

PHBS WORKING PAPER SERIES

Bargaining and Dynamic Competition

Shanglyu Deng
University of Macau

Dun Jia
Peking University

Mario Leccese
University of Maryland

Andrew Sweeting
University of Maryland

June 2024

Working Paper 20240601

Abstract

Industries with significant scale economies or learning-by-doing may come to be dominated by a single firm. Economists have studied how likely this is to happen, and whether it is efficient, using models where buyers are price or quantity takers, even though these industries are often also characterized by buyer-seller negotiations. We extend the dynamic “learning-by-doing and forgetting” model of Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) to allow for Nashin-Nash bargaining over prices. Price-taking and the social planner solution are captured as special cases. We show that sellers’ dynamic incentives, market concentration and welfare can change sharply, and non-monotonically, as one moves away from the price-taking assumption. We study the implications of buyer bargaining power for the existence of multiple equilibria, the design of subsidy policies and the welfare effects of policies designed to increase competition.

JEL Classification: C73, D21, D43, L13, L41

Peking University HSBC Business School
University Town, Nanshan District
Shenzhen 518055, China



PHBS 商界学校
北京大学汇丰商学院

Bargaining and Dynamic Competition
Shanglyu Deng, Dun Jia, Mario Leccese, and Andrew Sweeting
NBER Working Paper No. 32360
April 2024
JEL No. C73,D21,D43,L13,L41

ABSTRACT

Industries with significant scale economies or learning-by-doing may come to be dominated by a single firm. Economists have studied how likely this is to happen, and whether it is efficient, using models where buyers are price or quantity takers, even though these industries are often also characterized by buyer-seller negotiations. We extend the dynamic “learning-by-doing and forgetting” model of Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) to allow for Nash-in-Nash bargaining over prices. Price-taking and the social planner solution are captured as special cases. We show that sellers’ dynamic incentives, market concentration and welfare can change sharply, and non-monotonically, as one moves away from the price-taking assumption. We study the implications of buyer bargaining power for the existence of multiple equilibria, the design of subsidy policies and the welfare effects of policies designed to increase competition.

Shanglyu Deng
Department of Economics
University of Macau
China
shanglyudeng@gmail.com

Dun Jia
Peking University HSBC Business School
University Town, Nanshan District
Shenzhen 518055
China
dun.jia@phbs.pku.edu.cn

Mario Leccese
University of Maryland - Department of Economics
Tydings Hall
7343 Preinkert Dr.
College Park, MD 20742
United States
leccese@umd.edu

Andrew Sweeting
Department of Economics
University of Maryland
Tydings Hall 3135
College Park, MD 20742
and NBER
sweeting@econ.umd.edu

An online appendix is available at <http://www.nber.org/data-appendix/w32360>

1 Introduction

Industries where some feature of demand or supply may allow a firm that achieves an initial advantage to attain a position of lasting market dominance often attract the attention of policy-makers. For example, industrial and trade policies have been implemented or proposed for industries making electric vehicle batteries, semiconductors, solar panels and aircraft, where scale economies or learning-by-doing (LBD) are important, while, in a different type of setting, antitrust agencies and regulators are increasingly active when addressing the largest players in digital markets characterized by network effects. Even when non-economic concerns motivate policies, models of dynamic competition are required to evaluate their effects.

Existing theoretical or empirical models of dynamic competition in industrial organization (for example, Fudenberg and Tirole (1983), Cabral and Riordan (1994), Benkard (2000), Besanko, Doraszelski, Kryukov, and Satterthwaite (2010), Besanko, Doraszelski, and Kryukov (2014)) and international trade (Dasgupta and Stiglitz (1988), Leahy and Neary (1999), Neary and Leahy (2000)) assume that sellers unilaterally set prices or quantities. These models therefore miss how prices are negotiated in many industries, especially for capital goods, where dynamic competition is important.

This paper examines how bargaining affects outcomes and optimal policies in a model of dynamic price competition, where buyers arrive repeatedly and market structure is endogenous. To study novel interactions between bargaining and dynamics, we extend the well-known dynamic duopoly seller “learning-by-doing (LBD) and forgetting” computational model of Besanko, Doraszelski, Kryukov, and Satterthwaite (2010) (BDKS). Specifically, we replace the assumption that sellers simultaneously set take-it-or-leave-it prices each period with a more general assumption that prices are determined by a form of “Nash-in-Nash” bargaining between the buyer and the sellers (Horn and Wolinsky (1988), Collard-Wexler, Gowrisankaran, and Lee (2019)). The Nash-in-Nash structure, embedded in static models, has become widely used

to understand business-to-business transactions in healthcare (e.g., Gowrisankaran, Nevo, and Town (2015), Ho and Lee (2017), Grennan (2013)), cable television (e.g., Crawford, Lee, Whinston, and Yurukoglu (2018))) and consumer packaged goods settings (e.g., Draganska, Klapper, and Villas-Boas (2010)).

Assuming that there is a myopic single-unit buyer every period, an assumption that we maintain, and that sellers set prices, BDKS compute symmetric Markov Perfect Nash equilibria, assuming thirty levels of know-how ($M = 30$) for each firm. Equilibrium prices equal a markup plus the seller's opportunity cost of sale. The opportunity cost equals the current production cost, which falls with know-how, less how much the seller's expected future profit increases when it makes a sale (dynamic incentives). BDKS search for equilibria using numerical homotopies and find that equilibria where sellers set low prices in states where they are symmetric, even when they have accumulated significant know-how, and one seller will typically tend to have a significant lead in equilibrium, often co-exist with equilibria where prices are less sensitive to know-how and more symmetric market structures will tend to emerge.

We focus on how the allocation of bargaining power, determined by a single parameter, affects market structure and welfare. Shifting bargaining power to buyers will lower seller markups and, holding the evolution of market structure fixed, will tend to reduce sellers' dynamic incentives. However, markup shrinkage will also lead to the seller with the larger markup, typically the leader, making more sales so that leads last longer. When sellers' bargaining power is not too limited and leads are initially short-lived, this lead lengthening can increase the incentives of a seller to attain and to preserve a lead. We show, for a wide range of empirically relevant learning parameters, that this effect can cause market structure and welfare outcomes to change significantly and quickly as we move away from the standard assumption that sellers make take-it-or-leave-it offers.

We explore how the allocation of bargaining power changes the design and the effects of stylized policies. We show, for example, that a subsidy scheme that would

implement the social planner outcome with price-setting can be worse than no scheme at all when sellers only have much greater bargaining power than buyers. We also consider policies that have been considered as likely to promote more symmetric market structures, possibly at the cost of softening competition. We find that the effects of these policies on concentration and welfare can also be sensitive to the assumed allocation of bargaining power.

We also make methodological contributions. We reformulate the equations defining symmetric equilibria in terms of the probability that a laggard makes a sale, rather than prices and values. This reduces the number of equations from $2M^2$ to $\frac{M(M-1)}{2}$, allowing us to provide some limited analytical results for the $M = 2$ model, to use a non-homotopy method to enumerate equilibria when $M = 3$ (allowing us to confirm homotopy results), and to formulate easy-to-solve equations for optimal subsidies even when M is large. We find that, for both $M = 3$ and $M = 30$, the multiplicity of equilibria that is common with price-setting disappears when buyers have more than limited bargaining power.

By extending a price-setting model to allow for bargaining, our paper contributes directly to the literature studying dynamic competition with LBD that spans the analytic insights of Cabral and Riordan (1994), the computational analyses of BDKS, Besanko, Doraszelski, and Kryukov (2014), Besanko, Doraszelski, and Kryukov (2019a), Besanko, Doraszelski, and Kryukov (2019b) and Sweeting, Jia, Hui, and Yao (2022) (SJHY), and the empirical analyses of mergers and industrial policies by Benkard (2004), Kim (2014), Kalouptsidi (2018), An and Zhao (2019), and Barwick, Kalouptsidi, and Zahur (2019).¹ Our policy examples are very stylized, but they suggest how allowing for bargaining could also change the policy recommendations from richer models. SJHY assume price-setting but extend the Besanko, Doraszelski, and Kryukov (2014) model to allow for forward-looking buyers who expect to purchase in

¹A distinct literature on industry dynamics (Abbring and Campbell (2010) and Abbring, Campbell, Tilly, and Yang (2018)) assumes that firms are symmetric to guarantee equilibrium uniqueness in order to simplify an analysis of how industries respond to aggregate shocks.

some share of future periods. We will maintain BDKS’s assumption that buyers are myopic in this paper, except when we solve the social planner’s problem by assuming that there is a long-lived buyer who captures all surplus.

The Besanko, Doraszelski and Kryukov papers use a model (“BDK model”) where sellers can enter or exit, and there is no forgetting. We focus on the BDKS model for three reasons. First, reallocating bargaining power away from sellers in a model with fixed costs (or entry costs and scrap values) could cause the industry to disappear entirely. This effect is no doubt important, but it is distinct from the effects of bargaining power on dynamic incentives and leadership on which we focus. Second, our reformulation has particular benefits in the BDKS model where there is no outside good, although we allow for an outside good as a robustness check. Third, competitive dynamics in the BDKS model are arguably richer because stochastic know-how depreciation implies that a firm can hope to weaken even a well-established rival by depriving the rival of sales.

While most applications of Nash-in-Nash bargaining have been within static models, a small number of papers consider dynamic settings. Lee and Fong (2013) assume period-by-period Nash-in-Nash bargaining in a game where hospitals and insurance companies form networks that change stochastically, while Dorn (2023) shows how Kalai (1977) proportional bargaining facilitates consideration of multiperiod hospital-insurer contracts with index clauses. Yang (2020) embeds Nash-in-Nash bargaining over prices into a dynamic model of innovation in the smartphone device industry and Tiew (2024) considers a dynamic game where duopoly newspapers may bargain over forming or continuing a joint operating agreement. In the current paper bargained prices, through their impact on sales, directly affect sellers’ competitiveness and market structure in future periods, which is the central issue in the dynamic competition literature.

The paper is structured as follows. Section 2 outlines the model. Section 3 details the two alternative formulation of the equilibrium equations, our outcome measures

and analytic results. Section 4 uses the small state space of the $M = 3$ model to illustrate how reallocating bargaining power changes incentives and strategies. Section 5 provides results and a wider set of policy analyses for the $M = 30$ model. Section 6 concludes. The online Appendices contain proofs, details of methods and additional results.

2 Model

This section presents the model, focusing on where we depart from BDKS. Readers should consult BDKS for additional motivation.

2.1 States and Costs.

There is an infinite horizon, discrete time, discrete state game. The common discount factor is $\beta = \frac{1}{1.05}$.²

Sellers. There are two long-lived ex-ante symmetric but differentiated sellers ($i = 1, 2$). i 's production cost is $c(e_i) = \kappa \rho^{\log_2(\min(e_i, m))}$, where $e_i = 1, \dots, M$ is a commonly observed state variable that tracks the seller's "know-how".³ The industry state is $\mathbf{e} = (e_1, e_2)$. We will call a seller with higher know-how than its rival the "leader", and its rival the "laggard". A lower value of the progress ratio $\rho \in [0, 1]$ implies stronger learning economies. BDKS assume $m = 15$ and $M = 30$. We will also consider $m = M = 2$ and $m = M = 3$, the smallest state space where a leader can have either a "small" or a "large" advantage.

As described below, one unit will be purchased from one of the sellers in each

²As BDKS discuss, one can interpret this discount factor as corresponding to a period length of 1 month, a monthly discount rate of 1 percent and a probability that the industry survives each month of 0.96.

³Asker, Fershtman, Jeon, and Pakes (2020), Sweeting, Roberts, and Gedge (2020) and Sweeting, Tao, and Yao (2023) consider dynamic models where serially correlated state variables are private information.

period. e_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 and only 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 $\delta \in [0, 1)$.⁴ The probability of forgetting (Δ) is increasing in both δ and e_i . The know-how of a firm that makes a sale can increase by one or stay the same, while the know-how of the firm that does not make the sale will stay the same or decrease by one.

Buyers. Every period a buyer arrives who will purchase exactly one unit from one of the firms. A buyer is assumed to live only one period, and to choose the seller that maximizes her indirect utility, $v - p_i + \sigma \varepsilon_i$, where p_i is the price paid, and the ε_i s are i.i.d. private information Type I extreme value payoff shocks. σ parameterizes seller product differentiation.

