ONLINE APPENDICES FOR "DYNAMIC PRICE COMPETITION, LEARNING-BY-DOING AND STRATEGIC BUYERS" BY SWEETING, HUI, JIA AND YAO

A Implementation of Homotopy Methods

This Appendix provides details of our implementation of the homotopy algorithm. Our description will focus on our analysis of the BDKS model where we use a sequence of homotopies to try to enumerate the number of equilibria that exist for different values of (ρ, δ). Our implementation of other homotopies in the paper is similar to a single step in this sequence.

A.1 Preliminaries

For the BDKS model 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 $\rho \in [0, 1]$. We perform our procedure for six discrete values of b^p: 0 (i.e., the actual BDKS model), 0.01, 0.05, 0.1, 0.2 and 1. In Appendix C we report results using b^p-homotopies that do not assume discrete values. The state space of the game is defined by an (30×30) grid of values of the know-how of each firm.

A.2 System of Equations Defining Equilibrium

A Markov Perfect Equilibrium will be a combination of value functions for the buyer and the sellers, and a set of prices, that satisfy a system of 2,265 equations, which, as a group, we will denote F.

$$\left[\begin{array}{c} V_1^S(\mathbf{e}) - D_1(p(\mathbf{e}), \mathbf{e})(p_1(\mathbf{e}) - c_1(\mathbf{e})) - \sum_{k=1,2} D_k(p(\mathbf{e}), \mathbf{e})\mu_1^S(\mathbf{e}'_k) \\ V^B(\mathbf{e}) - \left\{ \begin{array}{c} b^p \log\left(\sum_{k=1,2} \exp\left(v_k - p_k + \mu^B(\mathbf{e}'_k)\right)\right) \\ +(1-b^p)\sum_{k=1,2} D_k(p(\mathbf{e}), \mathbf{e})\mu^B(\mathbf{e}'_k) \end{array} \right\} \\ D_1(p(\mathbf{e}), \mathbf{e}) + \left(p_1(\mathbf{e}) - c_1(\mathbf{e}) + \left[\mu_1^S(\mathbf{e}'_1) - \mu_1^S(\mathbf{e}'_2)\right]\right) \frac{\partial D_1(p(\mathbf{e}),\mathbf{e})}{\partial p_1} \end{array}\right] = 0.$$

where

$$D_k(p, \mathbf{e}) = \frac{\exp(v - p_k + \mu^B(\mathbf{e}'_k))}{\sum_{j=1,2} \exp(v - p_j + \mu^B(\mathbf{e}'_j))},$$

and $\mathbf{e}'_1 = (\min(e_1+1, M), e_2)$ and $\mathbf{e}'_2 = (e_1, \min(e_2+1, M))$, and $V^B(e_1, e_2) = V^B(e_2, e_1)$. The μ_ks are the continuation values of the players when a purchase has been made from firm k, the state has evolved to reflect the change in know-how due to the purchase but the realization of forgetting is yet to be realized.

$$\mu^B(e'_1, e'_2) = \beta \sum_{e''_{1,t+1}=\max(1,e'_{1,t}-1),e'_{1,t}} \sum_{e''_{2,t+1}=\max(1,e'_{2,t}-1),e'_{2,t}} V^B(e''_{1,t+1}, e''_{2,t+1}) \Pr(e''_{1,t+1}|e'_{1,t}) \Pr(e''_{2,t+1}|e'_{2,t}), \quad (10)$$

where the probabilities reflect the value of Δ given the state. $\mu^S(\mathbf{e_k})$, is the seller's continuation value when the buyer chooses to buy from seller k

$$\mu^S(e'_1, e'_2) = \beta \sum_{e''_{1,t+1}=\max(1,e'_{1,t}-1),e'_{1,t}} \sum_{e''_{2,t+1}=\max(1,e'_{2,t}-1),e'_{2,t}} V^S(e''_{1,t+1}, e''_{2,t+1}) \Pr(e''_{1,t+1}|e'_{1,t}) \Pr(e''_{2,t+1}|e'_{2,t}). \quad (11)$$

A.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, for our BDKS analysis, 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.

A.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.[36]

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

[36]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.

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 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 will be δ, ρ or b^p). 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 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.[37,38]

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.[39]

[37] 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.

[38] 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`.

[39] 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}.

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.

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.[40] 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. It happens after four rounds for $b^p = 0.1$.

While we do not use the b^p-homotopies in our calculation of the number of equilibria, we do use them as an additional check on whether we are missing equilibria.

[40] 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.

B Recursive Algorithms for Establishing Existence and Uniqueness of Accommodative and SELPM Equilibria in the BDK Model

Our second approach to finding equilibria in the BDK model uses backwards induction to exploit the fact that two common types of equilibria have absorbing terminal states.

Definition *An equilibrium is* ***accommodative*** *if* $\lambda_1(e_1, e_2) = \lambda_2(e_1, e_2) = 1$ *for all states* (e_1, e_2) *where* $e_1 > 0$ *and* $e_2 > 0$.

In an accommodative equilibrium there is no exit by active firms. If the industry starts off in state (1,1), it is guaranteed to arrive in state (M, M) in an accommodative equilibrium. This definition is the same as in BDK (2019), Appendix B.

Definition *An equilibrium has the "****S****ome* ***E****xit* ***L****eads to* ***P****ermanent* ***M****onopoly" (SELPM) property if there is some* $\mathbf{e} = (e_1, e_2)$ *where* $e_1 > e_2 > 0$*, and (i)* $\lambda_2(\mathbf{e}) < 1$*, (ii)* $\lambda_2(e_1', 0) = 0$ *and* $\lambda_1(e_1', 0) = 1$ *for all* $e_1' \geq e_1$*, and (iii)* $\lambda_1(e_1', e_2') = 1$ *for all* $e_1' \geq e_1, e_2' \geq e_2$.

In a SELPM equilibrium, once the game has reached a state $\mathbf{e}$ that satisfies the requirements the game will not return to that state once it has left it, and the game must eventually end up in one of the states (M, M), $(M, 0)$ or $(0, M)$. Exit followed by entry can occur in lower know-how states.

B.1 Existence of an Accommodative Equilibrium

We establish whether an accommodative equilibrium exists by solving, using `fsolve` in MATLAB with analytic derivatives, for equilibrium prices and values assuming that there is no exit from any duopoly state, and then verifying that it is always optimal for each duopolist to continue in every duopoly state by checking that $\beta V^S(e_1, e_2)$ is greater than the highest possible scrap value.

B.2 Existence of a SELPM Equilibrium and Uniqueness of Accommodative Equilibria

We now describe how we identify whether a SELPM equilibria exists, before noting how we can use a closely related procedure to check whether an accommodative equilibrium is unique. As we explain, our conclusions that an accommodative equilibrium is unique or that no SELPM equilibria exist rely on our ability to find **all** equilibria in a given state, holding fixed what will happen if and when the state transitions. We explain why we are confident that we can do this.

B.2.1 Algorithm for Establishing whether a SELPM Equilibrium Exists

Overview. The objective of our algorithm is to identify whether a state **e** which satisfies the SELPM conditions (i.e., no exit by the leader, some probability of exit by the laggard without any re-entry) exists. This is done by following equilibrium paths backwards through the state space from (M, M) until such a state is found (in which case the algorithm terminates with a success), or one of the conditions for a SELPM equilibrium is violated (e.g., re-entry can occur when the leader has a higher know-how state) or state (1,1) is reached without finding an e. In the latter two cases, the algorithm returns a failure for the path and another path is tried. If no state **e** has been found once all paths have been tried, we know no SELPM equilibrium exists.

We present our routine assuming that $\sigma = 1$ and $\beta^b = \beta$ to reduce notation. Our examples assume the illustrative parameters $\rho = 0.75$ and $\sigma = 1$ unless otherwise stated.

Step 1: Solving for Outcomes with Permanent Monopoly. In a SELPM equilibrium, behavior in duopoly states with higher know-how than state **e** will reflects players' assumption that there will be not be re-entry following exit. The first step in our procedure is to solve for equilibria in states with monopoly, $(1, 0)$ to $(M, 0)$, using backwards induction, assuming that the monopolist and the strategic buyers expect no entry by the potential entrant in these states, so that we can store these solutions in memory. Of course, once we have solved for strategies in the duopoly states, we can verify that no entry is actually optimal.

The equilibrium in state $(M, 0)$ will be unique for all values of b^p. With $b^p = 0$ this follows from the fact that, in state $(M, 0)$ the game is simply a repeated monopoly pricing problem. However, demand, and therefore the seller's problem, is identical in this state for any value of b^p, because the strategic buyer cannot affect the evolution of the state, so $b^p = 0$ pricing will be an equilibrium.

