

Martin, Simon; Normann, Hans-Theo; Pülpichhuisen, Paul; Werner, Tobias

Working Paper

The Spoils of Algorithmic Collusion: Profit Allocation Among Asymmetric Firms

CESifo Working Paper, No. 11629

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Martin, Simon; Normann, Hans-Theo; Pülpichhuisen, Paul; Werner, Tobias (2025) : The Spoils of Algorithmic Collusion: Profit Allocation Among Asymmetric Firms, CESifo Working Paper, No. 11629, CESifo GmbH, Munich

This Version is available at:

<https://hdl.handle.net/10419/314668>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

The Spoils of Algorithmic Collusion: Profit Allocation Among Asymmetric Firms

Simon Martin, Hans-Theo Normann, Paul Püplichhuisen, Tobias Werner

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

The Spoils of Algorithmic Collusion: Profit Allocation Among Asymmetric Firms

Abstract

We study the propensity of independent algorithms to collude in repeated Cournot duopoly games. Specifically, we investigate the predictive power of different oligopoly and bargaining solutions regarding the effect of asymmetry between firms. We find that both consumers and firms can benefit from asymmetry. Algorithms produce more competitive outcomes when firms are symmetric, but less when they are very asymmetric. Although the static Nash equilibrium underestimates the effect on total quantity and overestimates the effect on profits, it delivers surprisingly accurate predictions in terms of total welfare. The best description of our results is provided by the equal relative gains solution. In particular, we find algorithms to agree on profits that are on or close to the Pareto frontier for all degrees of asymmetry. Our results suggest that the common belief that symmetric industries are more prone to collusion may no longer hold when algorithms increasingly drive managerial decisions.

JEL-Codes: C730, D430, L130.

Keywords: algorithmic collusion, Cournot duopoly, asymmetric firms.

Simon Martin
Department of Economics
University of Vienna / Austria
simon.martin@univie.ac.at

Hans-Theo Normann
Düsseldorf Institute for Competition
Economics (DICE), Heinrich-Heine-University
Düsseldorf / Germany
normann@hhu.de

Paul Püplichhuisen
Düsseldorf Institute for Competition
Economics (DICE), Heinrich-Heine-University
Düsseldorf / Germany
pueplichhuisen@dice.hhu.de

Tobias Werner
Center for Humans and Machines at the Max
Planck Institute for Human Development
Berlin / Germany
werner@mpib-berlin.mpg.de

January 13, 2025

Computational infrastructure and support were provided by the Centre for Information and Media Technology at Heinrich Heine University Düsseldorf.

1 Introduction

It is well known that firms find it in their interest to (tacitly) collude, but that asymmetries are an obstacle to the successful implementation of collusive schemes. While firms may suppress output or set higher prices to achieve higher profits than those achieved under competition, they often differ along several dimensions, such as their cost structure, capacities, logistics, production networks, sales force, or product quality. Indeed, asymmetries seem to be the rule rather than the exception¹, and the conventional wisdom is that such asymmetries make (tacit) collusion more difficult (Ivaldi et al., 2003).

The challenge of collusion in asymmetric industries is two-dimensional. On the one hand, firms must (tacitly) agree on total output, and on the other hand, they must decide how to share the spoils of such collusion. In the case of a symmetric industry, collusion is easier because firms have the same preferences regarding total output, and sharing profits equally seems an intuitive and focal solution to the bargaining problem. In contrast, when firms are asymmetric, they disagree about the collusive level of total output, and profit sharing may prove cumbersome: Joint payoff maximization may require the inefficient firms to shut down. However, if all firms produce positive quantities, the outcome is inefficient, and a convex bargaining frontier emerges (Bishop, 1960, Schmalensee, 1987, Tirole, 1988). Coordinating on the two dimensions without explicit communication is usually considered to be difficult.²

The existing literature provides some insights into how asymmetric firms might collude, but there are limits to how the bargaining problem can be addressed. The theoretical literature mainly answers how asymmetric firms collude *optimally* (Harrington, 1991, Miklos-Thal, 2011), with or without side payments. However, the conditions under which asymmetric firms may be expected to find these optimal solutions to the bargaining problem, or to reach any alternative bargaining outcome, remain largely unexplored. The empirical literature on how firms maintain collusion in asymmetric industries, and in particular how they divide the spoils of collusion without side payments³ suffers from the availability of appropriate data. Assessing the extent of tacit collusion is difficult without robust counterfactuals, and cartel studies may be biased because they are based on detected antitrust cases. Several laboratory experiments (Fonseca and Normann, 2008, Argenton and Müller, 2012, Harrington et al., 2016, Fischer and Normann, 2019) have addressed this issue. A key finding of these experiments is that tacit collusion fails in asymmetric settings, but they cannot explain what successful collusion might look like.

¹For collusion studies on this point, see Grout and Sonderegger (2005), Levenstein and Suslow (2006), and Davies et al. (2011).

²In oligopoly models with asymmetric capacities, firms usually do not disagree on total output (or price), but the bargaining problem persists. With differentiated price competition, joint payoff maximization would require side payments, but the inefficient firms do not necessarily have to shut down.

³Clark and Houde (2013) find that the Canadian gasoline cartel implemented transfers between firms by adjusting delays during price changes.

In this paper, we propose a novel approach to this gap in the literature. Specifically, we analyze how self-learning algorithms share the profits of potentially colluding asymmetric firms. Methodologically, we solve the two-dimensional collusion problem by conducting a series of simulations with Q-learning algorithms (Watkins, 1989, Watkins and Dayan, 1992).

How an algorithm solves the simultaneous coordination and bargaining problems among asymmetric firms, is a subject of growing importance. Firms’ decisions are increasingly being outsourced to algorithms across a variety of industries.⁴ Furthermore, even in situations where human decision-makers remain in control, they frequently receive algorithmic recommendations (see, for instance, Garcia et al., 2022, Huelden et al., 2024).

We run several different parameterizations of the repeated asymmetric Cournot duopoly. Starting from a symmetric baseline variant, we explore six different degrees of asymmetry between firms, and we consider two setups that differ in their parameterization. The main parameterization setup explores increasingly asymmetric firms while keeping the total output of the (static) Cournot-Nash equilibrium constant. In the second parameterization setup, the marginal cost parameter of the efficient firm is kept constant, so that monopoly output remains invariant in this setting. We run Q-learning simulations for these fourteen variants for different exploration and learning parameters, and for both algorithms with and without memory.

Our main findings are as follows. Somewhat surprisingly, both consumers and firms can benefit from asymmetry. As firms become more asymmetric, the more efficient firm produces additional output. Some of these efficiency gains are indirectly shared with consumers. However, the algorithms in our setting do not implement fully efficient outcomes, according to which the inefficient firm is entirely inactive. Relative to joint profit maximization, algorithms produce more competitive outcomes when firms are symmetric, but less when they are asymmetric. Intuitively, the extreme benchmark of one firm staying inactive would require some form of compensation, for example, through side payments, which cannot be captured in our environment.

On the other extreme, Nash equilibrium (which assumes absence of coordination) underestimates the effect of asymmetry on total quantity but overestimates the effect on profits. Nevertheless, the two effects are similar in magnitude, and hence Nash equilibrium delivers surprisingly accurate predictions in terms of total welfare.

Equal relative gains best describe our results, a solution criterion by Roth and Malouf (1979). The idea is that firms profit from collusion proportionally to their respective disagreement profits, that is, the profits they would obtain under competition. Intuitively, the firm that expects to make the largest profits under competition should experience a

⁴For instance, similar self-learning algorithms to the ones we consider are used in pricing (Liu et al., 2019, Chen et al., 2023), ride sharing optimization (Qin et al., 2022), or inventory management in e-commerce (Madeka et al., 2022).

proportional increase in profits under collusion. Although the Q-learning agents in our simulation tend to produce more, on average, than predicted by equal relative gains, this appears to be primarily a level effect. In particular, the comparative statics of increasing asymmetry are well captured by equal relative gains.

Additionally, our results show that the profit allocation for all degrees of asymmetry lies on the Pareto frontier. In contrast, the Nash equilibrium does not. We interpret this as an indication that the outcome that algorithms converge to is well explained by bargaining solutions such as equal relative gains.

Our results have important implications for competition policy and the regulation of digital markets. Namely, our findings suggest that the common belief that symmetric industries are more prone to collusion may no longer hold when algorithms increasingly drive managerial decisions (Leisten, 2021, Garcia et al., 2022, Hunold and Werner, 2023).

On top of providing us with the possibility of addressing long-standing questions regarding asymmetry in collusion, this setup also allows us to address current debates about algorithmic collusion (Ezrachi and Stucke, 2017, Harrington, 2018, 2022, OECD, 2023a,b) as supra-competitive prices are detrimental to consumers and economic welfare.⁵ Supporting the concerns, recent research by Calvano et al. (2020) and Klein (2021) reveals that self-learning algorithms can learn to collude in repeated pricing games. This not only concerns supra-competitive prices, but algorithms can also devise strategies that use punishments to discourage deviations from agreed-upon behavior. So, price algorithms may tacitly collude. This makes this a challenge for competition policy because these self-learning algorithms were never instructed to do so. Their work has been extended by Hettich (2016) and Johnson et al. (2023), among others. Notably, most of these papers study price competition (with a few exceptions, as elaborated on below).

The remainder of this paper is structured as follows. We provide an overview of how our work relates to other literature in the next section. We explain our model and experimental setup in Section 3 and Section 4, respectively. Our main results, mechanism, and several robustness checks are presented in Section 5. Section 6 concludes.

2 Related Literature

Relative to the existing theoretical literature on collusion of asymmetric firms (Harrington, 1991, Miklos-Thal, 2011), our work offers novel insights into how the bargaining problem of asymmetric firms is resolved. We show that outcomes tend to be well explained by the concept of equal relative gains, in particular, the effect of increasing asymmetry. Similarly, the experimental literature (Fonseca and Normann, 2008, Argenton and Müller, 2012,

⁵Pricing algorithms can also have a positive impact on economic welfare: Algorithms are well equipped to handle the vast amount of data available online about competitors and customers. They can adjust prices dynamically, support consistent pricing strategies, and respond immediately to changes in the market environment (OECD, 2023b). These efficiencies could ultimately benefit consumers.

