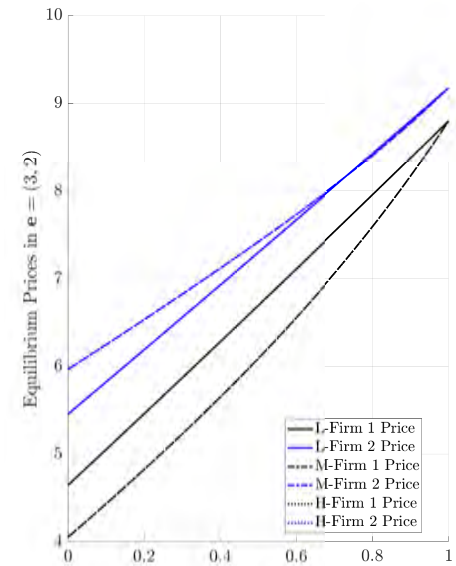

(c) $\tau = 0$ Prices in (3,2) with τ Homotopy Continuation Values

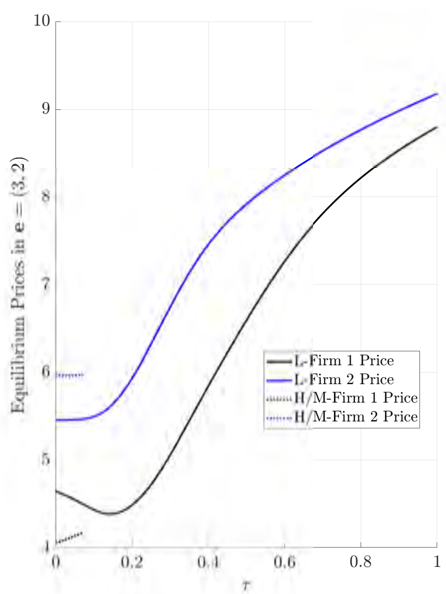

(d) Equilibrium Prices in (3,2): τ Homotopies

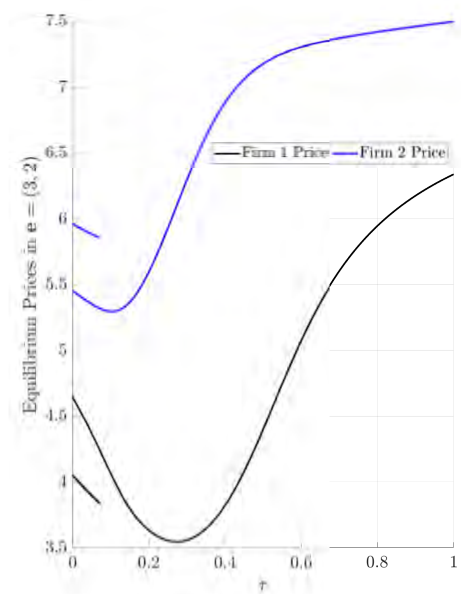

However, increasing τ also tends to lower sellers' expected profits from future play, by redistributing future surplus to the buyer. This will tend to reduce dynamic incentives so that opportunity costs move towards marginal production costs. Dynamic incentives are eliminated completely when $\tau = 1$. Figure 11(b) shows the (3,2) prices that sellers would set (i.e., if the in-period τ was zero) when the baseline ($\tau = 0$) equilibrium continuation values are multiplied by $(1 - \tau)$, where τ is on the horizontal axis. In this case, $\tau = 1$ is associated with static duopoly pricing. The changes in prices in both panels (a) and (b) are approximately linear in τ, but they are in different directions.

There is an additional effect caused by how changes in bargaining power affect which states the industry will visit in the future. In particular, the reduction in margins will have the greatest effect on the firm with the largest margin, so that its probability of sale will tend to increase. This firm will tend to be the leader, so that an increase in τ will tend to increase the leader's cost advantage, and, thereby, tend to increase a leader's future profits, and the leader's incentive to preserve its leadership and the incentives of both firms to become a leader (illustrated by the changes in sale probabilities shown in Figure 10(d) and to AB and AD incentives on the homotopy paths in Figure 10(e)). By lowering (the leader's) opportunity costs, this effect can tend to lower prices. Figure 11(c) shows what happens to the $\tau = 0$ prices that sellers would set in state (3,2) when we use the continuation values implied by the continuation values from the τ-homotopy paths. The fact that these prices are flat or initially declining as τ increases from zero, compared with the sharp increases in panel (b), reflects how the increase in a leader's expected future profits, due to its likely increasing advantage can offset or dominate how buyer bargaining power would otherwise tend to reduce dynamic incentives. As a result, firms' AB and AD incentives are non-monotonic in τ in Figure 10(e). Putting the three forces together produces the equilibrium (3,2) price paths shown in Figure 11(d) (identical to Figure 10(c)).

Of course, the balance of these three factors can vary across states. However, the fact that, on average, states where AB and, especially, AD incentives are large or growing are visited more as τ increases from zero to 0.2 is very clear in Figure 10(f), and is consistent with how AD incentives tend to be strong for a firm that has a small but significant lead.

These different effects also provide intuition for why the Mid-HHI and High-HHI equilibria are eliminated as τ increases. In these equilibria, the probability that the leader would sustain its advantage was already higher, because of the diagonal trench, which was sustained by the anticipation that the leader would make larger margins in states where it had a significant advantage. As τ increases, the larger effect is therefore through the reduction in the leader's expected future surplus, so that AB and AD incentives tend to fall, and it becomes impossible to sustain the low diagonal trench prices.

4.2 Effects of Bargaining Power on Counterfactual Policies for the Illustrative Parameters.

The fact that bargaining power has nonlinear effects on outcomes suggests that the effects of policies may also vary with bargaining power in subtle ways. This is also suggested by some of the effects of policies on outcomes when $\tau = 0.5$ are not close to the average of the effects when $\tau = 0$ and $\tau = 1$ (Table 2).

Figure 12 shows the relationship between τ and concentration, PDV CS and PDV TS under the baseline and the alternative policies (for the threshold rule we also consider a policy with a lower compliance cost).

The effects of all policies on medium-run concentration are greatest when $\tau \approx 0.2$, reflecting how, absent any policy, concentration rises quickly as buyers gain bargaining power, while the policies lead to greater seller symmetry. However, the no leader AD incentive policy raises HHI^{32} for high τ. The effects on surplus are more policy-specific. The threshold policies lower total and consumer surplus, with the largest effects for $\tau \approx 0.5$ for the milder policy, and $\tau \approx 0.7$ for the stricter policy. On the other hand, the incentive policies reduce surplus the most when τ is relatively low, but the elimination of leader AD incentives can raise PDV CS and TS surplus when buyers have most of the bargaining power.[18,19]

To help explain how removing the leader's AD incentives can increase surplus for high τ, Figure 13 shows no policy prices and the changes in prices caused by this policy, when $\tau = 0.5$ (when there is a loss in PDV CS and TS) (panel (c)) and $\tau = 0.9$ (small gains)

[18] The removal of the leader's AB/F and AD/F incentives can also lead to small gains in CS for high τ, and to small gains in the PDV of producer surplus for $\tau \leq 0.45$.

[19] We have also analyzed policies where, under the policy, the incentives of symmetric firms are also restricted. The effects of these policies are generally similar, although, under the removal of AD incentives policy, there is only a gain in PDV CS for $\tau \geq 0.94$.

Figure 12: Effects of Threshold and Incentive Policies on Medium-Run Concentration and Surplus for the Illustrative Parameters as a Function of τ. $b^p = 0$.

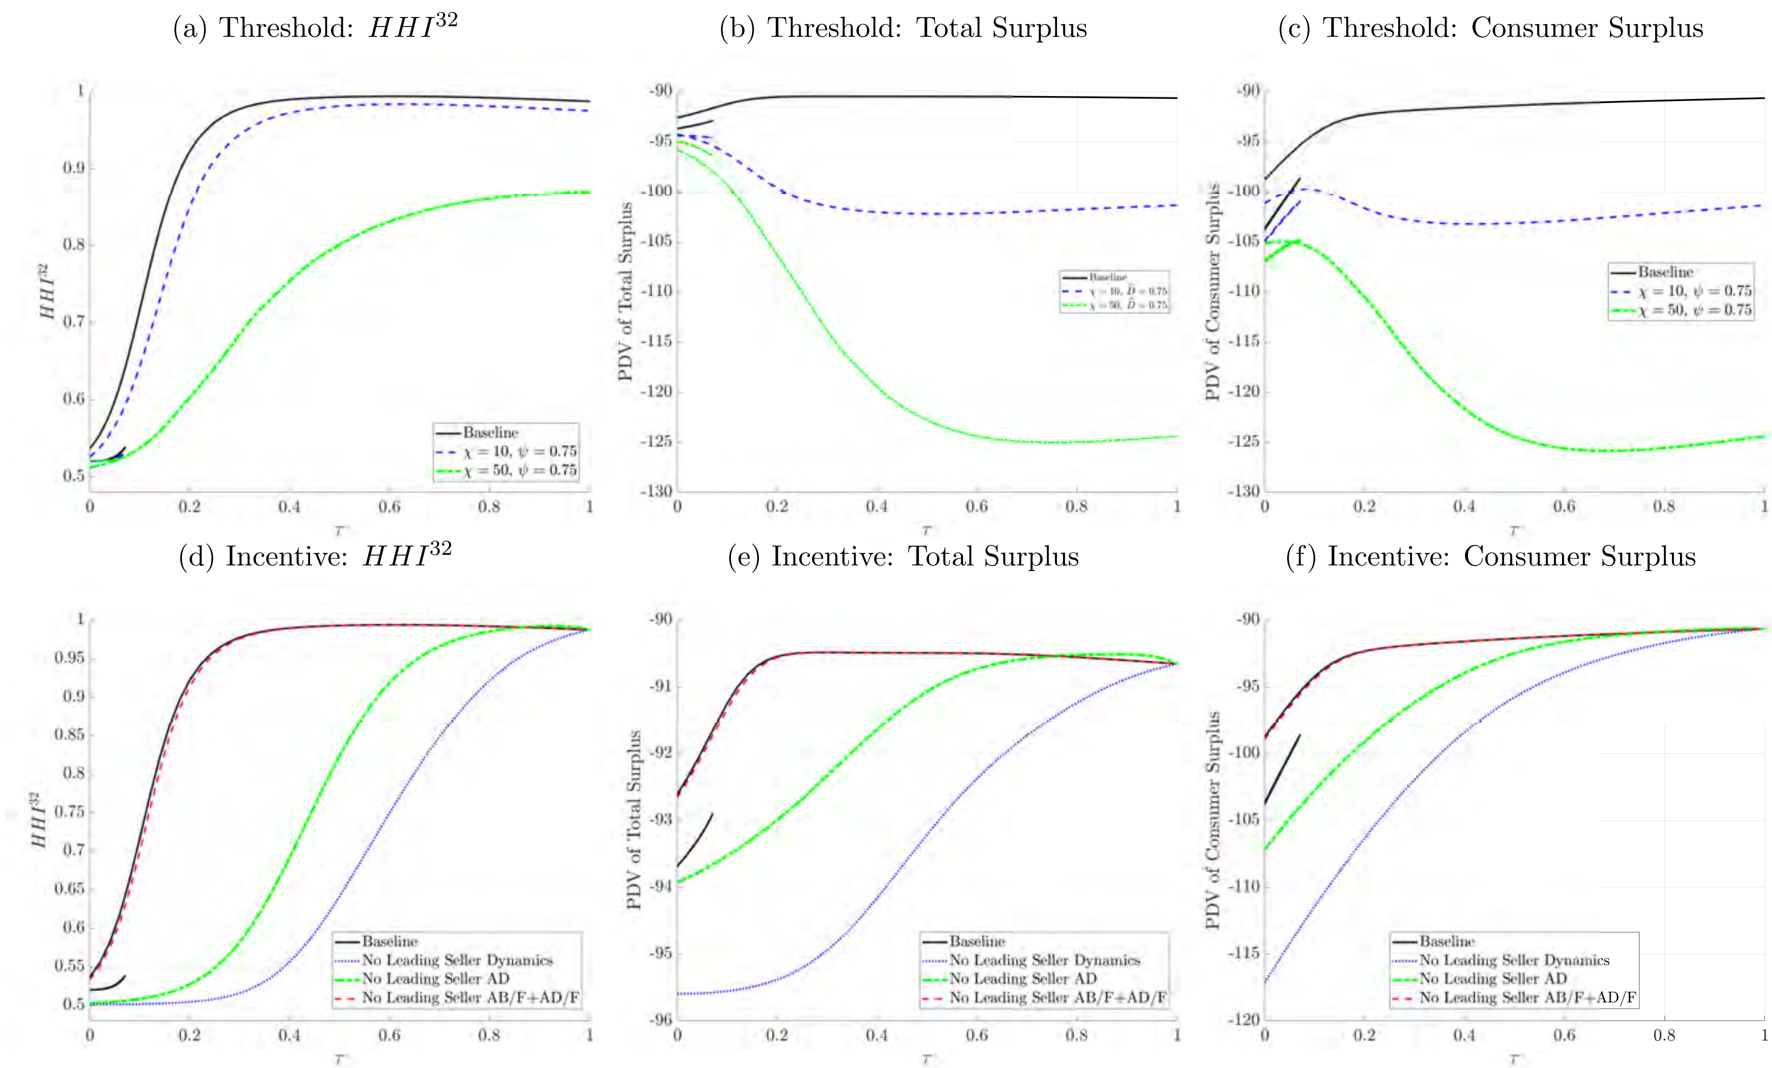

(panel (d)). With no policy in place, the variation in prices across states reflects marginal production costs, plus a small markup, quite closely and there is no pronounced well for the initial states, indicating that there is limited competition to become a leader.

However, when the leader cannot consider AD incentives, prices fall in the states at the top of the cost curves. The drop in prices covers a broader set of states when $\tau = 0.9$, and, in this case, the price decreases at the start of the game more than offset the increase in prices in later periods where one of the firms has a large lead. As noted, the policy actually increases HHI^{32} when $\tau = 0.9$, so that the benefit of the policy comes not from bringing sustained symmetric competition (the likely intention of such a policy), but because competition to become a leader becomes more intense.

Given the role of AD incentives in driving low pricing when $\tau = 0$, one might be surprised that the elimination of a leader's AD incentives could cause initial prices to fall relative to the no policy case. This happens because, in equilibrium, other incentives, including AB incentives, can increase when AD incentives are eliminated. As an illustration, Figure 14 shows equilibrium AB and AD incentives across states when $e_1 = 1$ and $\tau = 0$, $\tau = 0.5$ or $\tau = 0.9$. Under the no leader AD incentive policy, the AB incentives switch from the black solid line to the blue solid line, and, for the follower (or symmetric firm) AD incentives switch from the black dashed to the blue dashed line. The red dashed line indicates what the leader's AD incentive would be if it was allowed to be considered (and prices take their policy-based equilibrium values). When $\tau = 0$, the policy has only small effects on AB incentives, whereas it can increase AB incentives significantly when $\tau = 0.5$ or $\tau = 0.9$.

Figure 12 assumed that the PDV of surplus provided the relevant metric for evaluating policies, recognizing that softening of competition later in the game plays an important role in pricing in the initial periods where leadership may be established. Bargaining power also affects the expected impact of policies in either the medium-run (e.g., 32 periods) or the long-run (e.g., 1,000 periods) in non-linear ways. As one illustration, the calculations for polar cases show that, when $b^p = 0$ and $\tau = 1$, long-run market concentration tends to be higher than the social planner would choose because future buyers would benefit (in terms of εs) from a second low-cost provider. This is also true when τ is moderately high: for example, Table 1 shows that HHI^{1000} is 0.877 when $b^p = 0$ and $\tau = 0.5$, compared to 0.613

Figure 13: No Policy Prices and Changes in Prices From a Policy where the Leader Cannot Consider AD Incentives for $\tau = 0.5$ and $\tau = 0.9$. $b^p = 0$. The shown changes are changes in prices from the equilibrium with the same τ but with no policy.

(a) No Policy Prices: $\tau = 0.5$

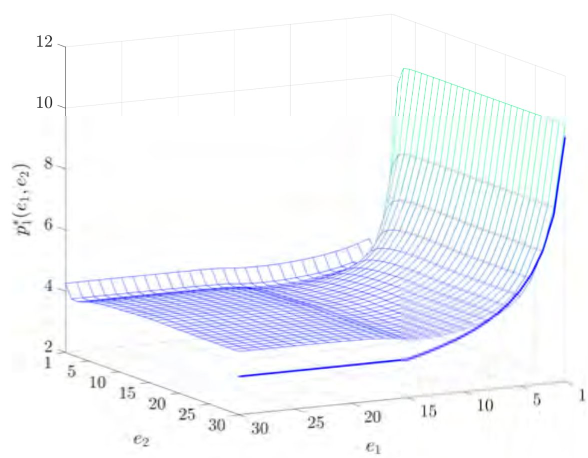

(b) No Policy Prices: $\tau = 0.9$

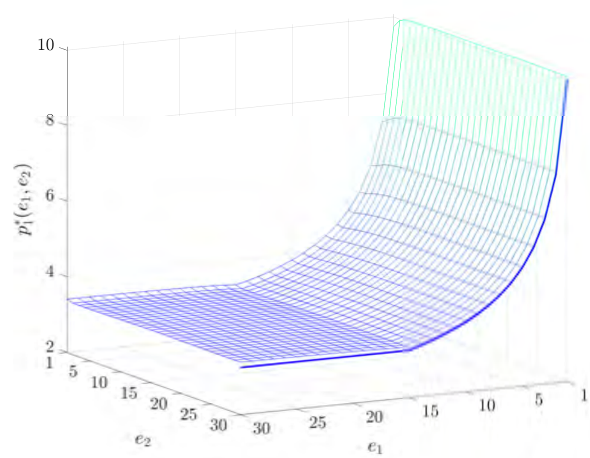

(c) Changes in Price: $\tau = 0.5$

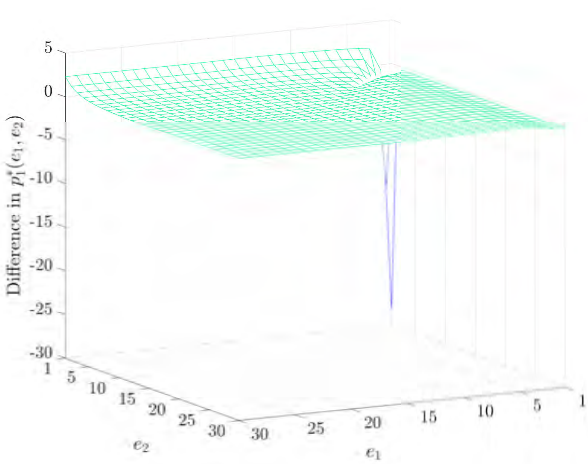

(d) Changes in Price: $\tau = 0.9$

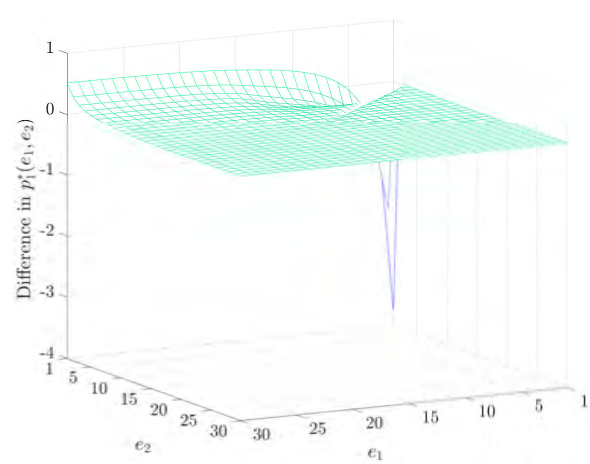

Figure 14: Sellers' AB and AD Incentives when $e_1 = 1$ for the Illustrative Parameters and $\tau = 0$, $\tau = 0.5$ or $\tau = 0.9$

(a) Firm 1 $\tau = 0$

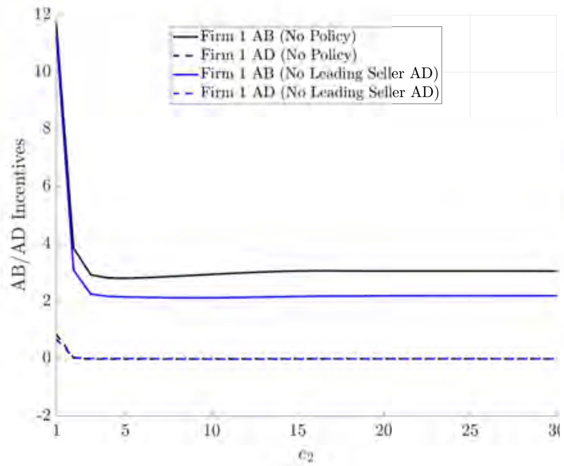

(b) Firm 1 $\tau = 0.5$

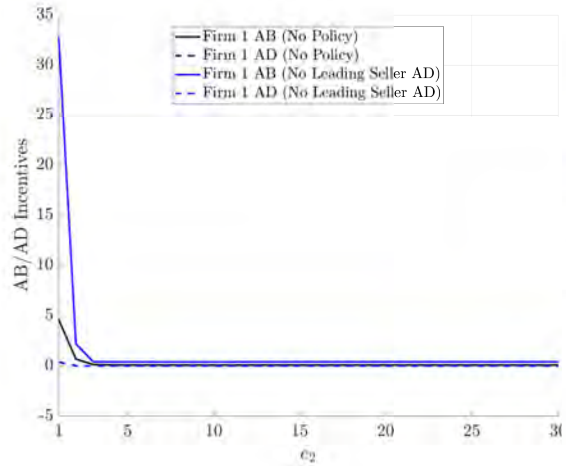

(c) Firm 1 $\tau = 0.9$

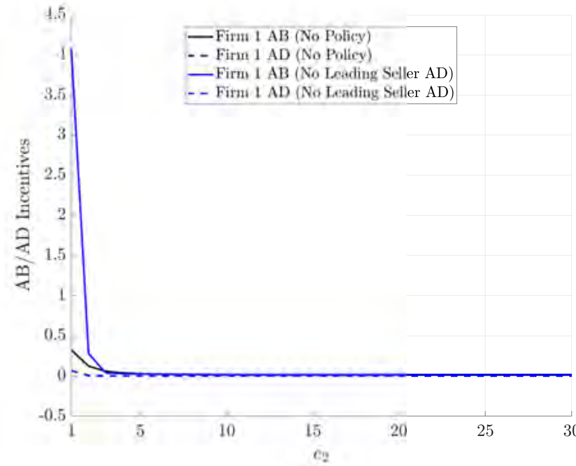

(d) Firm 2 $\tau = 0$

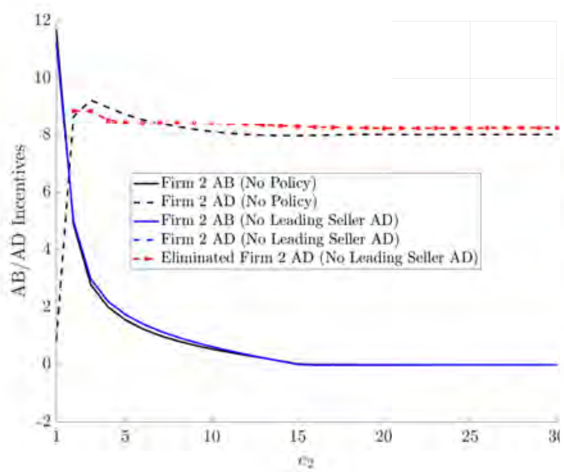

(e) Firm 2 $\tau = 0.5$

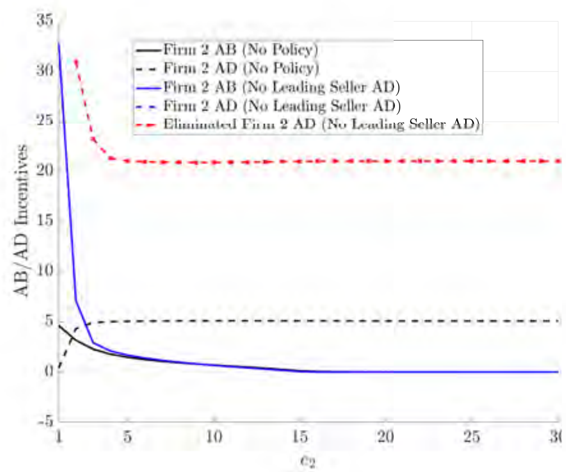

(f) Firm 2 $\tau = 0.9$

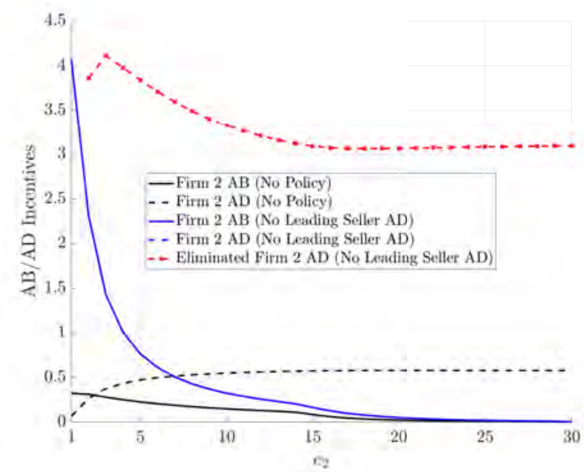

under the social planner. Therefore, policies that make it more likely that there will be two low-cost providers in the long-run, as the threshold policies do for all τ and the incentive policies do when $\tau < 1$, could potentially raise long-run welfare quite significantly.