Assuming that a monopolist will not exit (a condition that can be verified and is always satisfied for the parameters that we consider), the equations that determine the equilibrium values of V^B, V_1^S and p_1 in state $(e_1 < M, 0)$, where firm 1's marginal cost is c, are

$$V^B = b^p \left(\ln\left(\exp(\beta V^B) + \exp(v - p_1 + \beta V^B(e_1 + 1, 0))\right)\right) + ... \qquad (12)$$
$$\beta(1 - b^p)(D_1 V^B(e_1 + 1, 0) + (1 - D_1)V^B)$$
$$V^S = (p_1 - c + \beta V^S(e_1 + 1, 0))D_1 + \beta V^S(1 - D_1) \qquad (13)$$
$$D_1 + (p_1 - c + \beta V^S(e_1 + 1, 0) - \beta V^S)\frac{\partial D_1}{\partial p_1} = 0 \qquad (14)$$

where $D_1 = \frac{\exp\left(v - p_1 + \beta V^B(e_1+1,0)\right)}{\exp(v - p_1 + \beta V^B(e_1+1,0)) + \exp(\beta V^B)}$ assuming, following BDK, that $v_0 = p_0$. V^B and

Figure B.1: Monopoly State Equations in State $(10, 0)$: black curve is the value of V^B as a function of p_1, red curve is the optimal p_1 given V^B. There is an equilibrium where the lines intersect.

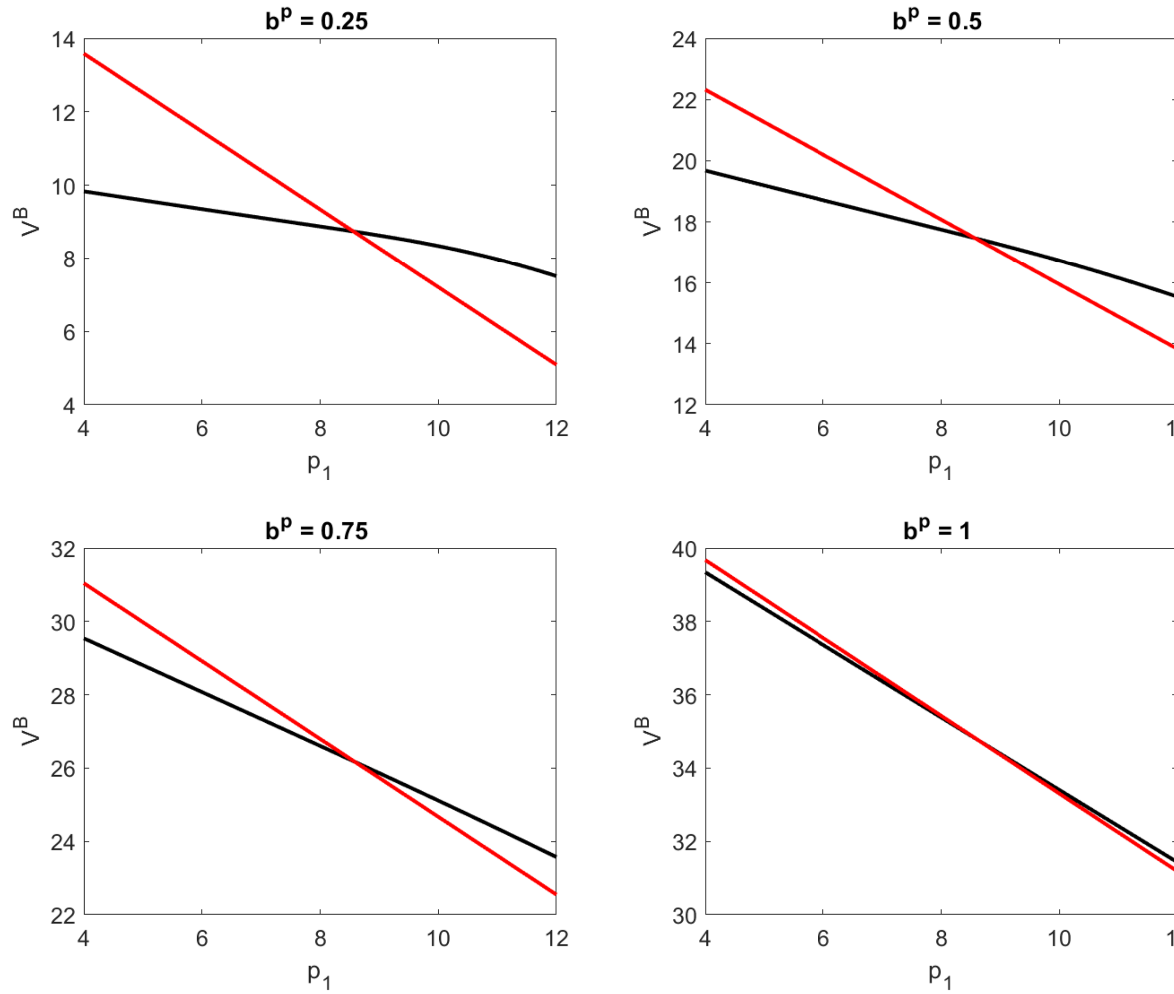

V^S denote $V^B(e_1, 0)$ and $V^S(e_1, 0)$ respectively. The last equation is the first-order condition for setting prices. When we are solving for the equilibrium in state $(e_1, 0)$, $V^B(e_1 + 1, 0)$ and $V^S(e_1 + 1, 0)$ are fixed.

We can check for evidence of multiplicity in states $(e_1 < M, 0)$ by identifying whether two equilibrium curves intersect more than once.[41] The first curve solves the value of V^B as a function of p_1, reflecting equation (12). The second curve solves for the value of p_1 that maximizes the seller's value, given V^B, as determined by equation (14).

Figure B.1 presents examples of what these curves look like for state $(10, 0)$ using the baseline parameters when $b^p = 0.25$, $0.5, 0.75$ and 1. The black curves denote the value of

[41] The standard proof of a unique pricing equilibrium in a game with logit demand is insufficient because prices affect continuation values.

V^B given p_1, and the red curves reflect the value-maximizing choices of p_1 given values of V^B. In each case the curves cross only once, consistent with a single equilibrium. We verify that there is generally a single equilibrium by repeating this exercise denoted in Figure 1 for a very large number of different values of ρ, σ, b^p, $V^S(e_1+1,0)$ and $V^B(e_1+1,0)$.[42] We find no values where there are multiple equilibria.

Step 2: Solving for the Equilibrium in State (M, M). The other possibile outcome in an accommodative or SELPM equilibrium is that the game ends up at state (M, M). As there can be no possibility of exit in state (M, M) in a SELPM equilibrium, and purchases cannot change the state, the equilibrium in this state must be unique and pricing strategies will be unaffected by b^p.[43] We can therefore find equilibrium prices by solving the standard pricing first-order conditions,

$$D_i + (p_i - c)\frac{\partial D_i}{\partial p_i} = 0,$$

and then calculating the implied seller values (V^S). We verify that βV^S is greater than the maximum possible scrap value, so that exit is not optimal. If exit could be optimal, there is no SELPM equilibrium.

Step 3: Recursive Procedure. We now use backwards induction to proceed backwards through the game, starting in state $(M, M-1)$ (of course, the state $(M-1, M)$ is symmetric). This procedure, for all equilibirum paths, is followed until a state $\mathbf{e}$ which satisfies the SELPM definition has been found, or we have confirmed that no such state exists. We first describe the procedure we apply to find all equilibria in state (e_1, e_2), with $e_1 \geq e_2$, before detailing the backwards induction routine in which this is embedded.

Finding Equilibria Consistent with SELPM in State $(e_1, e_2), e_1 \geq e_2$. In a state that can follow a state $\mathbf{e}$ in a SELPM equilibria, there may be some probability that the laggard firm exits unless $e_1 = e_2$, but the leader must continue with probability 1. We are therefore trying to solve for all equilibria with different combinations of

- prices (p_1, p_2)
- values (V_1^S, V_2^S, V^B)

[42] Specifically, we use b^p values on a grid [0.2,0.4,0.6,0.8,1], ρ values [0,0.1,0.2,..,0.9,1], σ values [0.5,0.6,..,1.1,1.2], $V^S(e_1+1,0)$ values [60,65,..,95,100] and $V^B(e_1+1,0)$ values b^p*[20,25,30,35,40]. This gives a total of 19,800 combinations that we check. We have also experimented with other values.

[43] As purchases cannot change the state, a strategic buyer will behave in the same way as a non-strategic buyer. There will be a single pricing Markov Perfect Nash equilibrium in a repeated game with multinomial logit demand and marginal costs that do not change.

