NBER WORKING PAPER SERIES

ESTIMATING DYNAMIC GAMES OF OLIGOPOLISTIC COMPETITION: AN EXPERIMENTAL INVESTIGATION

Tobias Salz
Emanuel Vespa

Working Paper 26765
http://www.nber.org/papers/w26765

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2020

For helpful discussions of this project we would like to thank John Asker, Isabelle Brocas, Colin Camerer, Juan Carrillo, Allan Collard-Wexler, Guillaume Frechette, Ali Hortacsu, Kei Kawai, Robin Lee, Alessandro Lizzeri, Ryan Oprea, Ariel Pakes, Tom Palfrey, Stephen Ryan, Andrew Schotter, Ralph Siebert, Matthew Shum, Anson Soderbery, Charles Sprenger, Severine Toussaert, Matan Tsur, Georg Weizsacker, Alistair Wilson, and Sevgi Yuksel. Vespa is grateful for financial support from the UCSB Academic Senate. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation
Tobias Salz and Emanuel Vespa
NBER Working Paper No. 26765
February 2020
JEL No. L10,L13

ABSTRACT

We evaluate dynamic oligopoly estimators with laboratory data. Using a stylized en-try/exit game, we estimate structural parameters under the assumption that the data are generated by a Markov-perfect equilibrium (MPE) and use the estimates to predict counterfactual behavior. The concern is that if the Markov assumption was violated one would mispredict counterfactual outcomes. The experimental method allows us to compare predicted behavior for counterfactuals to true counterfactuals implemented as treatments. Our main finding is that counterfactual prediction errors due to collusion are in most cases only modest in size.

Tobias Salz
MIT Deparatment of Economics
77 Massachusetts Avenue, E52-404
Cambridge, MA 02139
and NBER
tsalz@mit.edu

Emanuel Vespa
2127 North Hall
University of California
Santa Barbara, CA 93106
vespa@ucsb.edu

A data appendix is available at http://www.nber.org/data-appendix/w26765

1 Introduction

Many empirical studies attempt to recover primitives of an economic model that are then used to evaluate *counterfactual* scenarios. Identification of such primitives typically requires assumptions (on functional forms, equilibrium selection, etc.), and if these are not met, parameter estimates and counterfactual policy recommendations can be inaccurate. In this paper, we illustrate how the laboratory can be used to evaluate the extent to which specific modeling assumptions generate counterfactual prediction errors if the model is misspecified. The exercise we propose involves four steps. First, we implement a model of interest in the laboratory and obtain data resulting from (subjects') play under those primitives. Second, *under standard identification assumptions*, we use the laboratory-generated data to structurally recover the primitives. Third, we compare the true implemented primitives to the estimates. Fourth, we use the estimated primitives to predict behavior in a counterfactual scenario. Crucially, for step four, we also run the counterfactual scenario directly in the laboratory so that we can compare the prediction to actual behavior in the counterfactual scenario. This comparison allows us to evaluate to what extent specific assumptions for estimation lead to model-misspecification and counterfactual prediction errors. In particular, we are interested in the bias that results when the econometrician does not account for collusion.

We study a dynamic game of oligopolistic competition. The primitives in these models are often related to investments or fixed costs and *counterfactual* policy scenarios might study merger guidelines or other market interventions. The basic environment is one of repeated interactions, with a state variable that evolves endogenously (e.g. the number of firms in the market in an entry/exit model). The set of subgame-perfect equilibria (SPE) in dynamic games with an infinite horizon can be large (Dutta, 1995) and often hard to characterize. Empirical studies often focus on a subset of SPE known as Markov-perfect equilibria (MPE), where attention is restricted to stationary Markov strategies. On the one hand, this restriction is extremely useful as it allows for dynamic programming tools to solve for MPE and makes the model tractable. On the other hand, there are circumstances where the assumption of Markov play may be too restrictive. In fact, when the gains from collusion are large, behavior may not be properly captured by an MPE. Support of collusion as an SPE typically requires the threat of credible punishments to deter parties from otherwise profitable deviations. Hence, agents need to keep track of past play and use history to condition their present choices. Stationary Markov strategies, however, condition behavior only on the state variable, ignoring the particular history that led to the current state. Consequently, collusive equilibria that are supported by a switch to a punishment phase upon deviation cannot be enforced with a Markov strategy. It is therefore possible that equilibrium Markov strategies are not the true data generating process when the environment provides strong incentives to collude.

The central goal of this study is to test how restrictive the Markov assumption is for counterfactual predictions. To that end, we implement a dynamic oligopoly model in the laboratory where a key treatment variable is a structural parameter that affects whether collusion can be supported as an SPE or not. To provide a stringent test, the gains from collusion are very high in some treatments.