Bargaining. BDKS assume sellers make simultaneous take-it-or-leave-it price offers. We generalize by assuming that, before the ε s are realized, the buyer sends separate agents to each seller. Each agent-seller pair negotiates the price at which the buyer would be able to purchase from that seller. After negotiations are complete, the buyer observes the ε s and makes their purchase choice based on the negotiated prices with a purchase possible only if a price was agreed. We make the Nash-in-Nash assumption (Collard-Wexler, Gowrisankaran, and Lee (2019)) that each buyer agent-seller pair takes the price agreed by the other pair as given. We assume that the bargaining weight of the buyer’s agent is a parameter τ , with the bargaining weight of the seller $1 - \tau$. If $\tau = 1$, buyers have “all of the bargaining power”, and will extract all of the expected surplus from an agreement, whereas we will replicate

⁴At 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 ($f_{i,t} = 1$).

BDKS's equilibria if $\tau = 0$.

The Social Planner. We also compute outcomes given the choice probabilities that would be used by a social planner buyer that maximizes discounted expected total surplus (i.e., buyer utility less production costs) and only knows the current values of the ε s.⁵

Parameters. In the text we will assume $\sigma = 1$, $\kappa = 10$ and $v = 10$, although, with no outside good, the value of v does not affect equilibrium choice probabilities. Our focus will be on the technology parameters ρ and δ , and, especially, on what happens as τ varies from 0 to 1.

If $\tau \neq 0$, what values of τ are “reasonable” or “relevant”? $\tau = 0.5$, equal buyer-seller bargaining power, is often assumed in applications. As we will illustrate, equilibrium prices, and outcomes, when $\tau = 0.5$ will often not be close to the averages of prices when $\tau = 0$ and $\tau = 1$, as they would often be in the static model, because opportunity costs are endogenous. Alternatively, if one wants to interpret τ as reflecting differences in buyer and seller patience in an alternating offer process taking place in a negotiation, one could view our assumption that buyers are short-lived and sellers long-lived as being most consistent with τ s close to zero. Under this interpretation, the main interest will be in how outcomes often change rapidly as τ increases from zero.

3 Equilibrium and Outcomes

The equilibrium concept is symmetric and stationary Markov Perfect Nash equilibrium (MPNE, Maskin and Tirole (2001), Ericson and Pakes (1995), Pakes and McGuire (1994)). We now describe two alternative characterizations of equilibria,

⁵One can view our social planner as analogous to a single long-lived buyer in SJHY's model who faces prices equal to production costs each period.

how we measure outcomes and incentives, and results that can be proven analytically.

3.1 Formulation of Equilibrium Conditions for Prices and Values.

BDKS specify equations for equilibrium seller values, $VS^*(\mathbf{e})$, defined at the start of each period, and prices, $p^*(\mathbf{e})$. We extend their formulation to allow for bargaining. Symmetry implies that we only need to define equations for the prices and values of seller 1 (i.e., $p_2^*(e_1, e_2) = p_1^*(e_2, e_1)$ and $VS_2^*(e_1, e_2) = VS_1^*(e_2, e_1)$).

Beginning of period value for firm 1 (VS):

$$VS_1^*(\mathbf{e}) - D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - c(e_1)) - \sum_{k=1,2} D_k^*(\mathbf{e})\mu_{1,k}^S(\mathbf{e}) = 0, \quad (2)$$

where $\mu_{1,k}^S(\mathbf{e})$ is seller 1's continuation value when seller k makes the sale,

$$\mu_{1,k}^S(\mathbf{e}) = \beta \sum_{\forall e'_{1,t+1}|e_{1,t}} \sum_{\forall e'_{2,t+1}|e_{2,t}} VS_1^*(e'_{1,t+1}, e'_{2,t+1}) \Pr(e'_{1,t+1}|e_{1,t}, k) \Pr(e'_{2,t+1}|e_{2,t}, k), \quad (3)$$

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 $q_{k,t} = 1$. Given prices, the probability that $q_{k,t} = 1$, $D_k(\mathbf{e})$, is

$$D_k(\mathbf{e}) = \frac{\exp\left(\frac{v-p_k(\mathbf{e})}{\sigma}\right)}{\exp\left(\frac{v-p_1(\mathbf{e})}{\sigma}\right) + \exp\left(\frac{v-p_2(\mathbf{e})}{\sigma}\right)}. \quad (4)$$

$D_k^*(\mathbf{e})$ is the choice probability given equilibrium prices.

Negotiated prices (p): Our bargaining assumptions imply that, given p_2 , $p_1^*(\mathbf{e})$ will be

determined as

$$p_1^*(\mathbf{e}) = \arg \max_{p_1} (CS(p_1, p_2, \mathbf{e}) - CS(p_2, \mathbf{e}))^\tau \times \dots$$

$$(D_1(\mathbf{e})(\mu_{1,1}^S(\mathbf{e}) + p_1 - c(e_1)) + (1 - D_1(\mathbf{e}))\mu_{1,2}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e}))^{(1-\tau)} \quad (5)$$

where $CS(p_1, p_2, \mathbf{e}) = \sigma \log \left(\sum_{k=1,2} \exp \left(\frac{v - p_k(\mathbf{e})}{\sigma} \right) \right)$ (i.e., the expected future surplus of the buyer when it is able to choose from both firms). $CS(p_2, \mathbf{e}) = v - p_2(\mathbf{e})$ is the buyer's expected surplus if there is no agreement with seller 1 and seller 2 is the buyer's only option. Equilibrium $p_1^*(\mathbf{e})$ will therefore solve the first-order condition

$$\tau \frac{\partial CS(p_1^*(\mathbf{e}), p_2, \mathbf{e})}{\partial p_1} (D_1^*(\mathbf{e})(\mu_{1,1}^S(\mathbf{e}) + p_1^*(\mathbf{e}) - c(e_1)) + (1 - D_1^*(\mathbf{e}))\mu_{1,2}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e})) + \dots$$

$$(1 - \tau) (CS(p_1^*(\mathbf{e}), p_2, \mathbf{e}) - CS(p_2, \mathbf{e})) \left(D_1^*(\mathbf{e}) + \frac{\partial D_1^*(\mathbf{e})}{\partial p_1} (p_1^*(\mathbf{e}) - c(e_1) + \mu_{1,1}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e})) \right) = 0, \quad (6)$$

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

$$-\tau D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - \widehat{c}_1(\mathbf{e})) + (1 - \tau) [\sigma - (1 - D_1^*(\mathbf{e}))(p_1^*(\mathbf{e}) - \widehat{c}_1(\mathbf{e}))] \log \frac{1}{1 - D_1^*(\mathbf{e})} = 0. \quad (7)$$

$\widehat{c}_1(\mathbf{e}) = c(e_1) - (\mu_{1,1}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e}))$ is seller 1's opportunity cost of a sale where $(\mu_{1,1}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e}))$ is the increase in the seller's continuation values (equation (3)) when it, rather than its rival, makes the sale.

This first-order condition is the same as BDKS's if $\tau = 0$. If $\tau = 1$, the μ^S s will equal zero, and $p_1^*(\mathbf{e}) = c(e_1)$ will be the only solution.

3.2 Formulation of Equilibrium Conditions in Terms of Buyer Choice Probabilities.

We can rewrite the equilibrium choice probability equations as

$$\sigma \log \left(\frac{1}{D_1^*(\mathbf{e})} - 1 \right) - p_1^*(\mathbf{e}) + p_2^*(\mathbf{e}) = 0. \quad (8)$$

We will now show that prices can be expressed as functions of D_1^* s and parameters, so that the equilibrium can be characterized by equations (8) only.

From equation (7), firm 1's markup over its opportunity cost is

$$p_1^*(\mathbf{e}) - \widehat{c}_1(\mathbf{e}) = \Phi(D_1^*(\mathbf{e})) = \frac{(1 - \tau)\sigma \log \frac{1}{1 - D_1^*(\mathbf{e})}}{\tau D_1^*(\mathbf{e}) + (1 - \tau)(1 - D_1^*(\mathbf{e})) \log \frac{1}{1 - D_1^*(\mathbf{e})}}. \quad (9)$$

$\Phi(\mathbf{D}_1)$ will denote the stacked vector of markups. Using \mathbf{Q}_k to denote the seller 1 state transition matrix when k makes a sale,

$$\widehat{\mathbf{c}}_1 = \mathbf{c}_1 - \beta(\mathbf{Q}_1 - \mathbf{Q}_2)\mathbf{V}\mathbf{S}_1, \quad (10)$$

and

$$\mathbf{V}\mathbf{S}_1 = (\mathbf{I} - \beta\mathbf{Q}_2)^{-1} [\mathbf{D}_1 \circ \Phi(\mathbf{D}_1)]. \quad (11)$$

where \circ denotes the element-wise product of two vectors. Therefore,

$$\mathbf{p}_1 = \Phi(\mathbf{D}_1) + \mathbf{c}_1 - \beta(\mathbf{Q}_1 - \mathbf{Q}_2)(\mathbf{I} - \beta\mathbf{Q}_2)^{-1} [\mathbf{D}_1 \circ \Phi(\mathbf{D}_1)]. \quad (12)$$

There are M^2 D_1^* s, but symmetry with no outside good implies $D_1^*(e, e) = \frac{1}{2}\forall e$ and $D_1^*(e, e') = 1 - D_1^*(e', e)$, so we can reduce the problem to $\frac{M(M-1)}{2}$ equations and unknowns. This compares to $2M^2$ equations for prices and seller values in the standard formulation. For $M = 30$, our new formulation reduces the problem from 1,800 equations to 435, and for $M = 3$ the reduction is from 18 equations to only 3.

3.3 Formulation of the Social Planner Problem.

We can find the choice probabilities that maximize the discounted value of expected surplus by thinking of an infinitely lived buyer, with the same discount factor as the

sellers, who pays the production cost associated with its choice. In this case,

$$D_1^{SP}(\mathbf{e}) = \frac{1}{1 + \exp\left(\frac{c_1(\mathbf{e}) - c_2(\mathbf{e}) + \mu_2^{SP}(\mathbf{e}) - \mu_1^{SP}(\mathbf{e})}{\sigma}\right)}$$

where μ_i^{SP} is the continuation discounted surplus when seller i is chosen. In vector form,

$$\mu_2^{SP} - \mu_1^{SP} = \beta(\mathbf{Q}_2 - \mathbf{Q}_1) \left(\mathbf{I} - \beta \sum_{k=1,2} \mathbf{D}_k^{SP} \circ \mathbf{Q}_k \right)^{-1} \sum_{k=1,2} \left[\mathbf{D}_k^{SP} \circ \left(\sigma \log \frac{1}{\mathbf{D}_k^{SP}} + v - \mathbf{c}_k \right) \right], \quad (13)$$

where the last term reflects the expected surplus in each state given the choice probabilities. We solve the social planner's choice probabilities using these two sets of equations.

3.4 Outcomes.

We calculate outcome measures assuming the state is (1,1) in $t = 1$.

Concentration. Following BDKS, expected market structure in period t is measured by

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

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 $\lambda^t(\mathbf{e})$ is the probability that a game will be in state \mathbf{e} after t periods. The minimum value of HHI^t is 0.5. One could look at concentration at many different points, but we will focus primarily on HHI^{32} as a measure of medium-run concentration.⁶

⁶For many parameters, HHI^t converges slowly to a long-run value. For example, when $M = 30$, $\rho = 0.75$ and $\delta = 0.023$, $HHI^t \approx 0.5$ for $t > 4,000$, but $HHI^{1,000} \approx 0.6$.

Surplus. With no outside good, v does not affect equilibrium strategies and outcomes, so we exclude it from our surplus measures. TS^t , expected total surplus in period t , is calculated as the expected ε of the purchased good less the production cost. This value will usually be negative. Our focus is on TS^{PDV} , the present discounted value of total surplus. TS^{SP} is the maximized present discounted surplus under a social planner. Consumer surplus is measured as the expected ε less the price paid, and producer surplus as expected price paid less production costs.

3.5 Incentives.

Recall that seller 1's opportunity cost of sale is $\widehat{c}_1(\mathbf{e}) = c(e_1) - (\mu_{1,1}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e}))$ where $c(e_1)$ is the current production cost. Besanko, Doraszelski, and Kryukov (2014) recognize that the dynamic component, $(\mu_{1,1}^S(\mathbf{e}) - \mu_{1,2}^S(\mathbf{e}))$, can be expressed as the sum of two incentives.

Definition 1 *The firm 1 “advantage building” (AB) incentive is $\mu_{1,1}^S - \mu_{1,0}^S$. The firm 1 “advantage denying” (AD) incentive is $\mu_{1,1}^S - \mu_{1,2}^S$.*

$\mu_{1,0}^S$ would be seller 1's continuation value if the buyer was to purchase from *neither* seller. This possibility is only hypothetical in the BDKS model, where there is no outside good, but, as we will show, some interesting differences between AB and AD incentives, defined in this way, still arise.