- continuation probability for firm 2 (λ_2)

that satisfy the equilibrium equations (15)-(18).

$$V_i^S - D_i(p_1, p_2, V^B)(p_i - c_i(e_i)) - \sum_{k=0,1,2} D_k(p_1, p_2, V^B) V_i^{S,INT}(\mathbf{e}'_k) = 0 \text{ for } i = 1, 2 \tag{15}$$

If $e_1 < M$,

$$V_1^{S,INT}(\mathbf{e}'_1) = \beta \left(\lambda_2(e_1+1, e_2) V_1^S(e_1+1, e_2) + (1 - \lambda_2(e_1+1, e_2)) V_1^S(e_1+1, 0) \right)$$

so its value is known when solving for the equilibrium. Similarly,

$$V_2^{S,INT}(\mathbf{e}'_1) = \beta \left(\lambda_2(e_1+1, e_2) V_2^S(e_1+1, e_2) + (1 - \lambda_2(e_1+1, e_2)) E(X|\lambda_2(e_1+1, e_2)) \right),$$

where $E(X|\lambda_2(e_1+1, e_2))$ is the expected scrap value if firm 2 exits with probability $1 - \lambda_2(e_1+1, e_2)$. Alternatively, if $e_1 = M$, $V_1^{S,INT}(\mathbf{e}'_1) = \beta \left(\lambda_2 V_1^S + (1-\lambda_2) V_1^S(M, 0) \right)$ and $V_2^{S,INT}(\mathbf{e}'_1) = \beta \left(\lambda_2 V_2^S + (1-\lambda_2) E(X|\lambda_2) \right)$, so they depend on the endogenous λ_2 and V_1^S, because a sale to firm 1 does not change the state.

For all e_1,

$$V_1^{S,INT}(\mathbf{e}'_2) = \beta \left(\lambda_2(e_1, e_2+1) V_1^S(e_1, e_2+1) + (1 - \lambda_2(e_1, e_2+1)) V_1^S(e_1, 0) \right),$$
$$V_2^{S,INT}(\mathbf{e}'_2) = \beta \left(\lambda_2(e_1, e_2+1) V_2^S(e_1, e_2+1) + (1 - \lambda_2(e_1, e_2+1)) E(X|\lambda_2(e_1, e_2+1)) \right),$$

$V_1^{S,INT}(\mathbf{e}'_0) = \beta \left(\lambda_2 V_1^S + (1-\lambda_2) V_1^S(e_1, 0) \right)$ and $V_2^{S,INT}(\mathbf{e}'_0) = \beta \left(\lambda_2 V_2^S + (1-\lambda_2) E(X|\lambda_2) \right)$.

The first-order condition for prices and the equation defining the probability that firm 2 continues are

$$D_i(p_1, p_2, V^B) + \sum_{k=0,1,2} \frac{\partial D_k(p_1, p_2, V^B)}{\partial p_i} V_i^{S,INT}(\mathbf{e}'_k) + (p_i - c_i(e_i)) \frac{\partial D_i(p_1, p_2, V^B)}{\partial p_i} = 0 \text{ for } i = 1, 2 \tag{16}$$

$$\lambda_2 - F_{\text{scrap}}(\beta V_2^S) = 0 \tag{17}$$

where

$$V^B - b^p \log \left(\sum_{k=0,1,2} \exp \left(v_k - p_k + V^{B,INT}(\mathbf{e}'_k) \right) \right) - (1-b^p) \sum_{k=0,1,2} D_k(p_1, p_2, V^B) V^{B,INT}(\mathbf{e}'_k) = 0, \tag{18}$$

$$V^{B,INT}(\mathbf{e}'_2) = \beta \left(\lambda_2(e_1, e_2+1) V^B(e_1, e_2+1) + (1 - \lambda_2(e_1, e_2+1)) V^B(e_1, 0) \right), \tag{19}$$

$$V^{B,INT}(\mathbf{e}_0') = \beta\left(\lambda_2 V^B + (1-\lambda_2)V^B(e_1,0)\right), \tag{20}$$

and

$$V^{B,INT}(\mathbf{e}_1') = \beta\left(\lambda_2(e_1+1,e_2)V^B(e_1+1,e_2) + (1-\lambda_2(e_1+1,e_2))V^B(e_1+1,0)\right) \tag{21}$$

if $e_1 < M$, and

$$V^{B,INT}(\mathbf{e}_1') = \beta\left(\lambda_2 V^B + (1-\lambda_2)V_2^B(e_1,0)\right) \tag{22}$$

if $e_1 = M$.

There can be multiple solutions to these equations. We proceed assuming that there is a single equilibrium for a given value of λ_2, an assumption that we will verify below, and we solve equations (15), (16) and (18) given this value.[44] Given the prices and values, we then calculate what the optimal exit probability will be using (17). When we use a fine grid of values of λ_2 we can construct graphs like the ones shown in Figure B.2, which show this best response λ_2 function, as a function of an arbitrary λ_2, for the illustrative parameters, for states (30,1) and (30,5), with $b^p = 0$ and $b^p = 0.2$.

There is an equilibrium where the function crosses the 45^0 line, and we identify the precise solutions to all of the equations using the solutions at the gridpoints that surround the intersection as starting points. We then verify whether or not the leader wants to continue with probability 1, which is necessary for a SELPM equilibrium. Assuming that our initial grid of λ_2 is sufficiently fine, this approach will find all of the equilibria in a given state that are consistent with the SELPM criteria.

Backwards Induction Algorithm. Our aim is to identify whether a SELPM equilibrium exists. Therefore we follow equilibrium paths through the state space until we can identify a state which satisfies all of the requirements for a state $\mathbf{e}$ in the definition of a SELPM equilibrium, at which point the algorithm terminates in success, or, alternatively, we have followed all equilibrium paths and found that there is no such state, because along each path it is case that either a leader may choose to exit, no exit occurs anywhere in the game or there is positive probability that re-entry may occur after any exit. In this alternative case, no SELPM equilibrium exists.

The order in which we loop through states is to go from $e_1 = 30$ to $e_1 = 1$ as the outer loop, and to go from $e_2 = e_1$ to $e_2 = 1$ on the inner loop, while performing an additional check on whether re-entry may be optimal if $e_2 = 0$ (i.e., this verifies that we can use the monopoly solution calculated earlier).

[44] Occasionally the equations do not solve using the starting values chosen, in which case we use a Pakes and McGuire (1994) type of routine to find alternative starting values.

Figure B.2: Best Response Continuation Probability Functions for Firm 2 Given Endogenous Pricing Choices by Both Firms. Intersections with the 45^0-degree line are equilibria.