There are two potential threats posed by a violation of the MPE assumption. First, it may lead to biased estimates. Standard Monte Carlo simulations in our environment show that estimates are strongly biased if the data are generated according to a collusive equilibrium. Second, it may lead to biased counterfactual predictions. Monte Carlo simulations, again, confirm that such counterfactual prediction error can be severe in our setting. Consider a baseline in which the incentives to collude are low, and the data is actually consistent with an MPE. If the incentives to collude are larger in the counterfactual scenario, the selected equilibrium may change. The counterfactual might therefore not only entail a change in primitives but also in conduct, violating the ceteris paribus assumption of counterfactual comparisons. Therefore, a prediction based on MPE play in the counterfactuals may lead to errors. A Monte Carlo exercise can help to determine the extent of biases and prediction errors under *specific* assumptions on behavior, but it cannot resolve which of the assumptions better capture human behavior. The experimental exercise in this paper allows us to study the consequences of human behavior without having to take an a priori stance on what such behavior consists of. In this sense it is akin to a Monte Carlo exercise, except that the data are generated by humans in a laboratory.

Our design is based on a model that builds on the seminal contribution of Ericson and Pakes (1995), an infinite-horizon entry/exit game.[1] In our model each of two firms can be in or out of the market in each period and their state (in/out) is publicly observable. At the beginning of the game, both firms start in the market and each period consists of two stages: the quantity stage and the entry/exit stage. When both firms are in the market they play a quantity-stage game. Each firm can either select a *low* or a *high* level of production, where

[1] To be precise, we implement a model that includes privately observed shocks to firms' decisions, like in Bajari et al. (2007) and Aguirregabiria and Mira (2007).

high is associated with the stage-game Nash equilibrium and *low* with collusion. A firm that is not in the market does not participate in the quantity game and makes zero profits. If a firm is alone in the market, the optimal action is to set the *high* quantity. Firms receive feedback on their quantity-stage payoffs and then face an entry/exit stage that determines whether they are in or out of the market for the next period. Firms that are in the market choose whether to stay in or to receive a scrap value and exit the market, while firms that are out decide whether to stay out or to pay an entry fee. Scrap values and entry fees are privately observed and randomly drawn each period from common-knowledge distributions. There is no absorbing state; a firm that exits the market can re-enter at a future date. Total payoffs in each period consist of the quantity-stage and the entry/exit payoff, and profits are discounted by $\delta \in (0, 1)$.[2]

A first set of treatments manipulates the incentives for quantity-stage collusion. The treatment variable (A) is a parameter that determines by how much own quantity-stage profits increase when a firm changes production from *low* to *high*. We use three different parameterizations for A, keeping the other parameters fixed. In all cases we characterize the unique symmetric MPE in which firms select *high* in the quantity stage and make entry and exit decisions consistent with the quantity choice. When the gain from increasing own production is not large (lower values of A), we show that selecting *low* in the quantity stage can be supported as an SPE, and profit gains relative to the MPE range from 75% to 450% in our parameterizations.

Overall, collusion in the quantity stage creates a pattern of entry and exit that is different from the one under the MPE. This means, that if firms collude, structural estimates under the assumption of Markov play will be biased. However, the pattern of entry/exit decisions may differ from the predictions of the MPE for reasons other than the possibility of collusion. To control for such discrepancies, we conduct three additional treatments –one for each value of A– without a quantity stage choice. Whenever both firms are in the market, they are assigned the stage-game Nash payment corresponding to both selecting *high*, and agents only make entry/exit decisions. These treatments therefore serve as a baseline to capture discrepancies from the MPE entry/exit predictions that are unrelated to quantity-stage collusion.

The main finding is that for the purpose of counterfactual predictions the restriction to

[2] The goal of the simplified version is to recreate an environment that retains the main tensions and allows us to focus on a test of Markov perfection. In doing so our version is stripped of aspects that are meaningful when dealing with non-laboratory data, but not central to the questions in this paper. For instance, we set the number of firms and the set of quantity-stage available actions to two because it simplifies coordination hurdles required for collusion, and hence provides a more demanding test for Markov perfection.

equilibrium Markov strategies leads to a relatively modest bias. It seems reasonable to expect that, in settings with high incentives to collude, the restriction to Markov strategies will lead to misspecification and bias. Our experimental design therefore provides a stress test of the Markov assumption in an environment where it is expected to fail. However, we find that, for most counterfactual computations, differential incentives for collusion are only a minor source of errors.

The study of quantity-stage choices suggests a mechanism for this finding. Collusion leads to errors only as long as agents manage to sustain it. We do find that subjects' quantity-stage choices respond significantly to the collusion incentive. When collusion cannot be supported as an SPE, modal quantity-stage behavior coincides with the Nash equilibrium. Contrarily, in treatments with high gains from collusion, many subjects try to collude at the beginning of the game. But successful collusion phases are rare as they break down quickly after the initial attempt. Moreover, enacted punishments are close to the MPE after collusion breaks down. Hence, discrepancies from MPE quantity-stage choices only last for the few periods of a supergame. Equilibrium Markov play can therefore still serve as a sensible approximation of behavior despite the fact that the choices of a significant proportion of subjects are better rationalized by a non-Markovian strategy.