Harrington et al., 2016, Fischer and Normann, 2019) typically only contrast symmetric with asymmetric settings, but does not vary the degree of asymmetry as pronounced as we are doing. We thus offer novel insights into the comparative statics of different degrees of asymmetry.

We also contribute to the literature on collusion of self-learning algorithms. Calvano et al. (2020) analyze Q-learning in markets with logistic demand, horizontal product differentiation, and simultaneous price competition. For a variation with asymmetric costs, the results imply that relative to the competitive Nash equilibrium, the less efficient firm gains more compared to the more efficient firm under price competition. Klein (2021) consider sequential price competition with symmetric Q-learning agents. They show that algorithms can learn to play supra-competitive price and punishment strategies that can rationalize the prices as a subgame perfect Nash equilibrium. Waltman and Kaymak (2008) consider competition in quantities but focus on memoryless algorithms, which limits their ability to learn collusive strategies. Similarly, Abada and Lambin (2023) analyze repeated Cournot games with symmetric costs for electricity storages using Q-learning algorithms. They find that algorithms were able to reach supra-competitive profits above the three-player Cournot equilibrium. While they consider agents that are asymmetric in terms of their storage capacity as a robustness check, they do not vary the degree of asymmetry in a systematic way. Also, Wang (2022) and Taywade et al. (2022) analyze, among other specifications, Cournot competition in duopolies with cost asymmetries, but without considering the degree of asymmetry. In our paper, we explicitly analyze how variations in the degree of asymmetry yield different market outcomes. Furthermore, we benchmark different theoretical predictions against each other to test which theory best predicts collusion among algorithms in this environment.

There is also increasing evidence that algorithmic adoption is linked to anti-competitive behavior from empirical studies. For instance, Assad et al. (2024) find that an increase in the adoption of pricing algorithms is linked to an increase in mark-ups in the German gasoline markets. Other empirical studies focus on e-commerce platforms (Hanspach et al., 2024, Musolff, 2022) and similarly find evidence of anti-competitive effects from algorithms. By focusing on a controlled environment in simulations, we can analyze the mechanisms behind those phenomena when firms are asymmetric in a structured way.

While there are several experimental studies that either compare algorithmic and human collusion (Kasberger et al., 2023, Werner, 2022) or study the interaction between algorithms and humans (Normann and Sternberg, 2023, Schauer and Schnurr, 2022), they do not consider Cournot competition.

Also in experimental studies without algorithms, the focus is often not on asymmetries among the firms. The closest related study is Fischer and Normann (2019), who conduct an experiment on Cournot competition with cost asymmetries. They find that participants could not sustain collusion in this asymmetric environment without communication, while

supra-competitive outcomes with a relative advantage for the less efficient firm could be achieved in the case of communication. In our paper, we follow a similar approach to Fischer and Normann (2019) and analyze how firms can sustain collusion, but we focus on Q-learning algorithms instead.

3 Model

We consider a standard Cournot duopoly, allowing for asymmetries in marginal costs. There are two firms $i \in \{L, H\}$ which simultaneously choose non-negative quantities q_L and q_H from the interval $[0, q^{max}]$. The firms only differ by their constant marginal production costs c_i , with $c_H \geq c_L \geq 0$. The products are homogeneous, and the inverse demand function is given by

$$p(Q) = \max \{a - bQ, 0\},$$

where a and b denote the intercept and slope parameters, respectively, and $Q = q_H + q_L$ denotes total quantity. Firm i 's profit is given by

$$\pi_i = p(Q)q_i - c_i q_i. \tag{1}$$

We will analyze producer surplus (PS), consumer surplus (CS), and total surplus (TS):

$$\begin{aligned} PS &= \pi_H + \pi_L \\ CS &= \frac{a - p(Q)}{2} Q = \frac{bQ^2}{2} \\ TS &= PS + CS \end{aligned}$$

Given this model, we consider possible market outcomes under different modes of competition in the following. These will serve as benchmarks, when discussing the results in Section 5.

Nash equilibrium. Taking first-order conditions in the profit expression (1) immediately yields the standard static unique Nash equilibrium of the game:

$$q_i^{NE} = \frac{a - 2c_i + c_j}{3b}$$

and equilibrium profits are $\pi_i^{NE} = b(q_i^{NE})^2$.

Joint Profit Maximization. For joint profit maximization, only the cost-efficient firm L is active and produces its monopoly quantity, whereas the inefficient firm H stays out

of the market. So, $i = L$ produces the monopoly quantity

$$q_i^M = \frac{a - c_i}{2b} \quad (2)$$

resulting in total industry profits $\Pi^M = \pi_L^M + 0 = b(q_L^M)^2$.

Alternating monopoly. For this benchmark, we assume that one firm becomes the monopolist with probability ω , and the other firm stays out of the market altogether. In our application, we assume that $\omega = 0.5$, that is, each of the firms is a monopolist with probability 50%, chooses q_i^M and earns π_i^M . For example, they may coordinate on being the monopolist alternately.

The previous three solutions are standard oligopoly benchmarks, but they may have infeasible profit allocation implications for colluding firms. The following bargaining solutions represent alternative predictions for the level and distribution of profits. Later in the paper, we sometimes refer to “oligopoly benchmarks” and “bargaining solutions” to distinguish the two sets of solutions.

Bargaining solutions assume that firms coordinate on the Pareto frontier. Each point on the Pareto frontier maximizes the profit that can be earned by one firm given a profit target for the other firm. To obtain the Pareto frontier, we rewrite the profit function as $q_i = \pi_i / (p - c_i)$, and then sum over both firms:

$$Q = \frac{\pi_i}{p - c_i} + \frac{\pi_j}{p - c_j}. \quad (3)$$

Solving for π_j with $Q = a - p$ provides the Pareto frontier

$$\bar{\pi}_j(\pi_i) = \max_p \left(a - p - \frac{\pi_i}{p - c_i} \right) (p - c_j). \quad (4)$$

That is, for any profit of firm i , a market price p can be found that maximizes the profit of firm j . Since costs are asymmetric and both firms produce positive quantities, the Pareto frontier is convex (Bishop, 1960, Schmalensee, 1987, Tirole, 1988). Although all bargaining concepts below assume that the solution lies on the Pareto frontier, they differ in the choice of a point on this frontier.

Kalai-Smorodinsky. Kalai and Smorodinsky (1975) assume a profit split relative to the maximum profit a firm can achieve (π_i^M) and the profit under disagreement (π_i^{dis}), that is:

$$\frac{\pi_i^{KS} - \pi_i^{dis}}{\pi_i^M - \pi_i^{dis}} = \frac{\pi_j^{KS} - \pi_i^{dis}}{\pi_j^M - \pi_i^{dis}}. \quad (5)$$

In this specification, the maximum profit achievable is the one-firm monopoly profit. As disagreement profits (π_i^{dis}), we consider (i) non-cooperative Nash profits (π_i^{NE}) and (ii) min-max disagreement profits ($\underline{\pi}_i$). The latter are defined as the maximum profit a firm can achieve given that the other firm produces the maximum quantity at the upper bound of the action space, q^{max} :

$$\underline{\pi}_i = \max_{q_i} \pi_i(q_i, q^{max}).$$

The Kalai-Smorodinsky solution to (5) is then implicitly defined by plugging (5) into the Pareto frontier expression in (4), and maximizing with respect to p .

Equal relative gains. This bargaining concept of Roth and Malouf (1979) also follows the idea of profit sharing relative to the point of disagreement. However, it ignores the maximum (monopoly) profits:

$$\frac{\pi_i^{ERG}}{\pi_i^{dis}} = \frac{\pi_j^{ERG}}{\pi_j^{dis}}.$$

As with the Kalai-Smorodinsky solution, we consider here as disagreement profits (i) Nash profits and (ii) min-max profits.

Equal split. An equal split solution according to Roth and Malouf (1979) implies that the bargaining point on the Pareto frontier ensures both firms receive the same payoff:

$$\pi_i^{ES} = \pi_j^{ES}.$$

4 Experimental setup

4.1 Main parameterization of the Cournot model

We study seven different parameterizations, starting with symmetric firms and increasing the level of cost asymmetry between the two firms. Our main parameterization is shown in Table 1.⁶ Throughout, we use an intercept $a = 91$ and a slope $b = 1$, which nests the cases studied in Fischer and Normann (2019). We start with a symmetric specification (Sym) in which both firms have marginal costs of $c = 19$, resulting in Nash quantities of $q_L^{NE} = q_H^{NE} = 24$ and hence total Nash output of $Q^{NE} = 48$. Under joint profit maximization, output would be limited to $Q^M = 36$.

In specifications Asym1 to Asym6, we gradually make firms more asymmetric by decreasing c_L and simultaneously increasing c_H so that total Nash quantities Q^{NE} remain unaffected. However, under Nash competition, firms' output quantities also become increasingly asymmetric. For example, in Asym6 almost all total output is produced by the

⁶As a robustness check, we also examine an alternative parameterization setup in Section 5.4, where the marginal cost of the efficient firm is held constant, ensuring that monopoly output stays unchanged as we increase the degree of asymmetry.

more efficient firm. Furthermore, total monopoly quantities Q^M increase as firms become more asymmetric, since the only active firm L has decreasingly lower marginal costs.

Table 1: Cost parameterization.

	Sym	Asym1	Asym2	Asym3	Asym4	Asym5	Asym6
c_L	19	16	13	10	7	4	1
c_H	19	22	25	28	31	34	37
q_L^{NE}	24	27	30	33	36	39	42
q_H^{NE}	24	21	18	15	12	9	6
Q^{NE}	48	48	48	48	48	48	48
Q^M	36	37.5	39	40.5	42	43.5	45

For each of our specifications, the solutions discussed in Section 3 have different predictions. We will investigate outcomes based on experiments with Q-learning algorithms.

4.2 Q-Learning

Similar to Calvano et al. (2020) and others, we utilize Q-learning algorithms. Each period, the Q-learning agents can choose from a finite and discrete action space $a_{i,t} \in A$. Further, $s_{i,t} \in S$ denotes the state of the environment in the period t as an element of the set of possible states. We assume that an agent i has bounded knowledge of the actions chosen by all other agents over the past k periods. Then a state $s_{i,t} = \{\mathbf{a}_{t-1}, \dots, \mathbf{a}_{t-k}\}$ contains the actions of all agents for the past k periods, where \mathbf{a}_t is a vector of the actions of all agents in period t .