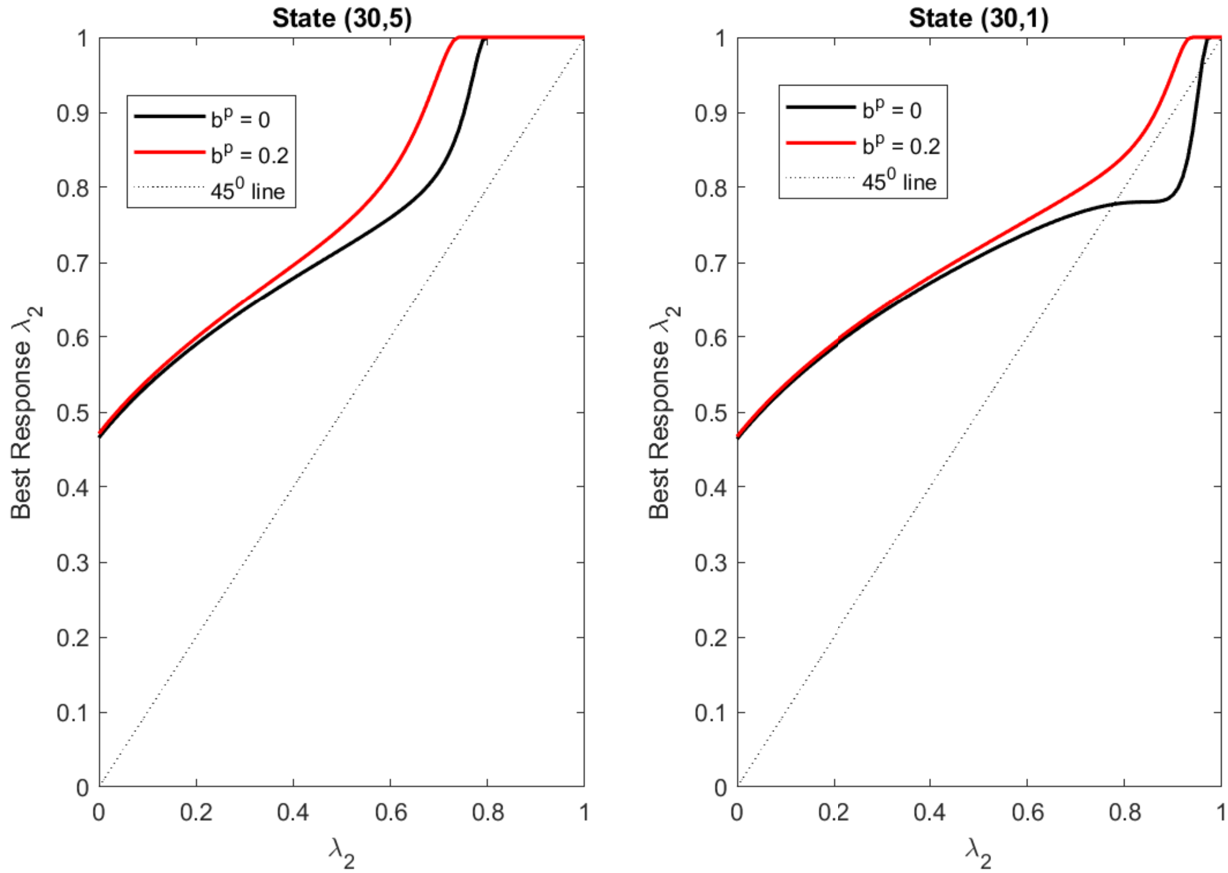

As an example, consider the illustrative parameters. With $b^p = 0$, the backwards induction algorithm identifies a single equilibrium path with no exit until state $(30, 1)$ where it identifies three equilibria with probabilities that firm 2 continues (λ_2) of 0.7777, 0.9577 and 1 (these correspond to the three intersections of the black line and the 45^0 line in the second panel of Figure B.2). If we take the first equilibrium, we can identify that it is never optimal for firm 2 to enter in state $(30, 0)$. This implies that $(30, 1)$ is a state that meets the criteria for a state **e** in a SELPM equilibria, so the algorithm terminates at this point. We would reach the same conclusion using the second equilibrium. On the other hand, when $b^p = 0.2$ we identify only a single equilibrium path where there is no exit in any state all of the way back to $(1, 1)$. In this case, we can conclude that no SELPM equilibrium exists.

Verifying the Uniqueness of Equilibrium in a State Given a Firm 2 Continuation Probability. The conclusion that there is no SELPM equilibrium for given parameters relies on us having identified **all** equilibrium paths that could meet the SELPM criteria. This will be the case as long as there cannot be multiple price equilibria for a given probability that the laggard firm continues.

There are two types of evidence that support this presumption. First, we have never identified an instance of multiple equilibria for any of the parameters that we have considered, even when using multiple different starting points or alternative solution algorithms. Second, we have tried to identify whether there could be multiple equilibria by using a reaction function-type of analysis.

Specifically, for a given value of λ_2 and the continuation values, we solve the equations for V^B, V_1^S and the first-order condition for p_1 for a grid of alternative values of p_2. We then solve the equations for V^B, V_2^S and the first-order condition for p_2 for a grid of alternative values of p_1. We can then draw curves $p_1^*(p_2)$ and $p_2^*(p_1)$, which reflect the best response behavior of strategic buyers to both prices. The intersections correspond to equilibria, and we can test whether they intersect more than once. Figures B.3 present some examples of these curves for the illustrative parameters, $b^p = 0$ or $b^p = 1$ and $e_1 = 30$ and $e_2 = 1$.

Recall that in the state (30,1), if the buyer purchases from firm 1, the state remains (30,1), whereas if firm 2 makes a sale, the state transitions to (30,2), where, for these parameters, there is always a unique equilibrium. If firm 2 sets a very low price, a strategic buyer will tend to shift demand towards firm 1 in order to keep the state the same in future periods. As a result, firm 1's optimal price is less sensitive to firm 2's price in this state when $b^p = 1$, which accounts for the change in the slope of the reaction functions. However, in all cases, the reaction functions only intersect once, and there is a single equilibrium.

In practice, it would be too computationally intensive to implement this check for every value of λ_2 in every state for all parameters. However, for all of the values of the parameters where our algorithm does not identify a SELPM equilibrium, but it does for nearby parameter values, we have performed the check for a grid of values of λ_2. In each case, we have found exactly one intersection (equilibrium) and we are therefore confident in our claim that we are not missing SELPM equilibria. This is also the case when we have solved games for many different sets of arbitrary continuation values and parameters.

B.2.2 Procedure for Identifying the Existence of AELPM Equilibria

Definition *An equilibrium has the "**A**ny **E**xit **L**eads to **P**ermanent **M**onopoly" (AELPM) property if (i) for any* $\mathbf{e} = (e_1, e_2)$ *where* $e_1 > e_2 > 0$*, and* $\lambda_2(\mathbf{e}) < 1$*,* $\lambda_2(e_1', 0) = 0$ *for all* $e_1' \geq e_1$*;(ii)* $\lambda_1(e, e) = 1$ *for all* $e > 0$*, and (iii)* $\lambda_1(e_1, e_2) = 1$ *for all* $e_1 > e_2$*.*

In an AELPM equilibrium, any exit by a laggard firm in duopoly will lead to permanent monopoly. We can establish whether an AELPM equilibrium exists using a similar equilibrium path following procedure to the one used for SELPM equilibria, except that we now need to

Figure B.3: Pricing Best Response Functions in State (30,1) for Different Assumed Continuation Probabilities for Firm 2 (λ_2).

(a) $b^p = 0$

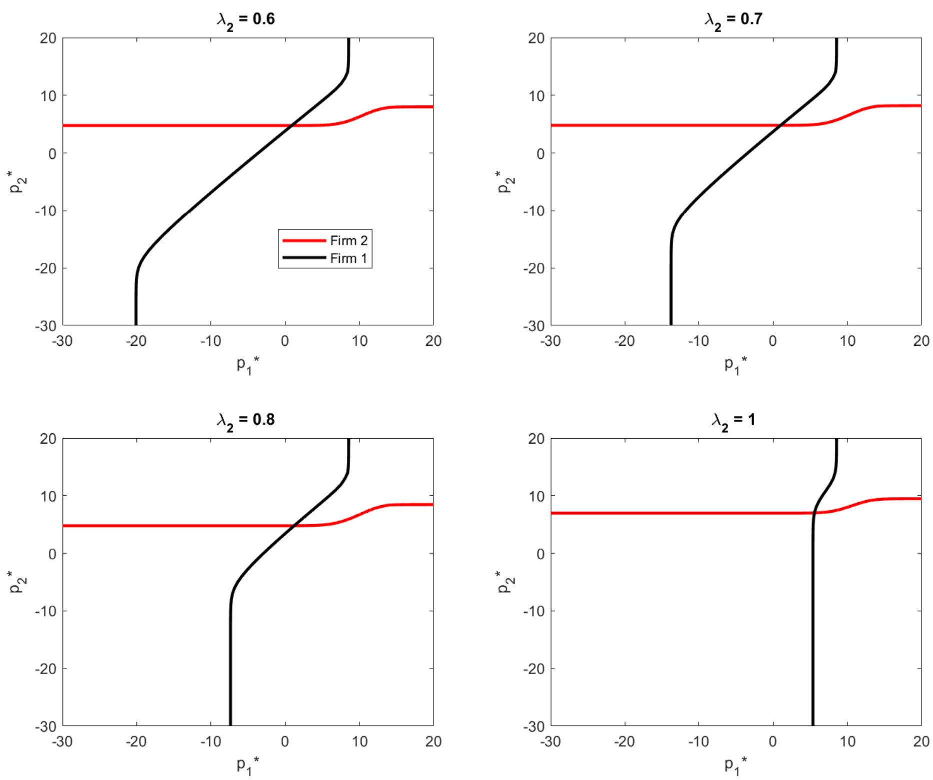

(b) $b^p = 1$

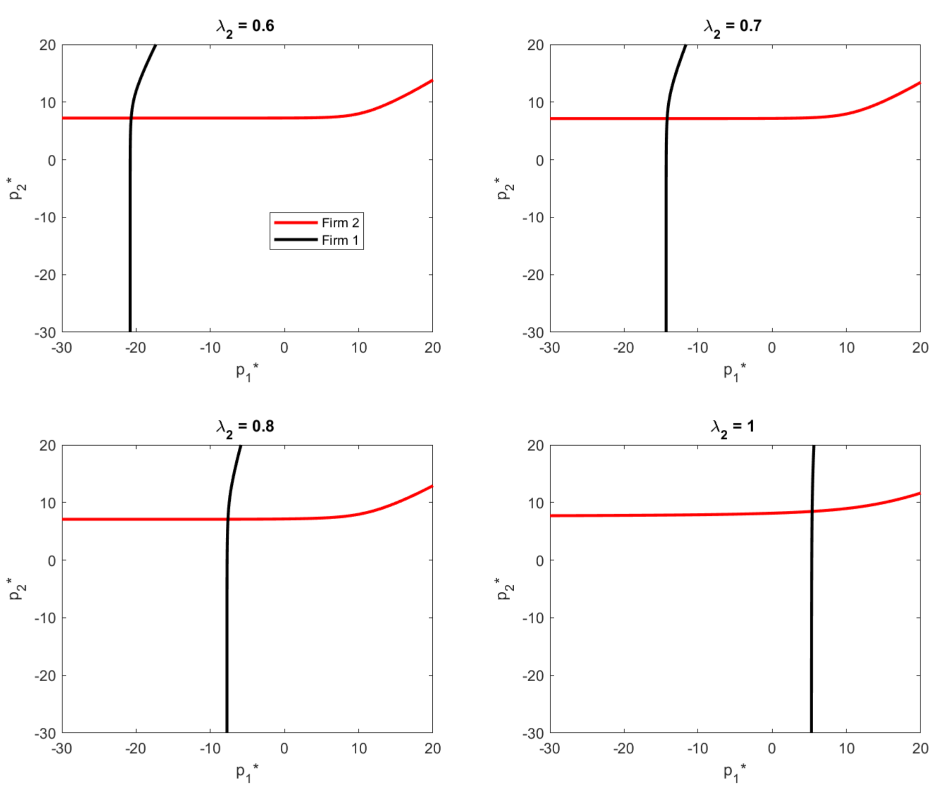

follow paths from $(30, 30)$ to $(1, 1)$ to make sure that no exit is followed by re-entry, and that symmetric firms will never exit.