3.6 Analytical Results.

Our focus is on cases where $M = 3$ or $M = 30$, and dynamics, the size of the state space and the possible existence of multiple equilibria make it difficult or impossible to derive analytic results. However, building off BDKS's results, one can show equilibria are unique in some situations.

Proposition 1 *In a model with any $m \leq M$,*

1. if $\tau = 1$, equilibrium prices will equal marginal production costs in all states for all ρ and δ .
2. there will be a unique symmetric MPNE when
 - (a) $\delta = 0$ for all ρ and τ , or
 - (b) $\tau = 1$ for all ρ and δ .

Proof. See online Appendix A.1. ■

The social planner problem will also, of course, have a unique solution. Uniqueness when $\delta = 0$ reflects how logit demand implies a unique price equilibrium when continuation values are fixed and how movements through the state space will be unidirectional until the game reaches absorbing state (M, M) so that backwards induction can be applied. When $\tau = 1$, buyers extract all of the expected surplus and prices must compensate sellers for their production costs.

We can prove more novel results when $m = M = 2$ and $\delta = 0$. The $M = 2$ model is tractable because, as $D_1^*(1, 1) = D_1^*(2, 2) = \frac{1}{2}$ in any symmetric equilibrium and the social planner solution, there is only one choice probability that can vary with τ .⁷

Proposition 2 *For $m = M = 2$ and $\delta = 0$, the unique symmetric equilibrium will have the following properties*

1. equilibrium $D_1^*(1, 2)(\tau) < \frac{1}{2}$ for all τ .
2. equilibrium $D_1^*(1, 2)(\tau)$ is strictly decreasing in τ .
3. for $t \geq 2$, HHI^t is strictly increasing in τ .
4. there exists a τ^* such that $TS^{PDV}(\tau^*) = TS^{SP}$, $TS^{PDV}(\tau^*)$ is strictly increasing in τ for $\tau \in (0, \tau^*)$ and strictly decreasing in τ for $\tau \in (\tau^*, 1)$.

⁷This also facilitates a one-dimensional search for equilibria. We have never found multiplicity with $M = 2$ for any δ .

Proof. See online Appendix A.2. ■

Social welfare will be maximized when $D_1^*(1, 2) = D_1^{SP}(1, 2)$. The key aspects of the result are that $D_1^*(1, 2)(\tau)$, the probability that the laggard will catch up, is monotonically decreasing in τ , so that, as τ increases, the state is more likely to be asymmetric, and that $D_1^*(1, 2)(\tau = 0)$ is more than $D_1^{SP}(1, 2)$ and $D_1^*(1, 2)(\tau = 1)$ is less than $D_1^{SP}(1, 2)$. Our computational analysis will show that asymmetric market structures become more likely and efficiency is maximized for τ s between 0 and 1 when we consider larger state spaces if LBD effects are significant and forgetting rates are not too high.⁸

3.7 Numerical Methods for Finding Multiple Equilibria.

For given parameters, a single equilibrium can be calculated by solving either formulation of the equilibrium equations, using the Pakes and McGuire (1994) algorithm or, if $\delta = 0$, using backwards induction. To track what happens to equilibrium outcomes when we vary a single parameter, we use BDKS’s numerical homotopy method to trace the equilibrium correspondence. We will call a homotopy where a parameter α is varied an “ α -homotopy”. When we want to identify all equilibria, we also follow BDKS by criss-crossing the parameter space using homotopies, starting new homotopies when additional equilibria are found.

The homotopy method is not guaranteed to find all equilibria in this model, and, in practice, many homotopies stall. For any given set of parameters, our homotopy results must carry the caveat that we could be missing equilibria. However, we are confident in our broad results. We develop an alternative approach for the $M = 3$ model to finding all equilibria, detailed in online Appendix B.2, that gives identical results. For the $M = 30$ version of the closely-related BDK model, SJHY show that a third approach, which identifies when certain types of equilibria exist, also produces

⁸If $M = 2$, Proposition 2 holds for any $\Delta(2)$ in a slightly changed model where $\Delta(1)$ is set equal to zero (so an $e_{i,t} = 1$ firm will always move to $e_{i,t+1} = 2$ if it makes a sale).

results consistent with an extensive homotopy search.

4 Bargaining, Concentration and Welfare with $M = 3$

A model with $m = M = 3$ is too simple to represent any industry, but it allows us to follow how bargaining power affects every price and choice probability.

4.1 Bargaining and the Existence of Multiple Equilibria.

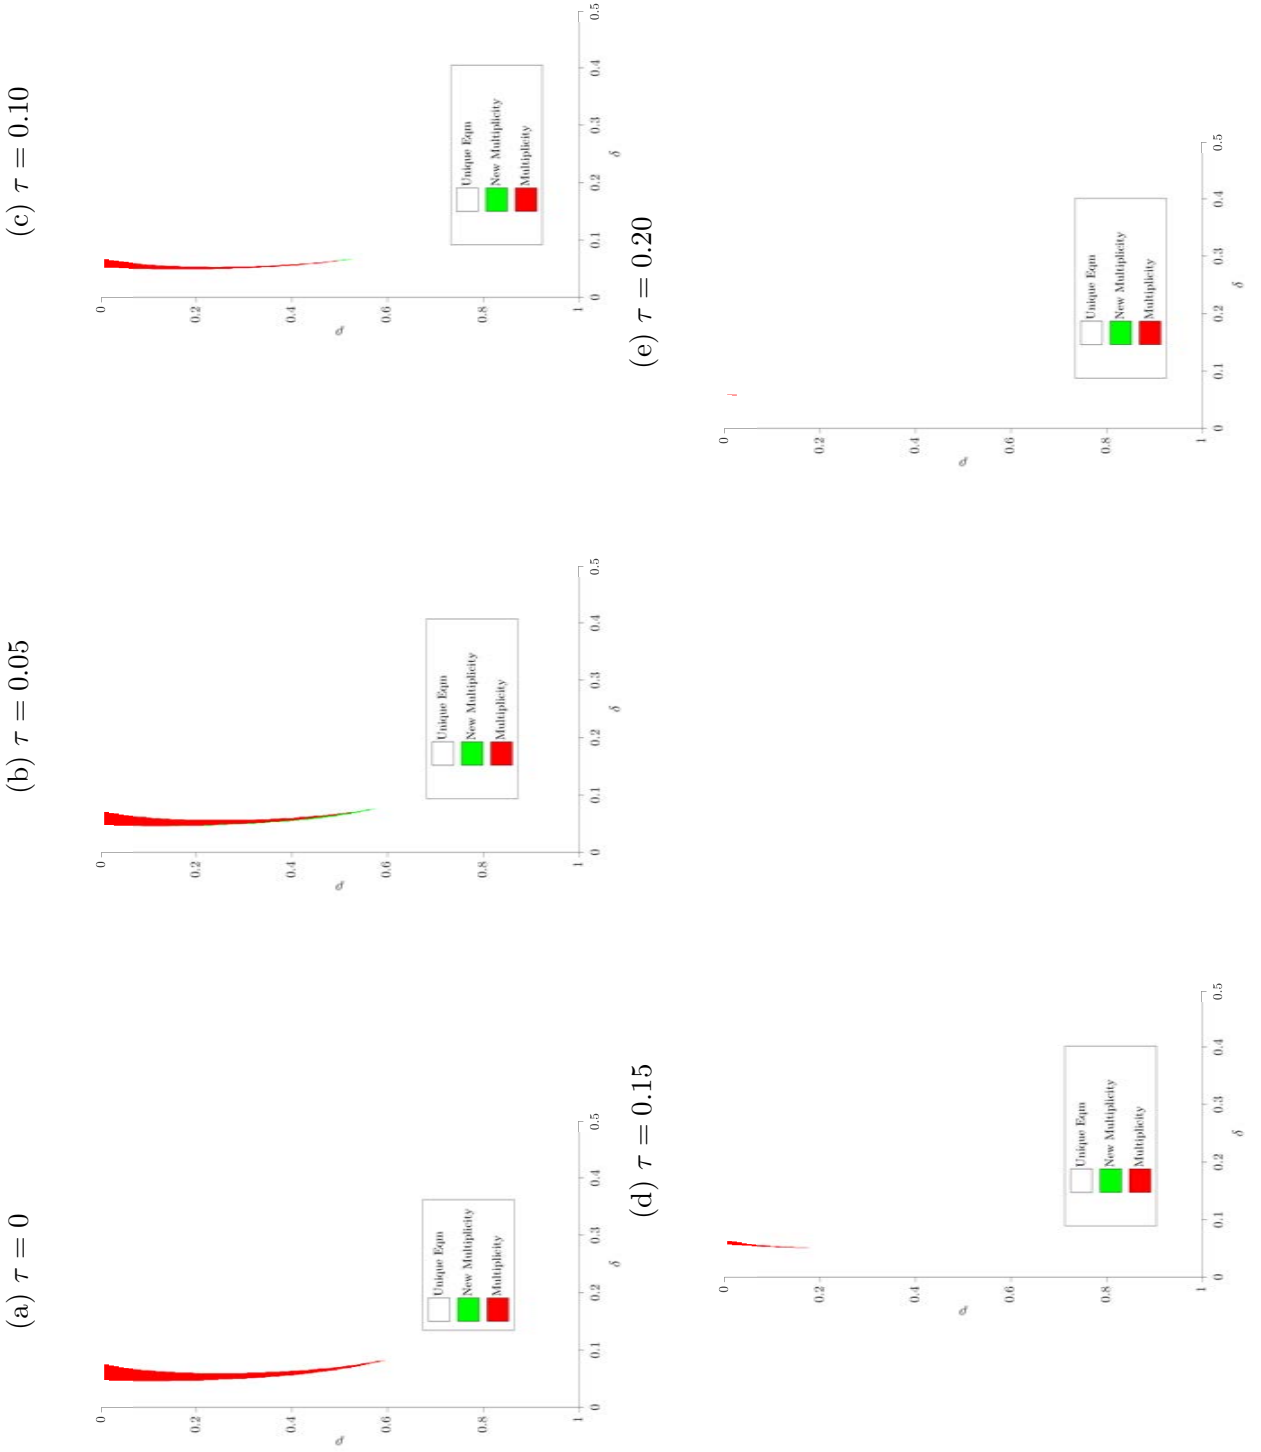
Multiple equilibria are widely understood to complicate analysis of comparative statics and counterfactuals. Uniqueness is guaranteed when $\tau = 1$, but it is unclear *a priori* whether multiple equilibria will be more common when sellers set prices or when bargaining power is split between buyers and sellers.

We identify equilibria by running sequences of ρ and δ homotopies for discrete values of τ in 0.05 steps. The set of identified equilibria are identical when we use our alternative method. The colored areas in Figure 1 indicate (ρ, δ) combinations with multiple equilibria (we never find more than three) for τ s up to 0.2. For $\tau \geq 0.25$ and $\delta \geq 0.1$ equilibria are always unique. While there are a few technology parameters, marked in green, where multiple equilibria are introduced when τ increases, the clear pattern is that allocating more bargaining power to buyers tends to eliminate multiplicity.

4.2 Bargaining, Equilibrium Prices and Outcomes for $\rho = 0.3$ and $\delta = 0.03$.

Figure 2 shows how choice probabilities, prices, concentration and surplus change when we trace equilibria using a τ -homotopy for $\rho = 0.3$ and $\delta = 0.03$. Equilibria are unique, for all τ , for these technology parameters. Online Appendix Table C.2 shows the exact prices and sale probabilities in each state, and the distributions of states after 4 and 32 periods, for $\tau = 0$, $\tau = 1$ and the social planner solution. $c(e_i) = 10$,

Figure 1: Multiplicity of Equilibria in the $M = m = 3$ Model. White = unique equilibrium. Green = multiple equilibria for τ and uniqueness for $\tau - 0.05$. Red = multiple equilibria for τ and (for $\tau \geq 0.05$) also for $\tau - 0.05$. For $\tau = 0.2$, there is a very small red area for $\rho < 0.05$ and $\delta \approx 0.058$. We do not identify any multiplicity for $\tau \geq 0.25$, using steps of 0.05, or for $\delta > 0.5$ for all (ρ, τ) . The figures were constructed using homotopies that criss-crossed the (ρ, δ) parameter space at least five times.



3 and 1.483 for $e_i = 1, 2, 3$ so that LBD effects are substantial. $\Delta(e_i) = 0.03, 0.059$ and 0.087 so that high know-how firms are unlikely to experience depreciation in any given period even if they do not make a sale.