Figure 15 shows the non-linearity of these effects for the incentive policies, plotting how HHI^t, CS^t and TS^t vary with τ for $t = 32$ and $t = 1000$ with no policies and the three leader incentive policies. Given the probabilities of forgetting implied by the illustrative δ, it is likely that both firms will be close to the bottom of their cost curves after 32 periods if they split sales relatively equal, which will be the efficient configuration looking at this period in isolation. This will also, obviously, be true after 1,000 periods. Therefore, when the policies lead to low expected concentration, but concentration would be high with no policy, the no leader dynamic and no leader AD policies can raise expected CS and TS significantly. For example, when $\tau = 0.3$, all four welfare measures are higher when the leader ignores dynamic incentives. On the other hand, when policies do not ensure symmetry and have only modest effects on concentration, as is the case for example when $\tau \geq 0.8$, consumers can be hurt, in the medium- and long-run, by the softening of competition that the restriction on incentives apply. Of course, when $\tau = 1$, incentive policies have no effect. Appendix Figure C.8 shows that these nonlinear and non-monotonic effects on welfare in the medium- and long-run are also seen for the threshold policies.

4.3 Effects of Bargaining Power on Outcomes for Alternative ρ and δ.

An obvious question is whether the effects that bargaining power has for the illustrative parameters are qualtiatively similar for alternative learning and forgetting parameters. In this sub-section, we vary ρ and δ in turn, before discussing what we see when we vary them together. Our figures will focus on values where $\rho \geq 0.6$ and $\delta \leq 0.2$ as these are the most empirically relevant ranges, although we will mention what happens for other values, with results in Appendix Figure C.2, to help illustrate the economic forces at play. When δ is greater than 0.0452 it is impossible for two firms to remain at the bottom of their cost curves in the long-run, in the sense that the probability of forgetting when $e_i = m$ is more than 0.5.

Figure 16 (spread over multiple pages) show ρ and δ homotopy paths, for alternative

Figure 15: Medium-Run and Long-Run Changes in Concentration and Expected Surplus as a Function of τ and Incentive Policies. Illustrative Parameters. $b^p = 0$.

(a) HHI^{32}

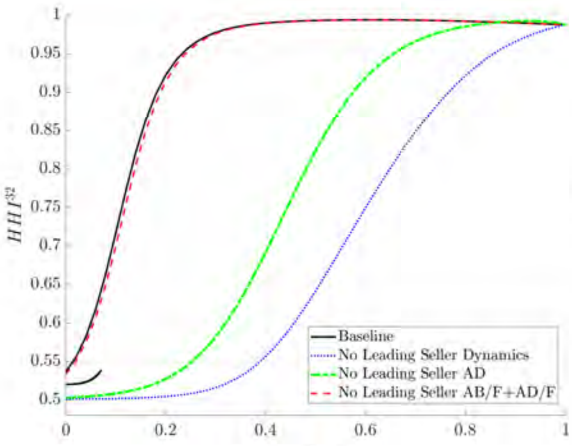

(b) CS^{32}

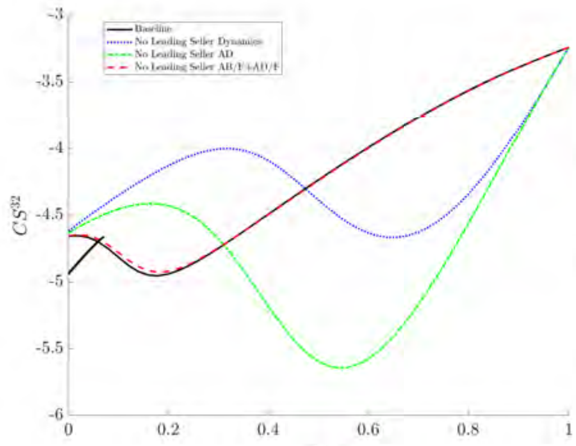

(c) TS^{32}

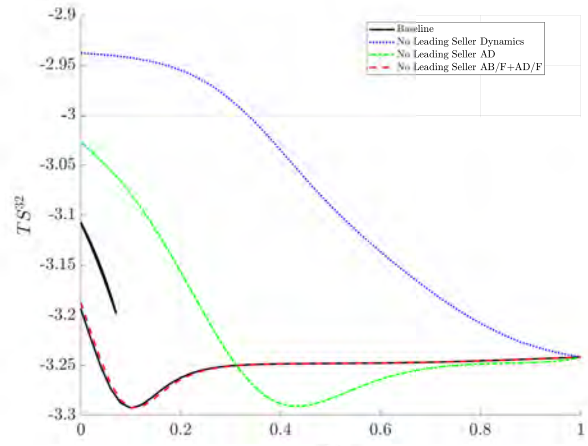

(d) HHI^{1000}

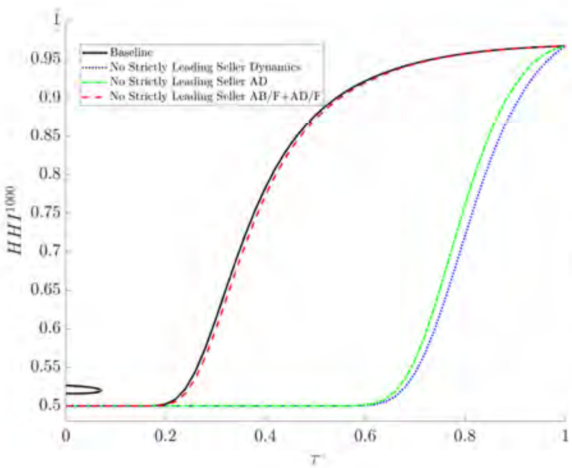

(e) CS^{1000}

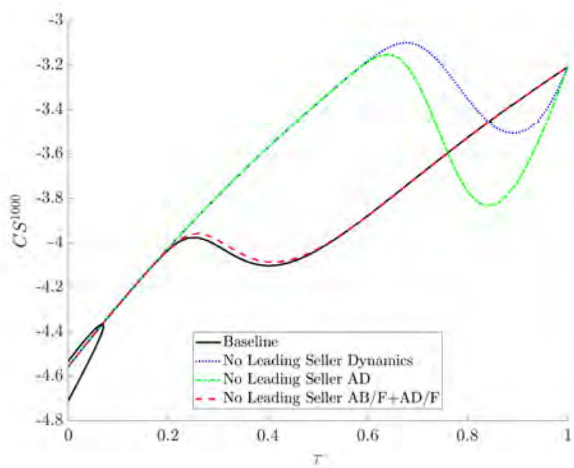

(f) TS^{1000}

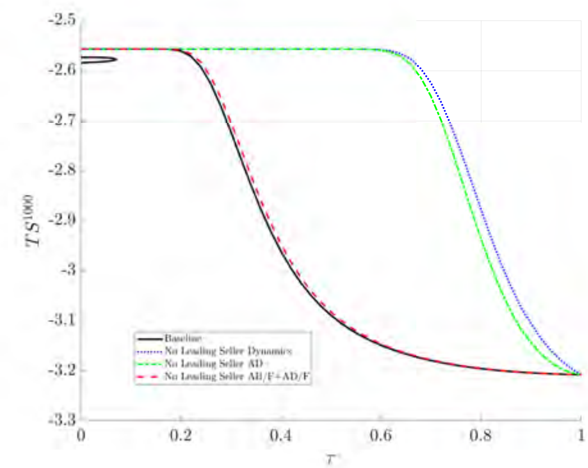

Figure 16: HHI^{32} for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. The thick black lines indicate HHIs implied by the social planner outcome.

(a) No Policy ρ Homotopies

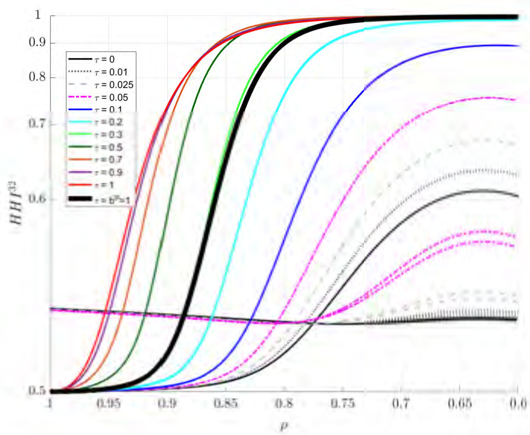

(b) No Policy δ Homotopies

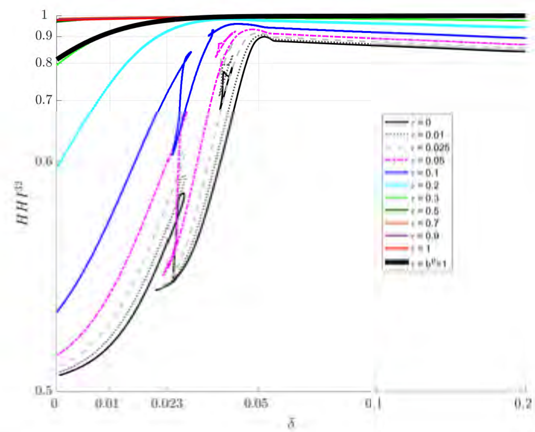

(c) $\chi = 50, \psi = 0.75$ ρ Homotopies

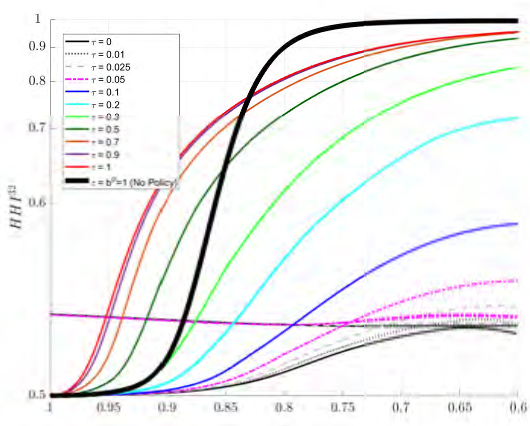

(d) $\chi = 50, \psi = 0.75$ δ Homotopies

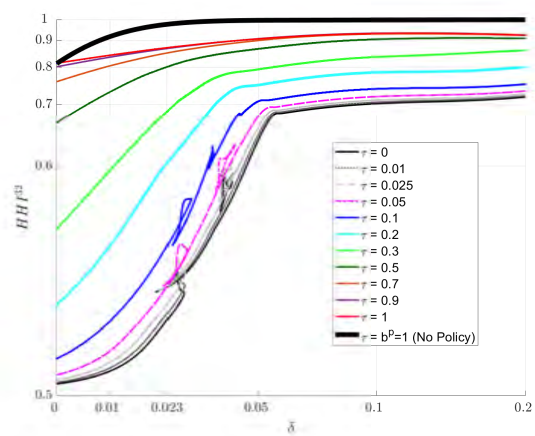

(e) No Leader AD Incentive ρ Homotopies

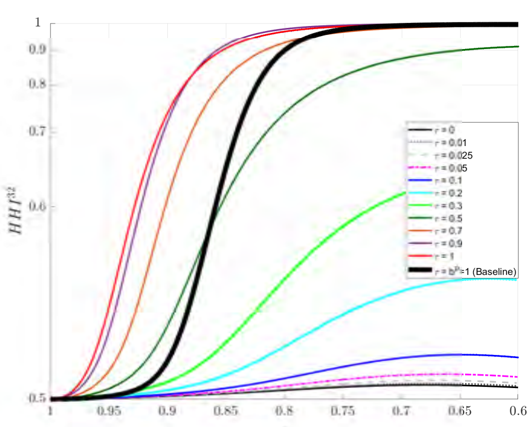

(f) No Leader AD Incentive δ Homotopies

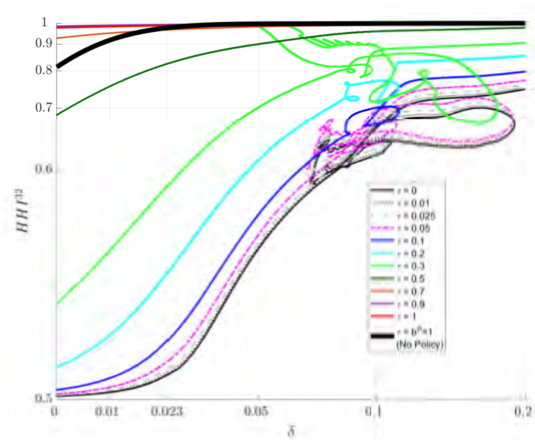

Figure 16: HHI^{1000} for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. The thick black lines indicate HHIs implied by the social planner outcome.

(g) No Policy ρ Homotopies

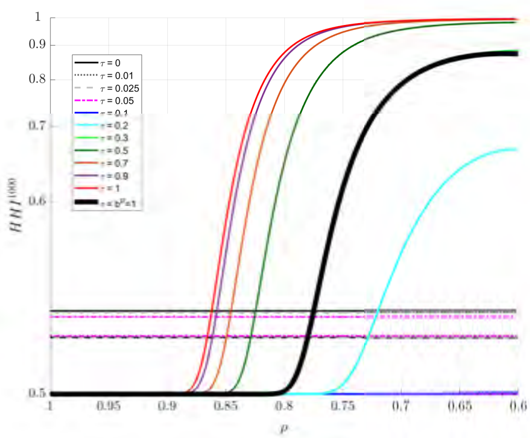

(h) No Policy δ Homotopies

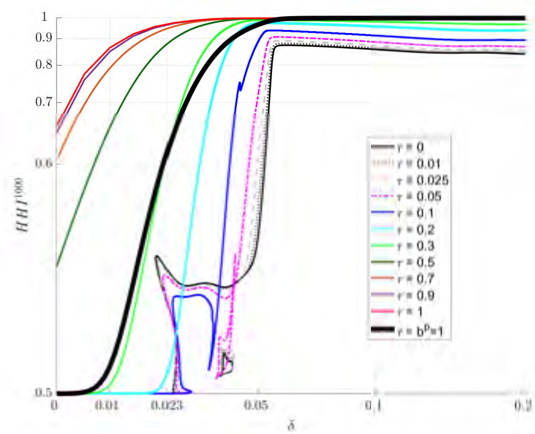

(i) $\chi = 50, \psi = 0.75$ ρ Homotopies

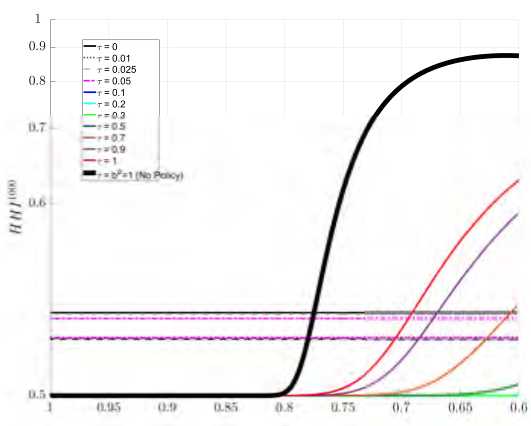

(j) $\chi = 50, \psi = 0.75$ δ Homotopies

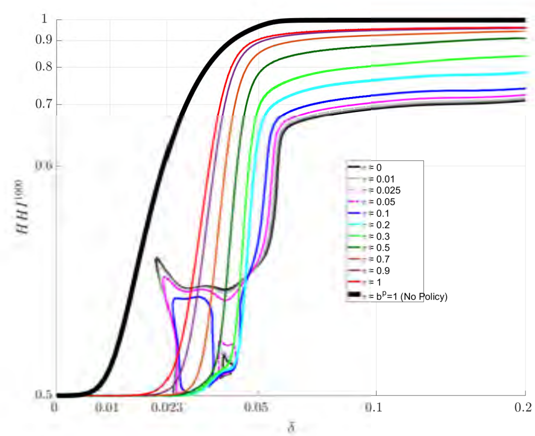

(k) No Leader AD Incentive ρ Homotopies

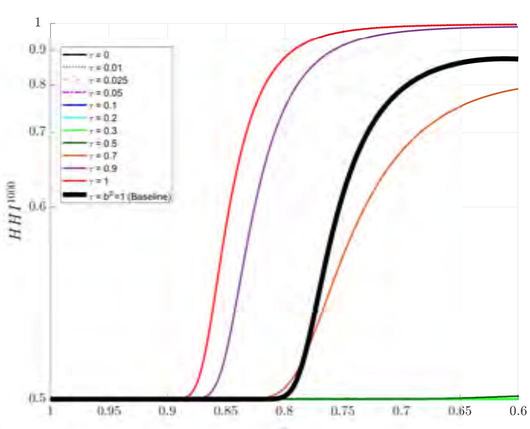

(l) No Leader AD Incentive δ Homotopies

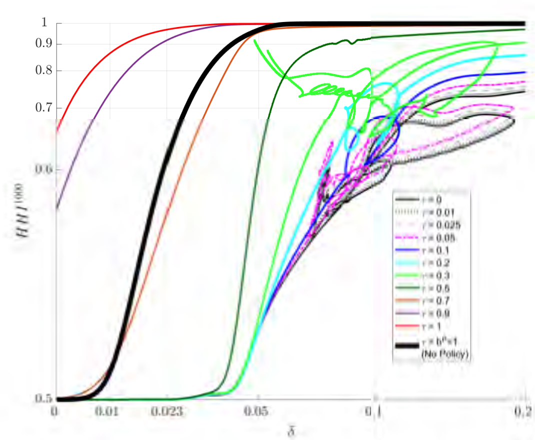

Figure 16: PDV CS for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. For the policies, we show the changes in CS caused by policy. The no policy plots show welfare relative to the social planner outcome for the same parameters.

(m) No Policy ρ Homotopies

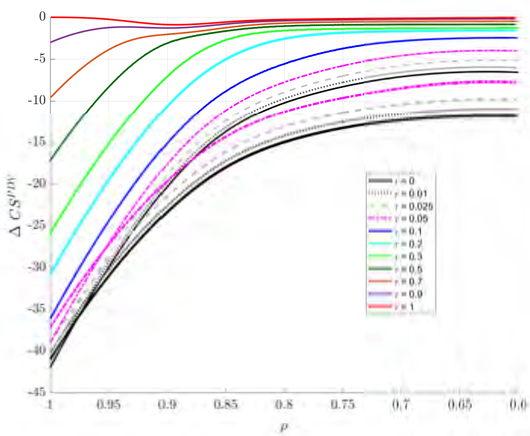

(n) No Policy δ Homotopies

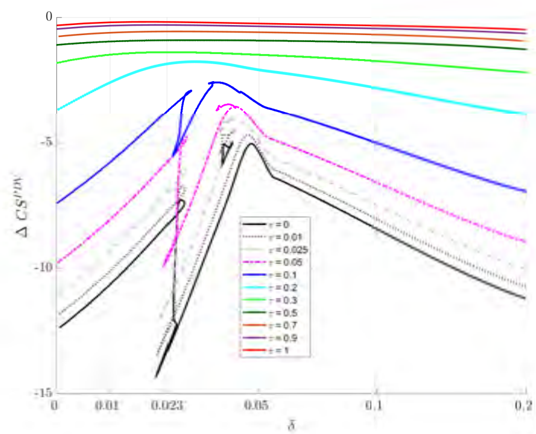

(o) $\chi = 50, \psi = 0.75$ ρ Homotopies

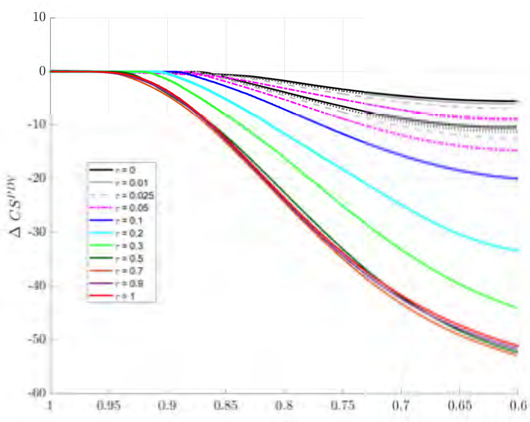

(p) $\chi = 50, \psi = 0.75$ δ Homotopies

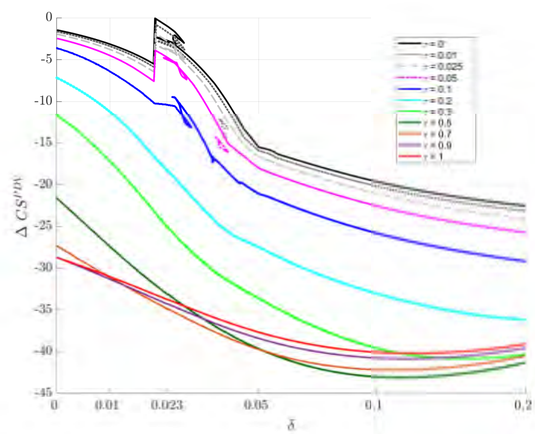

(q) No Leader AD Incentive ρ Homotopies

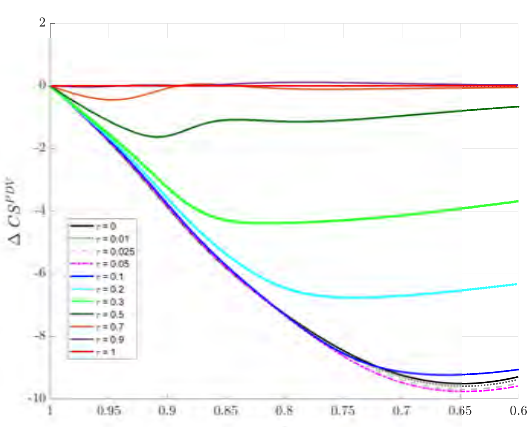

(r) No Leader AD Incentive δ Homotopies

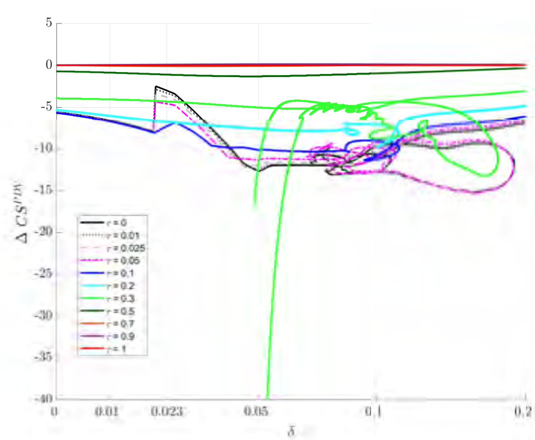

Figure 16: PDV TS for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. For the policies, we show the changes in TS caused by policy. The no policy plots show welfare relative to the social planner outcome for the same parameters.