For the illustrative baseline parameters with $b^p = 0$, the equilibrium path that contains $\lambda_2(30, 1) = 0.9577$ does find an AELPM equilibrium (which corresponds to the Mid-HHI equilibrium in Table 1). On the other hand, the equilibrium path that contains $\lambda_2(30, 1) = 0.7777$ fails the AELPM test when it reaches (1,1) as, in the corresponding High-HHI equilibrium, there is a small probability that both of the firms will exit, in which case re-entry may occur.

The computational burden involved in identifying whether AELPM equilibria exist is much higher than the one that identifies if SELPM equilibria exist, because of the need to follow a path all of the way through the state space and the possibility that there are a large number of paths, all of which fail when both firms have low-levels of know-how. Therefore, we focus on the SELPM results, and discuss AELPM equilibria in the context of classifying the equilibria identified by the homotopy approach.

B.2.3 Establishing the Uniqueness of Accommodative Equilibria

Our evidence that accommodative equilibria are unique comes from the applying the procedure described in the "Verifying the Uniqueness of Equilibrium in a State Given a Firm 2 Continuation Probability" subsection, but considering only $\lambda_2 = 1$. We have not found more than one intersection (equilibrium) in more than one state for any set of parameters.

C Additional Results

C.1 Additional Welfare Results for the Illustrative Parameters in the BDK Model

The results in Figure 3 in the main text show that, for $b^p = 0$, the NPV of consumer surplus is highest in the Mid-HHI equilibrium and lowest in the High-HHI equilibrium, whereas total surplus is highest in the accommodative equilibrium and lowest in the High-HHI equilibrium. As b^p increrases, both measures of surplus fall in the accommodative equilibrium as prices tend to increase (Figure 2(b)).

The reason why consumer surplus is higher in the Mid-HHI equilibrium is that initial prices are low, and the probability that the industry evolves into a long-run duopoly is not too much lower than in an accommodative equilibrium. To illustrate what happens to welfare in the first part of the game, Figure C.1(a) and (b) shows the expected surplus measures for the first ten periods of a game beginning at (1,1). Note that the reported numbers are sums and there is no discounting.

Figure C.1: Extended BDK Model: Equilibrium Expected Consumer Surplus, Total Surplus and Production Costs Over the First 10 Periods of the Game (starting at (1,1)) as a Function of b^p for the Illustrative Parameters. The black line traces the homotopy path from the Accommodative (A) baseline equilibrium. The red line traces the overlapping paths from the High-HHI (H) and Mid-HHI (M) baseline equilibria.

(a) Consumer Surplus

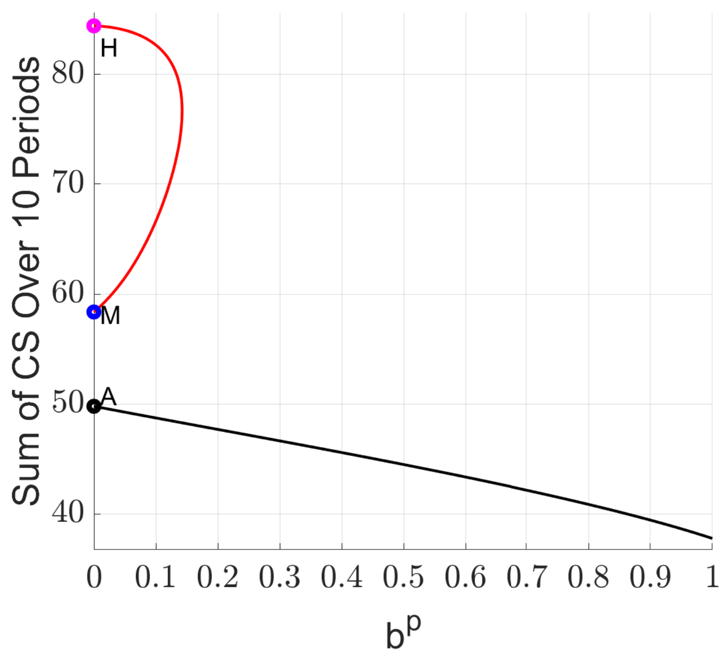

(b) Total Surplus

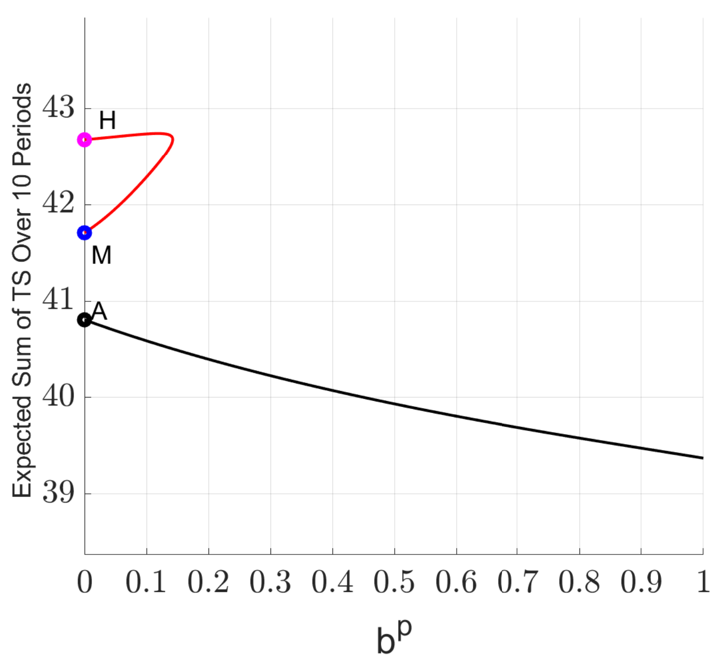

Figure C.1: Equilibrium Expected Consumer Surplus, Total Surplus and Production Costs Over the First 10 Periods of the Game (starting at (1,1)) as a Function of b^p for the Illustrative Parameters. The black line traces the homotopy path from the Accommodative (A) baseline equilibrium. The red line traces the overlapping paths from the High-HHI (H) and Mid-HHI (M) baseline equilibria.

(c) Expected Production Costs Over the First 10 Periods of the Game

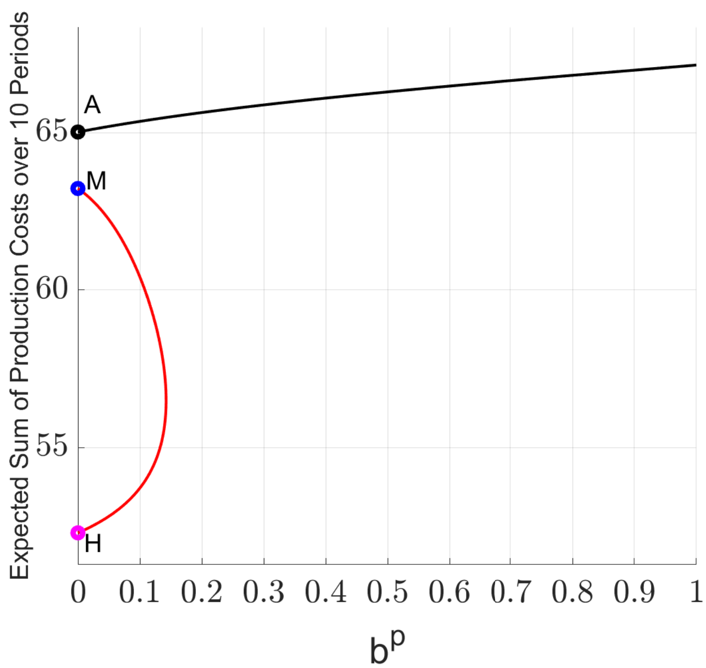

Figure C.2: Extended BDK Model: Expected NPV Per-Sale Production Costs and the Expected NPV of the Number of Sales (starting at (1,1)) for the Illustrative Parameters and Varying b^p.