Each agent aims to maximize the discounted expected future profits over the action set A , given the current state $s_{i,t}$. For a given period, the reward from picking a certain action in a certain state is denoted by $\pi_{i,t}$. For Q-learning algorithms, we can write

$$Q_{i,t}(s_{i,t}, a_{i,t}) = \mathbb{E}[\pi_{i,t} | s_{i,t}, a_{i,t}] + \delta \mathbb{E}[\max_a Q(s_{i,t+1}, a) | s_{i,t}, a_{i,t}] \quad (6)$$

Since the set of states and the set of actions are finite in our setup, the Q-function is a $|S| \times |A|$ matrix in which $Q_{i,t}(s_{i,t}, a_{i,t})$ is the current approximation of the stream of future discounted rewards associated with choosing action $a_{i,t}$ in state $s_{i,t}$.

When choosing action $a_{i,t}$ in state $s_{i,t}$ and receiving a reward $\pi_{i,t}$, the Q-matrix then gets updated as follows:

$$Q_{i,t+1}(s_{i,t}, a_{i,t}) = (1 - \alpha)Q_{i,t}(s_{i,t}, a_{i,t}) + \alpha \left(\pi_{i,t} + \delta \max_a Q_{i,t}(s_{i,t+1}, a) \right) \quad (7)$$

where $\alpha \in (0, 1)$ denotes the learning parameter. It governs how much weight the algorithm gives to the information it already has about the environment compared to newly arriving information in the current round.

When choosing an action a_t , an agent faces a trade-off between relying on past experiences (*exploitation*) versus exploring the action field for each state to find the optimal action-state combination (*exploration*). We incorporate experimentation in an ϵ -greedy process. That is, an agent chooses the current optimal action a_t^* for state s_t with probability $(1 - \epsilon_t)$. An agent chooses a random action with probability $\epsilon_t = e^{-\beta t}$, where $0 < \beta < 1$ is the exploration decay.

4.3 Methodology

We implement our study by having two independent Q-learning agents play a repeated game (with cost parameterization from Table 1) against each other. The same approach is taken, for example, in Calvano et al. (2020), Klein (2021), Johnson et al. (2023) and Kasberger et al. (2023).

Both agents use the same “technology”, that is, a learning rate α , experimentation with exponential decay at rate β , and a common discount factor $\delta = 0.95$. In our baseline specification with *moderate exploration*, we let $\alpha = 0.15$ and $\beta = 3.41 \times 10^{-6}$ as in Calvano et al. (2019).⁷

Both agents may condition their actions on the previous round’s actions, that is, the memory length $k = 1$. We discretize the action space in steps of three, that is, $q_i \in [0, 3, 6, \dots, 45]$, so $m = 16$, and we initialize the Q-matrix with random values from the uniform distribution on $[0, 1 \times 10^{-7}]$. We run the simulation until convergence, defined as not changing the best action for each state $s \in S$ over 100,000 subsequent periods. Upon convergence, we let the agents play 1,000 rounds. For each specification and parameterization, we repeat 1,000 simulations, and then consider the average across those simulations runs.

Given the results we obtain from our simulations, we compute a range of distance measures (goodness of fit) for several key outcome variables. Specifically, we are interested in the following four outcome variables $y \in \{Q, PS, CS, TS\}$, where Q denotes total quantity, CS the respective consumer surplus, $PS = \pi_L + \pi_H$ the producer surplus, and $TS = CS + PS$ the total surplus.

For each of these outcome variables, we consider the seven specifications $\sigma \in \Sigma = \{sym, asym1, asym2, asym3, asym4, asym5, asym6\}$. For each y and σ , there is a simulation outcome, y_σ^s , as well as eight theoretical bargaining solutions y_σ^b , for $b \in \{n, m, am, erg, es, ks, erg_n, ks_n\}$, denoting (static) Nash, monopoly, alternating monopoly, equal relative gains, equal split, Kalai-Smorodinsky, equal relative gains with Nash disagreement profit and Kalai-Smorodinsky with Nash disagreement profit, respectively.

⁷We chose a baseline β that provides the same expected times a cell would be visited by pure random exploration, that is, $\nu = ((m - 1)^n) / (m^{kn+n+1} \cdot (1 - e^{-\beta(n+1)})) \approx 21$, which yields $\beta = 3.41 \times 10^{-6}$. A slight adjustment is needed relative to Calvano et al. (2019) as our experimental setup involves $m = 16$ choices instead of $m = 15$ choices. We provide several robustness checks in these dimensions in Section 5.4.

For each y and b , we compute both an *average squared distance* to the simulation result (capturing the overall goodness of fit), given by

$$\overline{y^b} = \frac{1}{|\Sigma|} \sum_{\sigma \in \Sigma} (y_{\sigma}^s - y_{\sigma}^b)^2$$

as well as a *average squared normalized distance* to the simulation results (capturing the goodness of fit of the comparative statics), given by

$$\widehat{y^b} = \frac{1}{|\Sigma|} \sum_{\sigma \in \Sigma} (\widehat{y_{\sigma}^s} - \widehat{y_{\sigma}^b})^2$$

where $\widehat{y_{\sigma}^b}$ denotes normalized outcome variables to the symmetric case:

$$\widehat{y_{\sigma}^b} = \frac{y_{\sigma}^b}{y_{sym}^b}.$$

5 Results

5.1 Simulation results

This section provides an overview of our simulation results, focusing on the collusive potential of algorithms. We also examine the respective welfare implications.

Evidence of collusive outcomes. As shown in Figure 1, panel (a), total quantities across asymmetry specifications center around $Q \approx 40$. This falls between the theoretical benchmarks of static Nash level (48) and the collusive benchmark of alternating monopoly. The output produced by algorithms is restricted relative to the competitive benchmark. It provides evidence that the outcomes obtained by algorithms are collusive

We next investigate the effect of increasing asymmetry between firms. According to our simulation results, total quantity *increases* in asymmetry. However, for low levels of asymmetry, the simulated quantities are *above* the joint profit-maximizing solution (monopoly quantities), that is, the algorithms are *less collusive* than they could be. For high levels of asymmetry, in contrast, the simulated quantities are *below* the monopoly benchmark. In other words, algorithms produce more competitive outcomes (than suggested by models of collusion) when firms are symmetric but less collusive outcomes when they are very asymmetric. Importantly, this implies that asymmetric markets are relatively more collusive than symmetric ones.

While these results indicate that algorithms yield collusive outcomes, they do not explain the underlying mechanisms. We abstain from investigating these forces for the time being and continue analyzing outcomes for the moment. We explore these collusive mechanisms further in Section 5.3