(s) No Policy ρ Homotopies

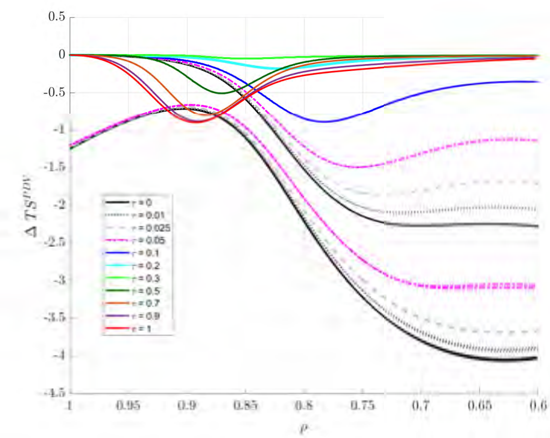

(t) No Policy δ Homotopies

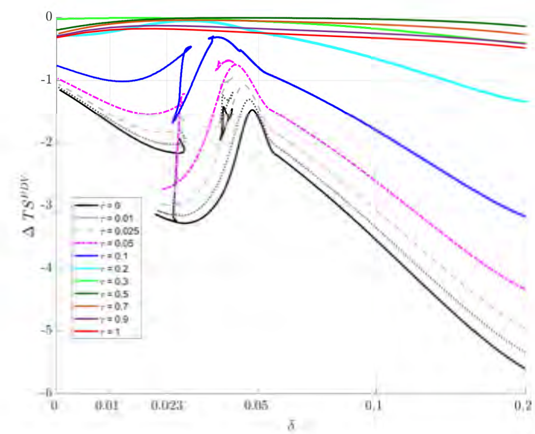

(u) $\chi = 50, \psi = 0.75$ ρ Homotopies

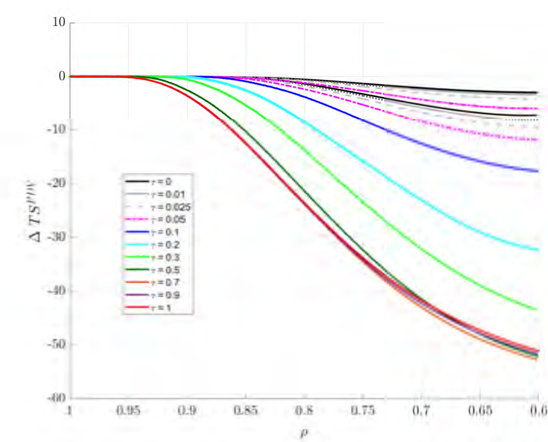

(v) $\chi = 50, \psi = 0.75$ ρ Homotopies

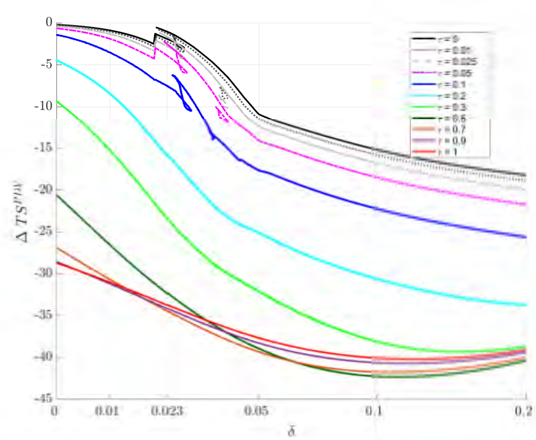

(w) No Leader AD Incentive ρ Homotopies

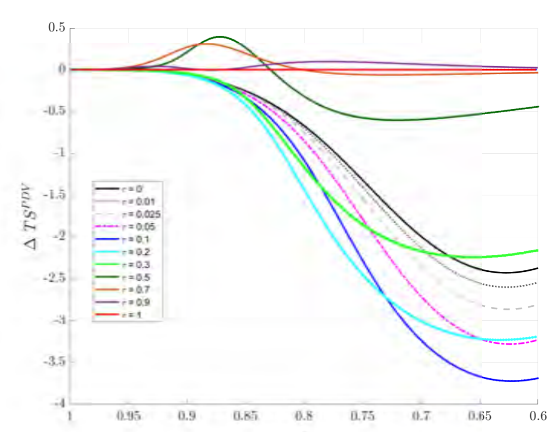

(x) No Leader AD Incentive δ Homotopies

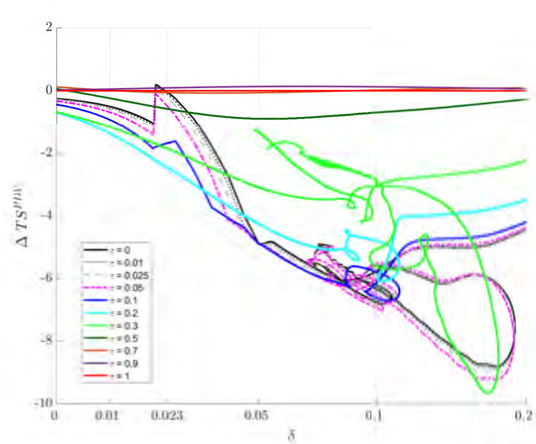

values of τ. The top figures on each page show outcomes when there are no policies, and the lower figures show outcomes (for HHIs) or changes in outcomes (CS and TS) under the threshold and removal of the leader AD's incentive policies.[20] We include some τ values close to zero so that we can understand how small deviations from a "sellers set prices" assumption affect the results. The no policy plots show welfare relative to the social planner outcome for the same parameters. The concentration plots also show the expected HHIs that the social planner solution would generate.

Examining the no policy figures, we see that there are always unique equilibria for $\tau \geq 0.2$, consistent with the illustrative parameter result. In this sense, buyer bargaining power eliminates multiplicity. However, close inspection of the δ homotopy plots reveals that there are small ranges of δ for which there is a single equilibrium when $\tau = 0$, but multiple equilibria when $\tau = 0.05$ or $\tau = 0.1$.[21]

A second feature of the HHI no policy plots is that, across δ and ρ values, concentration when $\tau = 0.3$ (and $b^p = 0$) is very similar to the level of concentration in the social planner solution, in both the medium-run and the long-run. While similar expected concentration does not imply socially optimal choices in every state, the PDV TS panels show that $\tau = 0.3$ also tends to produce the smallest, or one of the smallest, levels of inefficiency. Considering parameters outside the empirically relevant rages, the efficiency of $\tau = 0.3$ also holds for lower values of ρ (more LBD), but it does not hold when we consider very high probabilities of forgetting.[22]

For the illustrative parameters, and ignoring the complicating issue of multiplicity, buyer bargaining power increases PDV CS and long-run concentration monotonically, while medium-run concentration and PDV TS increases as τ rose from 0 to intermediate values of τ before falling slightly. We see most of the same patterns for concentration across values of δ and ρ,

[20]For the δ homotopies, the reported changes are based on taking the lowest value of PDV CS and TS under no policy as otherwise the figures become too messy. In this sense, the figures reflect the "best" case for policies.

[21]For example, there are several switchbacks in the blue $\tau = 0.1$ line for values of δ where there are no loops or switchbacks for smaller, or higher, values of τ.

[22]With $\delta > 0.5$, the social planner would focus its sales on a single firm in order to try to keep that firm at least one position below the top of the cost curve, whereas $\tau = 0.3$ tends to produce less concentration as a firm with a small lead will charge a significant markup in order to extract value from its (likely temporary) leadership position.

but there are some differences in the welfare effects.

For the very highest values of ρ, where supply-side dynamics are not so important, the relationship between PDV CS and τ is fairly linear, as one would expect in a static model. However, for ρs between 0.85 and 0.9 (which may be quite common empirically), atomistic buyers that have most of the bargaining power are too likely to buy from the leader in the sense that, even though this lowers production costs, by a relatively small amount because LBD is limited, the expected value of future εs is reduced by the resulting asymmetry in firm costs. This logic also applies when ρ is small because only one or two purchases from a laggard may be needed to produce a second supplier with a very low cost (costs fall very quickly with know-how because LBD is so high), but atomistic buyers ignore the benefits that this would create, and with low margins, because of bargaining power, they buy from the leader.

On the other hand, for intermediate ρ, the large markups set by the leading seller when the buyer lacks bargaining power tend to lead to inefficiently slow industry cost reductions, creating a welfare loss that more than offsets the gain in εs from more equal sales probabilities.[23] For these intermediate values of ρ, outcomes when τ is very high come close to being socially efficient. Similar to what was seen for the illustrative parameters, when $\tau = 0.5$, outcomes are (usually) quite similar to those when $\tau = 1$.

For very low values of δ, PDV TS increases rapidly in τ for low τ, but it is maximized around $\tau = 0.3$. The logic is that, as when ρ is very high, atomistic buyers are too likely to buy from the leader, given that, with a very low forgetting rate, a second low-cost provider can easily be sustained. On the other hand, with high δ, it is socially efficient to concentrate all purchases on the leader, and atomistic buyers are too likely to buy from the laggard. For these values, the relationship between τ and PDV TS becomes convex rather than the concave pattern observed for more plausible δ.

The lower panels on each page can be used to see the effects of threshold and the AD leader incentive policies. When LBD is very limited (high ρ), the threshold policy has little effect on outcomes for the simple reason that concentration is low for all τ, so that expected

[23]For example, when $\rho = 0.6$, the PDV of expected εs is 8-10.5 higher (depending on the equilibrium) when $\tau = 0$ than when $\tau = 1$, but expected production costs are 11-14.5 higher.

Figure 17: Expected Compliance Costs δ Homotopies for Different τs. $b^p = 0$.

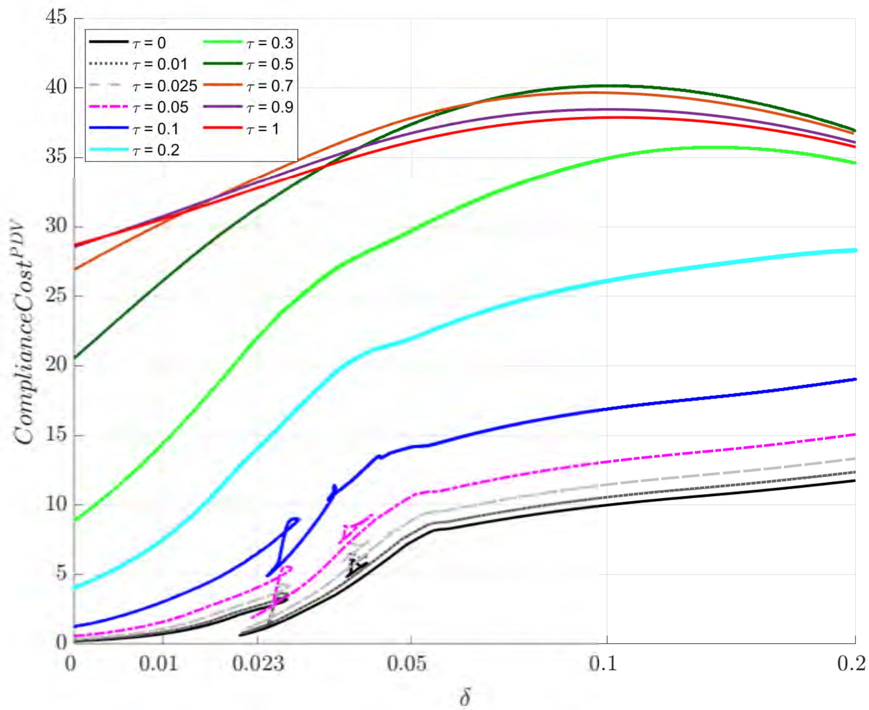

compliance costs are low and the effect on prices in realized states is minimal. This is not true when δ is very high because, for all τ, stochastic forgetting may lead the industry into states where one firm has a substantial advantage. Apart from when ρ is very high, the threshold policies are socially costly and costly to consumers. This partly reflects the level of compliance costs that are realized. Figure 17 shows how the PDV of compliance costs varies with δ and τ. For the illustrative parameters, compliance costs increase in τ up to $\tau = 0.5$, and are then fairly constant for higher values. When the rate of forgetting is higher, compliance costs are actually maximized for $\tau = 0.5$, declining modestly as additional bargaining power shifts to the buyer.

The no leader AD incentive policy figures show that, while the policy produces uniqueness for the illustrative parameters, it can actually lead to increased multiplicity for higher values of δ, for low τ. The other striking result is that for $\tau \approx 0.5$, the policy could raise the PDV TS significantly when $0.82 \leq \rho \leq 0.92$, even though it would lower surplus if τ was equal to 0.3 and essentially have no effect if τ was close to 1. Therefore, the non-monotonicity observed for the illustrative parameters applies more broadly, but the values of bargaining power that can make policies beneficial vary with the technological parameters.

Changing ρ and δ Together. We can also vary ρ and δ together. Here we will focus on $0.6 \leq \rho \leq 1$ and $0 \leq \delta \leq 0.2$ as combinations that are potentially empirically relevant and assess (i) how the existence of multiple equilibria varies with τ, and, for a few select values, (ii) the relationship between τ, social efficiency and the effects of policy. $\delta = 0.2$ implies that, if sales were split equally, sellers would expect to remain around $e_i = 3$ in the long-run.[24]

We identify multiple equilibria by running a sequence of ρ and δ homotopies across the parameter space on a grid of values of τ with 0.05 steps using the equilibria identified in the previous step as starting points. While this approach is not guaranteed to find all equilibria, our results when $\tau = 0$ are almost perfectly consistent with those in BDKS (there is a very small area of the parameter space where we identify additional equilibria). Figure 18(a) shows the smallest values of τ for which we identify multiple equilibria (for that τ), and (b) shows the smallest values of τ for which we identify uniqueness both for that τ and all higher values on the grid.

Consistent with what we noted about the switchbacks in the homotopy paths in Figure 16(b) and (d), there are some limited areas of the parameter space where there is only one equilibrium when $\tau = 0$ but there are multiple equilibria when $\tau = 0.05$, $\tau = 0.1$ or $\tau = 0.15$. However, the pattern that multiplicity exists only for low τ is clear. For example, multiplicity only exists for $\tau = 0.2$ when LBD is limited and forgetting probabilities are quite high (for example, when $\delta = 0.075$ a firm with know-how of 10 forgets with probability 0.54, and $\rho = 0.93$ implies a minimum marginal cost above 7.5 given $\kappa = 10$).

Appendix Figures C.4-C.7 show how the incentive policies affect learning for combinations of $\rho = 0.65$ and 0.9 (faster and slower learning than the illustrative ρ) and $\delta = 0.01$ and 0.05 (less and more likely forgetting than the illustrative δ). The results are broadly consistent with the patterns for the illustrative parameters. The incentive policies lower concentration when the expected HHI without any policy would be greater than 0.5 and less than 1. Faster learning, more likely forgetting (at least up to $\delta = 0.05$) and higher buyer bargaining power

[24]The implied forgetting probabilities in states 2, 3 and 4 are 0.36, 0.49 and 0.59.

Figure 18: Minimum and Maximum Values of τ with Multiple Equilibria when $b^p = 0$. Results Based on a Grid of τ Values with 0.05 steps. White space in (a) implies that multiplicity is never identified.

(a) The Smallest Value of τ where Multiple Equilibria Are Identified.

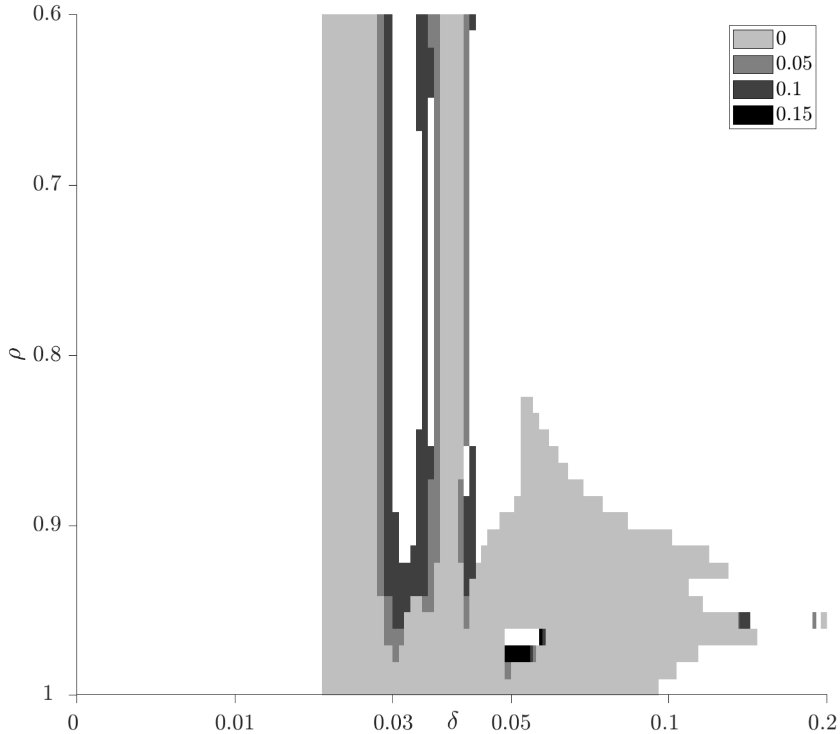

(b) The Smallest Value of τ' where Uniqueness is Identified For All $\tau \geq \tau'$.

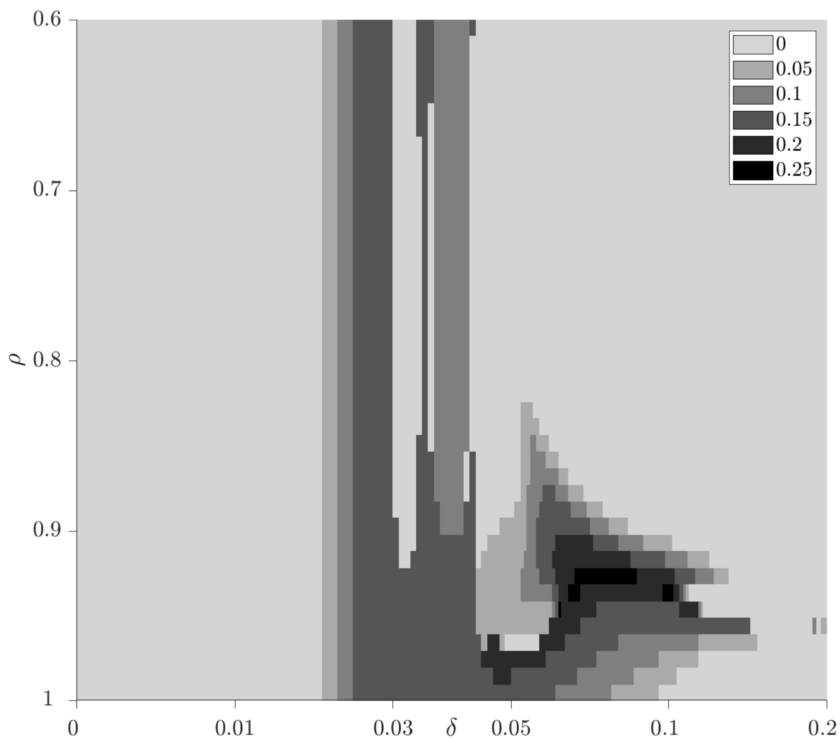

all tend to increase equilibrium concentration. The incentive policies tend to lower PDV CS, although there are exceptions when buyers have almost all of the bargaining power (atomistic buyers may be too likely to buy from the leader). The no leader dynamic incentives and no leader AD incentive policies can also raise PDV TS if learning is slow. Effects on long-run CS (CS^{1000}) are more dependent on the parameters: when $\rho = 0.65$, these incentive policies can raise long-run CS for low or intermediate τ, whereas for $\rho = 0.9$ and $\delta = 0.05$, this happens for very low τ and high τ. The effects on long-run and medium-run CS can also be different: for example, when $\rho = 0.9$ and $\delta = 0.05$, and τ is high, the incentive policies lower CS^{32} while increasing CS^{1000}.

5 Forward-Looking Buyer Behavior.

In this section we briefly consider how allowing for $b^p > 0$ affects outcomes when $\tau = 0$. While the BDKS and BDK models are not the same, the effects that we identify are similar to those in SJHY and readers should consult that paper for additional explanation. Our primary focus here is on introducing the effects that forward-looking buyer behavior has, so that we can discuss the interaction between b^p and τ in the next section.

Figure 19 shows no policy ρ and δ homotopy plots for HHI^{32}, PDV CS and TS for various values of b^p, with $\tau = 0$. Figure 20 shows the range of b^p values for which multiple equilibria exist for empirically plausible ρ and δ combinations.

The main pattern, which holds unless $\delta > 0.6$, is that increasing b^p tends to lower concentration, i.e., the sellers will be more symmetric, and that any multiplicity of equilibria is eliminated as b^p increases.[25] In these examples, we have uniqueness if $b^p \geq 0.1$. Unless $\rho > 0.9$, in which case this greater symmetry can increase total surplus by raising the expected value of εs (as discussed previously in the context of τ), increasing b^p tends to reduce PDV TS, by slowing learning, and tends to reduce PDV CS, by softening competition, which is consistent with the analytical results of Lewis and Yildirim (2002) for a model with

[25]For the illustrative parameters, multiplicity is eliminated when $b^p \geq 0.026$ (once again, it is the paths from the Mid-HHI and High-HHI baseline equilibria that form a loop). However, as with τ, there are parameter values for which multiplicity exists for very low b^p where it does not exist for $b^p = 0$.

Figure 19: Market Outcomes for Alternative b^p as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$). $\tau = 0$ (sellers set prices).

(a) HHI^{32} ρ Homotopies

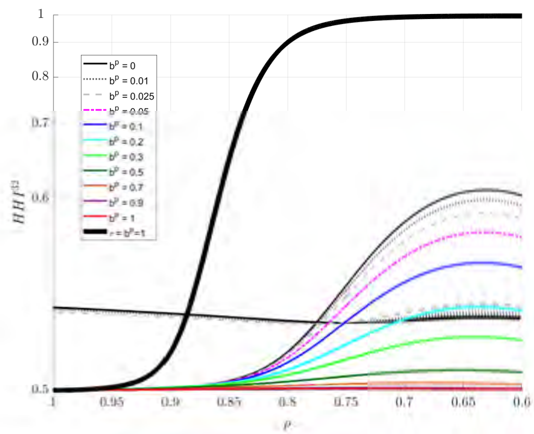

(b) HHI^{32} δ Homotopies

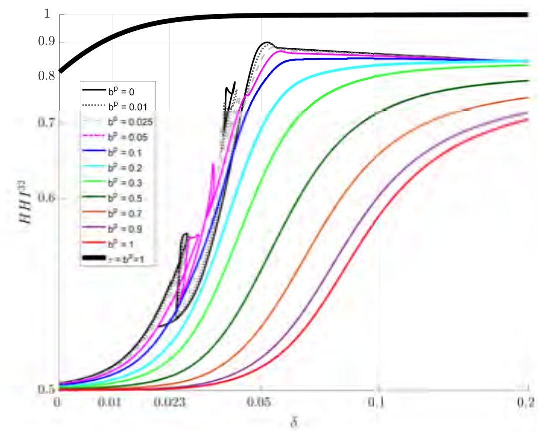

(c) Total Surplus ρ Homotopies (Relative to TS When $\tau = 1$ and $b^p = 1$)

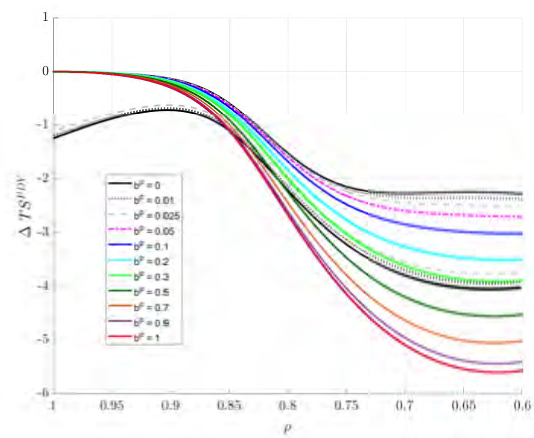

(d) Total Surplus δ Homotopies (Relative to TS When $\tau = 1$ and $b^p = 1$)

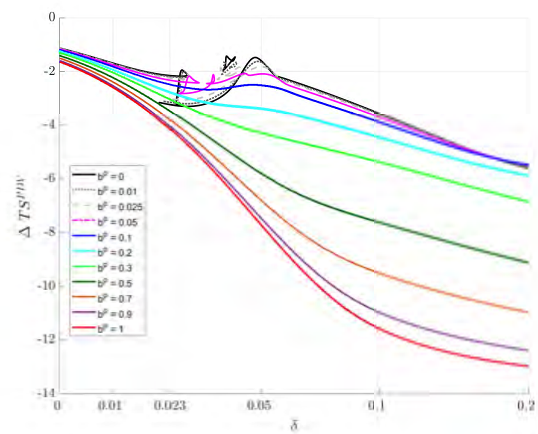

(e) Consumer Surplus ρ Homotopies (Relative to CS When $\tau = 1$ and $b^p = 1$)

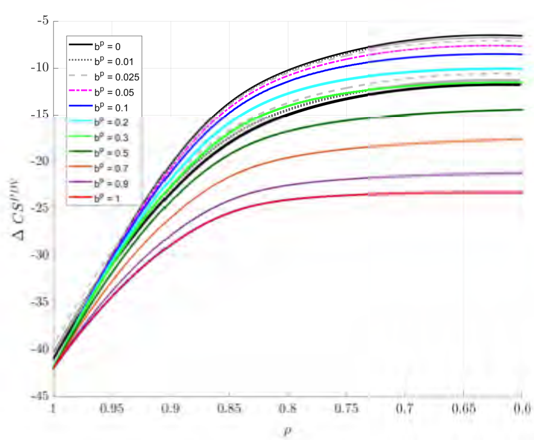

(f) Consumer Surplus δ Homotopies (Relative to CS When $\tau = 1$ and $b^p = 1$)

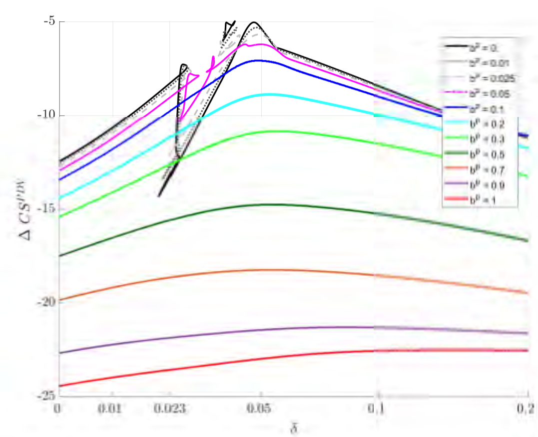

Figure 20: Minimum and Maximum Values of b^p with Multiple Equilibria. Results Based on a Combination of Using ρ and δ homotopies for different values of b^p and Using b^p homotopies. White space in (a) implies that multiplicity is never identified.

(a) The Smallest Value of b^p where Multiple Equilibria Are Identified. $\tau = 0$.