(a) Expected Discounted Per-Sale Production Costs

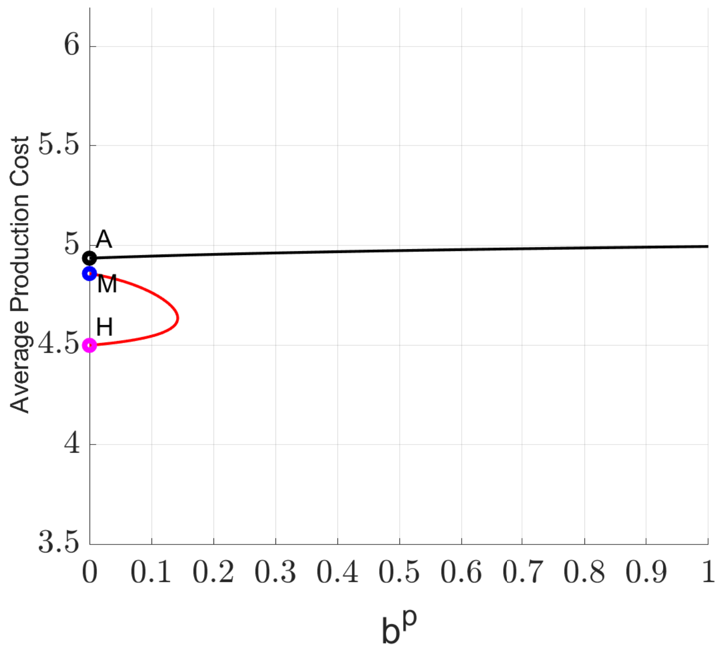

(b) Expected Discounted Number of Sales

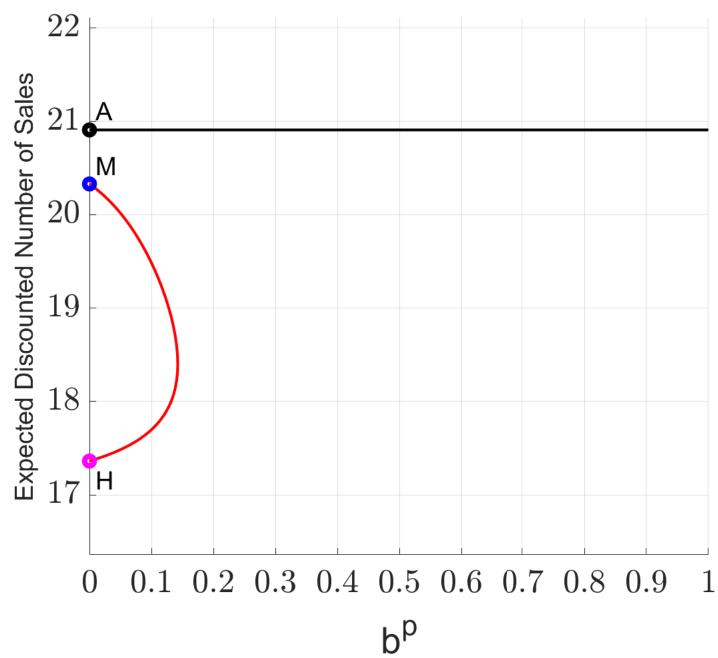

Consumer surplus is highest in the High-HHI equilibrium due to the very low initial prices. This also tends to increase total surplus. Total surplus is also increased by the reduction in production costs which results from one firm tending to make most of the sales. This is illustrated in Figure C.1(c), which shows the sum of production costs over the first ten periods. The effect that strategic buyer behavior increases prices in the accommodative equilibrium causes both measures of surplus to fall in the accommodative equilibrium as b^p is increased.

The NPV of total surplus is affected by the number of sales that are made and the costs of production. Figure C.2(a) shows that the expected discounted production cost per sale is highest in the accommodative equilibrium, due to slower early learning, and it is lowest in the High-HHI equilibrium, where learning will tend to be quickest.[45] Figure C.2(b) shows the discounted total number of sales that are made. Even though low prices mean that more sales are made at the very beginning of the game in the High-HHI and Mid-HHI equilibria, the discounted number of sales is highest in the accommodative model as, despite higher average production costs, long-run margins are low.

C.2 Classification of σ-Homotopy Equilibria in the Extended BDK Model

Figure 7(b) in the text shows the classification of the equilibria identified by ρ-homotopies into accommodative, AELPM, non-AELPM SELPM and BDK's aggressive equilibrium category. Figure C.3 shows a similar classification for the σ-homotopies. As with the ρ analysis, we observe that the accommodative/aggressive classification is not exhaustive, while the accommodative/SELPM (which includes AELPM) classification is exhaustive, and that AELPM equilibria are typically associated with lower HHI^∞ values than non-AELPM SELPM equilibria on a homotopy path for the same b^p.

C.3 Equilibrium Prices in the BDKS Model for $b^p = 0$, $\rho = 0.75$ and $\delta = 0.0275$.

Figure C.4 plots all of the equilibrium prices in the BDKS model when $b^p = 0$, $\rho = 0.75$ and $\delta = 0.0275$, which are the parameters considered in the illustrative example in Section 5.2.

[45] The reported number is the expected discounted total sum of production costs divided by the expected discounted total number of sales.

Figure C.3: Extended BDK Model: Classification of Equilibria Identified by σ-Homotopies for Full Range of b^p with Other Parameters at their Illustrative Values.

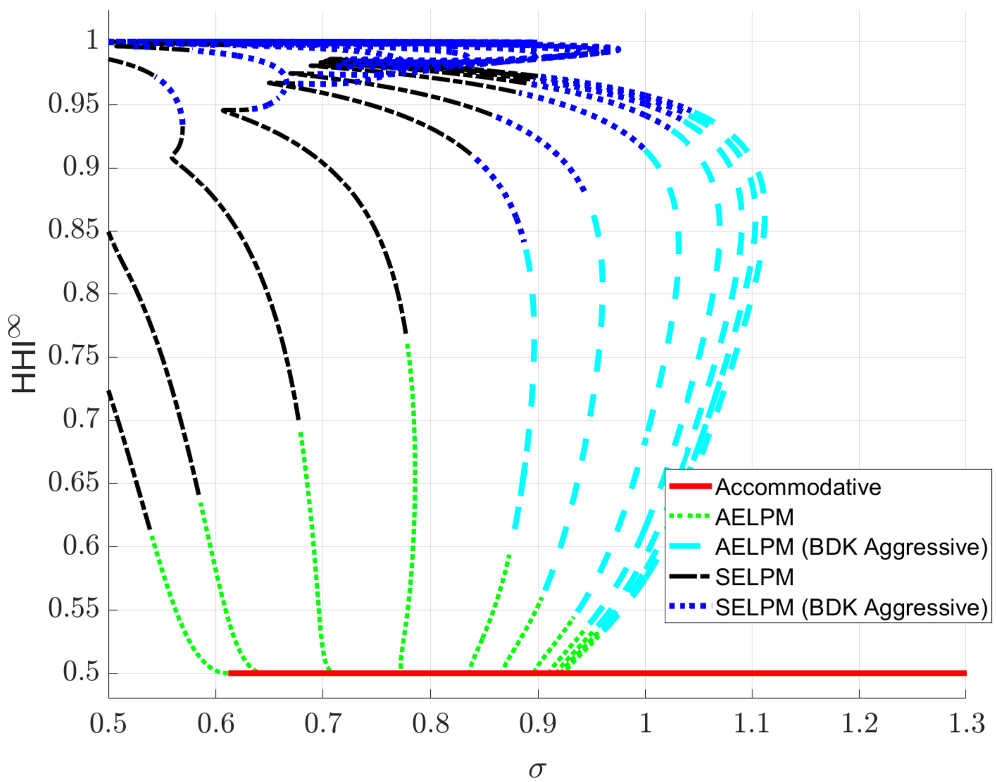

Figure C.4: $b^p = 0$ Equilibrium Prices for the Illustrative Parameters in the BDKS Model.

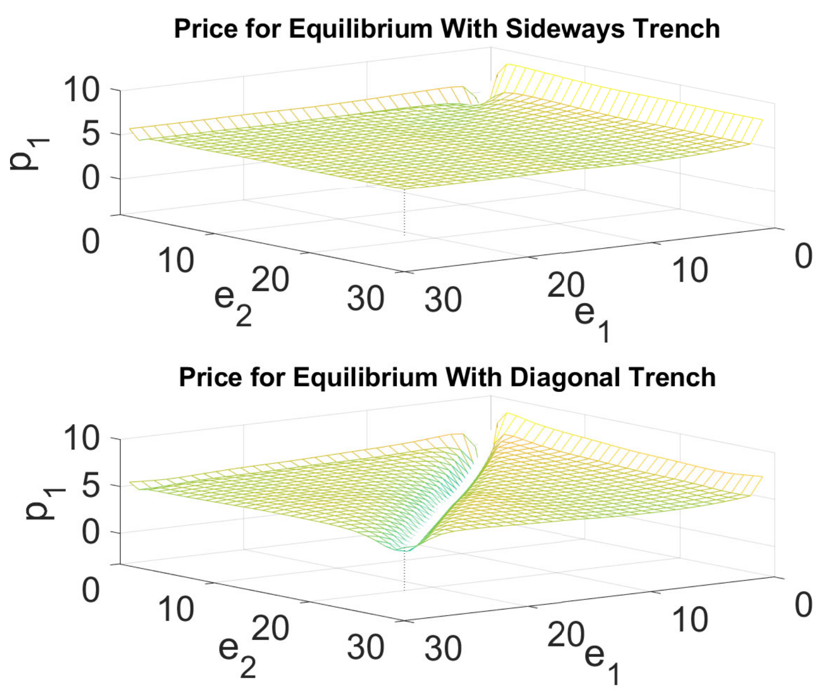

C.4 b^p-Homotopies Illustrating the Results for the Number of Equilibria in the BDKS Model

The main results from the analysis of the number of equilibria in the BDKS model are that (i) the number of equilibria typically declines quite quickly as b^p is increased, so that, when $b^p \geq 0.1$ there is typically a unique equilibrium, but (ii) there are some parameters where the number of equilibria increases when we move from $b^p = 0$ to $b^p = 0.01$ or 0.05. These results are based on δ- and ρ-homotopies. In this Appendix, we illustrate that the results are consistent with what we find using b^p-homotopies. While we only present results for two pairs of parameters, we see similar results for many other parameters that we have looked at.