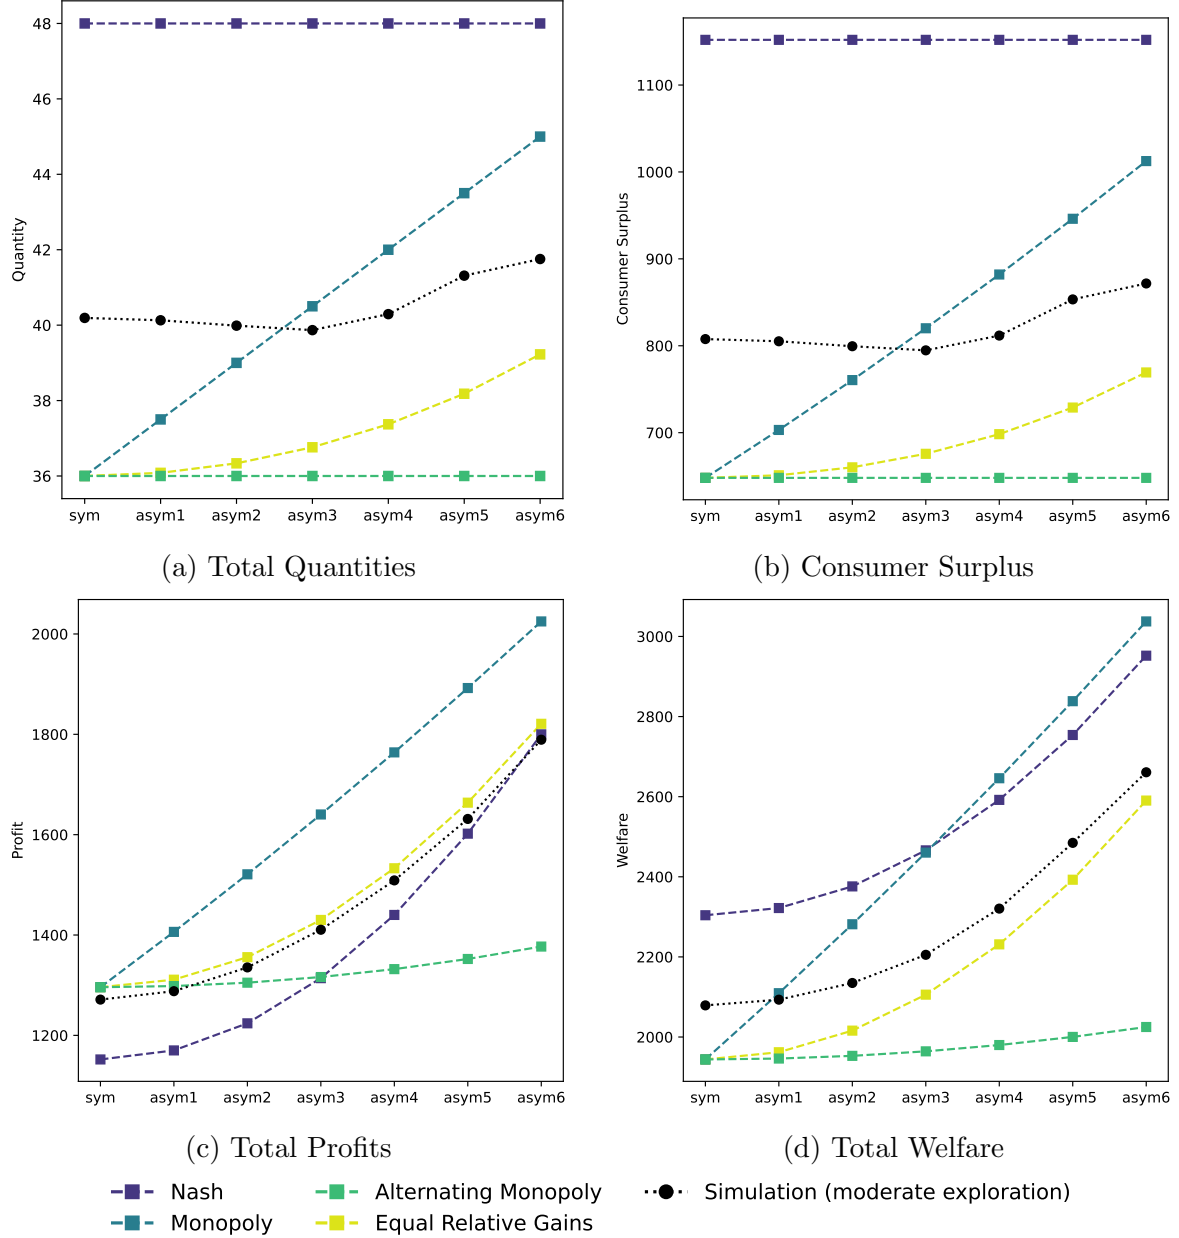


Figure 1: Simulation results for different degrees of asymmetry across different outcome variables. The figure compares the simulation results to four theoretical benchmarks: Nash equilibrium, monopoly, alternating monopoly, and equal relative gains. For comparisons with additional theoretical benchmarks, see Figure B.1 in Appendix B.

Welfare Effects. We next investigate the welfare effects of market asymmetry on the outcomes obtained by algorithms. In panel (b) of Figure 1, we depict the consumer surplus. In the homogeneous goods market we are investigating, consumer surplus is proportional, so the previous statements regarding total quantity translate directly into the effect on consumer surplus. From a consumer’s point of view, algorithms lead to outcomes that are not as favorable as those under (static Nash) competition, but better than those that would emerge from joint profit maximization.

The comparison of total profits, shown in panel (c), is less straightforward. In the presence of cost asymmetries, it not only matter which quantity is produced, but also which of the two firms produces which output. We find that total profits under simulated algorithms are smaller than under joint profit maximization (monopoly), but notably higher than those under competition. Thus, again, the outcome under simulated algorithms lies between the benchmarks of different modes of competition. For all cases, the monopoly solution constitutes the upper bound of what total profit could be achieved at maximum by the two firms as it implies only the low-cost firm provides positive quantities. The lower bound is the Nash solution for lower degrees of asymmetry and the alternating monopoly solution for higher degrees of asymmetry (*asym4*, *asym5*, *asym6*).

For the symmetric case, the solutions are in the interval [1152, 1296], where the lower bound corresponds to the Nash solution and the upper bound to the monopoly solution. With total profits of about 1275, our simulation results are close to the upper bound of the interval. In contrast, for the most asymmetric case, we observe total profits around 1800, which corresponds to the noncooperative Nash equilibrium. For all other cases in between, the algorithms achieve total profits well above the Nash equilibrium. In line with all theoretical benchmarks except equal split, total profits rise with the degree of asymmetry.

Given these results, the findings on total welfare, shown in panel (d), follow naturally. First, algorithms do not manage to appropriate welfare gains entirely. With strong cost asymmetry, total welfare would dictate the inefficient firm exiting the market. However, this outcome is unattainable without side payments, which are outside our model. Second, we find that total welfare can increase in cost asymmetry. Taken together, these results suggest that when algorithms take over managerial decisions, both consumers and firms can be better off in more asymmetric environments.

5.2 Explanatory power of oligopoly and bargaining concepts

We now explore to what extent the various concepts suggested in the literature can explain the outcomes achieved by the algorithms. To this end, we investigate which solution concept best describes the results using the methodology introduced in Section 4.3.

Level effects. We first investigate level effects, that is, which solution concept is closest to our simulated results in absolute terms. We may consider the relative distance of each concept to the simulation outcomes in Figure 1, separately for total quantities, consumer surplus, profits, and total welfare. Table 3 shows a systematic evaluation by computing squared distances.

As asymmetry increases, bargaining solutions align increasingly well with simulated output, particularly equal relative gains (for both disagreement profits). Interestingly, the simulation results in terms of quantity are closest, on average, to those predicted by a

monopoly (joint profit maximization). However, this averaging masks an important tendency: The monopoly prediction rises steeper in asymmetry than our simulation results. That is, the monopoly solution provides predictions significantly below the actual level for the symmetric and low asymmetry cases. However, its predictions rise well above the simulation results for higher degrees of asymmetry.

The equal split solution performs just as good or bad as the other bargaining solutions for the symmetric case. However, it performs increasingly worse for higher degrees of asymmetry. This highlights how an equal profit split might be a feasible solution when players are symmetric but is rather unlikely in the presence of (cost) asymmetries.

Total profits are predicted spot on by the equal relative gains solutions and Kalai-Smorodinsky with Nash disagreement profits for all degrees of asymmetry. Consequently, the average squared distance in Table 3 is small. In particular, they perform remarkably better than predicted by the oligopoly benchmarks. Equal relative gains and Kalai-Smorodinsky also provide relatively good predictions for total welfare.

Overall, concepts implicitly embedding a bargaining idea such as equal relative gains and Kalai-Smorodinsky provide a strong explanation for the outcomes generated by algorithms, despite underestimating total quantity. Similarly, while the static Nash equilibrium underestimates the impact of asymmetry on total quantity, it overestimates the effect on profits. These opposing biases nearly offset each other, resulting in relatively accurate predictions for total welfare

Table 3: Average Squared Distances from Simulation Results (moderate exploration)

	Q	CS	Π	W
Nash	57	110,618	7,945	65,716
Monopoly	6	10,208	41,506	68,265
Alternating Monopoly	21	30,526	41,379	126,583
Kalai-Smorodinsky	34	48,309	40,517	163,873
Equal Split	58	76,802	145,750	415,075
Equal Relative Gains	12	17,384	654	11,558
Kalai-Smorodinsky (Nash)	12	17,134	823	11,355
Equal Relative Gains (Nash)	9	13,241	4,428	7,914

Comparative statics. We are particularly interested in the differential effect of increasing asymmetry on various market outcomes. Therefore, we now abstract from different levels across theories by considering outcomes relative to the benchmark of a symmetric specification. This allows us to isolate the comparative statics effect of changes in asymmetry. The results are shown in Figure 2, with a corresponding average squared distance calculation in Table A.1 in the Appendix.

Note that we designed the simulations so that the total quantities under Nash remain constant across all degrees of (a)symmetry, resulting in a horizontal line at value 1 in Fig-