It is useful to begin by understanding the difference between the social planner's purchase strategy (choice probabilities indicated by the horizontal dashed lines in Figure 2(a)) and the choices of a myopic buyer facing prices equal to production costs, the equilibrium when $\tau = 1$.

When $\tau = 1$ a myopic buyer makes choices based on current production costs and current ε s. An $e_i = 1$ laggard has such a large cost disadvantage in asymmetric states that it is very unlikely to make a sale. This results in the industry state being (3,1) with probability 0.996 after 32 periods, and the seller with the lower ε making the sale with probability close to $\frac{1}{2}$ from the second period onwards.

The social planner recognizes that a second low cost seller will increase expected future ε s. The social planner's willingness to invest in a second low cost option is illustrated by how $D_1^{SP}(1, 3) > D_1^{SP}(1, 2) > D_1^*(1, 2)(\tau = 1)$, even though $c_2(1, 3) < c_2(1, 2)$, and $D_1^{SP}(2, 3) > \frac{1}{2} \gg D_1^*(2, 3)(\tau = 1)$. The social planner solution implies moderate concentration after four periods, but the industry will be in the highest surplus state, (3,3), after 32 periods with probability 0.789.

Besanko, Doraszelski, and Kryukov (2019a) decompose efficiency losses into a loss of surplus in each state relative to the social planner outcome, which they label the "PR distortion", and the loss because states are not visited with the probabilities that the social planner would choose ("MS distortion").⁹ Figure 2(e) (right axis) shows

⁹Define

$$DWL_\beta^{PR} = \sum_{t=0}^{\infty} \beta^t \sum_{\mathbf{e}} \lambda_t(\mathbf{e}) [TS^{SP}(\mathbf{e}) - TS(\mathbf{e})]$$

and

$$DWL_\beta^{MS} = \sum_{t=0}^{\infty} \beta^t \sum_{\mathbf{e}} [\lambda_t^{SP}(\mathbf{e}) - \lambda_t(\mathbf{e})] TS^{SP}(\mathbf{e}).$$

where $TS(\mathbf{e})$ and $TS^{SP}(\mathbf{e})$ represent the value of expected ε s less production costs given equilibrium and social planner choice probabilities in state \mathbf{e} respectively, and $\lambda_t(\mathbf{e})$ and $\lambda_t^{SP}(\mathbf{e})$ are the probabilities the game is in state \mathbf{e} after t periods. $DWL_\beta^{PR} + DWL_\beta^{MS}$ is the difference between

the values of these two distortions. When $\tau = 1$, the buyer maximizes surplus in each state, whereas the social planner's choice may forgo some current surplus. Therefore, DWL_{β}^{PR} is negative, but DWL_{β}^{MS} is positive and large as state (3,3) is rarely visited. The sum of these distortions indicates that TS^{PDV} is lowered by around 2.4 compared to the social planner optimum.

Now consider the equilibrium when $\tau = 0$. Sellers charge markups over their opportunity costs and their opportunity costs reflect dynamic incentives (i.e., expected future profits), so that DWL_{β}^{PR} is positive. However, relative to $\tau = 1$, these effects may move sale probabilities and state transitions in a socially desirable direction. For example, $D_1^*(1, 2)$ is increased by the laggard having larger dynamic incentives (shown below), lowering its opportunity cost of sale, and the leader having a larger markup. In fact, $D_1^*(1, 2)(\tau = 0) > D_1^{SP}(1, 2)$ and $D_1^*(1, 3)(\tau = 0) > D_1^{SP}(1, 3)$, and $D_1^*(2, 3)(\tau = 0) > \frac{1}{2}$ even though it is slightly lower than $D_1^{SP}(2, 3)$. These D_1^* s imply that the state is more likely to move from (2,1) to (2,2) (rather than (3,1)) than the social planner would choose, but they also imply that state (3,3) will be reached more quickly. As a result DWL_{β}^{MS} is negative, and the decrease in TS^{PDV} from the social planner optimum is small.


Looking between these polar cases, we see that HHI^4 and HHI^{32} (concentration) increase monotonically in τ , and that TS^{PDV} is maximized for $\tau \approx 0.25$. The maximized surplus is slightly lower than the social planner achieves as there is no τ where all three asymmetric state D_1^* s match their social planner values (in particular, no equilibrium gives the laggard a high enough probability of selling in state (2,3)). Concentration and consumer surplus (panel (f)) increase significantly between 0 and 0.5, with producer surplus falling sharply, and several prices, including $p^*(1, 1)$ and $p^*(2, 2)$, are non-monotonic in τ .

Cabral and Riordan (1994) reason that one firm will emerge as dominant when prices that satisfy the properties of “increasing dominance” (ID, lower cost firms


social planner and equilibrium surplus. Besanko, Doraszelski, and Kryukov (2019a) also consider a third component related to entry and exit which is not relevant in the BDKS model.

Figure 2: The Effects of Bargaining Power Parameter (τ) on Equilibrium Outcomes and Welfare for $\rho = 0.3$ and $\delta = 0.03$. SP = social planner outcome, $c(e_i)$ are production costs for a firm with know-how e_i . Homotopies and our alternative method establish that there is a unique equilibrium for every τ . DWL = Deadweight loss.


(a) D_1 in Asymmetric States.




(b) Prices in Symmetric States.




(c) Prices in Asymmetric States.




(d) Concentration.



(e) TS^{PDV} and DWL Decomposition.



(f) Producer and Consumer Surplus.



always set lower prices) and “increasing increasing dominance” (IID, $p_1^*(\mathbf{e}) - p_2^*(\mathbf{e})$ decreases in e_1). Prices equal to costs satisfy these properties, and concentration is excessive when $\tau = 1$, whereas $\tau = 0$ equilibrium prices do not satisfy these properties (e.g., $p_1^*(2, 1) < p_1^*(3, 1)$) and concentration is too low. However, consistent with an observation of BDKS, ID and IID are not necessary for concentration to be either high or excessive: for example, for $0.25 < \tau < 0.445$, prices do not satisfy these properties even though HHI^4 is close to one and much larger than the social planner would choose.

4.3 Bargaining Power on Dynamic Incentives.

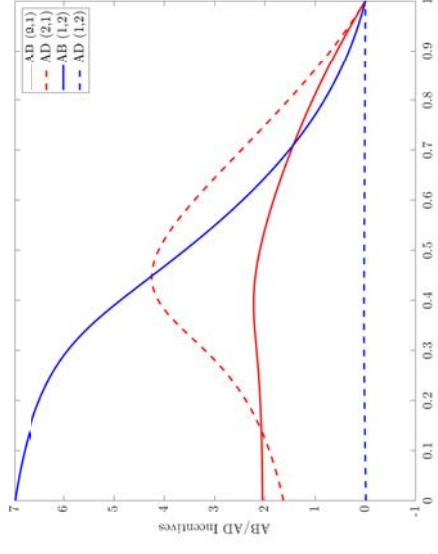


We now explore the relationship between τ and the equilibrium values of dynamic incentives. This relationship will explain why several prices have U-shaped relationships with τ and why concentration and some measures of surplus (including TS^{PDV} in the $M = 30$ model) change quickly as τ increases from zero.

Figure 3 illustrates different effects of varying τ . We focus primarily on state (1,1) as there is only one price. Panel (a) shows how $p(1, 1)$ would change with τ when continuation values are held fixed at their $\tau = 0$ values. This isolates the markup reduction effect that occurs in static models. $p(1, 1)$ falls monotonically, and close to linearly, to equal opportunity costs when $\tau = 1$.




The source of the non-monotonicity therefore lies in how continuation values vary with τ . One can consider τ as having two effects on continuation values. First, fixing industry evolution, increasing τ reduces seller continuation values, and shrinks the difference between the value of making a sale and the rival making a sale, as the share of surplus in every future negotiation is reduced. Panel (b) isolates this effect by showing the equilibrium price that sellers would set (i.e., $\tau = 0$ static game) when $\tau = 0$ continuation values are multiplied by one minus the τ value on the x-axis. As future profits shrink, a seller’s opportunity cost rises towards its production cost, and the take-it-or-leave-it $p(1, 1)$ rises monotonically.

Figure 3: Prices and Incentives as a Function of τ for $\rho = 0.3$ and $\delta = 0.03$ in the $M = m = 3$ Model.

- (a) Eqm. (1,1) Price Given $\tau = 0$ Continuation Values
 (b) $\tau = 0$ Eqm. (1,1) Price Given $(1-\tau) \times \tau = 0$ Continuation Values
 (c) AB and AD Incentives for $\mathbf{e} = (2,1)$ and (1,2)



- (d) Eqm. (1,1) Continuation Values and AB/AD Incentives
 (e) $\tau = 0$ Eqm. (1,1) Price Given τ Eqm. Continuation Values..



(f) Eqm. PDV AB/AD Incentives

Second, reallocating bargaining power will change how market structure evolves. In particular, the static markup reduction effect will increase the sale probability of the firm with the larger markup. If this firm is the leader, which it is in all asymmetric states when $\tau = 0$ for our parameters, then this will tend to make leads last longer. If leads would be short-lived when $\tau = 0$, this lead-lengthening effect can cause the relative value of acquiring or preserving a lead to increase in τ , so the leader's opportunity cost falls, lowering the leader's price.

To illustrate, consider what happens in state (2,1), which is the state that seller 1 in state (1,1) wants to reach. A firm in state (2,1) expects its lead to last for 5.4 periods when $\tau = 0$. This increases to 11.3 and 196 periods when $\tau = 0.25$ or 0.5 respectively. Correspondingly, a firm in state (2,1) loses its lead within three periods with probability 0.378 when $\tau = 0$, and with probabilities 0.168 and 0.009 when $\tau = 0.25$ and $\tau = 0.5$ respectively.

Panel (c) shows AB and AD incentives in state (2,1) for both sellers. This panel also reveals that changes in the laggard's dynamic incentives also make the leader more likely to sell in (2,1). Adding its AB and AD incentives together, the laggard has a larger dynamic incentive than the leader when $\tau = 0$, partially offsetting the laggard's production cost disadvantage. However, as τ increases, the leader's AD incentive increases as its lead is expected to last longer if its rival does not make a sale, while the reduction in equilibrium $p^*(2,2)$ reduces laggard's AB incentive. The leader's dynamic incentive exceeds the laggard's for $\tau \geq 0.25$.

Panel (d) shows that all three equilibrium $\mu_1^S(1,1)$ continuation values decline monotonically in τ increases but that the equilibrium AB incentive increases until $\tau \approx 0.4$, and the opportunity cost of sale decreases until $\tau \approx 0.35$. Panel (e) repeats the exercise in panel (b) but using these equilibrium continuation values. The changes in continuation values (i.e., dynamic effects) now lower the price, reinforcing the static markup effect, until $\tau \approx 0.37$. This explains why $p^*(1,1)$ initially declines in τ (Figure 2(b)). For larger τ , the shrinkage of continuation values, and the implied increase in

opportunity costs, is the stronger dynamic effect of increasing τ . $p^*(1, 1)$ therefore increases. While changes in $p^*(1, 1)$ do not change concentration, the reductions in the leader's markup and its opportunity cost increase the leader's probability of making sales in (2,1) and (3,1) (Figure 2(a)) which does cause concentration to increase.

The solid lines in panel (f) show that the discounted expected value of sellers' combined AB and AD incentives given equilibrium play. Both AB and AD incentives increase and then decrease in τ . The dotted lines show what the discounted values would be if we use the equilibrium values of the incentives in each state but assume $\tau = 0$ equilibrium state transitions. The differences between the solid and dotted lines highlight how the non-monotonicity in dynamic incentives is accentuated by how play shifts towards states such as (3,1), where dynamic incentives are larger, as τ increases from zero.

Generalization to Other Technology Parameters. Online Appendix C.2 provides an analysis of how these patterns generalize to other (ρ, δ) technology parameters. The pattern that concentration increases with τ holds for most parameters, but is not universal. Online Appendix C.3 shows an example with very limited LBD, where a twist in the equilibrium correspondence associated with multiple equilibria causes concentration to fall, at a low level, as τ increases from zero. Other examples, with significant LBD, involve very high concentration which declines slightly as τ increases from 0.5 to 1.