Figure C.5(a) shows an example where there are seven equilibria identified by the δ- and ρ-homotopies when $b^p = 0$. When we start the b^p-homotopies from each of these equilibria we find that there are 3 pairs that are connected by semi-circular loops, none of which extend past $b^p = 0.03$ and, from the equilibrium with the lowest concentration, a fairly flat path (with respect to the concentration outcome) that continues through higher values of b^p.

Figure C.5(b) shows an example where there is a single equilibrium identified by the δ- and ρ-homotopies when $b^p = 0$. The b^p-homotopy from this equilibrium forms a loop, but, when it reaches a value of b^p just above 0.01, it reverses itself and then continues out through higher values of b^p. In this case we find multiplicity of equilibria only for $b^p = 0.05$ out of the discrete b^p values that we consider.

C.5 Outcomes in the Extended BDKS Model as a Function of b^p for $\delta = 0.0275$ and $\rho = 0.85$ and $b^p = 0$ Seller Strategies

In the text we report what happens to a set of outcomes when we hold seller strategies fixed and increase b^p for the parameters $\delta = 0.0275$ and $\rho = 0.75$. We choose $\rho = 0.75$ so that the progress ratio is the same as in our illustrative example for the BDK model. In this Appendix, we shows what happens to firm 1's demand and its dynamic incentives in state (3,1), holding seller strategies fixed at their baseline equilibrium values, when $\delta = 0.0275$ and $\rho = 0.85$, as these are parameter that BDKS use. For these parameters, there are two baseline equilibria, which is unusual. One is a diagonal trench and the other is "flat with well" where there is only an obvious dip in prices in state (1,1).

The most obvious difference between Figure C.6 and the ones reported in the text is that in the flat with well equilibrium, demand is less sensitive to b^p, which makes sense as buying from the laggard is not much more likely to return the game to the well. The advantage-denying incentives are also somewhat smaller for both equilibria, although the advantage-building incentives retain a non-monotonic pattern.

Figure C.5: Extended BDKS Model: Examples of b^p-Homotopies From $b^p = 0$ Equilibria.

(a) $\delta = 0.066$, $\rho = 0.31$

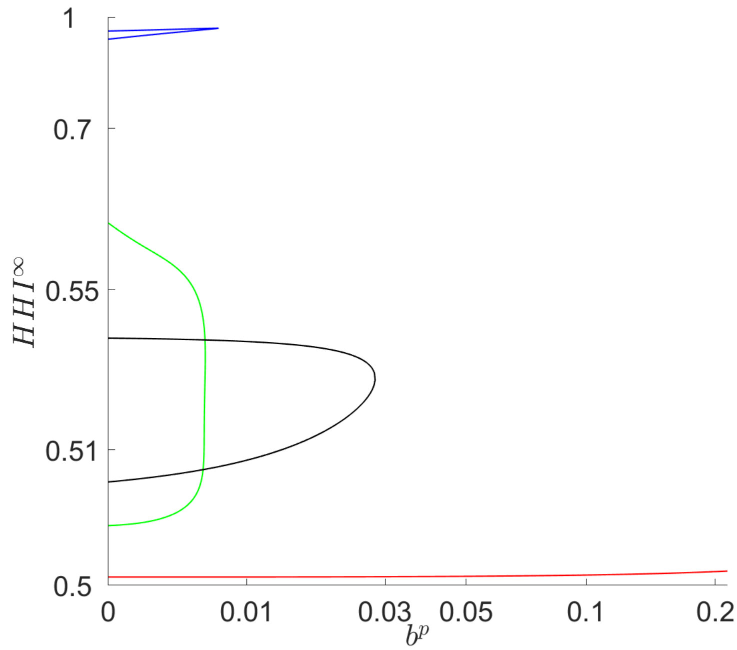

(b) $\delta = 0.028$, $\rho = 0.85$

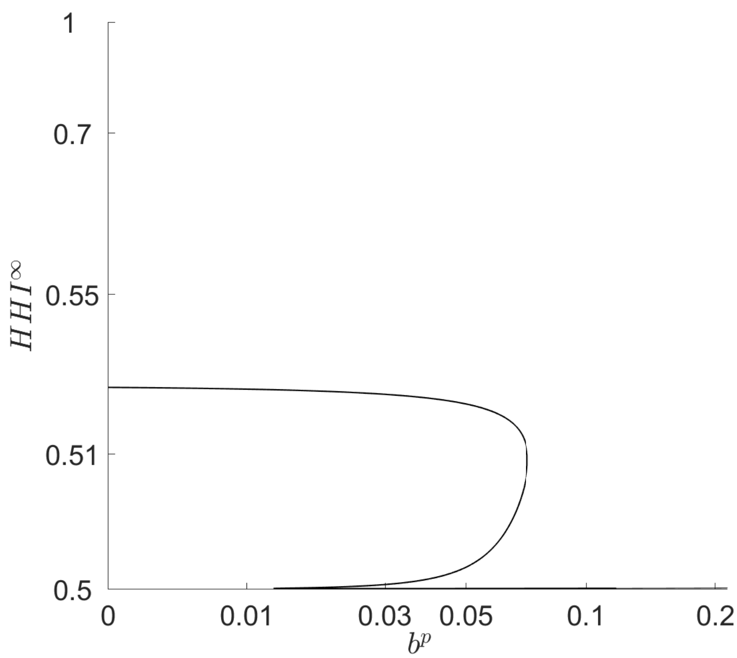

Figure C.6: Extended BDKS Model: Firm 1 Inverse Demand and Dynamic Incentives as a Function of b^p for $\delta = 0.0275$ and $\rho = 0.85$, Holding Seller Strategies Fixed at their $b^p = 0$ Values.

(a) Inverse Demand Curves For Seller 1 in State (3,1)

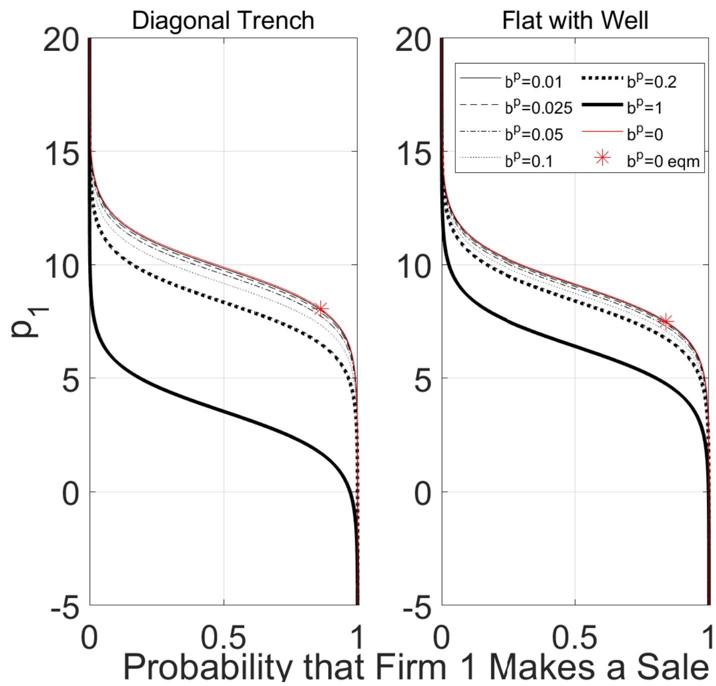

(b) Firm 1 Advantage-Building and Advantage-Denying Incentives in State (3,1): F =Flat with Well, DT = Diagonal Trench