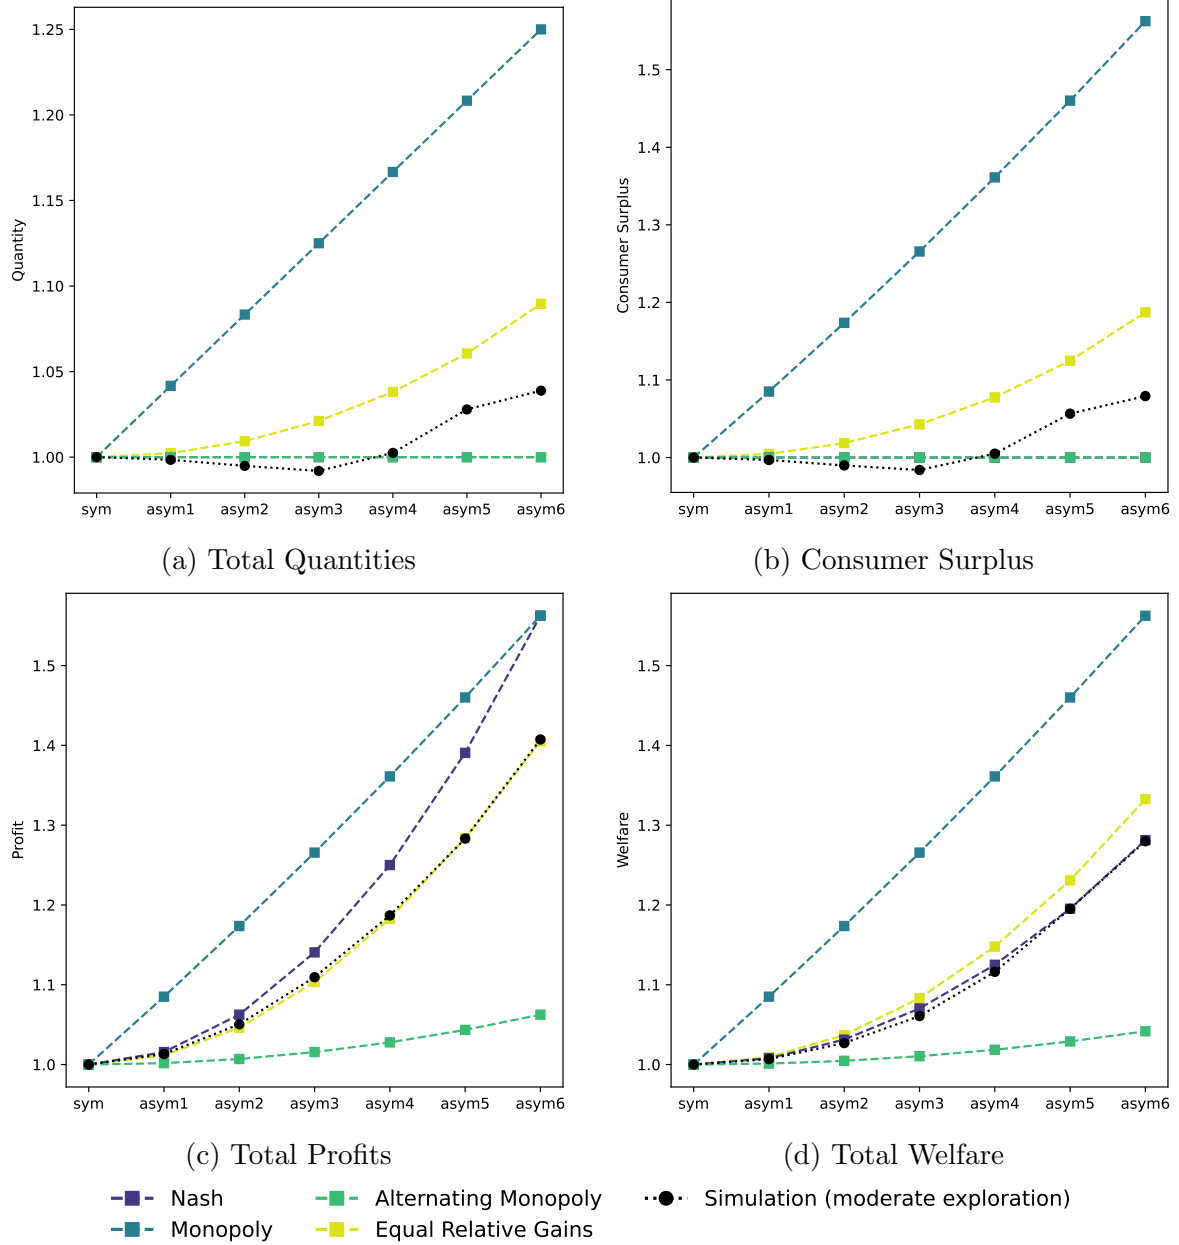


Figure 2: Simulation results (comparative statics; symmetric case = 1) for different degrees of asymmetry across various outcome variables. The figure compares the simulation results to four theoretical benchmarks: Nash equilibrium, monopoly, alternating monopoly, and equal relative gains. For comparisons with additional theoretical benchmarks, see Figure B.2 in Appendix B.

ure 2, panel (a), for the Nash equilibrium. The same is true for the alternating monopoly. Again, the trajectory of the simulation results is predicted best by equal relative gains and Kalai-Smorodinsky (with Nash deviation profits). Thus, these bargaining solutions do not only provide a best fit in terms of the overall (level) prediction accuracy but also a reasonably good fit for the relative change.

For total profits, equal relative gains and Kalai-Smorodinsky perfectly capture the

comparative statics of asymmetry perfectly. Similarly, the effect of asymmetry on total welfare is captured extremely well by the Nash prediction, as well as the bargaining concepts of equal relative gains and Kalai-Smorodinsky (with Nash). However, the good fit is masked by slight over- and underestimation of consumer surplus and total profits, respectively. In particular, the Nash prediction underestimates the effect on consumer surplus and overestimates the impact on total profits, while it is the opposite for the bargaining solutions.

Taken together, we conclude that the best explanation for our results is provided by equal relative gains, and bargaining concepts more generally. Algorithms seem to “agree” on profits that are on or close to the Pareto frontier for all degrees of asymmetry.

5.3 Mechanisms

A central problem for collusion of asymmetric firms concerns profit allocation. In contrast to collusion with symmetric firms where equal profit sharing is a natural candidate, profits of asymmetric firms under competition differ, making it unclear how profits should be split under collusion. This issue ultimately boils down to resolving an upfront bargaining problem over the spoils of collusion.

In Figure 3, we shed light on how algorithms solve the allocation bargaining problem. Panel (a) shows the outcomes for our main specification, where algorithms have a memory length of $k = 1$. The x-axis depicts the profits of the more efficient low-cost firm (Π_L) while the y-axis represents the profits of the inefficient high-cost firm (Π_H).

The solid and dotted lines show the Pareto frontiers of profits for different degrees of asymmetry. The steepest solid line corresponds to the Pareto frontier of the symmetric case, while the solid line with the flattest slope corresponds to the most asymmetric case (*asym6*). The dotted lines depict the Pareto frontiers of *asym1* to *asym5*. Additionally, the figure depicts the profits obtained in our simulations (black dots), under equal relative gains (yellow squares), Nash equilibrium (purple squares), and min-max-disagreement (grey circle) for different degrees of asymmetry (0 = symmetry, 6 = most asymmetric).

For all degrees of asymmetry, our results lie close to the Pareto frontier, with profits significantly exceeding those of the Nash equilibrium for both firms. Profits achieved by the algorithms are more than double those under min-max-disagreement. What is particularly striking is that the profits are close to equal relative gains for all degrees of asymmetry, confirming our results from Section 5.2. This suggests that the idea underlying equal relative gains are a good description of how algorithms reconcile the allocation problem: The spoils of collusion are divided such that both firms gain from collusion equally much *in relative terms*, and, importantly, such that the total profits are close to the Pareto frontier.

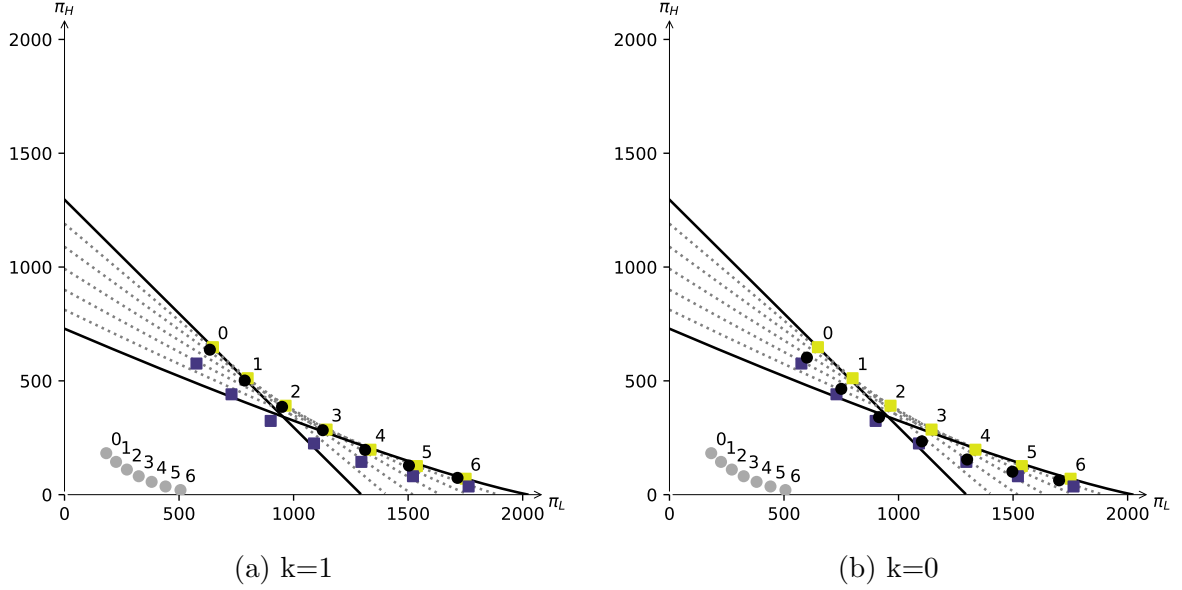


Figure 3: Pareto frontier and firm profits for our simulation results (black circles), equal relative gains (yellow squares), Nash equilibrium (purple squares), and min-max-disagreement (grey circles).

Algorithmic sophistication. A common concern in the analysis of self-learning algorithms regards the question of algorithmic sophistication, especially when algorithms obtain supra-competitive outcomes as in our setting (see, for instance, Calvano et al., 2023, Epivent and Lambin, 2024). Do the algorithms obtain high profits either because they are sufficiently sophisticated to develop strategies that support these outcomes, or are these outcomes rather a consequence precisely of the *lack of sophistication*, leaving gains of exploiting an innocent cooperator on the table?

To address these concerns, we provide two pieces of evidence. First, we investigate the outcomes that the algorithms would obtain if we do not endow them with any memory ($k = 0$). This mechanically prevents the emergence of any reward-and-punishment schemes typically encountered in collusive strategies. Not allowing the firms to condition their outcomes on previous periods' actions effectively renders the game static, so that the Nash outcomes become the unique (subgame-perfect) equilibrium of the game.

The results of these simulations are shown in panel (b) of Figure 3. For all degrees of asymmetry ($0 = \text{symmetry}$, $6 = \text{most asymmetric}$), our simulation results (black circles) are very close to the Nash equilibrium outcomes (grey circles) and well below the Pareto frontier. This is in sharp contrast to the results we obtained when we endowed the agents with memory ($k = 1$ in panel (a)).

To further investigate algorithmic sophistication, we provide a second piece of evidence by examining the incentives to deviate from the strategy prescribed by the Q-matrix upon convergence. Following Calvano et al. (2020), for each simulation, agent i , and parameterization of the game, we compute the static best response to the strategy of

the other player. Note that this computation is beyond the knowledge available to the algorithms and hence accessible only to us, setting the bar for algorithmic sophistication very high. It allows the deviating agent to gain the maximum one-period deviation profit.⁸

To evaluate the profitability, for both approaches, we simulate that agent i deviates to the respective action once. Subsequently, we simulate outcomes for 40 periods using the Q-matrices of both agents upon convergence.⁹ We then compute the fraction of simulations in which such a one-shot deviation was profitable. We consider a deviation profitable when the sum of discounted profits is higher under deviation compared to no deviation. For discounting, we consider the deviation period as $t = 0$ and $\delta = 0.95$ as in the initial Q-learning simulation.




Figure 4: Incentives to deviate based on one-period best response (moderate exploration)

The results are shown in Figure 4. For all degrees of asymmetry, in more than 60% of cases, none of the agents has any incentive to deviate. Only rarely do both agents want to deviate, but in this specification, the cost-efficient c_L firm might want to deviate in around 15% of the simulations.

Taken together, we interpret the two separate considerations of memory and deviation as very suggestive of the fact that collusive outcomes observed in our setting are indeed the result of sophisticated algorithms that learned to operate effectively in their environment.

⁸As an alternative deviation analysis, we let the agent choose a strictly higher quantity with the highest Q-value as their deviation action. The results of this robustness check are provided in Figure B.4 in the Appendix and are overall similar.

⁹As shown by Calvano et al. (2020), following a deviation, after a certain number of periods the algorithms usually return to playing the actions upon convergence. Thus, by choosing 40 simulation periods, we ensure that the deviation circle is completed.