The other patterns that we have described (e.g., TS^{PDV} and the discounted value of dynamic incentives having inverted-U relationships with τ) are observed when $\rho < 0.8$ (i.e., costs when $e_i = 3$ are no more than 70% of the costs when $e_i = 1$) and $\delta \leq 0.05$. For $\delta > 0.05$, forgetting is likely and a laggard with infrequent sales is unlikely to catch up, so that concentration is high and leads tend to be long-lasting even when $\tau = 0$. The effect of shrinking continuation values therefore tends to dominate the lead-lengthening effect and the values of dynamic incentives are

maximized when $\tau = 0$. For $\delta > 0.1$, it is too costly for the social planner to maintain two low cost sellers, and the planner's strategy shifts towards maximizing current surplus. As a result, efficiency is maximized for τ s close to, or equal to, 1.

Outside Option and Product Differentiation. We have assumed that $\sigma = 1$ and that buyers have no outside option. Online Appendix C.4 relaxes these assumptions. Outcomes are quite robust to moderate variation in these assumptions, such as allowing an outside good that is only really competitive when the sellers have the lowest know-how. However, significant reductions in product differentiation (lower σ) result in more concentrated equilibrium outcomes, and leads that are expected to last longer, for all τ . Consistent with the logic described so far, this tends to mean that dynamic incentives will be strongest when $\tau = 0$.

4.4 Subsidies.

We consider schemes that implement socially optimal choice probabilities by taxing or subsidizing sales by the laggard.¹⁰ Given socially optimal choice probabilities (D_1^{SP}), we can solve for M^2 prices and $\frac{M(M-1)}{2}$ subsidies (s_1 , negative for a tax) using M^2 equations

$$p_1(\mathbf{e}) - p_2(\mathbf{e}) = \sigma \log \left(\frac{1}{D_1^{SP}(\mathbf{e})} - 1 \right), \quad (14)$$

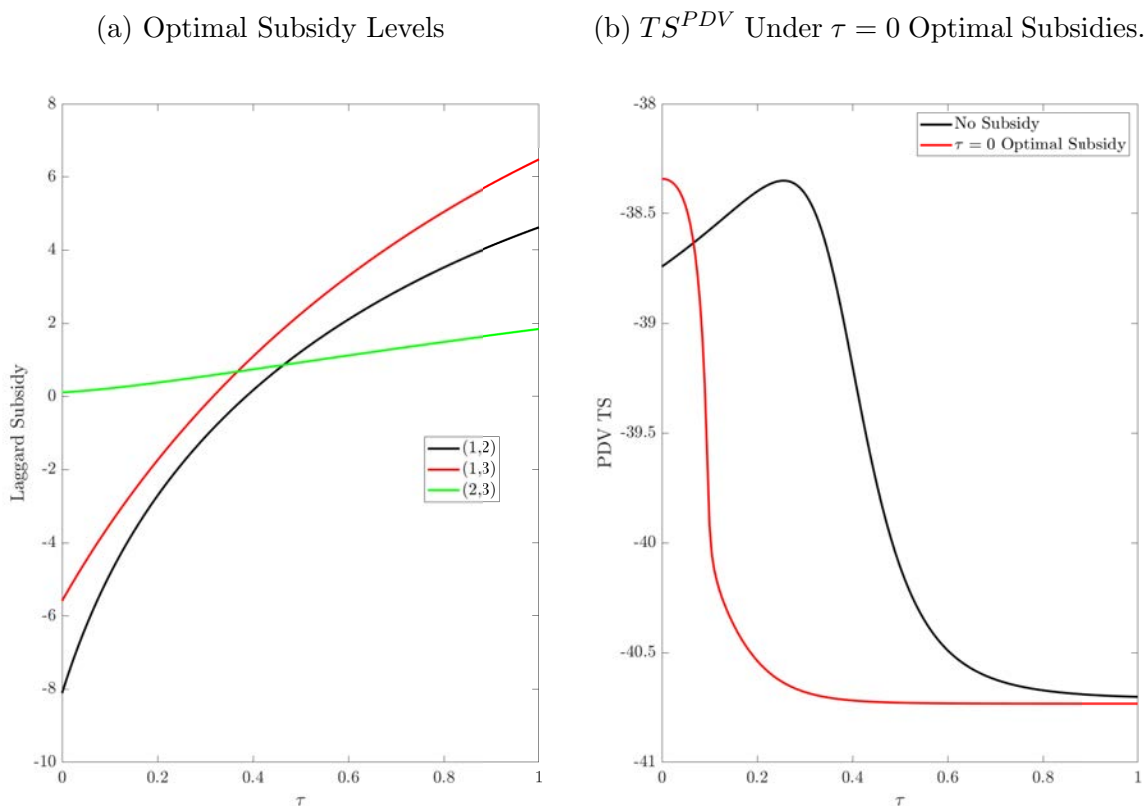
and $\frac{M(M-1)}{2}$ stacked equations

$$\mathbf{p}_1 + \mathbf{s}_1 = \Phi(\mathbf{D}_1^{SP}) + \mathbf{c}_1 - \beta(\mathbf{Q}_1 - \mathbf{Q}_2)(\mathbf{I} - \beta\mathbf{Q}_2)^{-1}[\mathbf{D}_1^{SP} \circ \Phi(\mathbf{D}_1^{SP})]. \quad (15)$$

Linearity in p and s implies that only one scheme can implement the social optimum but this scheme could support equilibria that are not optimal (see $M = 30$ example

¹⁰With no outside good, sale probabilities in symmetric states are efficient so it is natural to consider schemes that impose taxes or subsidies in asymmetric states. Taxes and subsidies are fully passed through to buyers when $\tau = 1$.

Figure 4: Subsidy Schemes and Welfare as a Function of τ for $\rho = 0.3$ and $\delta = 0.03$ in the $M = m = 3$ Model.



below).

Figure 4(a) shows the optimal subsidies for our example technology parameters. Laggard sales are taxed in states (1, 2) and (1, 3) when $\tau = 0$, but they are subsidized in all states when $\tau > 0.4$. Further calculations show that the optimal scheme can have non-monotonic effects on some outcomes: for example, relative to a no subsidy scenario, discounted consumer surplus increases when $\tau \approx 0$ and $\tau > 0.6$, but it falls slightly for $\tau \approx 0.3$.

We can also ask how welfare would change for different τ s if a subsidy scheme is designed assuming $\tau = 0$. $\tau = 0$ would be the natural assumption for a policy designer to make based on the existing literature. Figure 4(b) shows that the scheme would lower welfare, relative to having no scheme at all, for any τ that is larger than

0.06, a result that implies that buyers only have to have limited bargaining power for the subsidy policy to be welfare-reducing. The intuition for this result is that the $\tau = 0$ scheme aims to increase industry concentration which, with no subsidies, would be close to optimal for τ s that are not too large and already excessive for higher τ . Online Appendix C.5 shows that $\tau = 0$ subsidies would also lower welfare if $\tau = 0.2$ for a wide range of technology parameters with low δ s. For higher δ s, some of the $\tau = 0$ subsidies are so large that we cannot solve for equilibria with alternative τ s.

5 $M = 30$ Model

BDKS use the $m = 15$ and $M = 30$ model as a potentially stylized representation of a real-world industry. We describe how the features that we observe in the $M = 3$ model carry over to this larger state space model before examining optimal subsidies and some alternative policy counterfactuals.

Our policy analysis will assume $\rho = 0.75$ (so that costs for $e_i \geq 15$ are just over 30% of their $e_i = 1$ value) and $\delta = 0.023$ (implying $\Delta(30) \approx 0.5$), which we will refer to as the illustrative technology parameters. When considering what happens with no policies, we look across $0.6 \leq \rho \leq 1$, reflecting the range of empirical progress ratios identified in the literature survey of Ghemawat (1985), and $0 \leq \delta \leq 0.2$. $\delta = 0.2$ implies almost certain depreciation of know-how for $e_i \geq 20$.


5.1 Equilibria with No Policies.

The qualitative results from our $M = 3$ analysis carry over to the larger state space. First, multiple equilibria are eliminated once τ is even moderately high. Based on a homotopy analysis to identify equilibria for τ values in 0.05 increments, Figure 5(a) shows the smallest τ s for which multiple equilibria are identified, and (b) shows the τ values for which equilibria are unique for all larger τ s. Multiplicity exists for a wider range of technologies than in the $M = 3$ model when $\tau = 0$, but we still find that all

Figure 5: Multiplicity in the $M = 30$ Model. In panel (a) the white regions correspond to parameters for which multiplicity is never identified.

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

(b) Smallest Values of τ' where Equilibria are Unique For All $\tau \geq \tau'$.

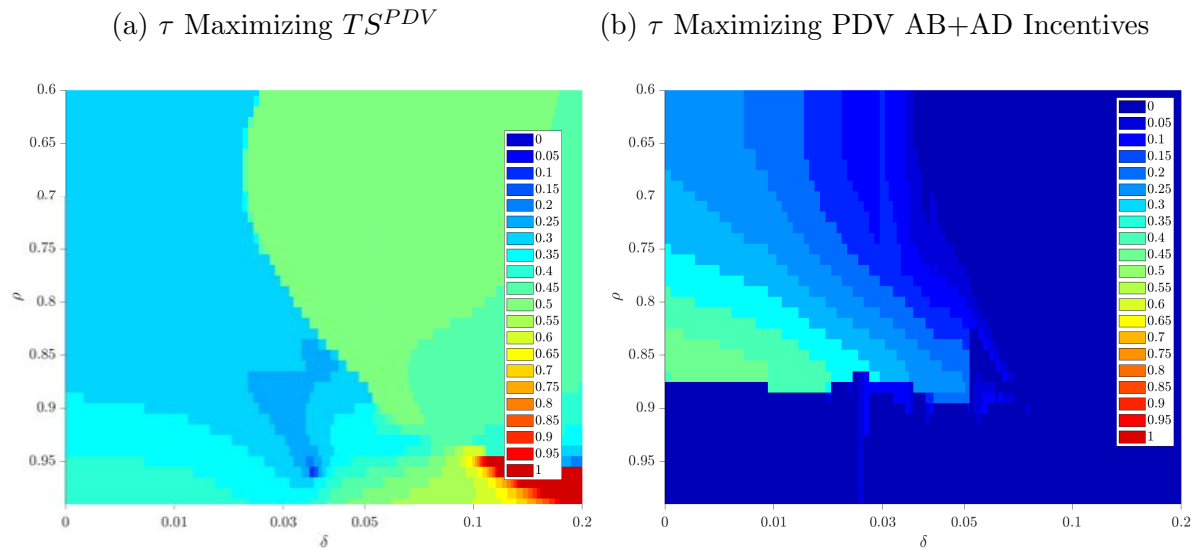


equilibria are unique once $\tau \geq 0.25$. We have undertaken further analysis to try to identify if there is a pattern as to the types of $\tau = 0$ equilibria (e.g., the nature of pricing strategies) that are on homotopy paths that ultimately reach $\tau = 1$ but we have not been able to identify a general pattern. See additional discussion in online Appendix C.3.

Second, as we detail in online Appendix D.2, the pattern that equilibrium concentration increases with τ , at least until it reaches a very high level, is fairly general when LBD effects are not too small.¹¹ HHI^{32} is lower than the social planner would choose when $\tau = 0$ for most technologies. For $\rho < 0.85$ and $\delta < 0.03$, equilibrium HHI^{32} is low when $\tau = 0$, very high when $\tau = 1$, and close to the level that the social planner would choose between $\tau = 0.25$ and $\tau = 0.5$. Consistent with this feature, Figure 6(a) shows that τ s between 0.2 and 0.5 maximize surplus for a wide range of technologies.

¹¹There are parameters with very limited LBD and moderate-to-high forgetting probabilities where HHI^{32} is above the social planner level, and where it falls when τ increases, before following an upward path. See online Appendix D.3 for an example.

Figure 6: Values of τ Maximizing TS^{PDV} and the PDV of Seller Dynamic Incentives (the sum of AB and AD incentives) in the $M = 30, m = 15$ model. We take the maximum value across equilibria when multiple equilibria exist for given τ .