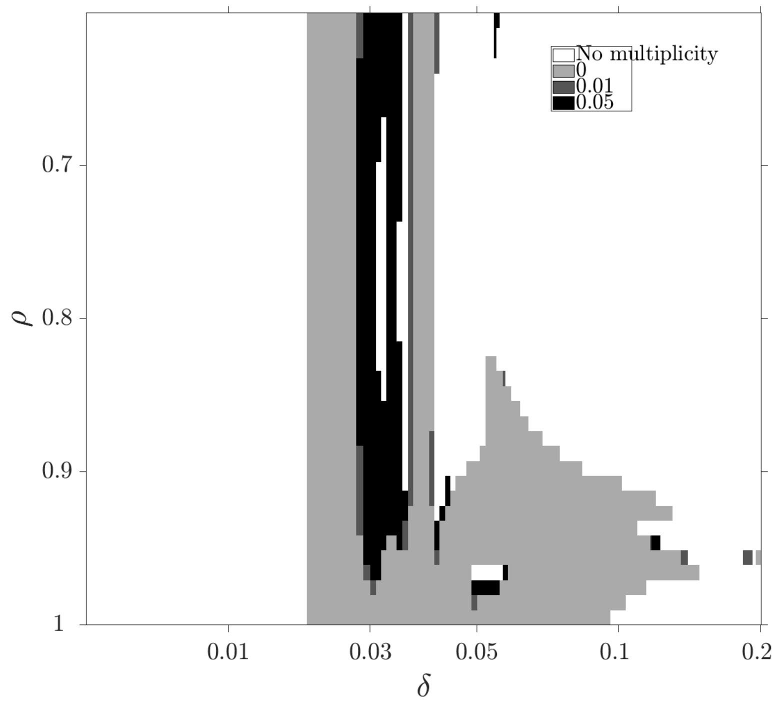

(b) The Smallest Value of b'^p where Uniqueness is Identified For All $b^p \geq b'^p$. $\tau = 0$.

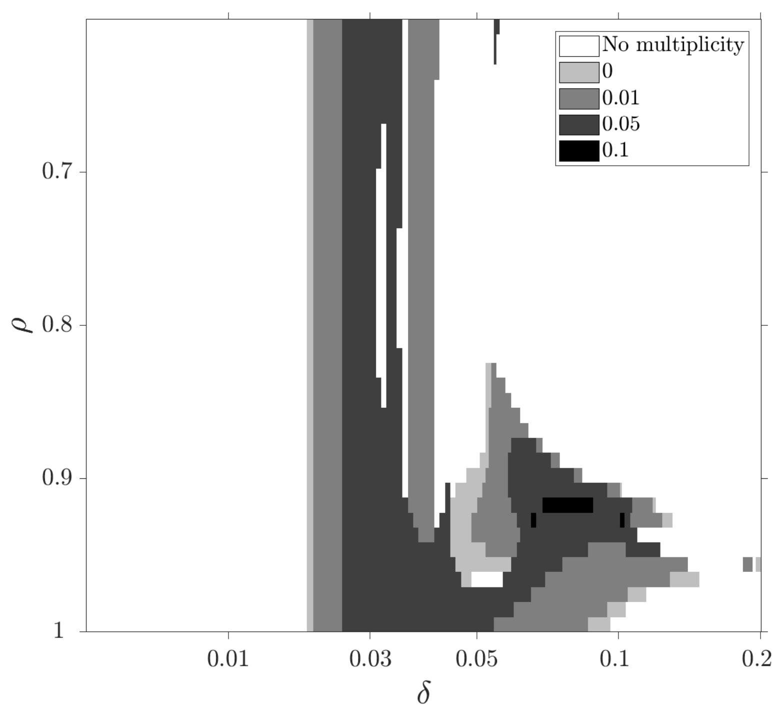

a monopsonist.[26]

The intuition for why forward-looking buyers aim for symmetry is fairly straightforward.[27] A forward-looking buyer will try to lower future prices, which can be done by lowering costs or lowering future margins. Unless δ is very high, the stronger incentive is to try to lower margins, as after a few periods at least one of the firms should have lower costs. With static pricing, seller symmetry would minimize margins. In a dynamic model, margins might be lower in some non-symmetric states, where AB and AD incentives are large. However, these incentives are large only when a seller expects that a sale will move the game to states with higher margins. But, a forward-looking seller will tend to move the industry away from high margin states so that in any equilibrium with forward-looking buyers, AB and AD incentives must be limited. Furthermore, the fact a seller will tend to have a lower future sale probability if it gets ahead can soften competition even relative to a static model, so that prices increase. The elimination of AB and AD incentives also provides why multiplicity, which, as BDKS describe, is rooted in the dynamics of the model, is eliminated as well.[28]

6 Interaction Between Bargaining Power and Forward-Looking Behavior

So far we have seen that, with atomistic buyers, one type of buyer sophistication, buyer bargaining power, tends to lead to greater industry concentration, at least when LBD is significant, and, typically, a redistribution of surplus from sellers to buyers, whereas, when sellers set prices, another type of sophistication, forward-looking behavior, tends to lead to symmetry and redistribution in the opposite direction. We also know that maximum sophistication of both types leads to socially optimal, and buyer optimal, outcomes. Therefore,

[26]Lewis and Yildirim (2002) assumes that sellers have some private information about their current marginal cost, in contrast to the current model.

[27]Appendix Figure C.3 provides more intuition. In particular, panel (a) shows how forward-looking buyers will shift demand towards the laggard in the state (3,2) for the baseline (illustrative parameter) equilibria (holding fixed prices in future states), particularly in the High-HHI and Mid-HHI equilibria where a purchase from laggard may move the industry into a symmetric state where deep trench strategies are very low. The knowledge that buyers will do this in the future weakens seller AB and AD incentives, as advantages are likely to be more temporary. This softens pricing so that state (3,2) prices rise monotonically in b^p (from the Low-HHI equilibrium).

[28]The multinomial logit form of demand in any period implies that price equilibria would be unique if we hold continuation values fixed.

it is interesting to ask what happens when these types of sophistication are introduced simultaneously, and what happens when we assume a value for one of the parameters that one might view as empirically plausible. Here we first look at what happens when $b^p = \tau$ and the parameters are changed together, and then we compare how the effects of varying b^p depend on whether $\tau = 0$ or $\tau = 0.5$, although we could also consider how the effects of varying τ vary with b^p.

Figure 21(a) and (b) show the b^p homotopy paths for the HHI measures and the components of welfare for the illustrative parameters when $\tau = 0$. Figure 21(c) and (d) show homotopy paths when we change b^p and τ together, so that the left-hand edge of each diagram coincides with the BDKS equilibria and the right-hand edge with the social optimum. Consistent with the results in the previous section, increasing b^p, with $\tau = 0$ lowers concentration, PDV CS and PDV TS, with the latter effects being more linear in b^p than when τ was changed, partly reflecting how concentration is low even in the BDKS case, so nonlinear changes resulting from changes in the distribution of future play are limited.

When b^p and τ increase together, concentration and the components of welfare can move in the opposite directions, consistent with how the social optimum is more concentrated than the BDKS outcomes. It is notable, however, that compared to when τ was changed, the PDV welfare components change relatively linearly. Multiplicity is eliminated for $b^p = \tau \geq 0.0196$ (roughly equivalent to 51 symmetric long-lived buyers that each have a tiny amount of bargaining power).

Figure 22 shows ρ homotopies for different b^p values when $\delta = 0.023$ and τ is equal to either 0 or 0.5 (Appendix Figure C.9 shows the δ homotopies when $\rho = 0.75$). When $\tau = 0$, forward-looking behavior by buyers lowers PDV CS and TS, except when its effect is to eliminate the lower surplus equilibira that exist when b^p is very small. When $\tau = 0.5$, increasing b^p tends to lower CS, although the magnitude of this effect is much smaller than when $\tau = 0$, but increasing b^p from 0 to 0.3 can raise PDV TS when $0.8 \leq \rho \leq 0.95$ as atomistic behavior can result in too much concentration.

Figure 23 illustrates the effects of the incentive policies on surplus, as a function of b^p, when $\tau = 0$ and $\tau = 0.5$. While the policies lower the PDV of CS and TS in both cases, for all b^p, the remaining panels show that there are policy-relevant interactions between τ

Figure 21: The Effects of Changing b^p with $\tau = 0$ and $b^p = \tau$ for the Illustrative Parameters.

(a) HHI^{32} and HHI^{1000}: b^p homotopy, $\tau = 0$

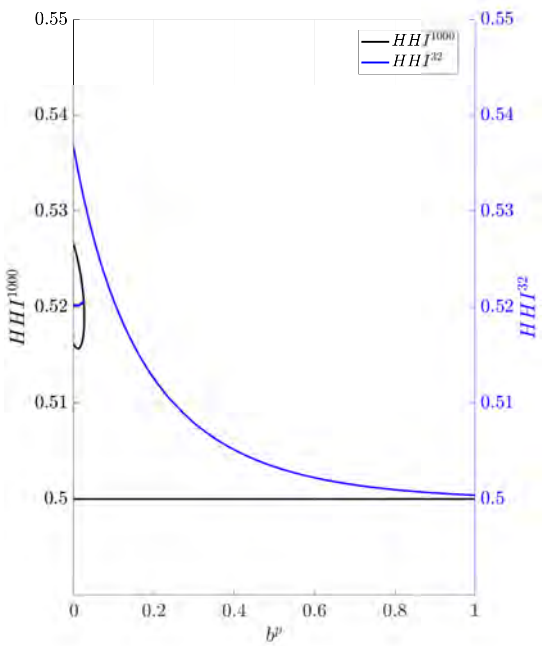

(b) PDV Welfare Components: b^p homotopy, $\tau = 0$

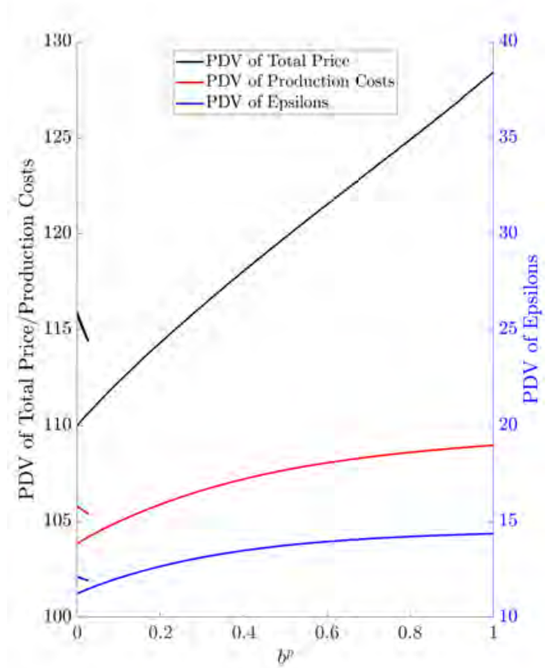

(c) HHI^{32} and HHI^{1000}: $b^p = \tau$ homotopy

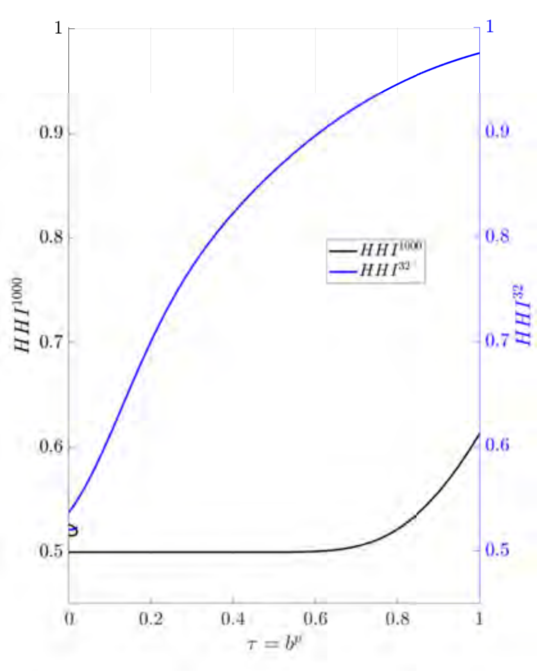

(d) PDV Welfare Components: $b^p = \tau$ homotopy

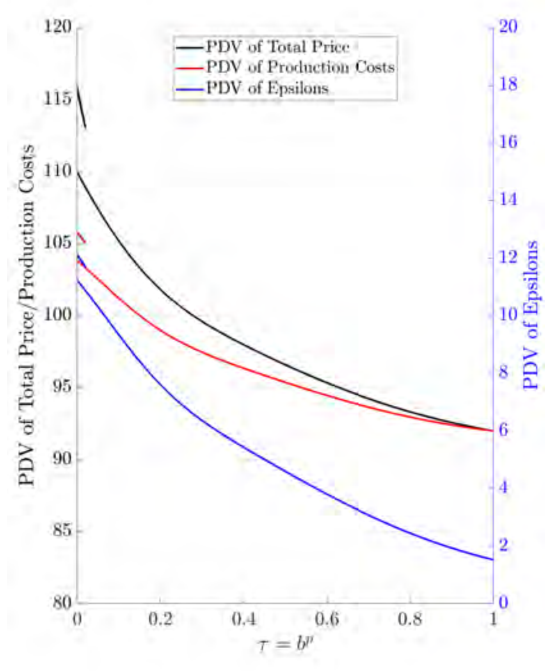

Figure 22: Market Outcomes for Alternative b^p as ρ Varies (with $\delta = 0.023$) for $\tau = 0$ and $\tau = 0.5$.

(a) HHI^{32} $\tau = 0$

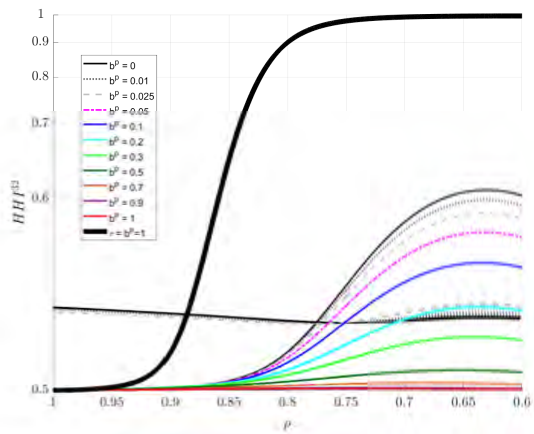

(b) HHI^{32} $\tau = 0.5$

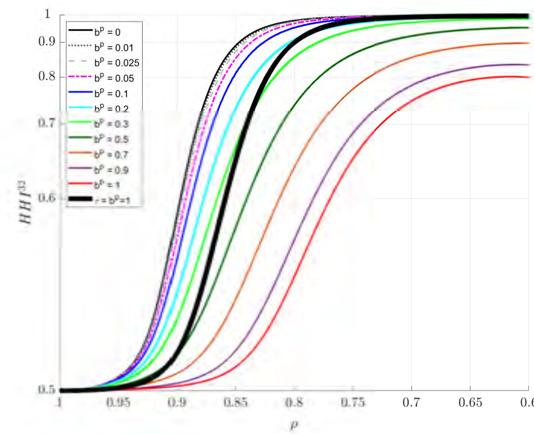

(c) Total Surplus (Relative to TS When $\tau = 1$ and $b^p = 1$) $\tau = 0$

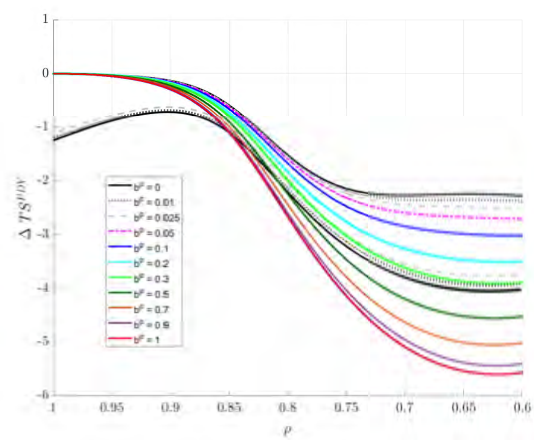

(d) Total Surplus (Relative to TS When $\tau = 1$ and $b^p = 1$) $\tau = 0.5$

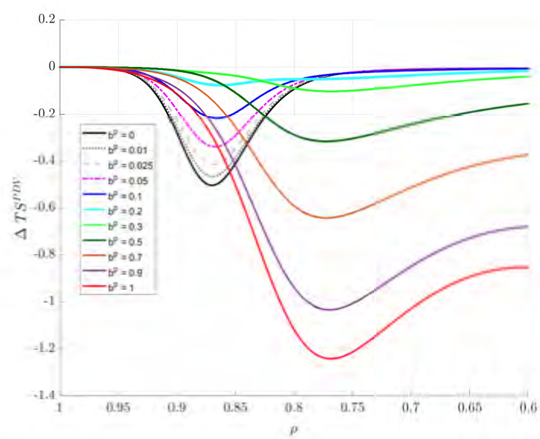

(e) Consumer Surplus (Relative to CS When $\tau = 1$ and $b^p = 1$)

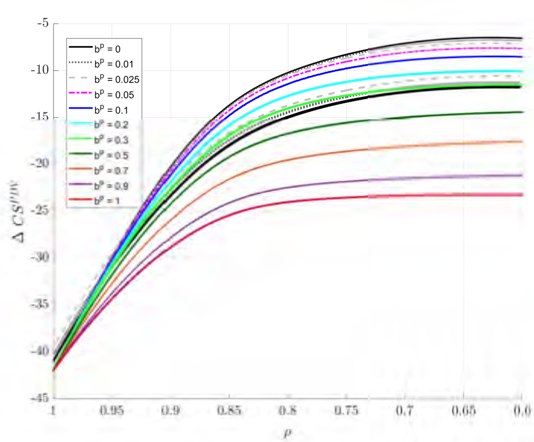

(f) Consumer Surplus (Relative to CS When $\tau = 1$ and $b^p = 1$)

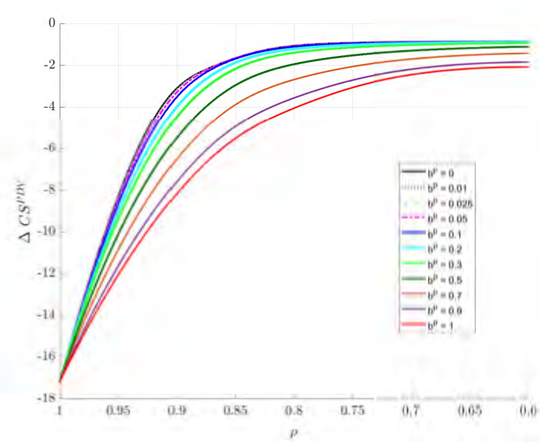

Figure 23: Comparison of the Effects of Incentive Policies on Consumer Surplus as b^p Varies for $\tau = 0$ and $\tau = 0.5$. Illustrative Parameters. Panels (a) and (b) show changes in PDV CS relative to the no policy outcome for the same parameters. The remaining panels show levels.

(a) PDV CS $\tau = 0$ (changes from no policy)

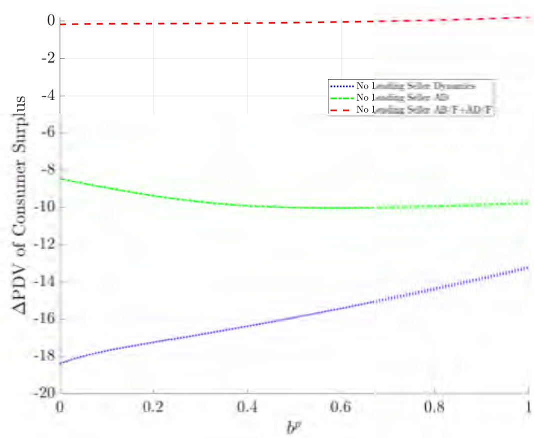

(b) PDV CS $\tau = 0.5$ (changes from no policy)

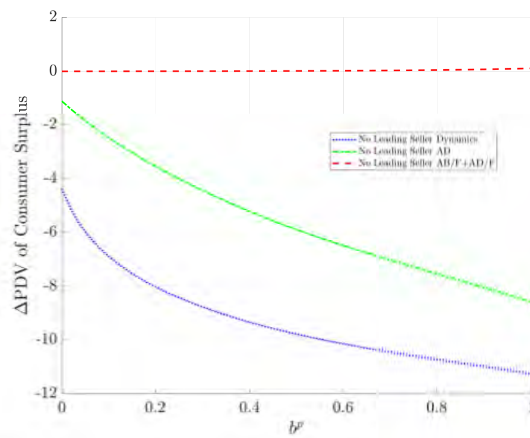

(c) CS^{32} $\tau = 0$ (levels)

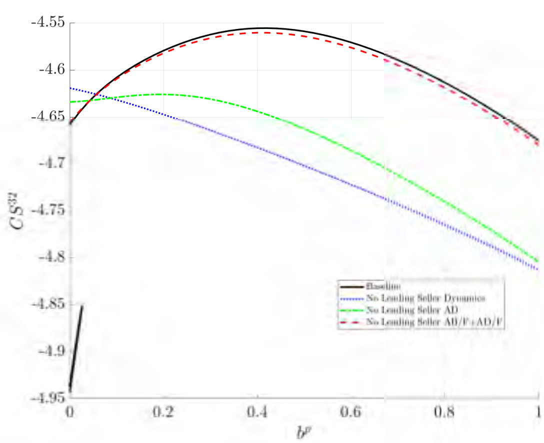

(d) CS^{32} $\tau = 0.5$ (levels)

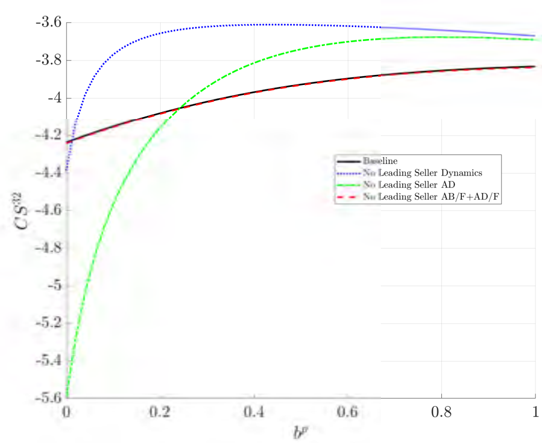

(e) CS^{1000} $\tau = 0$ (levels)

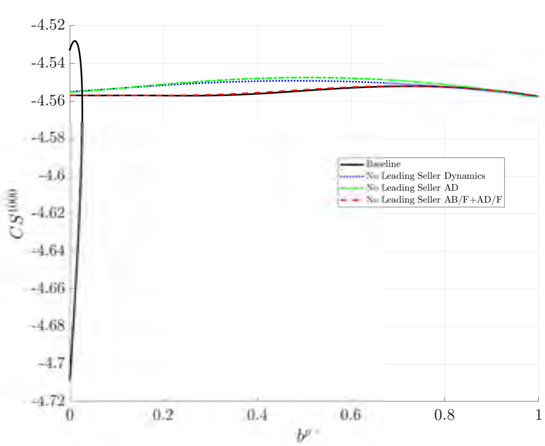

(f) CS^{1000} $\tau = 0.5$ (levels)

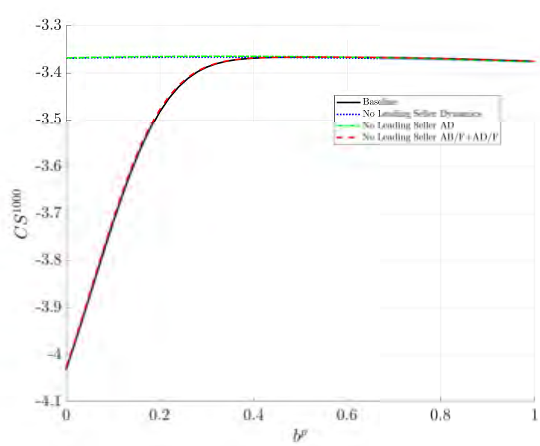

Figure 23: Comparison of the Effects of Incentive Policies on Total Surplus as b^p Varies for $\tau = 0$ and $\tau = 0.5$. Illustrative Parameters. Panels (g) and (h) show changes in PDV TS relative to the no policy outcome for the same parameters. The remaining panels show levels.