5.4 Robustness checks

Alternative cost asymmetry. The cost parameterizations of our main specification generally only depicts one of infinitely many possible cost asymmetries. Although chosen thoughtfully by keeping constant total Nash quantities, we provide an alternative cost parameterization in Appendix C. Under the alternative cost parameterization, the one-firm monopoly output is kept constant (that is, $c_L = 19$) while we increase the costs of the high-cost firm c_H . For this exercise, we focus on moderate exploration ($\nu = 21$). Qualitatively, the results are similar to our main specification. Firms reduce total quantities well below the Nash prediction and the profit split lies on the Pareto frontier for all cases. The profit split aligns closely with the equal relative gains prediction using min-max deviation profits but exhibits some deviations for high degrees of asymmetry. As in our main specification, Nash and equal relative gains (particularly with Nash deviation profits) effectively capture the relative changes compared to the symmetric case.

Different α, β parameters. In our main specification, we consider moderate exploration such that the expected times of random visits per cell in the Q-matrix is $\nu = 21$. We assumed an $\alpha = 0.15$. As a robustness test, we ran simulations with high exploration ($\nu = 100$), leading to qualitatively similar results, as shown in Figures B.1, B.2, B.3, and B.4 in the Appendix. For simulation results with high exploration, we only observe a slightly stronger increase of total quantities for higher degrees of asymmetry and an even lower incentive to deviate compared to moderate exploration. As a robustness check, we further investigate the outcomes when we set the learning parameter $\alpha = 0.2$ instead for moderate and high exploration. Our results remain robust to this adjustment as shown in Figure B.5 and B.6 in the Appendix.

Algorithmic sophistication. Our deviation analysis in Figure 4 shows an increase in the share of simulation runs where the deviation was not profitable. Given that we also observe an increase in total quantities for higher degrees of asymmetry, one might suspect the share of profitable deviations to be the main driver of this effect. To address this, we evaluated total quantities separately for two subsamples: (i) incentive-compatible simulations and (ii) simulations where at least one firm could profitably deviate. The robustness check reveals the same trend for both subsamples, with average total quantities in incentive-compatible simulations approximately 0.5 units higher than in simulations where at least one firm could profitably deviate, see Figure B.7.

6 Conclusion

Sustaining collusion in asymmetric industries poses a significant challenge, as firms must tacitly agree on both total output and profit allocation. While an equal profit split is an

intuitive focal point in symmetric markets, asymmetries introduce complexities, raising questions about whether collusion can be sustained and how firms allocate profits.

This paper complements the existing theoretical, empirical, and laboratory literature on collusion in asymmetric markets by conducting several simulation experiments with self-learning algorithms. In particular, we let the self-learning algorithms compete in a Cournot duopoly and gradually increase the degree of cost asymmetry while keeping the total (static) Nash quantity constant. The approach allows us to benchmark the different degrees of asymmetry against a robust counterfactual to illustrate what a collusive outcome might look like.

Our results reveal that algorithms consistently reduce total output below the static Nash equilibrium across all degrees of asymmetry, indicating tacit collusion. For a higher degree of asymmetry, the results imply an increase in total quantities, and thus consumer surplus. While total profits in our simulation exceed total Nash profits for lower degrees of asymmetry, the effect diminishes for higher degrees of asymmetry. Nevertheless, this does not necessarily imply a lower degree of (tacit) collusion, as firm profits for all degrees of asymmetry lie on the Pareto frontier. Thus, the firm algorithms robustly established cooperation.

We evaluate the simulation results against several different theoretical predictions. The concept of equal relative gains is the bargaining idea that best captures the effect of asymmetry and overall results. Firms allocate cooperative profits in *relative* terms to a disagreement point.

Our findings have important implications for competition policy. Typically, symmetry among firms is viewed as potentially facilitating collusion and, *ceteris paribus*, demanding the competition authority's closest attention. However, according to our results, the opposite may hold in industries where decisions are increasingly delegated to autonomous algorithms: precisely *asymmetric* firms might be more prone to dampening competition and yielding inefficient market outcomes.

References

- Abada, Ibrahim and Xavier Lambin**, “Artificial intelligence: Can seemingly collusive outcomes be avoided?,” *Management Science*, 2023.
- Argenton, Cédric and Wieland Müller**, “Collusion in experimental Bertrand duopolies with convex costs: The role of cost asymmetry,” *International Journal of Industrial Organization*, 2012, 30 (6), 508–517.
- Assad, Stephanie, Robert Clark, Daniel Ershov, and Lei Xu**, “Algorithmic pricing and competition: Empirical evidence from the German retail gasoline market,” *Journal of Political Economy*, 2024, 132 (3), 723–771.
- Bishop, Robert L**, “Duopoly: Collusion or Warfare?,” *The American Economic Review*, 1960, 50 (5), 933–961.
- Calvano, Emilio, Giacomo Calzolari, Vincenzo Denicolò, and Sergio Pastorello**, “Algorithmic pricing what implications for competition policy?,” *Review of industrial organization*, 2019, 55 (1), 155–171.
- , – , **Vincenzo Denicolo, and Sergio Pastorello**, “Artificial intelligence, algorithmic pricing, and collusion,” *American Economic Review*, 2020, 110 (10), 3267–3297.
- , – , **Vincenzo Denicolò, and Sergio Pastorello**, “Algorithmic collusion: Genuine or spurious?,” *International Journal of Industrial Organization*, 2023, 90, 102973.
- Chen, Ji, Yifan Xu, Peiwen Yu, and Jun Zhang**, “A reinforcement learning approach for hotel revenue management with evidence from field experiments,” *Journal of Operations Management*, 2023, 69 (7), 1176–1201.
- Clark, Robert and Jean-François Houde**, “Collusion with asymmetric retailers: Evidence from a gasoline price-fixing case,” *American Economic Journal: Microeconomics*, 2013, 5 (3), 97–123.
- Davies, Stephen, Matthew Olczak, and Heather Coles**, “Tacit collusion, firm asymmetries and numbers: evidence from EC merger cases,” *International Journal of Industrial Organization*, 2011, 29 (2), 221–231.
- Epivent, Andréa and Xavier Lambin**, “On algorithmic collusion and reward–punishment schemes,” *Economics Letters*, 2024, 237, 111661.
- Ezrachi, Ariel and Maurice E Stucke**, “Artificial intelligence & collusion: When computers inhibit competition,” *U. Ill. L. Rev.*, 2017, p. 1775.
- Fischer, Christian and Hans-Theo Normann**, “Collusion and bargaining in asymmetric Cournot duopoly—An experiment,” *European Economic Review*, 2019, 111, 360–379.
- Fonseca, Miguel A and Hans-Theo Normann**, “Mergers, asymmetries and collusion: Experimental evidence,” *The Economic Journal*, 2008, 118 (527), 387–400.
- Garcia, Daniel, Juha Tolvanen, and Alexander K Wagner**, “Demand estimation using managerial responses to automated price recommendations,” *Management Science*, 2022, 68 (11), 7918–7939.

- Grout, Paul and Silvia Sonderegger**, *Predicting cartels*, Office of Fair Trading United Kingdom, 2005.
- Hanspach, Philip, Geza Sapi, and Marcel Wieting**, “Algorithms in the Market-place: An Empirical Analysis of Automated Pricing in E-Commerce,” *Information Economics and Policy*, 2024, p. 101111.
- Harrington, Joseph E**, “The determination of price and output quotas in a heterogeneous cartel,” *International Economic Review*, 1991, pp. 767–792.
- , “Developing competition law for collusion by autonomous artificial agents,” *Journal of Competition Law & Economics*, 2018, 14 (3), 331–363.
- , “The effect of outsourcing pricing algorithms on market competition,” *Management Science*, 2022, 68 (9), 6889–6906.
- , **Roberto Hernan Gonzalez, and Praveen Kujal**, “The relative efficacy of price announcements and express communication for collusion: Experimental findings,” *Journal of Economic Behavior & Organization*, 2016, 128, 251–264.
- Hettich, Matthias**, “Algorithmic Pricing and Tacit Collusion,” 2016.
- Huelden, Tobias, Vitalijs Jascisens, Lars Roemheld, and Tobias Werner**, “Human-machine interactions in pricing: Evidence from two large-scale field experiments,” *Available at SSRN 4763132*, 2024.
- Hunold, Matthias and Tobias Werner**, “Algorithmic price recommendations and collusion: Experimental evidence,” *Available at SSRN*, 2023.
- Ivaldi, Marc, Bruno Jullien, Patrick Rey, Paul Seabright, and Jean Tirole**, “The economics of tacit collusion,” 2003.
- Johnson, Justin P, Andrew Rhodes, and Matthijs Wildenbeest**, “Platform design when sellers use pricing algorithms,” *Econometrica*, 2023, 91 (5), 1841–1879.
- Kalai, Ehud and Meir Smorodinsky**, “Other solutions to Nash’s bargaining problem,” *Econometrica: Journal of the Econometric Society*, 1975, pp. 513–518.
- Kasberger, Bernhard, Simon Martin, Hans-Theo Normann, and Tobias Werner**, “Algorithmic Cooperation,” *Available at SSRN 4389647*, 2023.
- Klein, Timo**, “Autonomous algorithmic collusion: Q-learning under sequential pricing,” *The RAND Journal of Economics*, 2021, 52 (3), 538–558.
- Leisten, Matthew**, “Algorithmic Competition, with Humans,” Technical Report, working paper, 1 September 2021.
- Levenstein, Margaret C and Valerie Y Suslow**, “What determines cartel success?,” *Journal of economic literature*, 2006, 44 (1), 43–95.
- Liu, Jiayi, Yidong Zhang, Xiaoqing Wang, Yuming Deng, and Xingyu Wu**, “Dynamic pricing on e-commerce platform with deep reinforcement learning: A field experiment,” *arXiv preprint arXiv:1912.02572*, 2019.