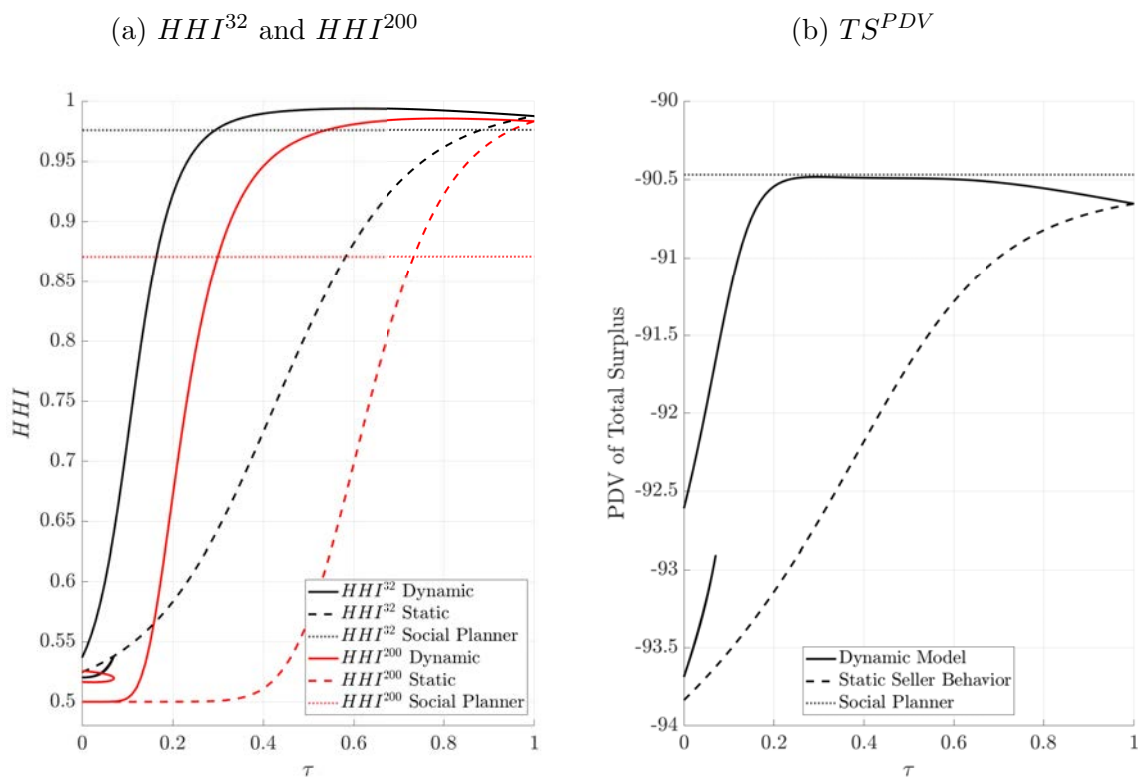


Third, for $\delta < 0.03$ and $\rho < 0.85$, we also see that τ s between 0.2 and 0.5 maximize sellers' dynamic incentives (Figure 6(b)). The increase in dynamic incentives as τ increases from zero is associated with leads lasting significantly longer: for example, for the illustrative technology parameters, the PDV of dynamic incentives increases from around 80 when $\tau = 0$ to around 170 when $\tau = 0.2$. This is associated with the number of periods that a firm in state (2,1) expects its lead to last increasing from 90 to 300, and the probability that it loses its lead within three periods falling from 0.119 to 0.009.¹²

The increases in dynamic incentives also leads to a divergence between predicted outcomes in the dynamic model and the outcomes that would be predicted by a model where sellers are assumed to behave statically even though know-how can evolve. Figure 7 illustrates this divergence by showing the values of HHI^{32} , HHI^{200} and TS^{PDV} along τ -homotopy paths for these two types of equilibrium given the illustrative technology parameters. Equilibria will always be unique when sellers

¹²These statistics are associated with the $\tau = 0$ equilibrium on the τ -homotopy path that stretches from $\tau = 0$ to $\tau = 1$ (see Figure 7).

Figure 7: Concentration and Welfare Outcomes with Dynamic and Static Equilibrium Seller Behavior as a Function of τ for $\rho = 0.75$ and $\delta = 0.023$ with no policies.

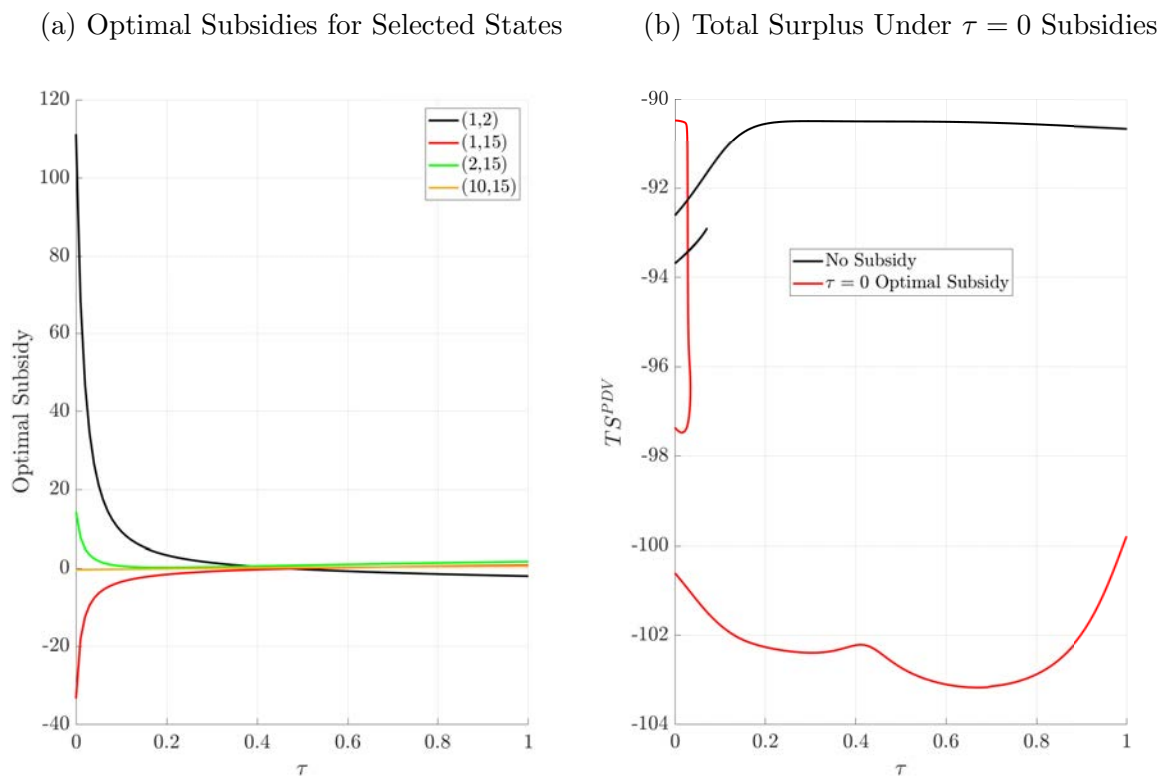


ignore dynamic incentives, and they will coincide with the dynamic equilibrium when $\tau = 1$. When $\tau = 0$, concentration levels in the three dynamic equilibria and the static equilibrium are similar, but equilibrium concentration increases much more quickly as τ increases in the dynamic model, and HHI^{32} actually falls, by a small amount, as τ increases from 0.5 to 1. The difference in TS^{PDV} given dynamic and static behavior is maximized for $\tau \approx 0.2$, where the dynamic equilibrium gets very close to the planner optimum.

5.2 Subsidy Policies.

We now turn to the policy analysis, by considering optimal subsidies in the $M = 30$ model, before considering alternative policies to promote competition. We will focus on the illustrative technology parameters, as possibly reflecting an industry with

Figure 8: Optimal Subsidies and Welfare under $\tau = 0$ Subsidies as a Function of τ for the Illustrative Technology Parameters. The subsidy is given to the laggard when it makes a sale, and a negative number is a tax.



fairly strong LBD effects and some depreciation. We have found that some other technologies support many equilibria for alternative policies. Online Appendix D.3 shows some policy results for technology parameters with limited LBD.

Figure 8(a) shows optimal laggard subsidies in four different states as a function of τ . The optimal subsidy in state (1,2) and the optimal tax on the laggard in state (1,15) change dramatically as τ increases from zero, illustrating how small deviations from the price-setting assumption would need to be accounted for in policy design.

The direction of the $\tau = 0$ subsidies may also appear counterintuitive. When $\tau = 0$, equilibrium concentration is lower than the social planner would choose. So why is it optimal to subsidize sales by the laggard?

The reason lies in how the future evolution of market structure affects dynamic

incentives. The social planner’s evolution, which is being implemented, will lead to a firm in (3,1) being likely to preserve its lead for a long-time. Under this evolution, the dynamic incentives of a $\tau = 0$ leader in state (2,1) would be extremely large (its AB incentive is 127.3 and its AD incentive is 80.7) without subsidies, implying the leader would make a sale with a probability that is actually much higher than the social planner wants. Therefore, a laggard subsidy is appropriate. Similarly, because the social planner’s choices would lead the laggard to make most of the sales once the leader has know-how above m , the leader’s dynamic incentives in (1,15) would be much smaller than the laggard’s so that a tax on the laggard is appropriate.

The red lines in Figure 8(b) show TS^{PDV} under the subsidies that would be optimal if $\tau = 0$. The black line shows TS^{PDV} in the no subsidy equilibrium. When τ is very small, $\tau = 0$ optimal subsidies support multiple equilibria with different welfare levels. The only equilibrium when $\tau \geq 0.04$ implies that $\tau = 0$ optimal subsidies lead to lower welfare than with no policy, further illustrating the importance of knowing the value of τ when determining policies.

5.3 Policies to Increase Competition.

We use the illustrative technology parameters to analyze policies that might be expected to “promote competition” (i.e., more symmetric market structures). Note that concentration is too low, at least initially, for these parameters when $\tau = 0$, so one might think of our analysis as primarily trying to understand the costs of the types of policies that have been suggested in industries where dynamic competition is important, whereas, if $\tau = 0.5$, one might read our analysis as assessing whether the welfare benefits of increased seller symmetry are offset by the distortions that the policies will introduce. We will also see that, for higher values of τ , some of the policies can also have unexpected effects on market structure.

5.3.1 Alternative Policies.

Restriction on Market Concentration. Benkard (2004), analyzing the market for wide-bodied commercial aircraft, considers a counterfactual where a limit is imposed on the market share of the largest firm in a given quarter.¹³ In our duopoly single-buyer-per-period model, we implement the share restriction as a soft constraint by assuming that a market leader has to pay a compliance penalty of $\chi \times \max\{0, D_i - \psi\}^2$ whenever its sale probability is above a threshold $\psi > 0.5$. As χ increases, it becomes more costly for a firm to have a high market share.¹⁴ We calculate TS^{PDV} assuming that compliance penalties are not costs to society. The results presented in the paper assume that $\psi = 0.75$, so that policy maker aims that the leader should be no more than three times as likely as the laggard to make a sale, and $\chi = 50$, a value that is large enough that the constraint is rarely breached in equilibrium. See online Appendix D.4 for an analysis of how concentration and TS^{PDV} change with the assumed level of χ .

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

$$\begin{aligned}
 & -\tau [D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - \hat{c}_1) - \chi \max\{0, D_1^*(\mathbf{e}) - \psi\}^2] + (1 - \tau) \log\left(\frac{1}{1 - D_1^*(\mathbf{e})}\right) \times \\
 & [\sigma - (1 - D_1^*(\mathbf{e}))(p_1^*(\mathbf{e}) - \hat{c}_1) + 2\chi(1 - D_1^*(\mathbf{e})) \max\{0, D_1^*(\mathbf{e}) - \psi\}] = 0,
 \end{aligned} \tag{16}$$

and the equation for the seller's value becomes

$$VS_1^*(\mathbf{e}) - D_1^*(\mathbf{e})(p_1^*(\mathbf{e}) - c(e_1)) - \sum_{k=1,2} D_k^*(\mathbf{e})\mu_{1,k}^S(\mathbf{e}) - \chi \max\{0, D_1^*(\mathbf{e}) - \psi\}^2 = 0. \tag{17}$$

¹³While absolute restrictions on market shares are rare, market shares can play an important role in determining potential liability for actions that agencies or rivals claim are anticompetitive.

¹⁴However, smaller values of χ can make it more likely that the game is in states where the leader has a large share, and actually increase expected penalty costs.

Restrictions on Pricing Incentives. Besanko, Doraszelski, and Kryukov (2014) and Besanko, Doraszelski, and Kryukov (2019b) consider the effects of alternative limitations on the dynamic incentives that firms are able to consider using the BDK model, motivated by alternative standards for assessing prices to be predatory.¹⁵ We consider how equilibrium outcomes change when a leader¹⁶:

- is unable to consider dynamic incentives at all (i.e., the leader has to use an opportunity cost that equals its current production cost).
- is unable to consider AD incentives, but can consider AB incentives, the level of which may change endogenously.

Restrictions on Pricing. An easier-to-implement restriction would prevent the leader from setting a price below its current production cost. Below-cost pricing is often viewed as a necessary, but not sufficient, condition for pricing to be viewed as predatory.¹⁷

We find equilibria using the homotopy method with an additional set of equations associated with Lagrangian constraints. We find an initial equilibrium using an iterative guess-and-verify approach to identify states where the constraints on the leader’s price bind.