(g) PDV TS $\tau = 0$ (changes from no policy)

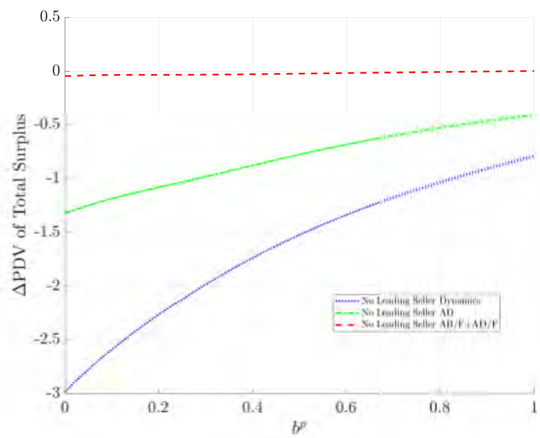

(h) PDV TS $\tau = 0.5$ (changes from no policy)

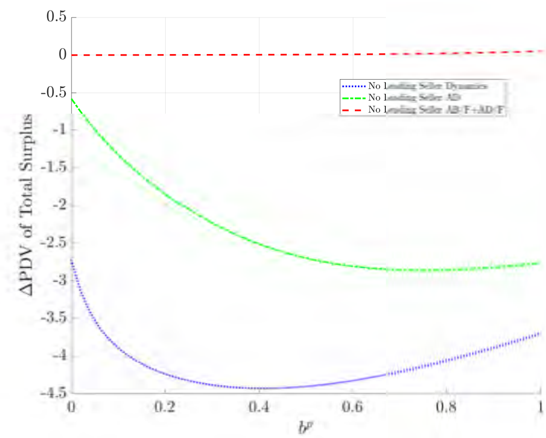

(i) TS^{32} $\tau = 0$ (levels)

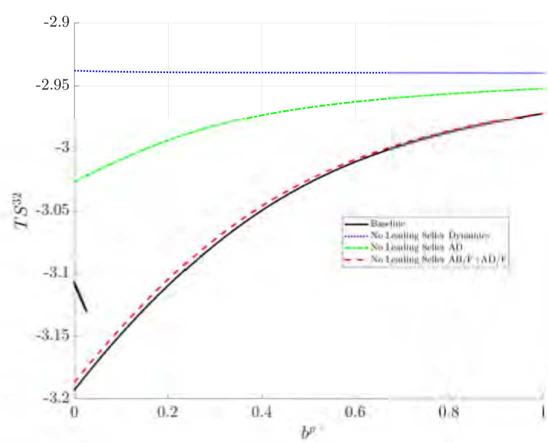

(j) TS^{32} $\tau = 0.5$ (levels)

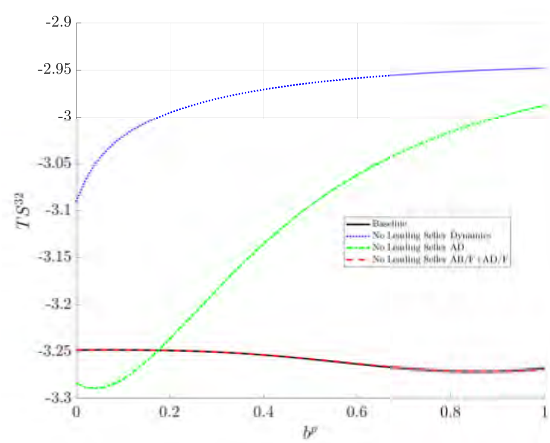

(k) TS^{1000} $\tau = 0$ (levels)

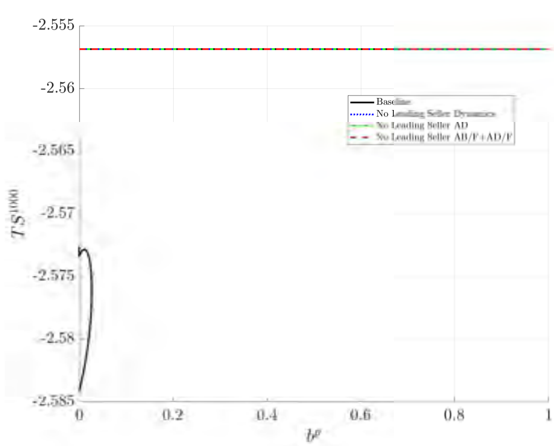

(l) TS^{1000} $\tau = 0.5$ (levels)

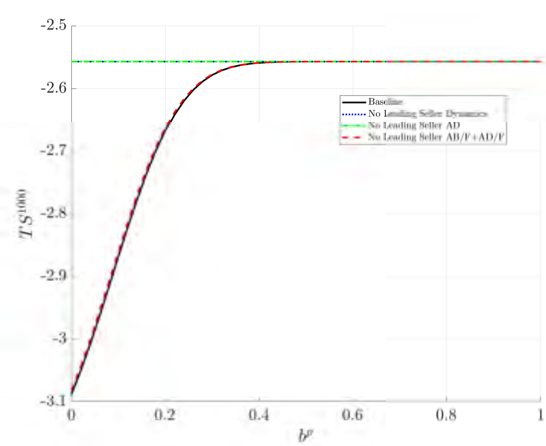

and b^p, especially for CS. For example, when $\tau = 0$, the no leader dynamics or no leader AD policies lower CS^{32} unless b^p is very close to zero, and the magnitude of effects are always small, whereas when $\tau = 0.5$, the policies raise CS^{32} when b^p is large enough. Long-run CS is increased significantly by these policies when $b^p \leq 0.4$ and $\tau = 0.5$, whereas, ignoring the multiplicity issue, effects are tiny when $\tau = 0$.

7 Conclusion

Models of dynamic competition have played an important role in economists' understanding of how the benefits of scale and LBD should be weighed against the dangers of market power when designing antitrust, industrial and international trade policies. This paper has extended one of the common models in this literature to understand how model predictions and policy outcomes may change when we allow for buyers and sellers to bargain over prices and for buyers, as well as sellers, to be forward-looking.

Our clearest result is that allowing for buyer-seller bargaining, rather than unilateral seller price-setting, can make a great deal of difference to both the efficiency of dynamic competition and the cost and benefits of stylized policies. Forward-looking behavior by buyers also affects dynamic competition. What also emerges, however, is that many of these ceffects are non-monotonic and non-linear in parameters that we use to characterize buyer bargaining power and forward-looking behavior, as well as in the parameters that characterize technology. These results reflect how atomistic buyers with bargaining power may tend to result in too much seller concentration and how forward-looking buyers without bargaining power may inadvertently soften competition, and how effects of future bargaining on continuation values can dominate the static effects with which the literature is familiar. These complications mean that unambiguous policy predictions will rarely exist even though allowing for bargaining and forward-looking behavior eliminates the multiplicity of equilibria (at least in the absence of distortionary policies) that has often been the focus of papers that have assumed atomistic buyers and sellers that set prices (e.g., Cabral and Riordan (1994), BDK and BDKS).

References

An, Y., and W. Zhao (2019): "Dynamic Efficiencies of the 1997 Boeing-McDonnell Douglas Merger," *RAND Journal of Economics*, 50(3), 666–694.

Anton, J. J., G. Biglaiser, and N. Vettas (2014): "Dynamic Price Competition with Capacity Constraints and a Strategic Buyer," *International Economic Review*, 55(3), 943–958.

Asker, J., C. Fershtman, J. Jeon, and A. Pakes (2020): "A Computational Framework for Analyzing Dynamic Auctions: The Market Impact of Information Sharing," *RAND Journal of Economics*, 51(3), 805–839.

Benkard, C. L. (2000): "Learning and Forgetting: The Dynamics of Aircraft Production," *American Economic Review*, 90(4), 1034–1054.

——— (2004): "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *Review of Economic Studies*, 71(3), 581–611.

Besanko, D., U. Doraszelski, and Y. Kryukov (2014): "The Economics of Predation: What Drives Pricing When There is Learning-by-Doing?," *American Economic Review*, 104(3), 868–97.

——— (2019): "Sacrifice Tests for Predation in a Dynamic Pricing Model: Ordover and Willig (1981) and Cabral and Riordan (1997) meet Ericson and Pakes (1995)," *International Journal of Industrial Organization*, 70, 102522.

Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2010): "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," *Econometrica*, 78(2), 453–508.

Cabral, L. M., and M. H. Riordan (1994): "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 62(5), 1115–1115.

Caves, R. E., and M. E. Porter (1977): "From Entry Barriers to Mobility Barriers: Conjectural Decisions and Contrived Deterrence to New Competition," *Quarterly Journal of Economics*, 91(2), 241–261.

Dafny, L. S. (2005): "Games Hospitals Play: Entry Deterrence in Hospital Procedure Markets," *Journal of Economics & Management Strategy*, 14(3), 513–542.

Dick, A. R. (1991): "Learning-By-Doing and Dumping in the Semiconductor Industry," *Journal of Law and Economics*, 34(1), 133–159.

Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62(1), 53–82.

Gaynor, M., H. Seider, and W. B. Vogt (2005): "The Volume-Outcome Effect, Scale Economies, and Learning-by-Doing," *American Economic Review*, 95(2), 243–247.

Ghemawat, P. (1985): "Building Strategy on the Experience Cirve," *Harvard Business Review*, 63(2), 143–149.

Horn, H., and A. Wolinsky (1988): "Bilateral Monopolies and Incentives for Merger," *RAND Journal of Economics*, 19(3), 408–419.

Irwin, D. A., and P. J. Klenow (1994): "Learning-by-Doing Spillovers in the Semiconductor Industry," *Journal of Political Economy*, 102(6), 1200–1227.

Joskow, P. L., and N. L. Rose (1985): "The Effects of Technological Change, Experience, and Environmental Regulation on the Construction Cost of Coal-Burning Generating Units," *RAND Journal of Economics*, 16(1), 1–27.

Lee, R. S., and K. Fong (2013): "Markov Perfect Network Formation: An Applied Framework for Bilateral Oligopoly and Bargaining in Buyer-Seller Networks," Discussion paper, New York University.

Lee, R. S., M. D. Whinston, and A. Yurukoglu (2021): "Structural Empirical Analysis of Contracting in Vertical Markets," in *Handbook of Industrial Organization*, ed. by K. Ho, A. Hortaçsu, and A. Lizzeri, vol. 4, pp. 673–742. Elsevier.

Lewis, T. R., and H. Yildirim (2002): "Managing Dynamic Competition," *American Economic Review*, 92(4), 779–797.

Lieberman, M. B. (1984): "The Learning Curve and Pricing in the Chemical Processing Industries," *RAND Journal of Economics*, 15(2), 213–228.

——— (1987a): "Patents, Learning by Doing, and Market Structure in the Chemical Processing Industries," *International Journal of Industrial Organization*, 5(3), 257–276.

——— (1987b): "Postentry Investment and Market Structure in the Chemical Processing Industries," *RAND Journal of Economics*, 18(4), 533–549.

Loertscher, S., and L. M. Marx (2022): "Incomplete Information Bargaining with Applications to Mergers, Investment, and Vertical Integration," *American Economic Review*, 112(2), 616–49.

Maskin, E., and J. Tirole (2001): "Markov Perfect Equilibrium: I. Observable Actions," *Journal of Economic Theory*, 100(2), 191–219.

Pakes, A., and P. McGuire (1994): "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, 25(4), 555–589.

Qiu, Y., and L. D. Anadon (2012): "The Price of Wind Power in China During its Expansion: Technology Adoption, Learning-by-Doing, Economies of Scale, and Manufacturing Localization," *Energy Economics*, 34(3), 772–785.

Saini, V. (2012): "Endogenous Asymmetry in a Dynamic Procurement Auction," *RAND Journal of Economics*, 43(4), 726–760.

Sweeting, A., D. Jia, S. Hui, and X. Yao (2022): "Dynamic Price Competition, Learning-by-Doing, and Strategic Buyers," *American Economic Review*, 112(4), 1311–33.

Sweeting, A., J. W. Roberts, and C. Gedge (2020): "A Model of Dynamic Limit Pricing with an Application to the Airline Industry," *Journal of Political Economy*, 128(3), 1148–1193.

Sweeting, A., X. Tao, and X. Yao (2019): "Dynamic Games with Asymmetric Information: Implications for Mergers," Discussion paper, University of Maryland.

Thompson, P. (2001): "How Much Did the Liberty Shipbuilders Learn? New Evidence for an Old Case Study," *Journal of Political Economy*, 109(1), 103–137.

Thornton, R. A., and P. Thompson (2001): "Learning from Experience and Learning from Others: An Exploration of Learning and Spillovers in Wartime Shipbuilding," *American Economic Review*, 91(5), 1350–1368.

Zimmerman, M. B. (1982): "Learning Effects and the Commercialization of New Energy Technologies: The Case of Nuclear Power," *Bell Journal of Economics*, 13(2), 297–310.

ONLINE APPENDICES FOR "Bargaining, Forward-Looking Buyers and Dynamic Competition" by Deng, Hui, Jia, Leccese and Sweeting

A Proofs of Proposition 1

Recall Proposition 1.

Proposition 1 *1. There will be a unique symmetric MPNE when*

(a) $\delta = 0$ *and* $b^p = 0$*, for all* ρ *and* τ*.*

(b) $\tau = 1$*, and either* $b^p = 0$ *or* $b^p = 1$ *for all* ρ *and* δ*.*

2. If $\tau = 1$*, prices will equal marginal production costs in all states, for all* b^p*,* ρ *and* δ*.*

Proof of Part1(a). Follows from the recursive proof of BDKS, and the fact there can only be one (static) MPNE in the absorbing state (M, M).[29]

Proof of Part 2. The structure of the proof is to show that VS must be zero in every state, which implies that price will equal marginal production costs. Since $\tau = 1$, the solution to the bargaining problem satisfies the following

$$(p_i(\mathbf{e}) - c_i(e_i))D_i(p(\mathbf{e}), \mathbf{e}) + \sum_{k=0,1,2} D_k(p(\mathbf{e}), \mathbf{e})\mu^S_{i,k}(\mathbf{e}) = \sum_{k\in\{0,1,2\}\setminus\{i\}} D^{-i}_k(p(\mathbf{e}), \mathbf{e})\mu^S_{i,k}(\mathbf{e}), \quad (17)$$

where $\mu^S_{i,k}(\mathbf{e})$ is firm i's continuation value after the buyer purchases from firm k, and $D^{-i}_k(p(\mathbf{e}), \mathbf{e})$ denotes the buyer's choice probability for product k when firm i's product

[29] Note that in state (M, M), the value of b^p has no effect on any buyer's strategy. However, this does not necessarily imply that there is uniqueness in earlier states even when movements through the state space are unidirectional. See Appendix D in SJHY for a discussion.

is not available. Specifically,

$$\mu_{i,k}^S(\mathbf{e}) = \beta \sum_{\mathbf{e}'} \Pr(\mathbf{e}'|\mathbf{e}, q_k) VS_i(\mathbf{e}'), \text{ and}$$

$$D_k^{-i}(p(\mathbf{e}), \mathbf{e}) = \frac{\exp\left(\frac{v_k - p_k(\mathbf{e}) + \mu_k^B(\mathbf{e})}{\sigma}\right)}{\sum_{j\in\{0,1,2\}\setminus\{i\}} \exp\left(\frac{v_j - p_j(\mathbf{e}) + \mu_j^B(\mathbf{e})}{\sigma}\right)},$$

where $\mu_k^B(\mathbf{e})$ denotes the buyer's continuation value after the buyer purchases from firm k.

Recall that firm i's beginning of period value satisfies the following equation (18),

$$VS_i(\mathbf{e}) = (p_i(\mathbf{e}) - c_i(e_i)) D_i(p(\mathbf{e}), \mathbf{e}) + \sum_{k=0,1,2} D_k(p(\mathbf{e}), \mathbf{e}) \mu_{i,k}^S(\mathbf{e}). \tag{18}$$

Plugging (17) into (18) yields that

$$VS_i(\mathbf{e}) = \sum_{k\in\{0,1,2\}\setminus\{i\}} D_k^{-i}(p(\mathbf{e}), \mathbf{e}) \mu_{i,k}^S(\mathbf{e}).$$

By definition of $\mu_{i,k}^S(\mathbf{e})$ and changing the order of summation, we have that

$$VS_i(\mathbf{e}) = \beta \sum_{\mathbf{e}'} \sum_{k\in\{0,1,2\}\setminus\{i\}} D_k^{-i}(p(\mathbf{e}), \mathbf{e}) \Pr(\mathbf{e}'|\mathbf{e}, q_k) VS_i(\mathbf{e}').$$

The above equation can be rewritten in a matrix form,

$$\mathbf{VS}_i = \beta \mathbf{Q} \mathbf{VS}_i,$$

where $\mathbf{VS}_i$ is an $M^2 \times 1$ vector, and $\mathbf{Q}$ is an $M^2 \times M^2$ Markov matrix. Therefore,

$$\mathbf{VS}_i = \left(\lim_{T\to\infty} \beta^T \mathbf{Q}^T\right) \mathbf{VS}_i = 0.$$

Plugging $VS_i(\mathbf{e}) = 0$ into (17) yields that $p_i(\mathbf{e}) = c_i(e_i)$.

Proof of Part 1(b). The previous proof shows that , if $\tau = 1$, $VS_i(\mathbf{e}) = 0$ and $p_i(\mathbf{e}) = c_i(e_i)$ in all states. As a result, the MPE is characterized by the equations concerning the buyer's value:

$$VB^*(\mathbf{e}) - b^p\sigma \log\left(\sum_{k=1,2} \exp\left(\frac{v_k - p_k^*(\mathbf{e}) + \mu_k^B(\mathbf{e})}{\sigma}\right)\right) - (1-b^p)\sum_{k=1,2} D_k^*(\mathbf{e})\mu_k^B(\mathbf{e}) = 0, \quad (19)$$

where

$$\mu_k^B(\mathbf{e}) = \beta \sum_{\mathbf{e}'} \Pr(\mathbf{e}'|\mathbf{e}, q_k) VB^*(\mathbf{e}').$$

If $b^p = 1$, (19) can be rewritten as

$$VB^*(\mathbf{e}) = \sigma \log\left(\sum_{k=1,2} \exp\left(\frac{v_k - p_k^*(\mathbf{e}) + \beta \sum_{\mathbf{e}'} \Pr(\mathbf{e}'|\mathbf{e}, q_k) VB^*(\mathbf{e}')}{\sigma}\right)\right).$$

One can view the right hand side as a functional of $VB^*(\cdot)$. We denote the functional by $\mathcal{T}$ and show next that $\mathcal{T}$ satisfies Blackwell's sufficient conditions for a contraction (see Theorem 3.3 in Stokey, Lucas and Prescott, 1989).

It is clear that $\mathcal{T}$ is monotone. That is, $[T(VB)](\mathbf{e}) \leq [T(\widehat{VB})](\mathbf{e})$ if $VB(\mathbf{e}) \leq \widehat{VB}(\mathbf{e})$. Note also that for any constant $a \geq 0$, $[T(VB + a)](\mathbf{e}) = \beta a + [T(VB)](\mathbf{e})$. Therefore, Blackwell's sufficient conditions are satisfied and $\mathcal{T}$ is a contraction mapping. By the Contraction Mapping Theorem, $\mathcal{T}$ has exactly one fixed point, which implies the existence and uniqueness of the MPE.

If $b^p = 0$, the buyer's problem is effectively static and (19) can be rewritten in a matrix form,

$$\mathbf{VB} = \beta\widehat{\mathbf{Q}}\mathbf{VB},$$

where $\mathbf{VB}$ is an $M^2 \times 1$ vector, and $\widehat{\mathbf{Q}}$ is an $M^2 \times M^2$ Markov matrix. Therefore, $\mathbf{VB} = (\mathbf{I} - \beta\widehat{\mathbf{Q}})^{-1}\mathbf{0} = \mathbf{0}$ and $VB^*(\mathbf{e}) = 0$ uniquely solves (19).

B Implementation of Homotopy Methods

This Appendix provides details of our implementation of the homotopy algorithm using the example of how we use a sequence of homotopies to try to enumerate the number of equilibria that exist for different values of (ρ, δ) for given values of b^p. Our implementation of other homotopies in the paper is similar to a single step in this sequence.

B.1 Preliminaries

We identify equilibria at particular gridpoints in (ρ, δ) space. We specify a 1000-point evenly-spaced grid for the forgetting rate $\delta \in [0, 1]$ and a 100-point evenly-spaced grid for the learning progress ratio $\rho \in [0, 1]$. The state space of the game is defined by an (30×30) grid of values of the know-how of each firm.

B.2 System of Equations Defining Equilibrium

An MPE is defined by a system of 2,265 equations (one VS^* equation (text equation (13)) for each of 900 states, one p^* equation (text equation (7)) for each of 900 states, and 465 VB^* equations (text equation (4)), reflecting our symmetry assumption. As a group, we will denote three equations as F.

B.3 Homotopy Algorithm: Overview

The idea of the homotopy is to trace out an equilibrium correspondance as one of the parameters of interest is changed, holding the others fixed. Starting from any equilibrium, the numerical algorithm traces a path where a parameter (such as δ), and the vectors $V^B(\mathbf{e})$, $V^S(\mathbf{e})$ and $p(\mathbf{e})$ are changed together so that the equations F continue to hold, by solving a system of differential equations. The differential equation solver does not return equilibria exactly at the gridpoints so it is necessary to interpolate between the solutions returned by the solver. Homotopies can be run starting from different equilibria and varying different parameters. When these different homotopies return solutions at the same gridpoint it is necessary to define a numerical rule for when two different solutions should be counted as different equilibria.

B.4 Procedure Details

Step 1: Finding Equilibria for $\delta = 0$. The first step is to find an equilibrium (i.e., a solution to the 2,265 equations) for $\delta = 0$ for each value of ρ on the grid. There will be a unique Markov Perfect equilibrium for $\delta = 0$, as, in this case, movements through the state space are unidirectional, so that the state will eventually end up in the state (M, M) where no more learning is possible.[30]

We solve for an equilibrium using the Levenberg-Marquardt algorithm implemented using `fsolve` in MATLAB, where we supply analytic gradients for each equation. The solution for the previous value of ρ are used as starting values. To ensure that the solutions are precise we use a tolerance of 10^{-7} for the sum of squared values of each equation, and a relative tolerance of 10^{-14} for the price and value variables that we are solving for.

Step 2: δ-Homotopies. Using the notation of BDKS, we explore the correspondence

$$F^{-1}(\rho) = \{(\mathbf{V}^*, \mathbf{p}^*, \delta) | F(\mathbf{V}^*, \mathbf{p}^*; \rho, \delta) = \mathbf{0}, \quad \delta \in [0, 1]\},$$

The homotopy approach follows the correspondence as a parameter, s, changes (in our analysis, s could be δ, ρ, τ or b^p, or we could set $\tau = b^p$ and change them together). Denoting $\mathbf{x} = (\mathbf{V}^*, \mathbf{p}^*)$, $F(\mathbf{x}(s), \delta(s), \rho) = \mathbf{0}$ can be implicitly differentiated to find how $\mathbf{x}$ and δ must change for the equations to continue to hold as s changes.

$$\frac{\partial F(\mathbf{x}(s), \delta(s), \rho)}{\partial \mathbf{x}} \mathbf{x}'(s) + \frac{\partial F(\mathbf{x}(s), \delta(s), \rho)}{\partial \delta} \delta'(s) = \mathbf{0}$$

where $\frac{\partial F(\mathbf{x}(s), \delta(s), \rho)}{\partial \mathbf{x}}$ is a (2,265 x 2,265) matrix, $\mathbf{x}'(s)$ and $\frac{\partial F(\mathbf{x}(s), \delta(s), \rho)}{\partial \delta}$ are both (2,265 x 1) vectors and $\delta'(s)$ is a scalar. The solution to these differential equations will have the

[30] BDKS discuss this result for $b^p = 0$. It will also hold for any higher value of b^p, as movements through the state space are unidirectional.

following form, where $y_i'(s)$ is the derivative of the ith element of $\mathbf{y}(s) = (\mathbf{x}(s), \delta(s))$,

$$y_i'(s) = (-1)^{i+1} \det \left(\left(\frac{\partial F(\mathbf{y}(s), \rho)}{\partial \mathbf{y}} \right)_{-i} \right)$$

where $_{-i}$ means that the ith column is removed from the (2,266 x 2,266) $\frac{\partial F(\mathbf{y}(s),\rho)}{\partial \mathbf{y}}$.

To implement the path-following procedure, we use the routine FORTRAN routine FIXPNS from HOMPACK90, with the ADIFOR 2.0D automatic differentiation package used to evaluate the sparse Jacobian $\frac{\partial F(\mathbf{y}(s),\rho)}{\partial \mathbf{y}}$ and the STEPNS routine is used to find the next point on the path.[31,32]

The FIXPNS routine will return solutions at values of δ that are not equal to the gridpoints. Therefore we adjust the code so that after *each* step, the algorithm checks whether a gridpoint has been passed and, if so, the routine ROOTNX is used to calculate the equilibrium at the gridpoint, using information on the solutions at either side.[33]

The time taken to run a homotopy is usually between one hour and seven hours, when it is run on UMD's BSWIFT cluster (a moderately sized cluster for the School of Behavioral and Social Sciences).

Step 3: Enumerating Equilibria. Once we have collected the solutions at each of the (ρ, δ) gridpoints we need to identify which solutions represent distinct equilibria, taking into account that small differences may arise because of numerical differences that are within our tolerances. For this paper, we use the rule that solutions count as different equilibria if at least some elements of the price vector differ by more than 0.001.

[31]STEPNS is a predictor-corrector algorithm where hermetic cubic interpolation is used to guess the next point, and an iterative procedure is then used to return to the path.

[32]For details of the HOMPACK subroutines, please consult manual of the algorithm at `https://users.wpi.edu/~walker/Papers/hompack90,ACM-TOMS_23,1997,514-549.pdf`.