- Madeka, Dhruv, Kari Torkkola, Carson Eisenach, Anna Luo, Dean P Foster, and Sham M Kakade**, “Deep inventory management,” *arXiv preprint arXiv:2210.03137*, 2022.
- Miklos-Thal, Jeanine**, “Optimal collusion under cost asymmetry,” *Economic Theory*, 2011, *46*, 99–125.
- Musolff, Leon**, “Algorithmic pricing facilitates tacit collusion: Evidence from e-commerce,” in “Proceedings of the 23rd ACM Conference on Economics and Computation” 2022, pp. 32–33.
- Normann, Hans-Theo and Martin Sternberg**, “Human-algorithm interaction: Algorithmic pricing in hybrid laboratory markets,” *European Economic Review*, 2023, *152*, 104347.
- OECD**, “Algorithmic competition – Note by Italy,” 2023.
- , “Algorithmic competition – OECD Competition Policy Roundtable Background Note,” 2023.
- Qin, Zhiwei Tony, Hongtu Zhu, and Jieping Ye**, “Reinforcement learning for ridesharing: An extended survey,” *Transportation Research Part C: Emerging Technologies*, 2022, *144*, 103852.
- Roth, Alvin E and Michael W Malouf**, “Game-theoretic models and the role of information in bargaining,” *Psychological review*, 1979, *86* (6), 574.
- Schauer, Andreas and Daniel Schnurr**, “Competition Between Human and Artificial Intelligence in Digital Markets: An Experimental Analysis,” *Available at SSRN 4139527*, 2022.
- Schmalensee, Richard**, “Competitive Advantage and Collusive Optima,” *International Journal of Industrial Organization*, 1987, *5* (4), 351–367.
- Taywade, Kshitija, Brent Harrison, and Judy Goldsmith**, “Using Non-Stationary Bandits for Learning in Repeated Cournot Games with Non-Stationary Demand,” *arXiv preprint arXiv:2201.00486*, 2022.
- Tirole, Jean**, *The Theory of Industrial Organization*, MIT Press, 1988.
- Waltman, Ludo and Uzay Kaymak**, “Q-learning agents in a Cournot oligopoly model,” *Journal of Economic Dynamics and Control*, 2008, *32* (10), 3275–3293.
- Wang, Yanwen**, “Will Algorithms Learn to Collude?,” 2022.
- Watkins, Christopher JCH and Peter Dayan**, “Q-learning,” *Machine learning*, 1992, *8*, 279–292.
- Watkins, Christopher John Cornish Hellaby**, “Learning from delayed rewards,” 1989.
- Werner, Tobias**, “Algorithmic and human collusion,” *Available at SSRN 3960738*, 2022.

Appendix

A Additional Tables

Table A.1: Average squared normalized distances (multiplied by 1,000) from simulation results (moderate exploration)

	Q	CS	Π	W
Nash	0.34	1.41	5.81	0.03
Monopoly	18.78	92.02	18.63	39.95
Alternating Monopoly	0.34	1.41	30.38	13.86
Kalai-Smorodinsky	3.44	13.21	29.80	20.60
Equal Split	12.30	43.38	97.45	71.79
Equal Relative Gains	0.86	3.71	0.01	0.81
Kalai-Smorodinsky (Nash)	0.98	4.31	0.04	1.05
Equal Relative Gains (Nash)	2.71	12.28	1.03	4.59

B Additional Figures

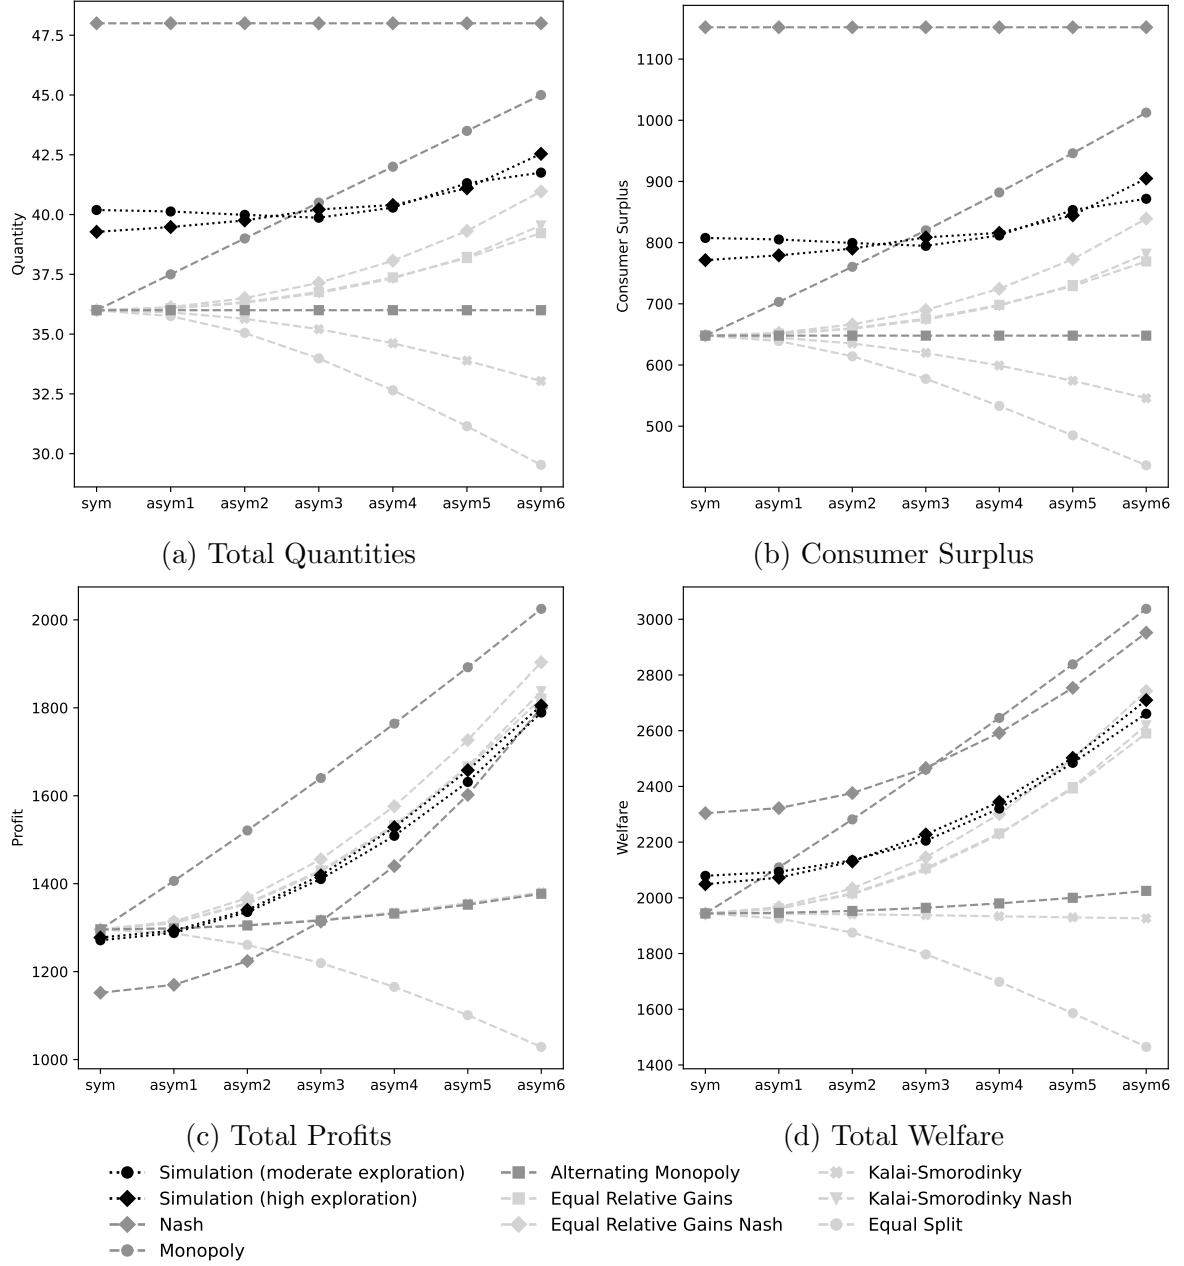





Figure B.1: Simulation results for different degrees of asymmetry (main results including high exploration; all theoretical predictions).




(a) Total Quantities



(b) Consumer Surplus



(c) Total Profits



(d) Total Welfare







Figure B.2: Simulation results (comparative statics; symmetric case = 1) for different degrees of asymmetry.




(a) Total Quantities



(b) Consumer Surplus



(c) Total Profits



(d) Total Welfare







Figure B.3: Simulation results (levels) for different degrees of asymmetry without memory.




(a) Best response (moderate exploration)



(b) Best response (high exploration)




(c) Q-Value (moderate exploration)




(d) Q-Value (high exploration)


Figure B.4: Incentives to deviate




(a) Total Quantities



(b) Consumer Surplus



(c) Total Profits



(d) Total Welfare




Figure B.5: Simulation results for different degrees of asymmetry ($\alpha = 0.2$ including high exploration; all theoretical predictions).




Figure B.6: Pareto frontier and simulation results for $\alpha = 0.2$ and $\nu = 21$.




Figure B.7: Total quantities for main specification ($\nu = 21$, $\alpha = 0.15$) overall and sub-samples divided based on whether deviation is profitable.