Trigger Policies. In practice, political pressure to promote competition may arise only once one firm is established. The social planner also wants to aim for symmetry

¹⁵While the lack of exit may appear to make the BDKS model a less attractive setting for considering rules motivated by the predation literature, the literature on anticompetitive conduct (e.g., Caves and Porter (1977), Lieberman (1987)) does not assume exit must be possible, and know-how depreciation allows a firm in the BDKS model to try to weaken its rival in a wide range of states.

¹⁶We have also computed some results imposing these restrictions on both firms, with the no dynamics case corresponding to the static equilibrium considered above. However, allegations of anticompetitive conduct usually focus on the market leader, so focusing on policies that target the leader seems more relevant.

¹⁷See discussion in the Department of Justice 2008 report “Competition and Monopoly: Single-Firm Conduct Under Section 2 of the Sherman Act”, <https://www.justice.gov/sites/default/files/atr/legacy/2009/05/11/236681.pdf>. Of course, application would entail many choices about the appropriate way to measure costs.

once one firm has lowered its costs. We therefore also compute equilibria under “trigger versions” of the policies listed above, meaning that the policy will be introduced as soon as one firm reaches know-how state e' and will then last forever. We assume players know the value of e' when the game begins.

Multiplicity of Equilibria. For the illustrative technology parameters, we find some multiplicity for small τ under the concentration restriction and Leader $p \geq mc$ policies, but the predictions across these equilibria, and the equilibria with no policies, are sufficiently similar that they do not complicate our conclusions. Based on our searches, which do identify multiplicity for other technology parameters, equilibria under the incentive policies are unique.

5.3.2 Predicted Policy Effects.

Figure 9(a)-(d) shows how the policies change HHI^{32} , TS^{PDV} and PS^{PDV} as a function of τ . The only policy to affect outcomes when $\tau = 1$ is the concentration restriction policy as the seller is able to avoid the compliance penalty by not agreeing a price.

Ignoring the slight complications in interpretation caused by multiple equilibria when τ is small, all policies weakly lower expected medium- and long-run concentration for $\tau < 0.8$. However, this reduction is associated with a softening of competition in low know-how states which leads to discounted seller surplus increasing, and a loss in social welfare as production costs are expected to fall more slowly. The policies decrease TS^{PDV} the most when $\tau \approx 0.2$, reflecting the efficiency of the no policy outcome. Comparing across policies, the $p \geq mc$ policy has the smallest effects, reflecting how, with no policies, prices are significantly below production costs in only a relatively small set of states.¹⁸

For $0.8 \leq \tau \leq 1$, the No Leader AD policy increases TS^{PDV} (and CS^{PDV} , not

¹⁸Online Appendix D.1 lists prices in a subset of states in the no policy $\tau = 0$ equilibria.

Figure 9: Effects of Policies Introduced at the Start of the Industry's Life on Concentration and Total and Producer Surplus, for the Illustrative Technology Parameters. The compliance costs of the Concentration Restriction policy are not counted as costs to society in total surplus calculations, although they are costs to the sellers.

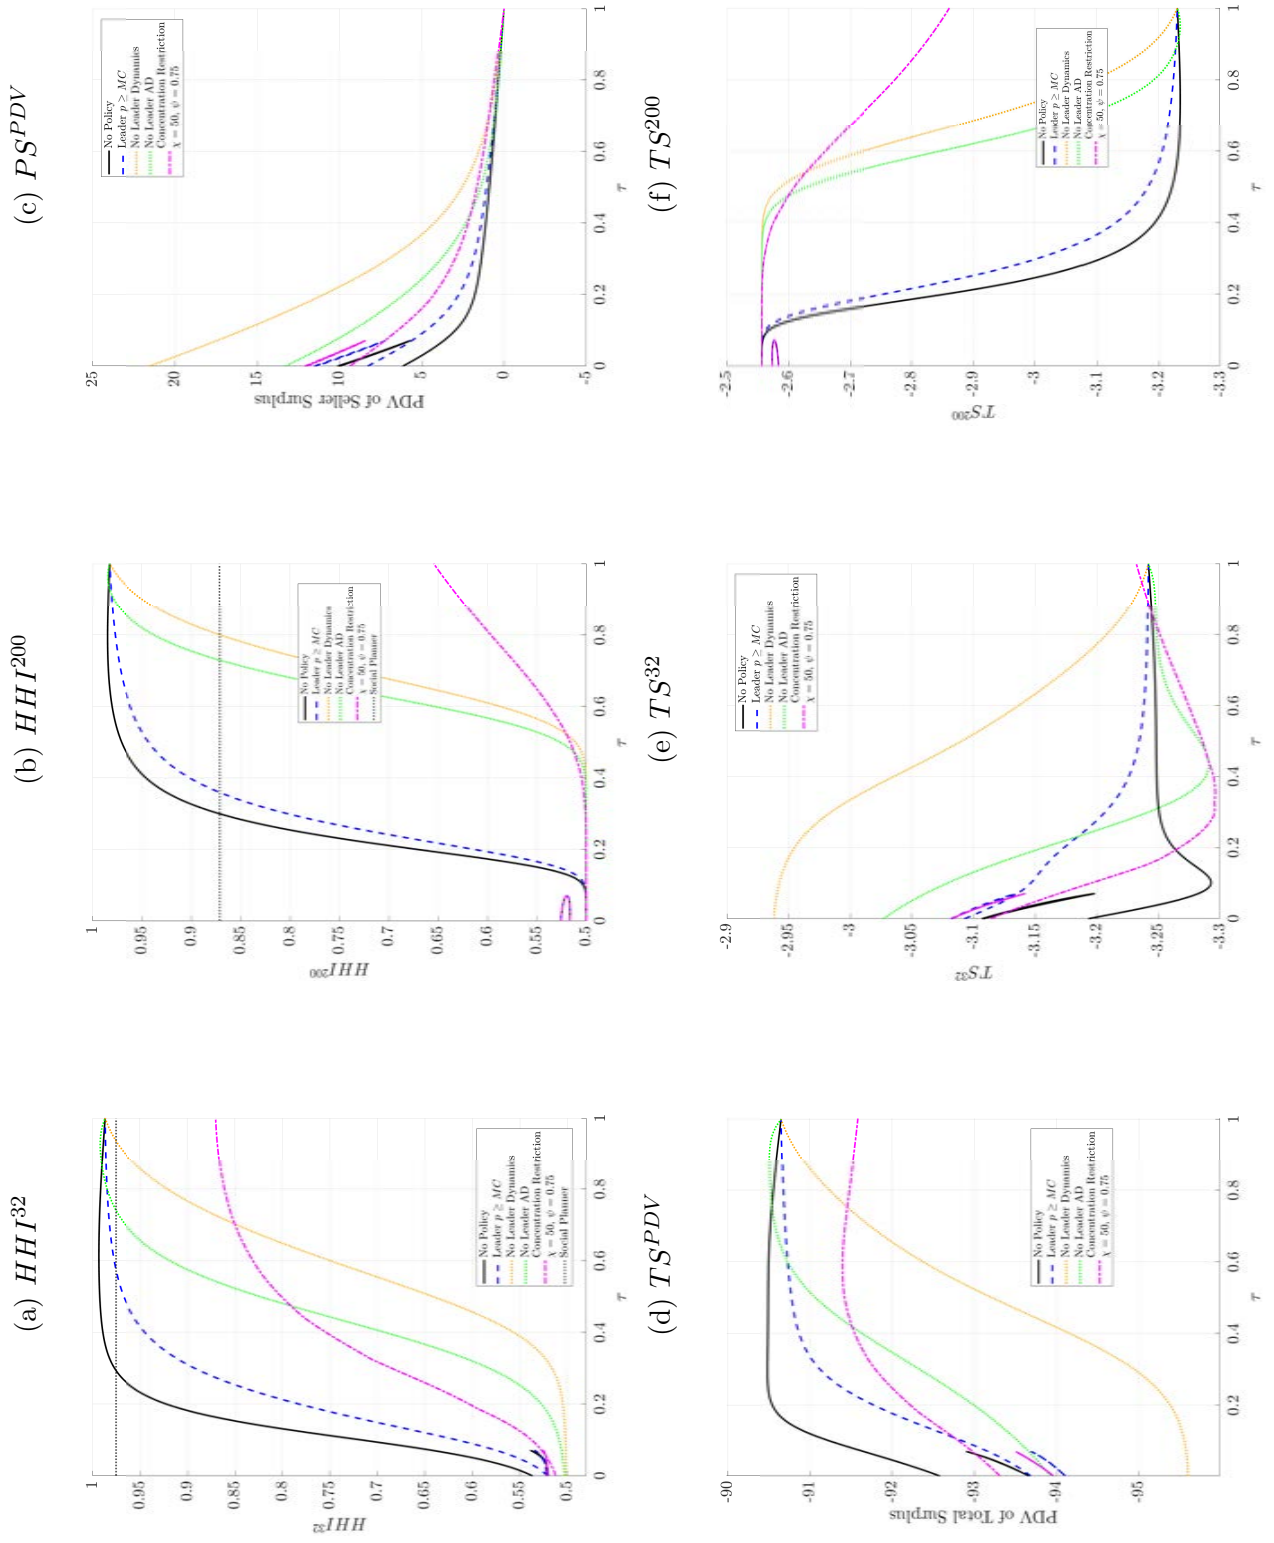



Figure 10: The Present Discounted Value of Sellers' AB and AD Incentives With and Without the No Leader AD Policy for the Illustrative Technology Parameters.