[33]It can happen that the ROOTNX routine stops prematurely so that the returned solution is not exactly at the gridpoint value of δ. We do not use the small proportion of solutions where the difference is more than 10^{-6}. Varying this threshold does not affect the reported results. We also need to decide whether the equations have been solved accurately enough so that the values and strategies can be treated as equilibria. The criteria that we use is that solutions where the value of each equation residual should be less than 10^{-10}. Otherwise, the solution is rejected. In practice, the rejected solutions typically have residuals that are much larger than 10^{-10}.

Step 4: ρ-Homotopies. With a set of equilibria from the δ-homotopies in hand, we can perform the next round of the criss-crossing procedure, using equilibria found in the last round as starting points.[34] From this round on, we run homotopies from starting points in both directions i.e., we follow paths where ρ is falling as well as paths where ρ is increasing. We have found that this is useful in identifying additional equilibria.

This second round of homotopies can also help us to deal with gridpoints where the first round δ-homotopies identify no equilibria because a homotopy run stops (or takes a long sequence of infinitesimally small steps). As noted by BDKS (p. 467), the homotopies may stop if they reach a point where the evaluated Jacobian $\frac{\partial F(\mathbf{y}(s),\rho)}{\partial \mathbf{y}}$ has less than full rank. Suppose, for example, that the δ-homotopy for $\rho = 0.4$ stops at $\delta = 0.3$, so we have no equilibria for δ values above 0.3. Homotopies that are run from gridpoints where we did find equilibria with $\delta = 0.350, .., 1$ and higher or lower values of ρ may fill in some of the missing equilibria.

Step 5: Repeat steps 3, 2 and 4 to Identify Additional Equilibria Using New Equilibria as Starting Points. We use the procedures described in Step 3 to identify new equilibria at the gridpoints. These new equilibria are used to start new sets of δ-homotopies, which in turn can identify equilibria that can be used for new sets of ρ-homotopies. This iterative process is continued until the number of additional equilibria that are identifed in a round has no noticeable effect on the heatmaps which show the number of equilibria. For the BDKS, $b^p = 0$ case, this happens after 8 rounds.

[34] In practice, using all new equilibria could be computationally prohibitive. We therefore use an algorithm that continues to add new groups of 10,000 starting points when we find that using additional starting points yields a significant number of equilibria that have not been identified before. We have experimented with different rules, and have found that alternative algorithms do not find noticeably more equilibria, across the parameter space, than the algorithm that we use.

C Additional Tables and Figures

Table C.1: Equilibria in the BDKS Model for $\delta = 0.0275$, $\rho = 0.75$, $b^p = 0$, $\tau = 0$.

						Low-HHI $HHI^{1000} = 0.500$		Mid-HHI $HHI^{1000} = 0.516$		High-HHI $HHI^{1000} = 0.527$	
e_1	e_2	c_1	c_2	Δ_1	Δ_2	p_1	p_2	p_1	p_2	p_1	p_2
1	1	10.00	10.00	0.0230	0.0230	-0.54	-0.54	-1.63	-1.63	-1.61	-1.61
2	1	7.50	10.00	0.0455	0.0230	4.91	7.21	5.16	7.60	5.15	7.60
2	2	7.50	7.50	0.0455	0.0455	4.22	4.22	0.77	0.77	0.77	0.77
3	1	6.34	10.00	0.0674	0.0230	5.82	8.18	6.56	8.71	6.55	8.70
3	2	6.34	7.50	0.0674	0.0455	4.65	5.46	4.06	5.97	4.05	5.97
3	3	6.34	6.34	0.0674	0.0674	5.11	5.11	1.49	1.49	1.47	1.47
4	1	5.62	10.00	0.0889	0.0230	5.95	8.29	6.67	8.55	6.67	8.54
4	2	5.62	7.50	0.0889	0.0455	4.86	5.85	5.46	7.08	5.46	7.08
4	3	5.62	6.34	0.0889	0.0674	5.08	5.42	3.81	5.53	3.80	5.54
4	4	5.62	5.62	0.0889	0.0889	5.22	5.22	1.72	1.72	1.70	1.70
10	1	3.85	10.00	0.2076	0.0230	5.89	8.16	6.14	7.71	6.13	7.69
10	2	3.85	7.50	0.2076	0.0455	5.05	6.06	5.75	6.43	5.75	6.42
10	3	3.85	6.34	0.2076	0.0674	5.20	5.81	5.80	6.31	5.80	6.31
10	8	3.85	4.22	0.2076	0.1699	5.10	5.20	4.49	5.85	4.49	5.86
10	9	3.85	4.02	0.2076	0.1889	5.11	5.15	3.26	4.56	3.25	4.55
10	10	3.85	3.85	0.2076	0.2076	5.12	5.12	2.47	2.47	2.43	2.43
15	1	3.25	10.00	0.2946	0.0230	5.79	8.05	5.98	7.36	5.97	7.38
15	2	3.25	7.5	0.2946	0.045	5.02	5.93	5.63	6.18	5.62	6.17
15	3	3.25	6.34	0.2946	0.0674	5.22	5.74	5.67	6.02	5.67	6.01
15	10	3.25	3.85	0.2946	0.2076	5.19	5.20	5.40	5.94	5.41	5.95
15	14	3.25	3.34	0.2946	0.2780	5.23	5.21	3.46	4.44	3.43	4.44
15	15	3.25	3.25	0.2946	0.2946	5.24	5.24	3.16	3.16	3.10	3.10
16	16	3.25	3.25	0.3109	0.3109	5.28	5.28	3.24	3.24	3.18	3.18
20	20	3.25	3.25	0.3721	0.3721	5.25	5.25	3.32	3.32	3.20	3.20
22	22	3.25	3.25	0.4007	0.4007	5.25	5.25	3.44	3.44	3.26	3.26
25	25	3.25	3.25	0.4411	0.4411	5.25	5.25	3.90	3.90	3.28	3.28
27	27	3.25	3.25	0.4665	0.4665	5.25	5.25	4.62	4.62	3.34	3.34
28	28	3.25	3.25	0.4787	0.4787	5.25	5.25	4.98	4.98	3.52	3.52
29	1	3.25	10.00	0.4907	0.0230	5.79	8.05	5.63	7.62	5.57	7.46
29	2	3.25	7.50	0.4907	0.0455	5.01	5.91	5.04	5.77	5.07	5.72
29	10	3.25	3.85	0.4907	0.2076	5.23	5.17	5.35	5.17	5.39	5.15
29	15	3.25	3.25	0.4907	0.2946	5.27	5.22	5.45	5.34	5.52	5.33
29	29	3.25	3.25	0.4907	0.4907	5.25	5.25	5.22	5.22	3.98	3.98
30	1	3.25	10.00	0.5024	0.0230	5.79	8.05	5.67	7.66	5.63	7.53
30	2	3.25	7.50	0.5024	0.0455	5.01	5.91	5.10	5.84	5.12	5.81
30	10	3.25	3.85	0.5024	0.2076	5.23	5.17	5.33	5.21	5.35	5.19
30	15	3.25	3.25	0.5024	0.2946	5.27	5.22	5.42	5.35	5.45	5.33
30	29	3.25	3.25	0.5024	0.4907	5.25	5.25	5.30	5.20	4.29	4.60
30	30	3.25	3.25	0.5024	0.5024	5.25	5.25	5.27	5.27	4.77	4.77

Notes: c_i, p_i, Δ_i are the marginal costs, equilibrium price and probability of forgetting for firm i. HHI^{∞} is the expected long-run value of the HHI.

Table C.2 complements text Table 2 by showing some additional policies, specifically, a share restriction with a lower compliance cost and policies that prevent all sellers from considering particular dynamic incentives.

Figure C.1: Expected Value of the HHI^t Measure for the Illustrative Parameters for Alternative Polar Values of b^p and τ. Note the different scales on the x-axis.

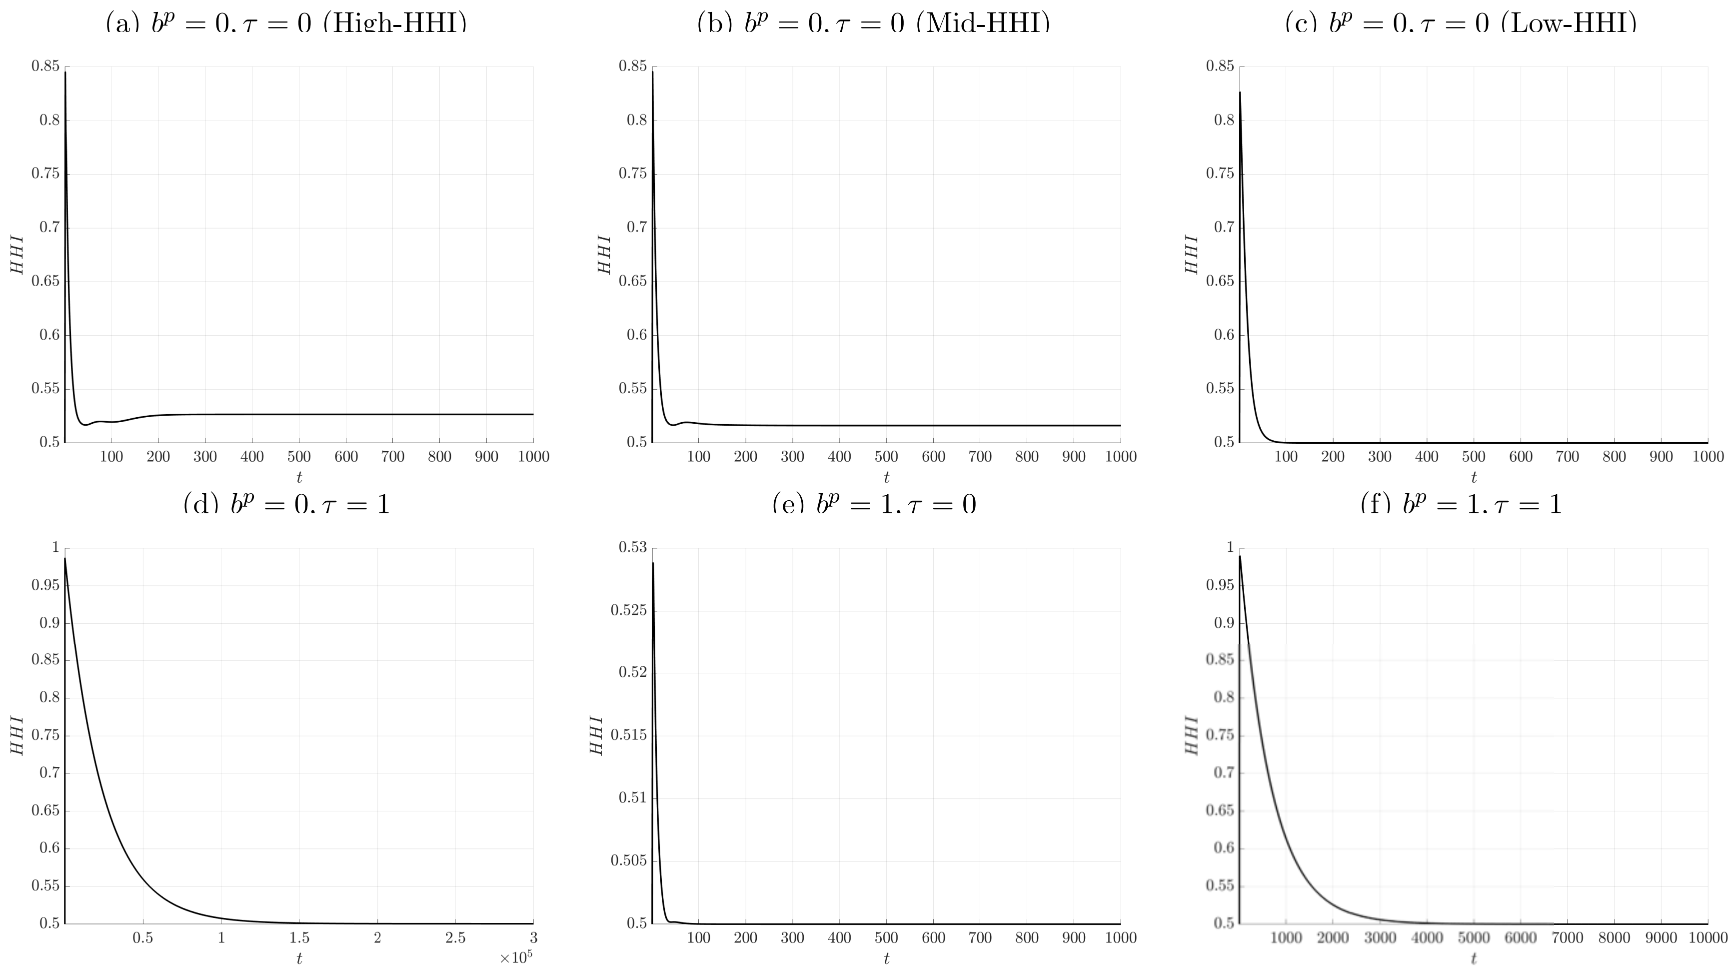

Table C.2: Effect of Alternative Policies to Restrict Dominance in the Baseline Cases For the Illustrative Parameters. Δ indicates changes from the values in the upper part of the table, where consumer surplus (CS) is the expected value of εs less prices, and producer surplus is expected prices less production costs. Low, Mid and HHI denote the relevant baseline equilibria. When a policy rule eliminates multipliicity the table shows changes from the Low-HHI baseline equilibrium only.

Parameters		$b^p = 0,\ \tau = 0$		$b^p = 1,\ \tau = 1$	$b^p = 0,\ \tau = 1$	$b^p = 1,\ \tau = 0$	$b^p = 0,\ \tau = 0.5$	$b^p = 0.05,\ \tau = 0.5$
No Policy Baseline	Low	Mid	High					
HHI^{32}	0.537	0.520	0.520	0.976	0.987	0.500	0.993	0.989
HHI^{1000}	0.500	0.516	0.527	0.613	0.967	0.500	0.877	0.802
PDV CS	-98.750	-103.659	-103.793	-90.469	-90.653	-114.043	-91.378	-91.394
PDV PS	6.149	9.975	10.105	0	0	19.437	0.887	0.903
CS^{1000}	-4.557	-4.533	-4.710	-2.723	-3.208	-4.557	-4.032	-3.877
$\mathrm{PS}^{\ 1000}$	2.000	1.960	2.2126	0	0	2.000	0.943	0.890
Threshold Rule, $\chi = 10$, $\psi = 0.75$								
# of equilibria		3		1	1	1	1	1
ΔHHI^{32}	-0.011	0.000	0.000	-0.399	-0.012	0.000	-0.012	-0.028
$\Delta\ HHI^{1000}$	0.000	-0.000	-0.000	-0.113	-0.078	0.000	-0.202	-0.255
Δ PDV CS	-2.413	-1.271	-1.251	-3.763	-10.705	0.000	-11.765	-11.161
Δ PDV PS	0.754	0.564	0.546	0	0	0.000	0.048	0.097
PDV Compliance Cost	1.458	0.622	0.628	1.689	10.598	0.000	11.706	11.010
Δ CS^{1000}	0.000	0.000	0.000	0.167	-0.375	0.000	0.144	0.370
Δ PS^{1000} (incl. compliance)	0.000	0.000	-0.001	0	0	0.000	-0.079	-0.068
Compliance Cost^{1000}	0.000	0.000	0.000	0.000	0.477	0.000	0.211	0.055
No Seller Dynamics								
# of equilibria		1		1	1	1	1	1
ΔHHI^{32}	-0.012	0.005	0.004	0	0	-0.000	-0.182	-0.259
$\Delta\ HHI^{1000}$	-0.000	-0.016	-0.027	0	0	0.000	-0.377	-0.302
Δ PDV CS	-48.520	-43.611	-43.479	0	0	-23.373	-45.251	-35.659
Δ PDV PS	47.294	43.459	43.330	0	0	23.089	44.056	33.883
Δ CS^{1000}	0.000	-0.024	0.153	0	0	0.000	0.656	0.502
Δ PS^{1000}	0.000	0.040	-0.126	0	0	0.000	-0.124	-0.071
No Seller AD Incentives								
# of equilibria		1		1	1	1	1	1
ΔHHI^{32}	-0.034	-0.017	-0.017	0	0	0.000	-0.200	-0.293
$\Delta\ HHI^{1000}$	0.000	-0.016	-0.027	0	0	0.000	-0.377	-0.302
Δ PDV CS	-9.829	-4.920	-4.787	0	0	-10.771	-1.834	-2.556
Δ PDV PS	8.510	4.675	4.545	0	0	10.335	1.261	1.571
Δ CS^{1000}	0.003	-0.021	0.156	0	0	-0.001	0.666	0.513
Δ PS^{1000}	-0.003	0.037	-0.129	0	0	0.001	-0.134	-0.0821
No Seller AB/Forgetting or AD/Forgetting Incentives								
# of equilibria		1		1	1	1	1	1
ΔHHI^{32}	-0.003	0.014	0.013	0	0	-0.000	-0.000	-0.000
$\Delta\ HHI^{1000}$	-0.000	-0.016	-0.027	0	0	-0.000	-0.005	-0.006
Δ PDV CS	-0.225	4.684	4.817	0	0	0.197	0.007	0.007
Δ PDV PS	0.181	-3.654	-3.784	0	0	-0.187	-0.007	-0.006
Δ CS^{1000}	-0.000	-0.024	0.153	0	0	-0.000	0.005	0.007
Δ PS^{1000}	0.000	0.040	-0.126	0	0	0.000	-0.002	-0.001

Figure C.3 illustrates this logic using the baseline equilibria. Figure C.3(a) shows what happens to demand for firm 1 (i.e., the probability that firm 1 makes a sale) as a result of forward-looking behavior by buyers as b^p changes, holding seller prices in all states fixed at their baseline equilibrium values.[35] In the equilibria with a diagonal trench, a forward-looking buyer has a strong incentive to favor the laggard, so that the industry is likely to become symmetric in the next period. Figure C.3(b) shows how the AB/AD incentives of each seller change due to buyer behavior, once again holding seller strategies fixed (the Mid- and High-HHI incentives are indistinguishable). Firm 1's incentive to prevent firm 2 from making the sale declines rapidly in b^p, reflecting how future strategic behavior by buyers means that a laggard is likely to catch-up soon even if it does not make the current sale. While the loss in demand in the High and Mid-HHI equilibrium provides firm 1 with an incentive to lower its price, the change in dynamic incentives means that charging a low price is unprofitable once b^p has increased by a small amount. This explains why, along the b^p homotopy path from the baseline equilibria, shown in Figure C.3(c), firm 1's equilibrium price falls before the path, which does not extend beyond $b^p = 0.026$, loops back. In contrast, on the surviving path, the weakening of both AB and AD dynamic incentives (Figure C.3(d)) shows what happens to these on the homotopy paths), dominates the small decrease in demand so that prices increase.

[35]The figure shows the probability of sale when firm 1's price is varied, assuming that firm 2's price in (3,2), the prices of both firms in other states and the price of firm 1 if (3,2) is visited in the future are at their baseline equilibrium values.

Figure C.2: HHI^{32} for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. The thick black lines indicate HHIs implied by the social planner outcome. Full ranges of ρ and δ.

(a) No Policy ρ Homotopies

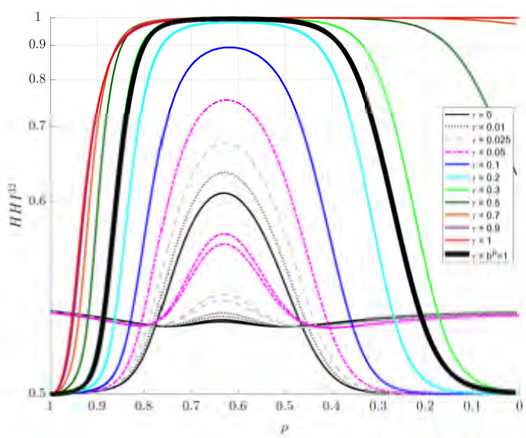

(b) No Policy δ Homotopies

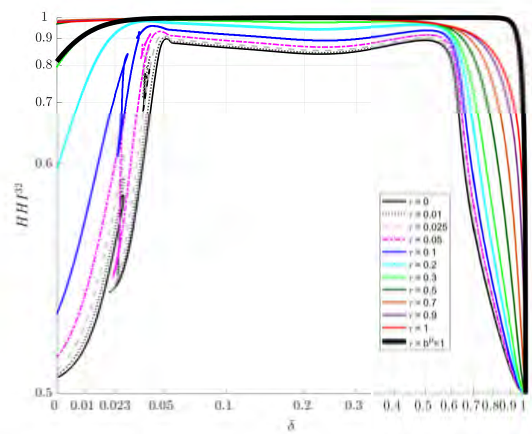

(c) $\chi = 50, \psi = 0.75$ ρ Homotopies

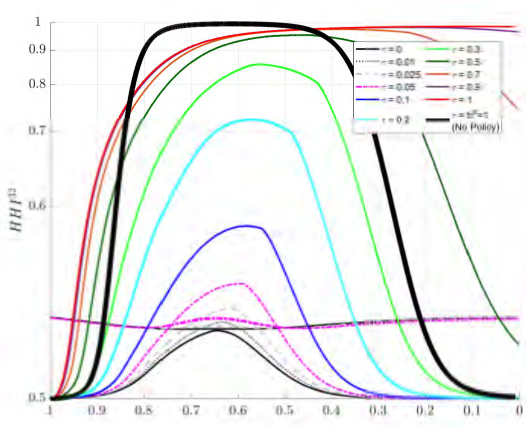

(d) $\chi = 50, \psi = 0.75$ δ Homotopies

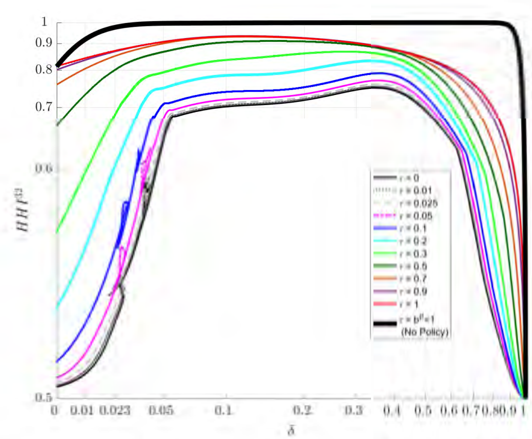

(e) No Leader AD Incentive ρ Homotopies

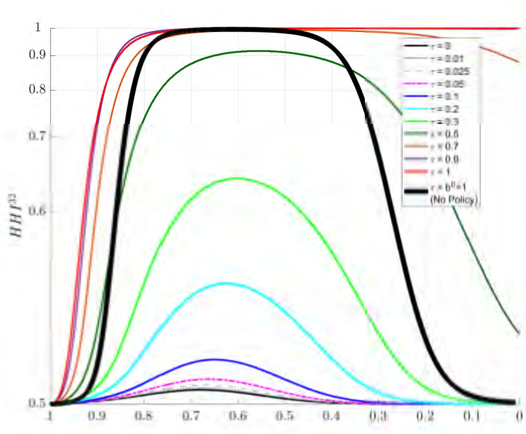

(f) No Leader AD Incentive δ Homotopies TBA

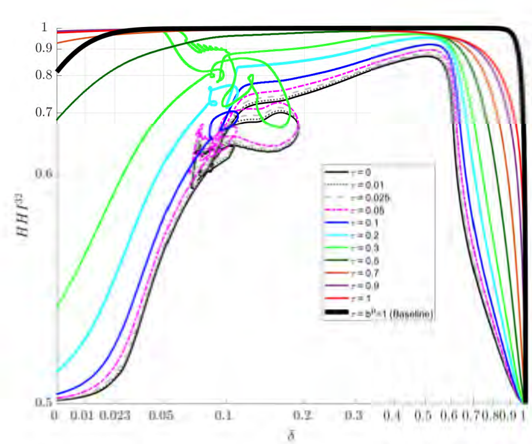

Figure C.2: HHI^{1000} for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. The thick black lines indicate HHIs implied by the social planner outcome.