C Robustness Check: Alternative Cost Asymmetry

Table C.1: Alternative cost parameterization

	Sym	Asym1	Asym2	Asym3	Asym4	Asym5	Asym6
c_L	19	19	19	19	19	19	19
c_H	19	22	25	28	31	34	37
q_L^{NE}	24	25	26	27	28	29	30
q_H^{NE}	24	22	20	18	16	14	12
Q^{NE}	48	47	46	45	44	43	42
Q^M	36	36	36	36	36	36	36

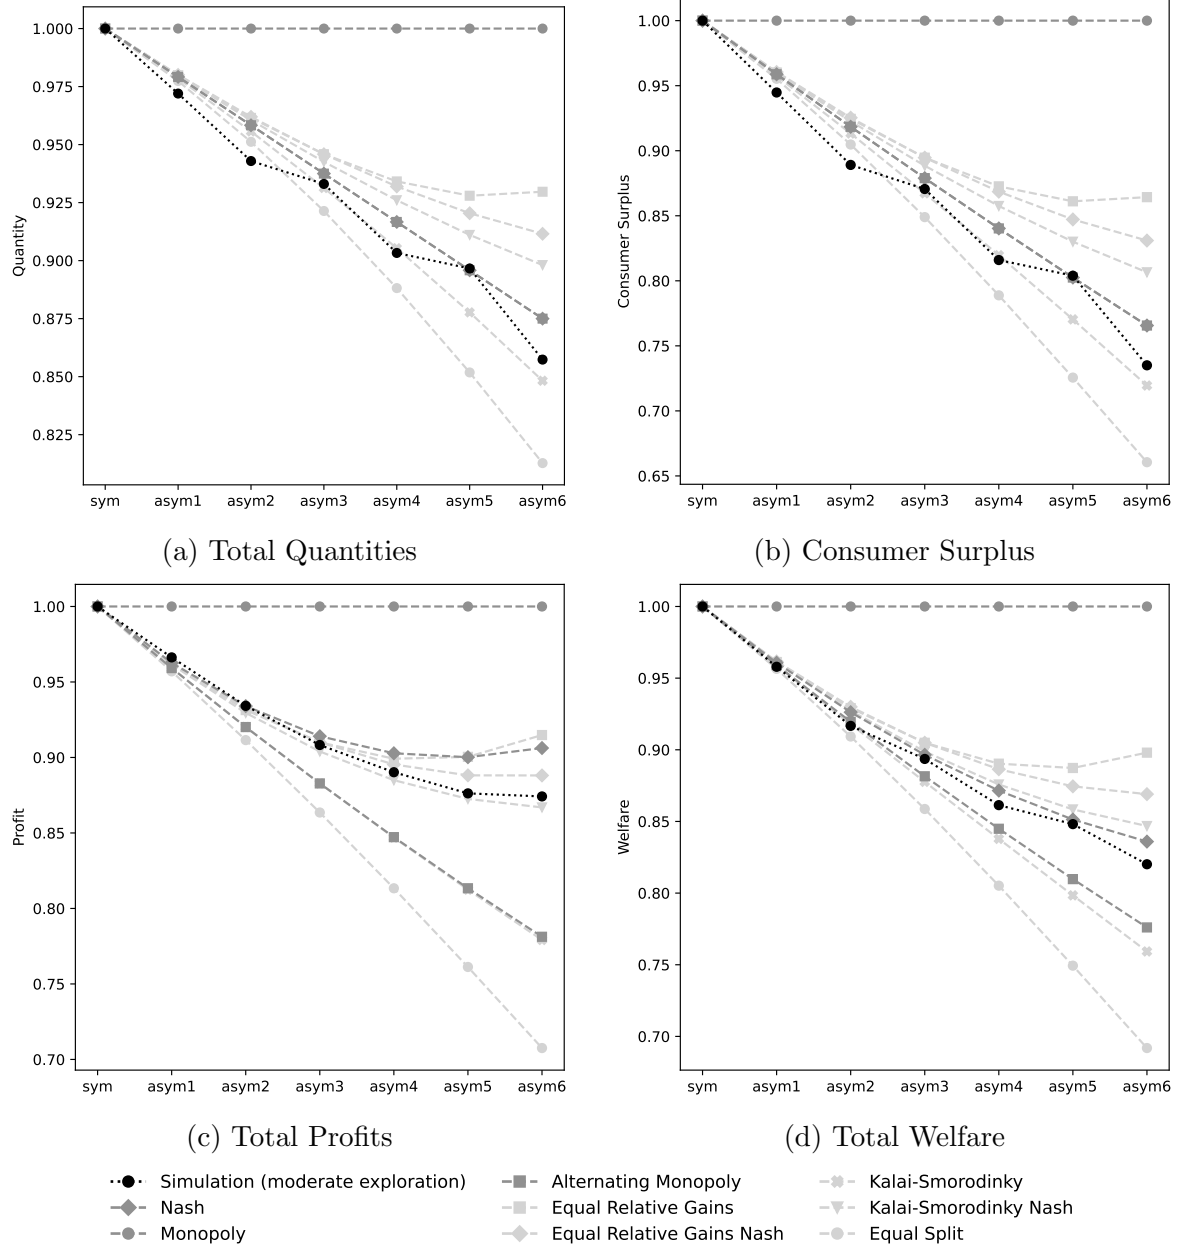


Figure C.1: Simulation results (comparative statics; symmetric case = 1) for alternative different degrees of asymmetry.

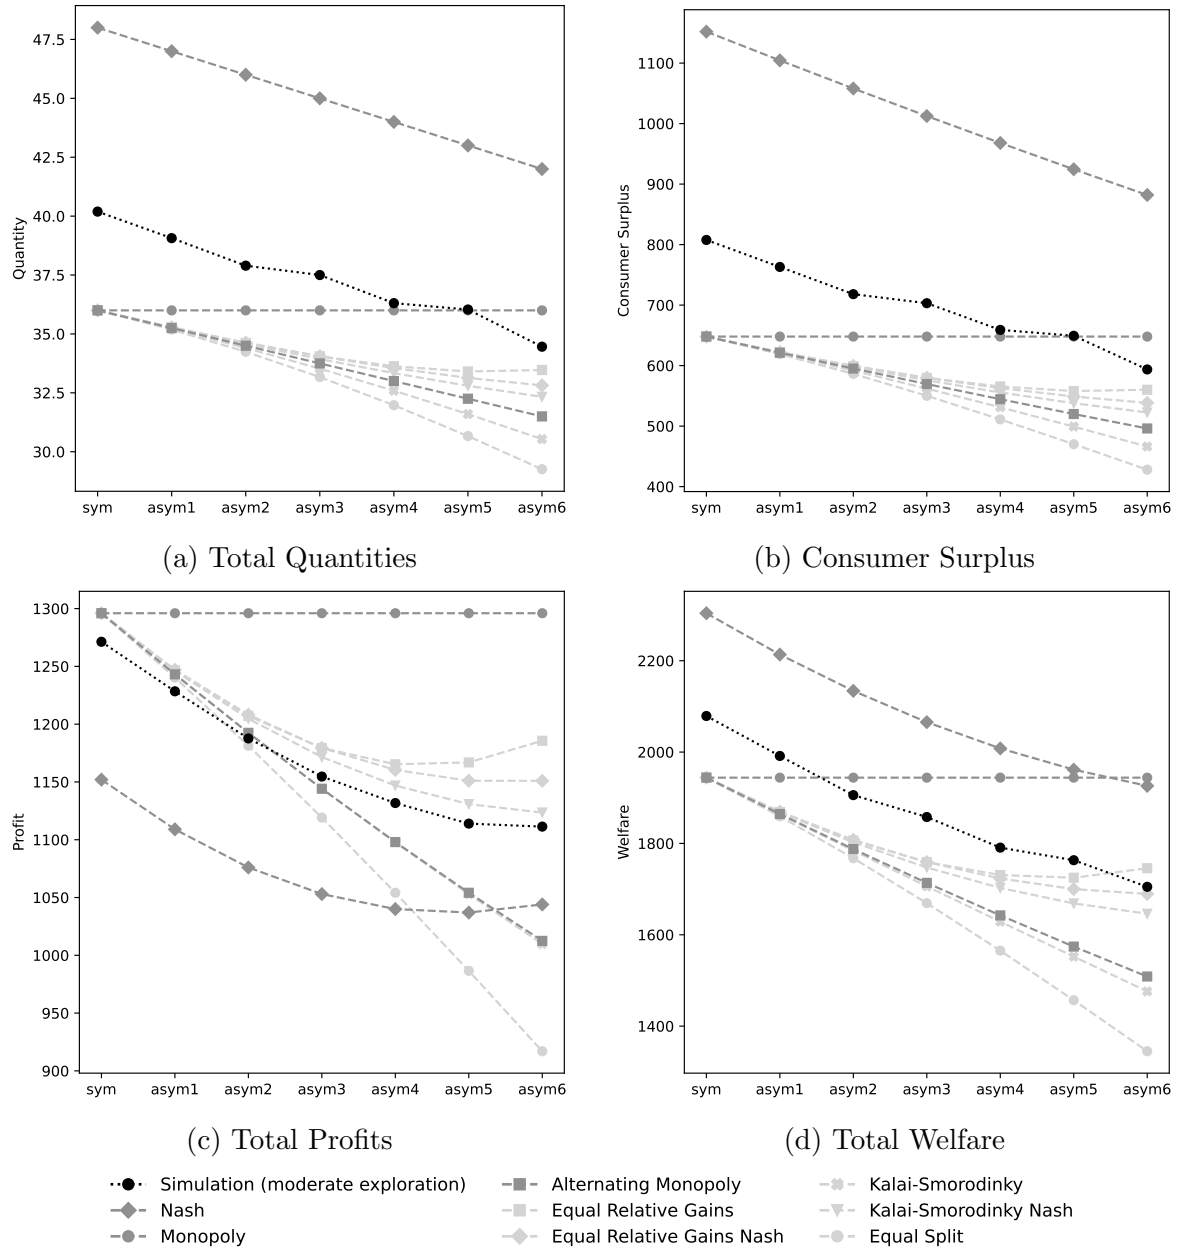


Figure C.2: Simulation results (levels) for alternative different degrees of asymmetry without memory.




Figure C.3: Pareto frontier and simulation results for alternative different degrees of asymmetry with (symmetric case: 0).

Table C.2: Average squared distances (alternative cost parameterization)

	Q	CS	Π	W
Nash	59	100,091	10,024	47,206
Monopoly	5	7,111	18,731	20,373
Alternating Monopoly	13	16,849	2,207	23,635
Kalai-Smorodinsky	16	19,550	2,310	27,943
Equal Split	20	24,063	8,868	52,103
Equal Relative Gains	10	13,285	1,626	8,528
Kalai-Smorodinsky (Nash)	11	14,896	309	10,931
Equal Relative Gains (Nash)	10	13,848	823	8,750

Table C.3: Average squared normalized distances (multiplied by 1,000)

	Q	CS	Π	W
Nash	0.11	0.38	0.26	0.07
Monopoly	6.99	24.95	8.15	13.52
Alternating Monopoly	0.11	0.38	2.19	0.55
Kalai-Smorodinsky	0.09	0.31	2.27	1.00
Equal Split	0.64	1.89	7.07	4.38
Equal Relative Gains	1.11	3.61	0.34	1.25
Kalai-Smorodinsky (Nash)	0.41	1.31	0.02	0.17
Equal Relative Gains (Nash)	0.70	2.28	0.06	0.58