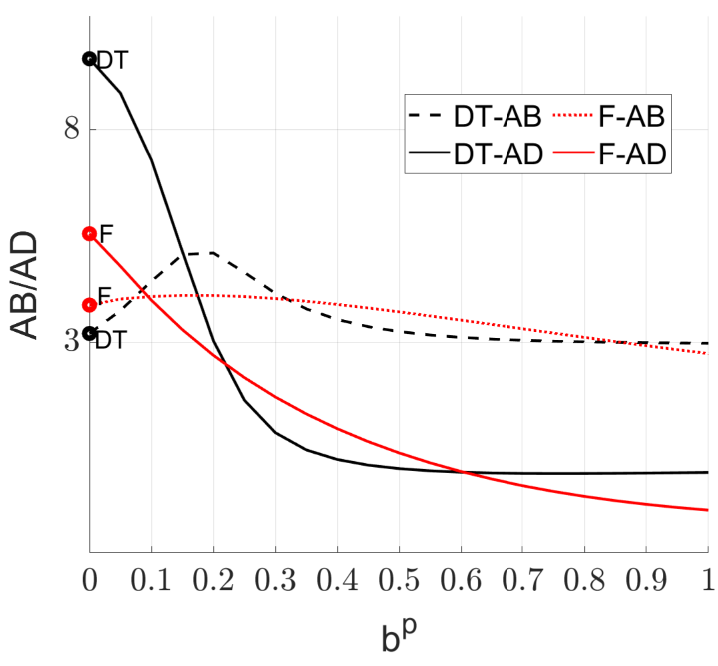

C.6 Extended BDKS Model: ρ-Homotopies with $\delta = 0.05$

Figures C.7 and C.8 show how expected long-run market concentration and average prices, and the NPV of consumer and total surplus, for a game starting at (1,1), vary with ρ and b^p, holding δ fixed at 0.05. The figures should be read with $\rho = 1$, at the left edge, corresponding to no LBD, with faster and more dramatic cost reductions from LBD as one moves to the right. $\delta = 0.05$ corresponds to a forgetting probability of 0.54 for $e_i = m = 15$, so, in expectation, it is not quite possible to sustain two firms at the bottom of their cost curves.

Consistent our results for the number of equilibria, there is multiplicity for low b^p for high and low values of ρ, but these are eliminated as b^p rises, so that there is a unique equilibrium for all ρ for $b^p \geq 0.1$. Consistent with the rest of our analyses, expected long-run concentration tends to be lower with higher b^p, and when $b^p = 1$ the firms are expected to be close to symmetric in the long-run even though there will be a continuous process where the firms gain and lose (forget) know-how. Concentration also tends to be low for very high and low ρ as, in both cases, marginal costs are close to their minimum possible values for low levels of know-how so that, even with forgetting, both firms can be expected to be close to the bottom of their cost curves, whereas for intermediate values of ρ only a firm that gains an advantage over its rival may be able to maintain low costs.

Figure C.7(b), which reports differences in prices from those when $b^p = 1$, shows that when ρ lies between 0.4 and 0.8, long-run prices are significantly higher when b^p is very low, reflecting the fact that the market is dominated by a single firm. On the other hand, for $b^p \geq 0.1$ the effect that increasing b^p tends to soften competition dominates, and prices increase. For more extreme values of ρ, where there can be multiple equilibria for low b^p, the prices in some equilibria are higher than those in the unique equilibrium when buyers are more strategic, and prices in other equilibria are lower.

Figure C.8 compare the NPV of total and consumer surplus as a function of ρ for different levels of b^p. In general, the differences in surplus across different values of b^p (they are measured relative to the unique outcome when $b^p = 1$) are small, but they indicate that even when more aggressive pricing results in more concentrated outcomes and higher long-run prices, the lower initial prices may raise the NPV of surplus.[46] Similar to a pattern in the BDK model (see Section 4.2), for very high ρ, strategic buyer behavior can slightly increase total surplus as they internalize how purchases can lower costs over a number of periods.

[46] As sales are always made in the BDKS model, the increase in total surplus reflects the reduction in production costs that comes from most sales being made by the lower cost firm.

Figure C.7: Extended BDKS Model: Equilibrium Long-Run Expected HHI and Average Prices with $\delta = 0.05$.

(a) Expected Long-Run HHI (HHI^{∞})

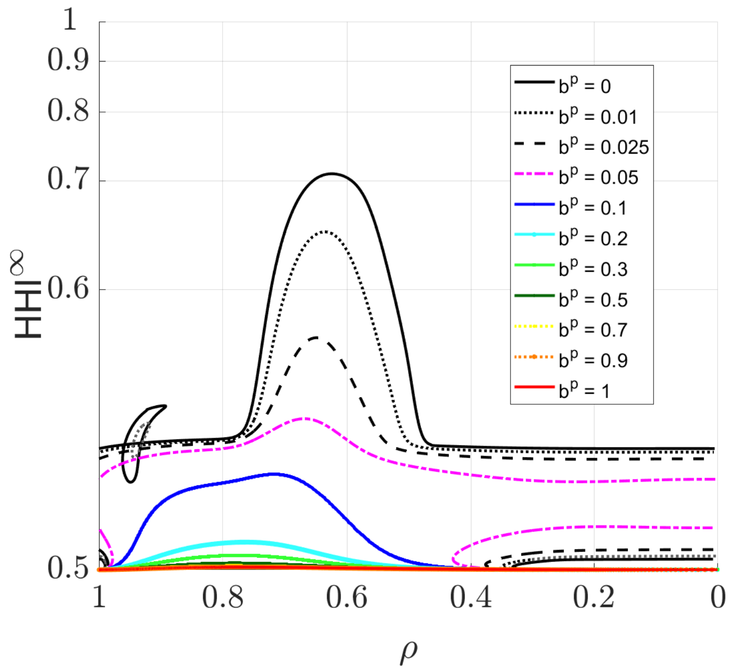

(b) Expected Long-Run Prices (Relative to Unique Outcome Where $b^p = 1$)

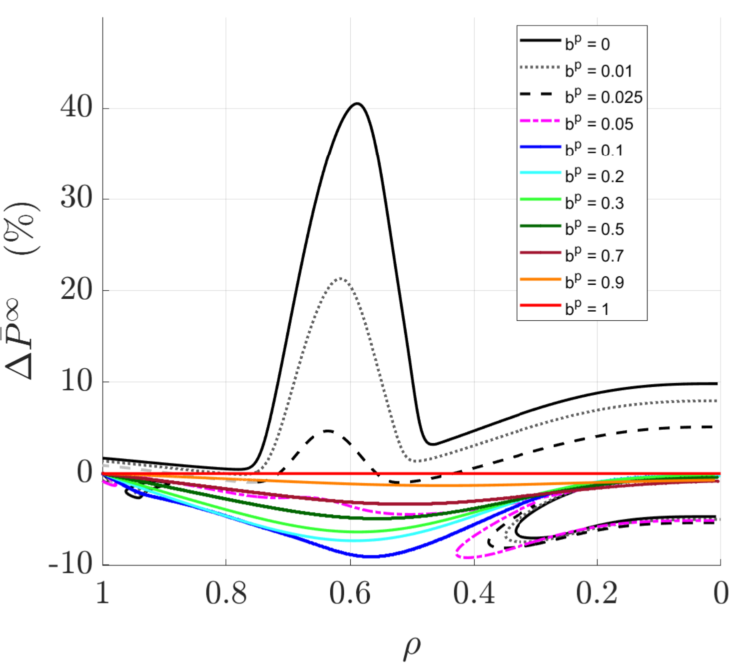

Figure C.8: Extended BDKS Model: Equilibrium NPV of Welfare with $\delta = 0.05$. Measured Relative to the Unique Outcome Where $b^p = 1$.

(a) Consumer Surplus

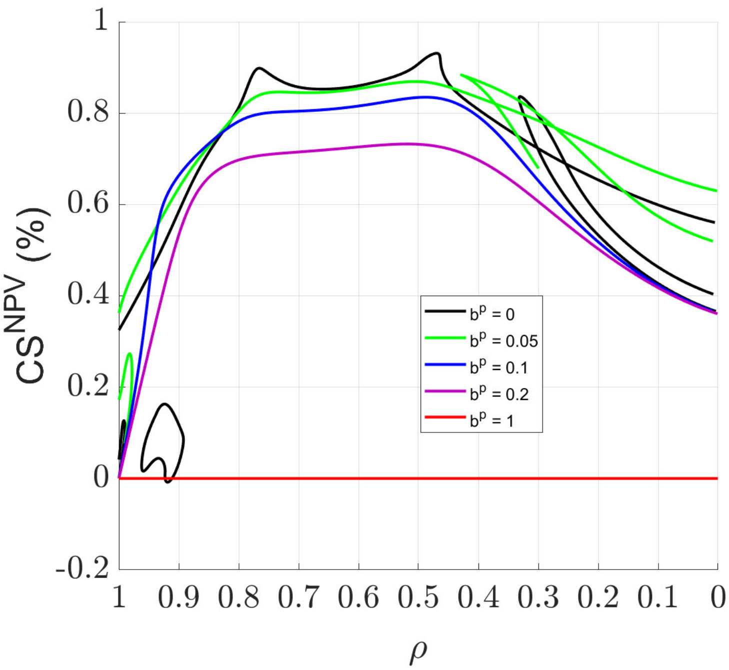

(b) Total Surplus