(g) No Policy ρ Homotopies

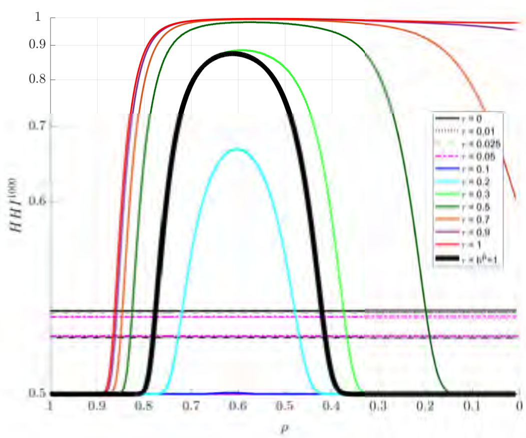

(h) No Policy δ Homotopies

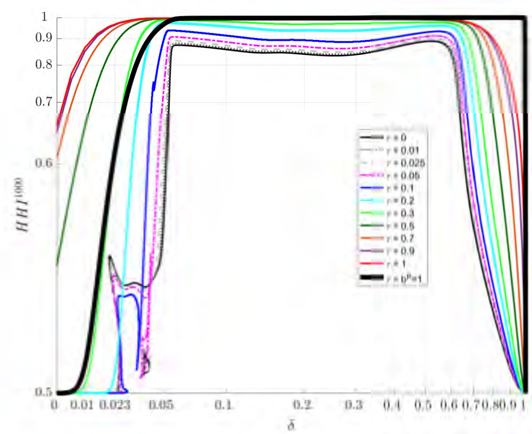

(i) $\chi = 50, \psi = 0.75$ ρ Homotopies

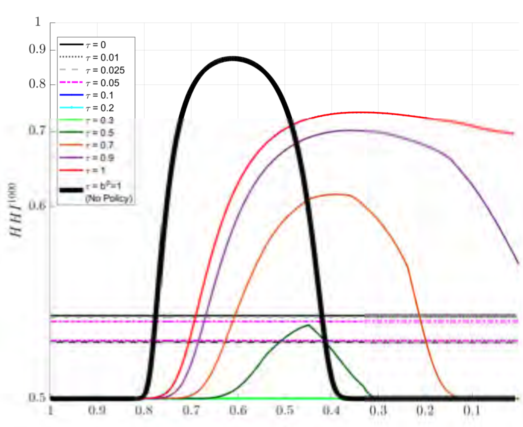

(j) $\chi = 50, \psi = 0.75$ δ Homotopies

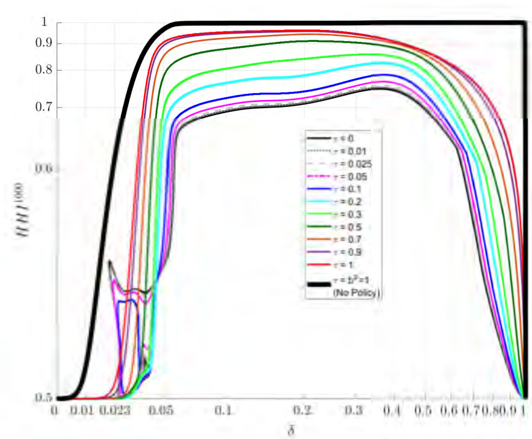

(k) No Leader AD Incentive ρ Homotopies

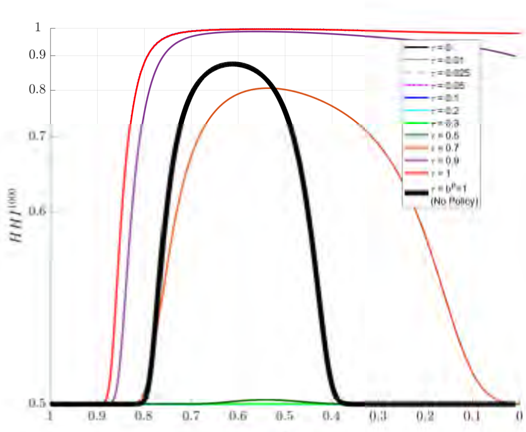

(l) No Leader AD Incentive δ Homotopies

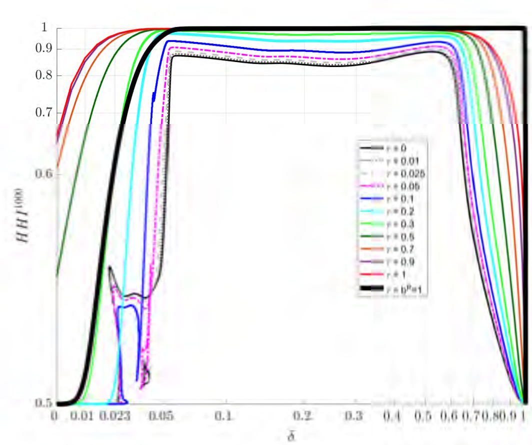

Figure C.2: PDV CS for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. For the policies, we show the changes in TS caused by policy.

(m) No Policy ρ Homotopies

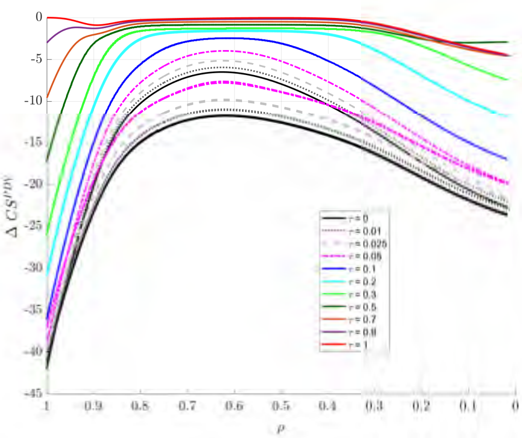

(n) No Policy δ Homotopies

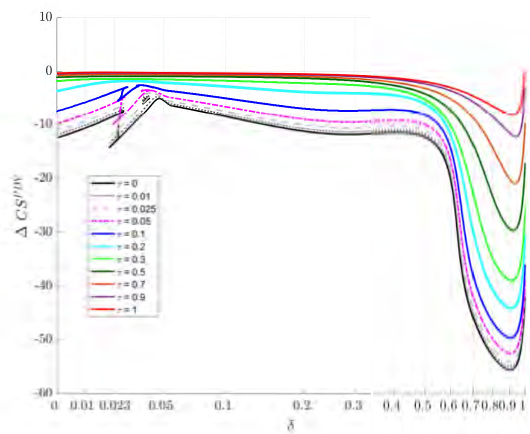

(o) $\chi = 50, \psi = 0.75$ ρ Homotopies

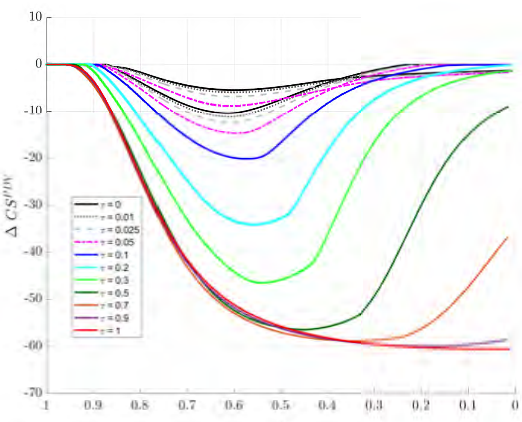

(p) $\chi = 50, \psi = 0.75$ δ Homotopies

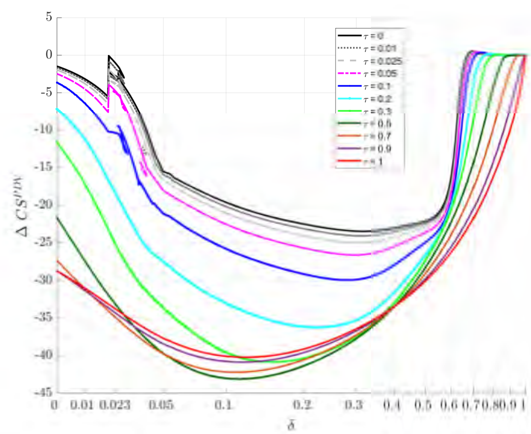

(q) No Leader AD Incentive ρ Homotopies

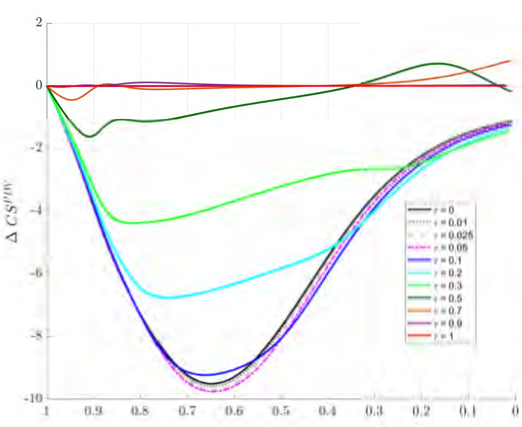

(r) No Leader AD Incentive δ Homotopies

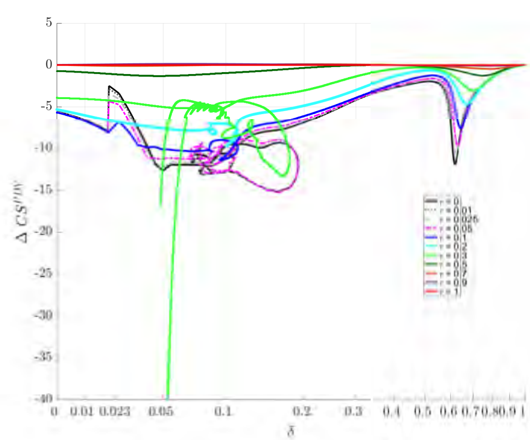

Figure C.2: PDV TS for Alternative τ as ρ Varies (with $\delta = 0.023$) and δ Varies (with $\rho = 0.75$) with No Policy and Under Threshold and Incentive Policies. $b^p = 0$. For the policies, we show the changes in TS caused by policy.

(s) No Policy ρ Homotopies

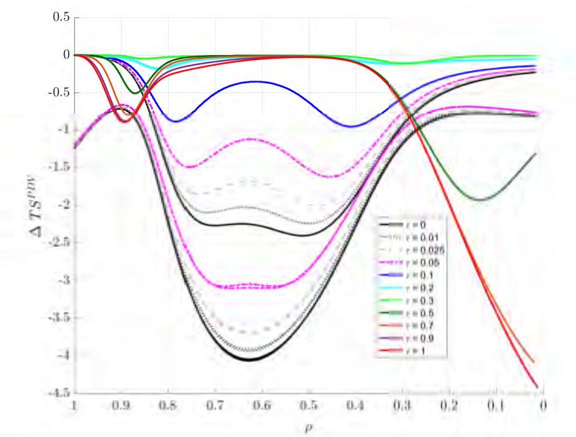

(t) No Policy δ Homotopies

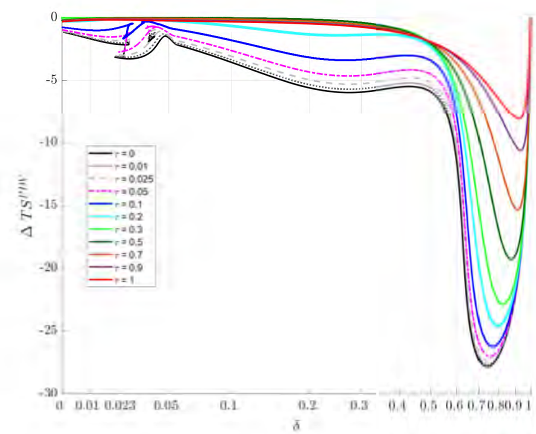

(u) $\chi = 50, \psi = 0.75$ ρ Homotopies

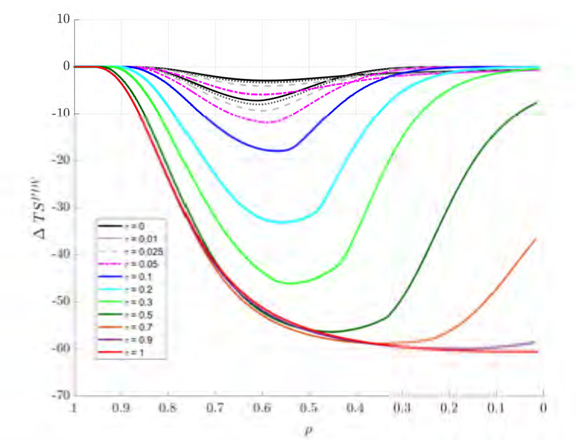

(v) $\chi = 50, \psi = 0.75$ ρ Homotopies

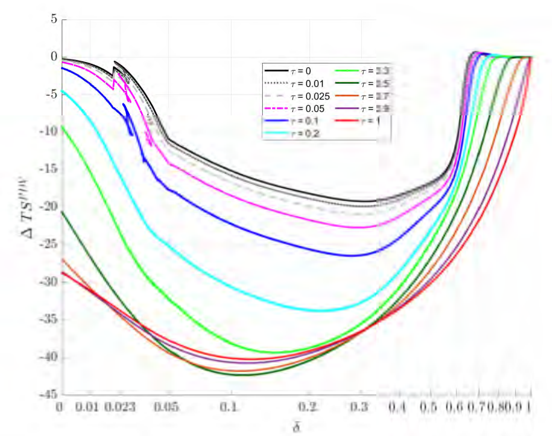

(w) No Leader AD Incentive ρ Homotopies

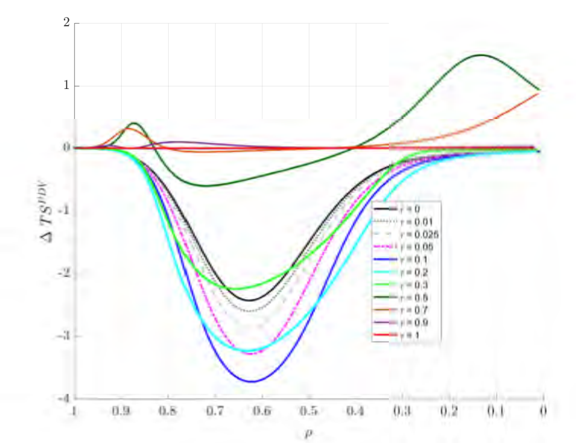

(x) No Leader AD Incentive δ Homotopies

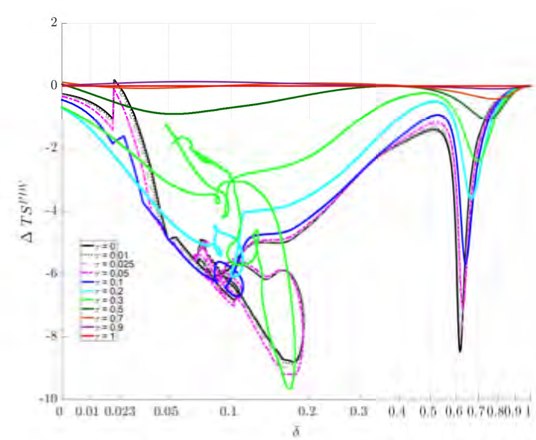

Figure C.3: The Effects of Changes in b^p with $\tau = 0$ For the Illustrative Parameters.

(a) Demand in State (3,2) with Seller Strategies Fixed at $b^p = 0$ Values

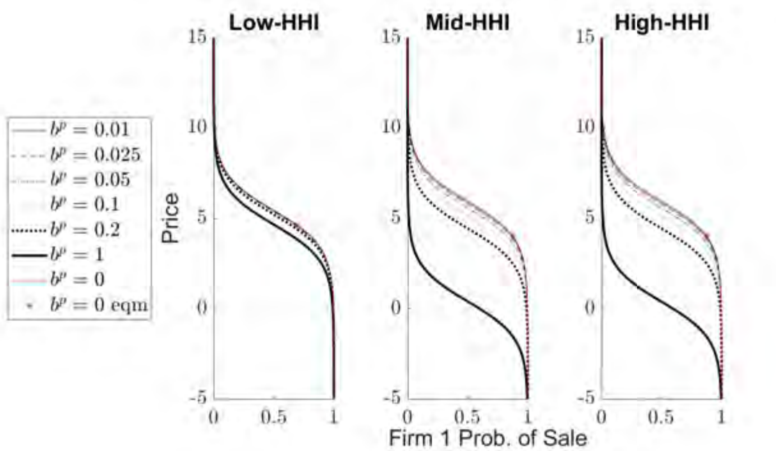

(b) Seller Incentives in State (3,2) with Seller Strategies Fixed at $b^p = 0$ Values