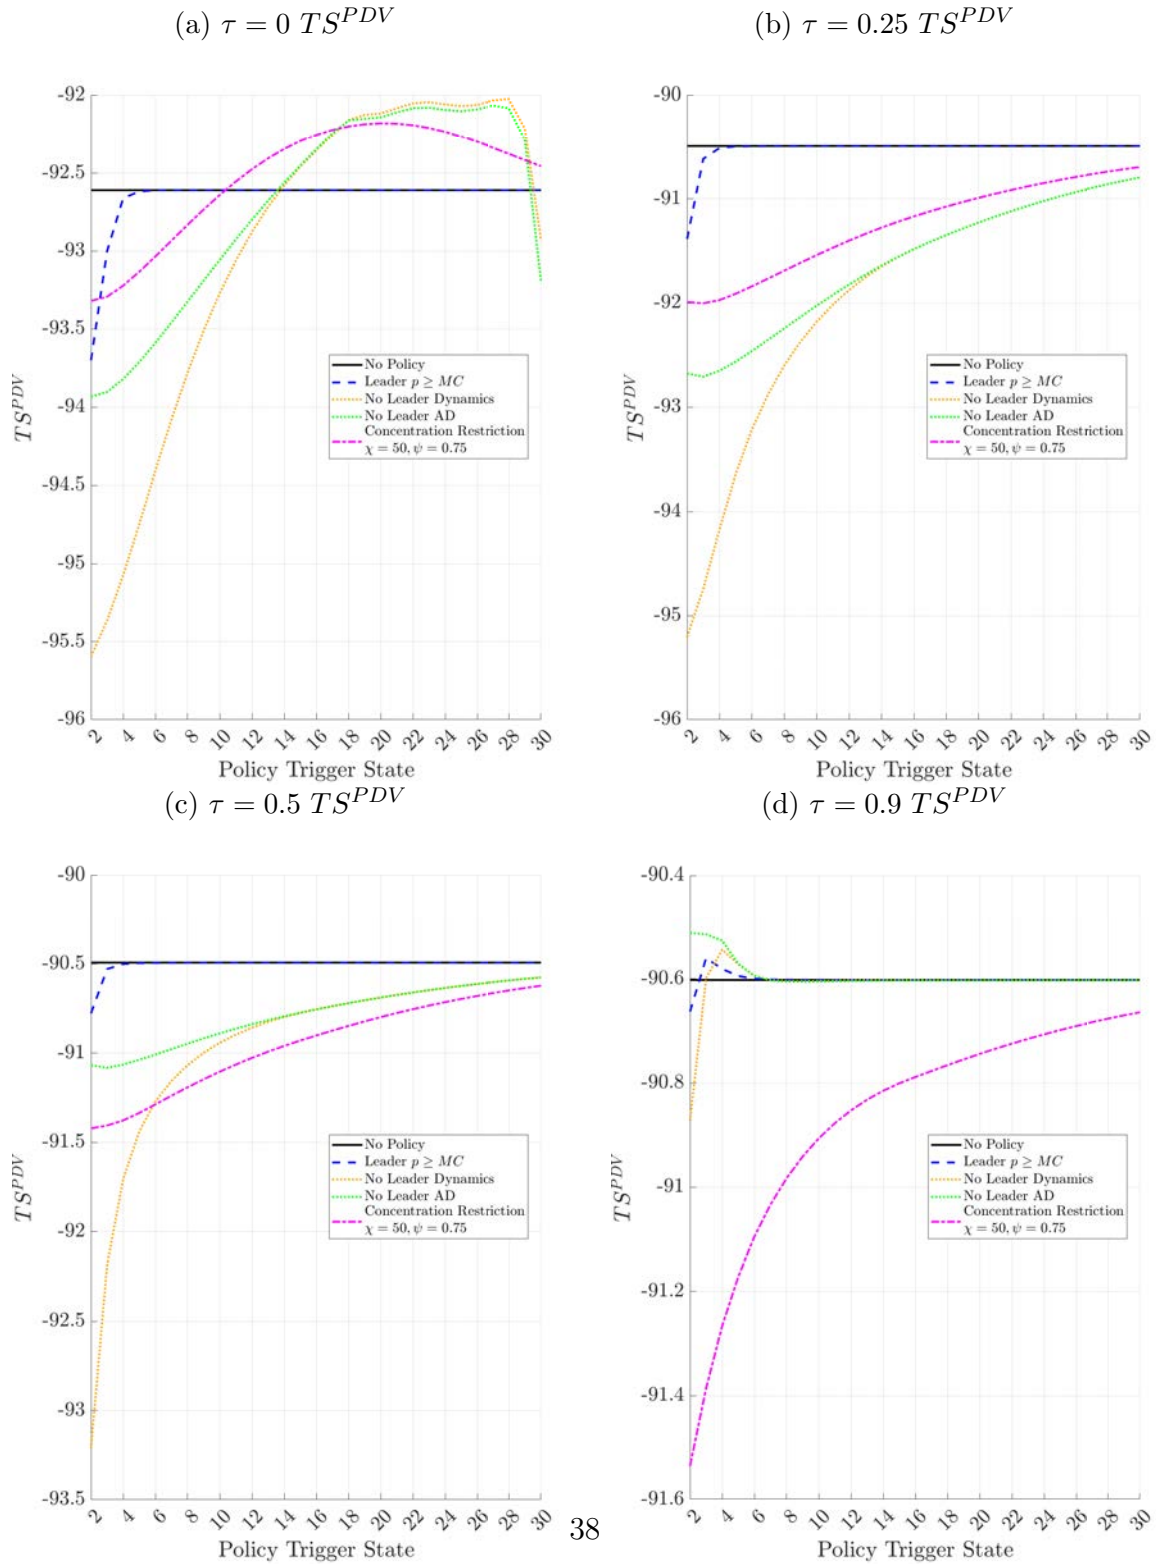


shown), and, for $0.9 \leq \tau < 1$, it also increases HHI^{32} . While applications would rarely assume that $\tau > 0.8$, these changes illustrate the subtle ways that dynamic incentives can interact. Figure 10 shows the PDV of dynamic incentives with no policies, and the value of dynamic incentives that sellers are allowed to consider under the No Leader AD policy. For small and moderate levels of τ , the policy significantly reduces the sum of considered dynamic incentives, but, for $\tau > 0.8$, the sum increases as AB incentives under the policy increase. Further analysis shows that, for $\tau > 0.8$, the No Leader AD policy increases the relative size of the leader's AB incentives in low know-how states, even though buyers have most of the bargaining power, so that the leader moves more quickly down its cost curve.

Figure 9(e) and (f) show that, while the policies lower TS^{PDV} for $\tau < 0.8$, they may increase TS^{32} and TS^{200} , reflecting how they increase the probability that there will be two sellers with low costs. This also suggests that trigger strategies, with appropriate e 's, might be able to enhance efficiency.

Figure 11 shows how TS^{PDV} changes for four different τ , as a function of the

Figure 11: Effects of Policies on TS^{PDV} as a Function of the Trigger State for the Illustrative Technology Parameters. The compliance costs of the Concentration Restriction policy are not counted as costs to society in total surplus calculations, although they are costs to sellers.



trigger state e' . As none of the policies directly constrain sellers in state (1,1), the outcomes when the trigger is $e' = 2$ are identical to those when the policy is introduced from the start of the industry. This explains why we observe the No Leader AD policy raising welfare when $\tau = 0.9$ and $e' = 2$.

The figures show that, in fact, the trigger strategies do not raise TS^{PDV} when $\tau = 0.25$ or $\tau = 0.5$. Further analysis shows that TS^{PDV} is reduced for τ s even closer to zero: for example, the incentive (concentration restriction) policies with triggers of $e' = 20$ lower TS^{PDV} if $\tau > 0.06$ (0.08), although the reductions in welfare are small if $\tau > 0.8$. Therefore, though equilibrium medium- or long-run concentration is too high with no policies for $\tau = 0.25$ or $\tau = 0.5$, none of the considered pro-competition policies raise efficiency.

The incentive and concentration restriction policies are, however, predicted to raise TS^{PDV} when $\tau = 0$ and e' is chosen so that the policy only applies once one firm has reached, or almost reached, the bottom of its cost-curve (recall $m = 15$). The increase in efficiency may seem surprising given that concentration is too low when $\tau = 0$ and there are no policies, but this is, in fact, exactly the reason why efficiency increases. The laggard anticipates that the policy will help it catch-up once the leader reaches e' . Therefore, the laggard's incentives to compete before this state is reached are reduced, and the leader moves down its cost-curve more quickly, increasing concentration and surplus. This can lower medium-run surplus and raise medium-run concentration: for example, TS^{32} is lower with the policies, than with no policy, if $e' \geq 13$ and HHI^{32} increases if $e' \geq 20$.

6 Conclusion

The dynamic competition literature has assumed that sellers compete by setting prices or quantities. The existing applied literature on bargaining has focused on settings where either market structure is fixed or prices and market shares do not have direct

effect on firms' future competitiveness. We extend the dynamic competition and applied bargaining literatures by embedding buyer-seller bargaining into an existing and tractable model of dynamic price competition, with the standard assumption that sellers make take-it-or-leave-it price offers nested as a special case. By doing so, our model allows the allocation of bargaining power to affect market structure, and endogenous market structure to affect sellers' outside options, their dynamic incentives and negotiated outcomes. These interactions are likely to be important in the types of capital goods industries where negotiations are common and which have motivated much of the policy-oriented dynamic competition literature in international trade and industrial organization.

We find that, for a wide-range of technology parameters, the mark-ups that leaders charge when they set prices causes leads to be relatively short-lived, so that sellers' dynamic incentives are relatively weak and, because sellers move down their cost curves relatively slowly, production costs are quite high. These patterns can change quickly when buyers have even quite limited bargaining power, as mark-ups are reduced and longer lasting leads incentivize sellers to try to get ahead. As a result, conclusions that would be drawn assuming sellers set prices may not generalize even to cases where sellers have most, but not quite all, of the bargaining power.

We show how the allocation of bargaining power affect the design of optimal subsidy policies and the effects of policies that might be expected to promote competition. For example, subsidies that would be optimal with seller price-setting only require small deviations from this assumption to be welfare-reducing. Of course, it is an empirical question how bargaining power is actually allocated between buyers and sellers, and our results suggest that the accurate estimation of bargaining power in models with dynamic price competition is an important direction for future research. While our framework naturally extends to considering forward-looking buyers (Sweeting, Jia, Hui, and Yao (2022)), a convincing empirical analysis would likely need to include additional real-world features, such as multi-unit and multi-period contracts,

that the current paper has abstracted away from.

References

- ABBRING, J. H., AND J. R. CAMPBELL (2010): “Last-In First-Out Oligopoly Dynamics,” *Econometrica*, 78(5), 1491–1527.
- ABBRING, J. H., J. R. CAMPBELL, J. TILLY, AND N. YANG (2018): “Very Simple Markov-Perfect Industry Dynamics: Theory,” *Econometrica*, 86(2), 721–735.
- AN, Y., AND W. ZHAO (2019): “Dynamic Efficiencies of the 1997 Boeing-McDonnell Douglas Merger,” *RAND Journal of Economics*, 50(3), 666–694.
- 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.
- BARWICK, P. J., M. KALOUPSTIDI, AND N. B. ZAHUR (2019): “China’s Industrial Policy: An Empirical Evaluation,” Discussion Paper 26075, National Bureau of Economic Research.
- 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.
- (2019a): “How Efficient is Dynamic Competition? The Case of Price as Investment,” *American Economic Review*, 109(9), 3339–64.

- (2019b): “Sacrifice Tests for Predation in a Dynamic Pricing Model: Ordoover 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. SATTERTHWAITTE (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.
- COLLARD-WEXLER, A., G. GOWRISANKARAN, AND R. S. LEE (2019): ““Nash-in-Nash” Bargaining: A Microfoundation for Applied Work,” *Journal of Political Economy*, 127(1), 163–195.
- CRAWFORD, G. S., R. S. LEE, M. D. WHINSTON, AND A. YURUKOGLU (2018): “The Welfare Effects of Vertical Integration in Multichannel Television markets,” *Econometrica*, 86(3), 891–954.
- DASGUPTA, P., AND J. STIGLITZ (1988): “Learning-By-Doing, Market Structure and Industrial and Trade Policies,” *Oxford Economic Papers*, 40(2), 246–268.
- DORN, J. (2023): “Dynamic Bargaining Between Hospitals and Insurers,” Discussion paper, Princeton University.
- DRAGANSKA, M., D. KLAPPER, AND S. B. VILLAS-BOAS (2010): “A Larger Slice or a Larger Pie? An Empirical Investigation of Bargaining Power in the Distribution Channel,” *Marketing Science*, 29(1), 57–74.

- ERICSON, R., AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62(1), 53–82.
- FUDENBERG, D., AND J. TIROLE (1983): “Learning-by-Doing and Market Performance,” *Bell Journal of Economics*, 14(2), 522–530.
- GHEMAWAT, P. (1985): “Building Strategy on the Experience Curve,” *Harvard Business Review*, 63(2), 143–149.
- GOWRISANKARAN, G., A. NEVO, AND R. TOWN (2015): “Mergers When Prices Are Negotiated: Evidence from the Hospital Industry,” *American Economic Review*, 105(1), 172–203.
- GRENNAN, M. (2013): “Price Discrimination and Bargaining: Empirical Evidence from Medical Devices,” *American Economic Review*, 103(1), 145–177.
- HO, K., AND R. S. LEE (2017): “Insurer Competition in Health Care Markets,” *Econometrica*, 85(2), 379–417.
- HORN, H., AND A. WOLINSKY (1988): “Bilateral Monopolies and Incentives for Merger,” *RAND Journal of Economics*, 19(3), 408–419.
- KALAI, E. (1977): “Proportional Solutions to Bargaining Situations: Interpersonal Utility Comparisons,” *Econometrica*, 45(7), 1623–1630.
- KALOUPTSIDI, M. (2018): “Detection and Impact of Industrial Subsidies: The Case of Chinese Shipbuilding,” *The Review of Economic Studies*, 85(2), 1111–1158.
- KIM, M. (2014): “Strategic Responses to Used Goods Markets: Airbus and Boeing,” Discussion paper, University of Oklahoma.
- LEAHY, D., AND J. P. NEARY (1999): “Learning by Doing, Precommitment and Infant-Industry Promotion,” *Review of Economic Studies*, 66(2), 447–474.

- 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.
- LIEBERMAN, M. B. (1987): “Postentry Investment and Market Structure in the Chemical Processing Industries,” *RAND Journal of Economics*, 18(4), 533–549.
- MASKIN, E., AND J. TIROLE (2001): “Markov Perfect Equilibrium: I. Observable Actions,” *Journal of Economic Theory*, 100(2), 191–219.
- NEARY, J. P., AND D. LEAHY (2000): “Strategic Trade and Industrial Policy Towards Dynamic Oligopolies,” *Economic Journal*, 110(463), 484–508.
- 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.
- 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 (2023): “Dynamic Games with Asymmetric Information: Implications for Mergers,” Discussion paper, University of Maryland.
- TIEW, A. (2024): “Flailing Firms and Joint Operating Agreements: An Application to US Local Daily Print Newspapers from 1932 to 1992,” Discussion paper, Harvard University.
- YANG, C. (2020): “Vertical Structure and Innovation: A Study of the SoC and Smartphone Industries,” *RAND Journal of Economics*, 51(3), 739–785.