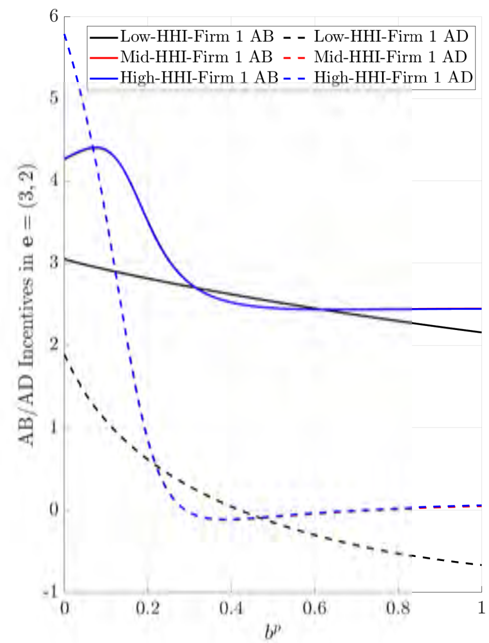

(c) (3,2) Prices: b^p homotopy

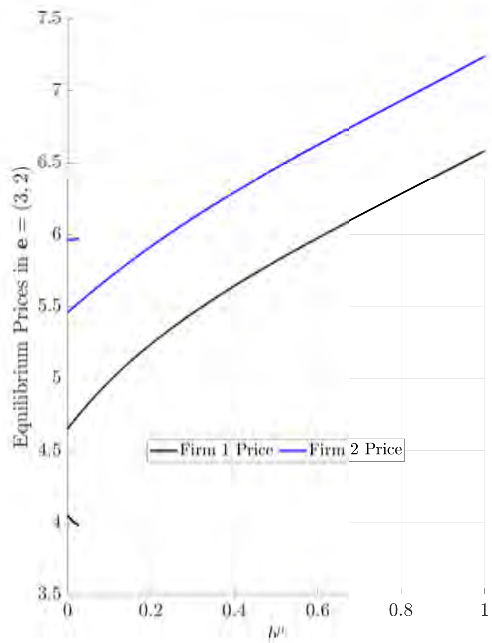

(d) Seller Incentives in State (3,2) : b^p homotopy

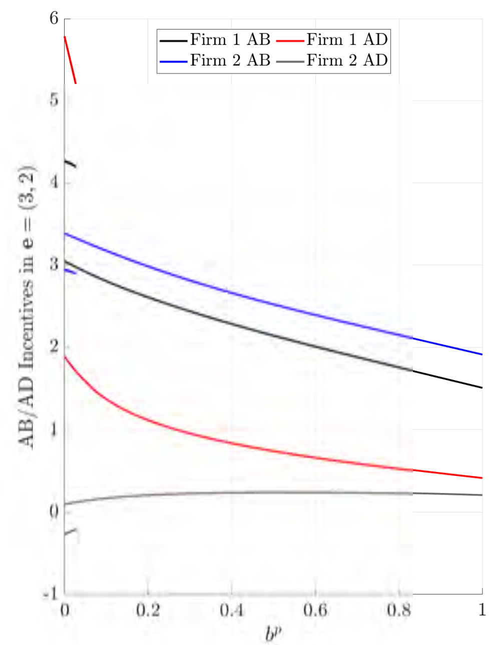

Figure C.4: Effects of Incentive Policies For $\rho = 0.9$ and $\delta = 0.01$

(a) HHI^{32}

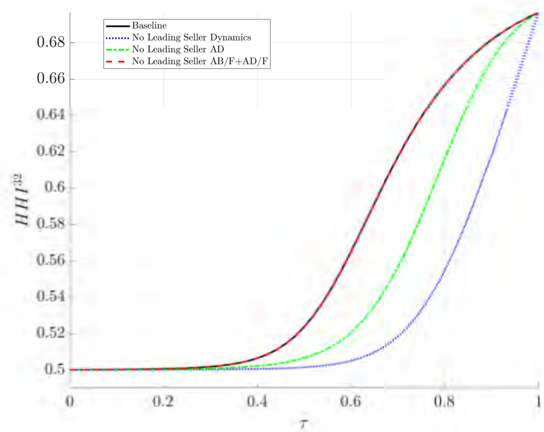

(b) HHI^{1000}

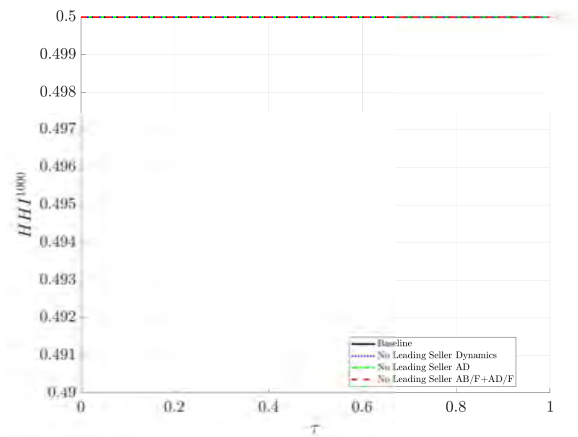

(c) PDV CS

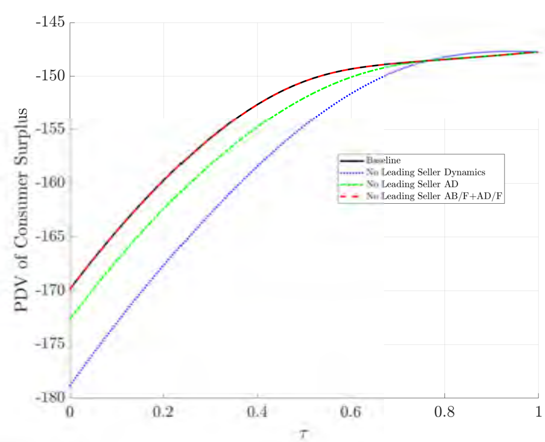

(d) PDV TS

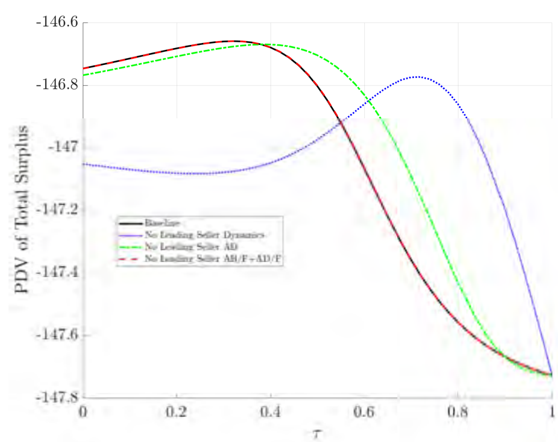

(e) CS^{32}

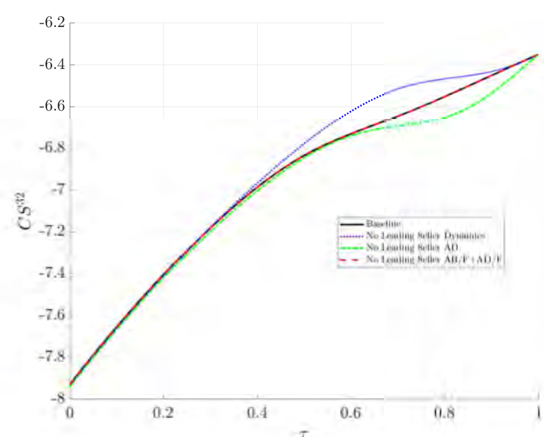

(f) CS^{1000}

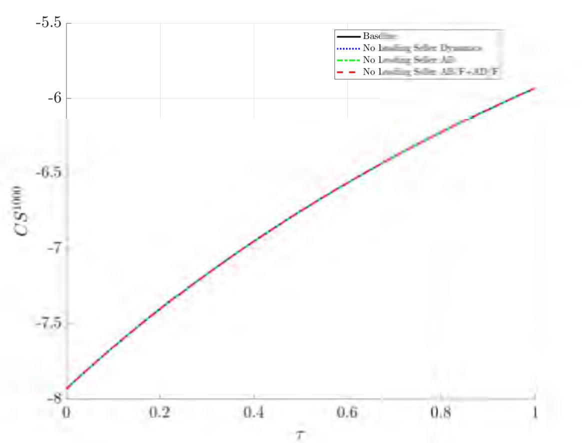

Figure C.5: Effects of Incentive Policies For $\rho = 0.65$ and $\delta = 0.01$

(a) HHI^{32}

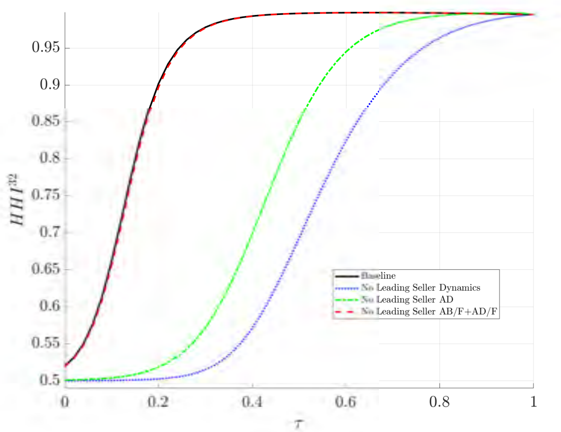

(b) HHI^{1000}

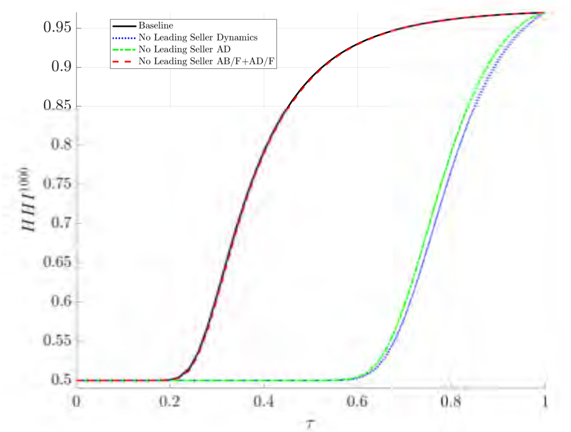

(c) PDV CS

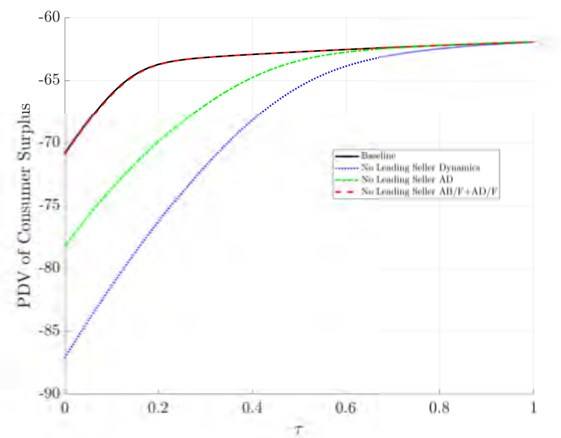

(d) PDV TS

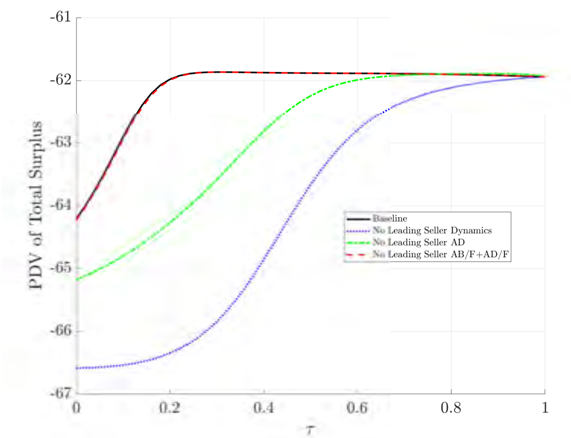

(e) CS^{32}

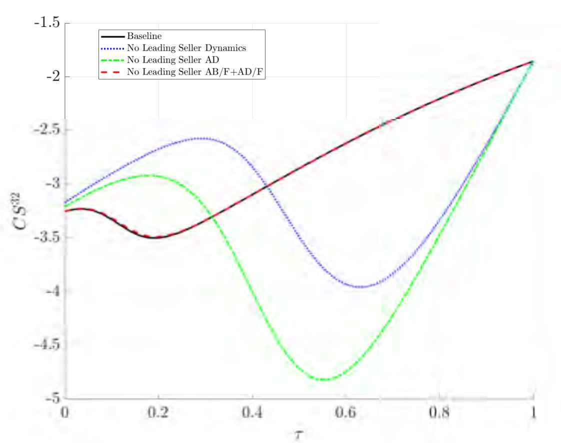

(f) CS^{1000}

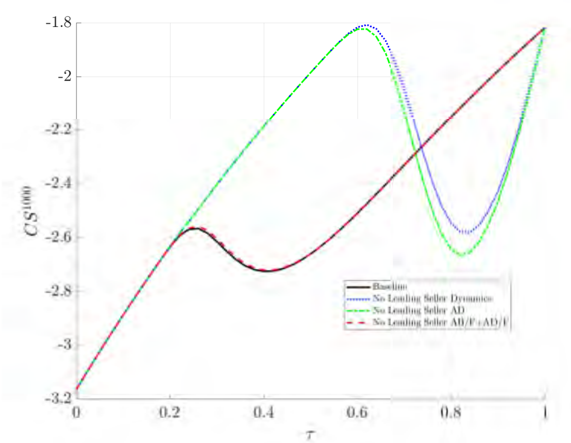

Figure C.6: Effects of Incentive Policies For $\rho = 0.9$ and $\delta = 0.05$

(a) HHI^{32}

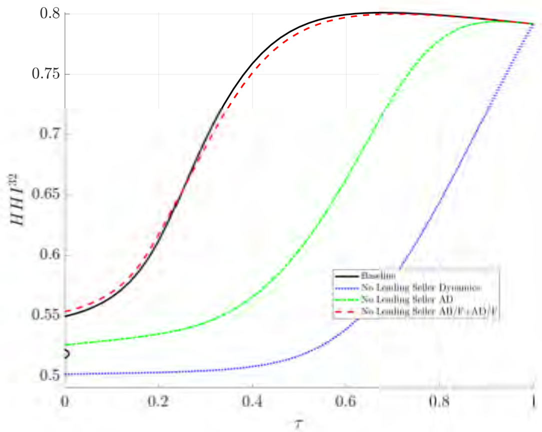

(b) HHI^{1000}

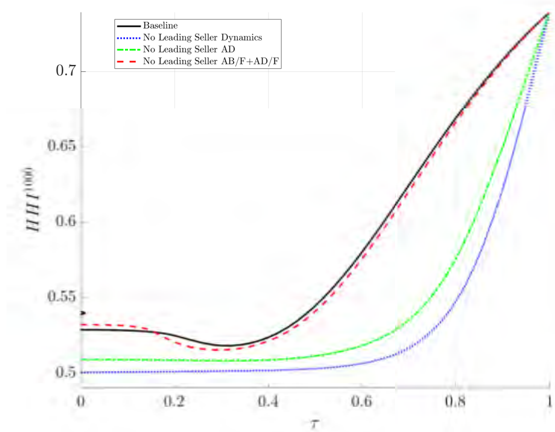

(c) PDV CS

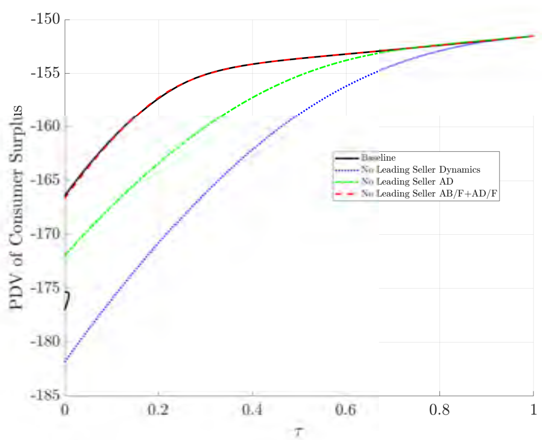

(d) PDV TS

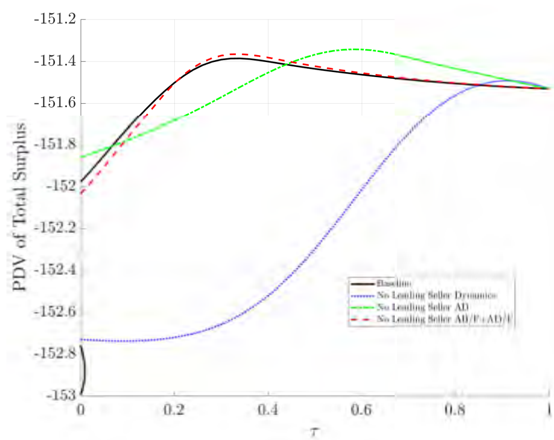

(e) CS^{32}

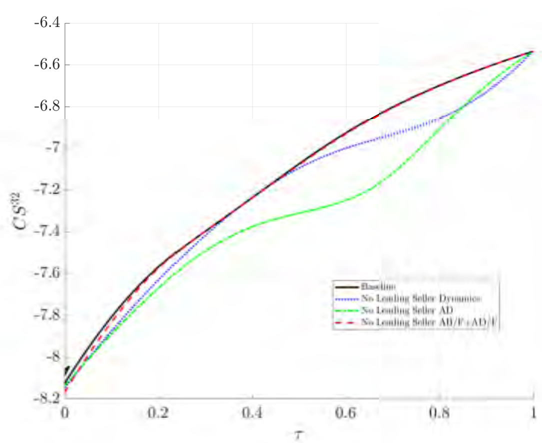

(f) CS^{1000}

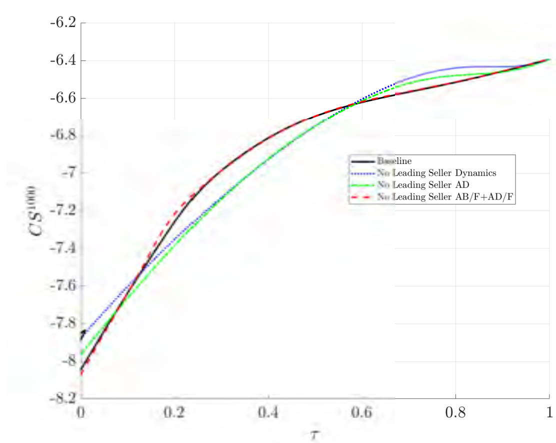

Figure C.7: Effects of Incentive Policies For $\rho = 0.65$ and $\delta = 0.05$

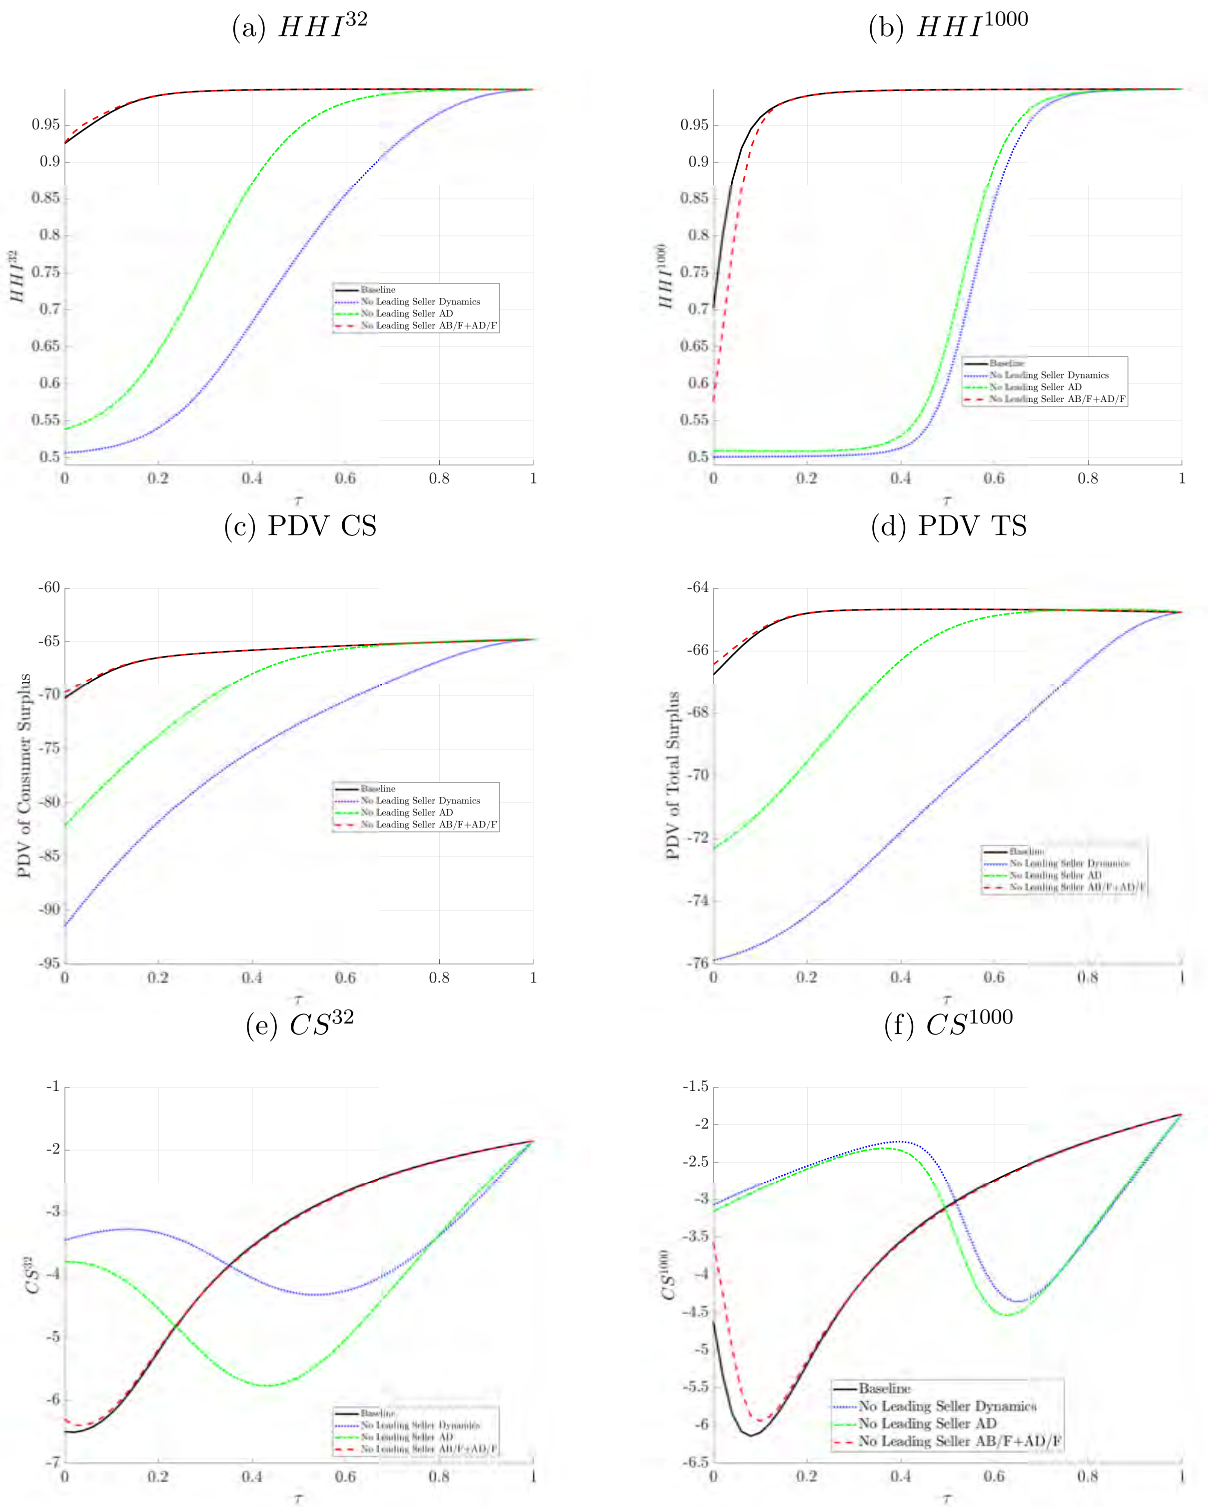

Figure C.8: Medium-Run and Long-Run Changes in Concentration and Expected Surplus as a Function of τ and Incentive Policies. Illustrative Parameters. $b^p = 0$.

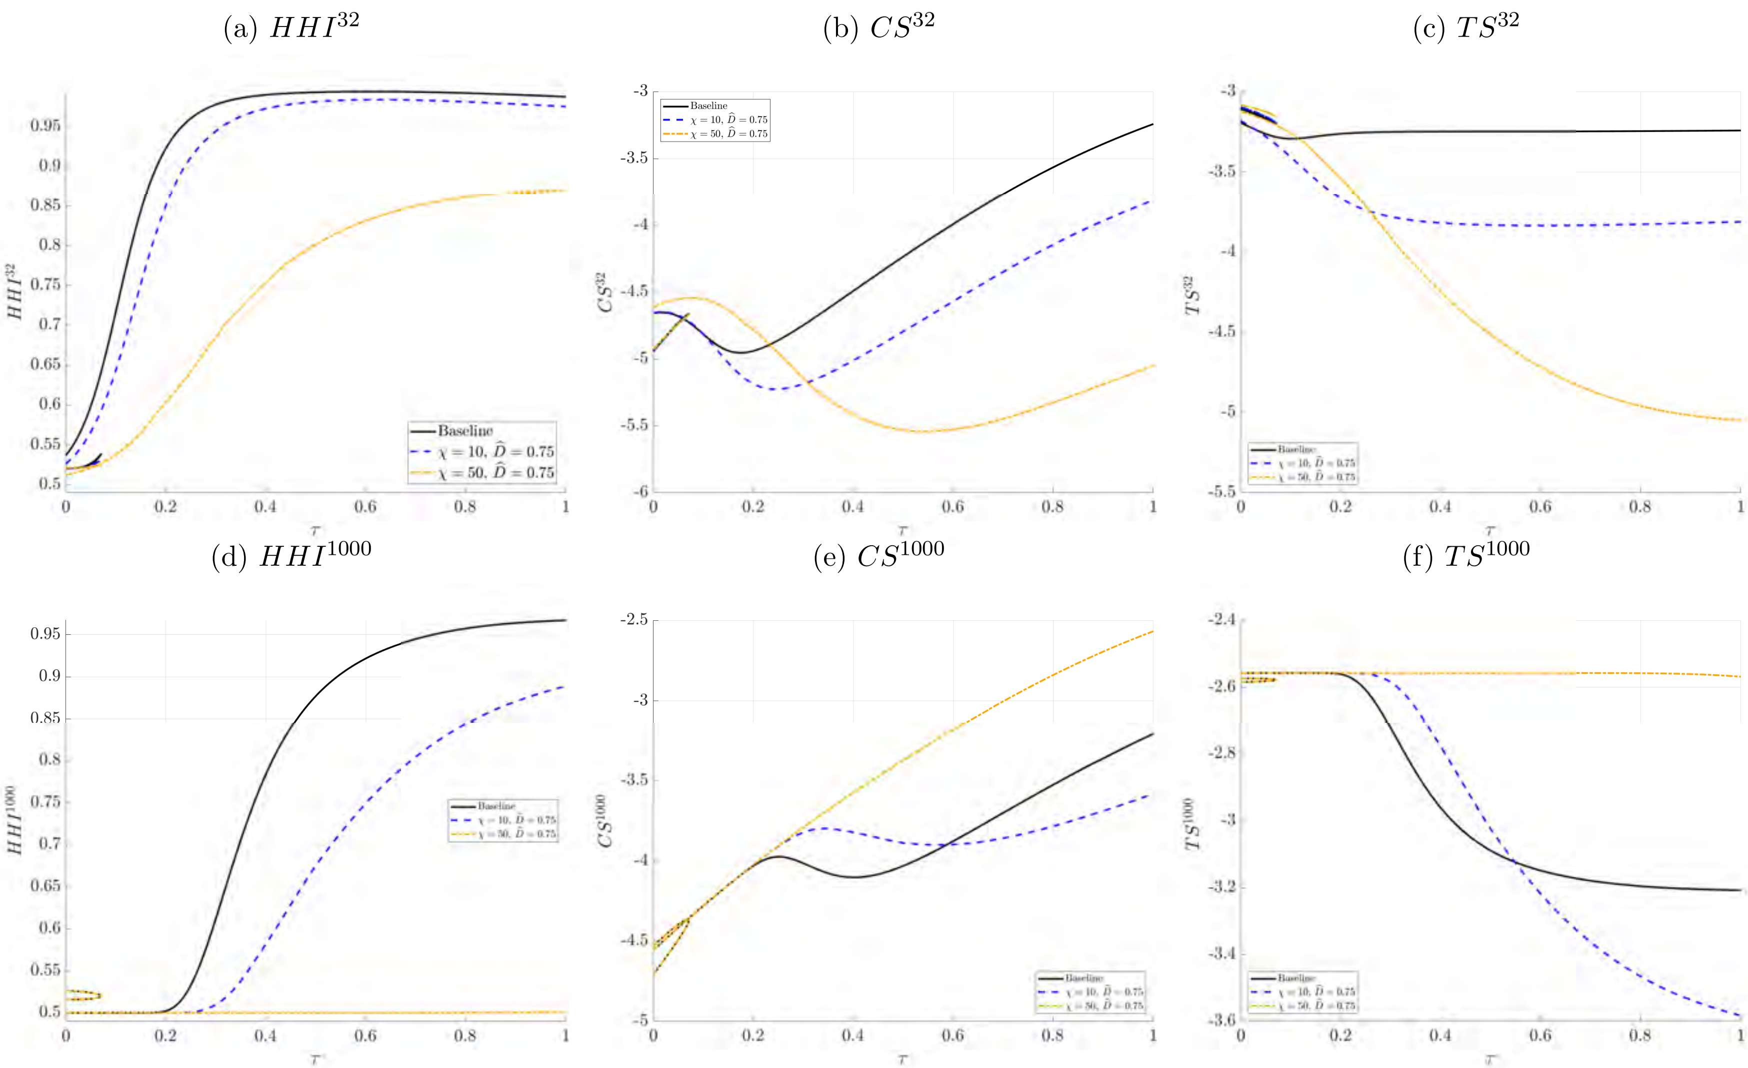

Figure C.9: Market Outcomes for Alternative b^p as δ Varies (with $\rho = 0.75$) for $\tau = 0$ and $\tau = 0.5$.

(a) HHI^{32} $\tau = 0$

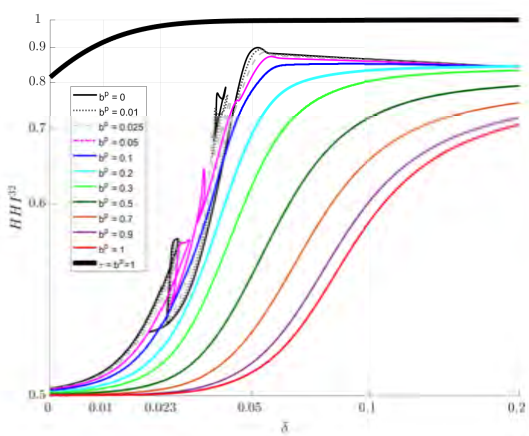

(b) HHI^{32} $\tau = 0.5$

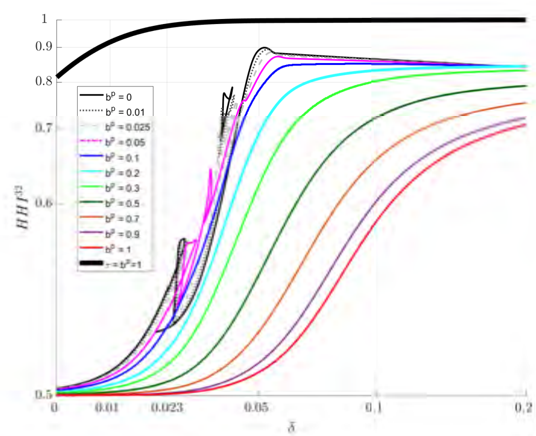

(c) Total Surplus (Relative to TS When $\tau = 1$ and $b^p = 1$) $\tau = 0$

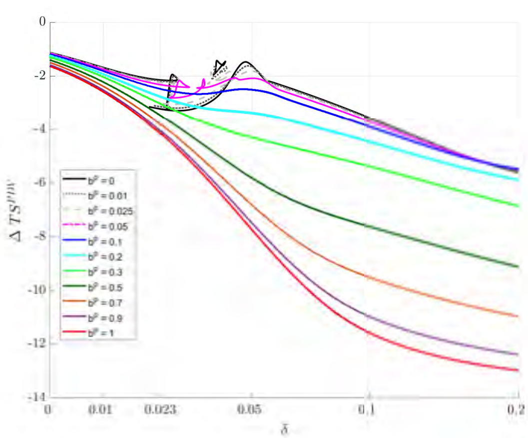

(d) Total Surplus (Relative to TS When $\tau = 1$ and $b^p = 1$) $\tau = 0.5$

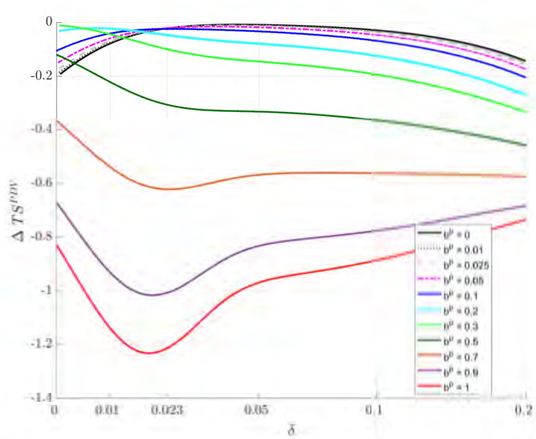

(e) Consumer Surplus (Relative to CS When $\tau = 1$ and $b^p = 1$) $\tau = 0$

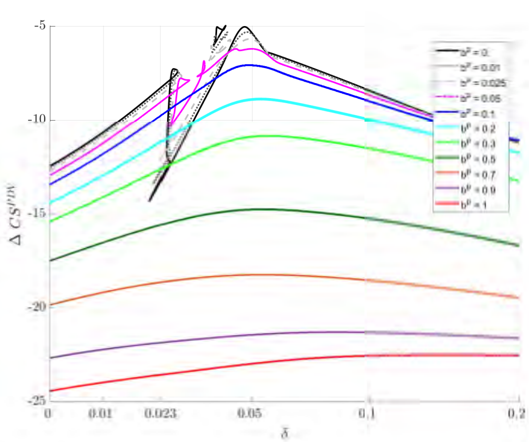

(f) Consumer Surplus (Relative to CS When $\tau = 1$ and $b^p = 1$) $\tau = 0.5$