The fact that collusion in the quantity stage does not succeed for long is consistent with patterns documented for actual firms outside of the laboratory. While the incentives to collude in our environment are high and coordination challenges are minimal, equilibrium collusion is tacit. However, empirical evidence suggests that successful cartels usually require inter-firm communication and/or transfers, referred to as *explicit* collusion. The Sherman Act, part of the US anti-trust law, forbids any explicit agreement to coordinate on quantities or prices. Our setting should, therefore, be well suited to capture situations where firms do not engage in illegal agreements to coordinate.[3] However, an important insight of the "new empirical industrial organization" is that each industry operates subject to a unique set of market forces and institutional constraints, often permitting an empirical one size fits all approach. This caveat of external validity certainly also applies to the results here. To what extent the Markov assumption is sensible should therefore be carefully judged on a case-by-case basis. However, we do believe that experiments can give useful auxiliary information on the validity of specific modeling assumptions.

[3]For example, according to Marshall and Marx (2012), page 3: "(...), it appears that repeated interaction is not enough in practice, at least not for many firms in many industries. Even for duopolies (...) explicit collusion was required to substantially elevate prices and profits." See Marshall et al. (2008) and Harrington (2008) for additional evidence.

Beyond these results on collusion, we also document that subjects' decisions deviate in other ways from the MPE, leading to less entry and exit than expected. We provide an analysis of this phenomenon, which we call *inertia*, and also suggest ways to correct for it in the estimation.

Our findings also offer avenues for future research. One could inquire what market structures or other conditions lead to sustained collusion and how those would impact counterfactual predictions. Those include direct (cheap-talk) communication and multi-market contacts. In this study, we focus on collusion, but one can imagine other reasons why the assumption of Markov perfection might be violated: in many dynamic environments an MPE can be a very demanding concept in terms of the information that agents ought to have and their capacity to process it. The literature discusses ways to relax these requirements (see Pakes, 2015), and the lab might be used as a tool to evaluate such alternatives.

Other possible extensions concern the definition of the state. Our exercise uses standard payoff-relevant states to define MPE, but another approach expands the definition of the state to include elements of the history that are not payoff relevant (e.g. Fershtman and Pakes, 2000). In this case, the researcher decides how to expand the state, since some elements of the history must be excluded (otherwise the set of MPE and the set of SPE would coincide). While in the environment that we study we find that expanding the definition of a state would not lead to meaningful improvements; further experimental research can provide guidance on when this could be the case.

Finally, in environments with multiple equilibria, the laboratory can help evaluate if there are systematic deviations from a particular assumption on equilibrium selection. One reason to assume sub-game perfection and a Markovian structure is often the lack of more compelling assumptions. Moreover, since there are many ways in which behavior can deviate from Markov perfection, exploring the consequences of possible deviations might not be feasible. Laboratory data, however, can provide some guidance. Based on the observation of human behavior, it can help to sort out which of all theoretically possible deviations from Markov perfection is most likely to occur.

Related Literature

The original framework for the game we study is formulated in Ericson and Pakes (1995).[4] Most empirical applications of Markov-perfect industry dynamics rest on the *two-stage* approach, which substantially reduces the computational burden of estimation relative to full solution methods.[5] We describe the application of the *two-stage* approach in section 4. The two-stage approach builds on the idea that the equilibrium law of motions, on which players base their actions, is directly observed in the data. There are also suggestions for estimators in which players believes are not in equilibrium (Aguirregabiria and Magesan, 2016).

Our paper is also related to the literature that studies the issue of equilibrium multiplicity in empirical models of strategic behavior. Note that there are, in fact, two issues related to equilibrium selection. First, a maintained solution concept, like Markov perfection, will rule out many plausible equilibria. Second, even the set of equilibria admitted by a solution concept may be large. This is also true for MPEs. Typical empirical applications (especially those with asymmetric firms) might allow for many potential MPE.[6] For suggestions on how to deal with the latter, see Borkovsky et al. (2014).[7] Our paper focuses on the former, and the environment is constructed to have just one symmetric MPE. However, a similar experimental design might be used to study equilibrium selection within the set of MPE. For a broader discussion of the issue of multiple equilibria in empirical models see De Paula (2013).

The literature is clearly aware that assuming Markov play can lead to biases under collusion. In fact, in their seminal paper on static entry games Bresnahan and Reiss (1990) have pointed out that collusion can affect estimates of market structure parameters that are obtained under the assumption of simple static Nash. More broadly, there is a literature on the detection of collusion in dynamic settings, of which Porter (1983) is a famous example. Harrington and Skrzypacz (2011) present theoretical work that builds on recent insights from the literature on repeated games to explain collusive practice in a dynamic context. The focus of our paper, however, is not on how to detect collusion but to quantify model mis-specification (and

[4]To give just some recent examples of applications see Collard-Wexler (2013) and Ryan (2012) for entry exit choices, and Goettler and Gordon (2011), Schmidt-Dengler (2006), Blonigen et al. (2013), and Sweeting (2013) for the introduction and development of new products. See Aguirregabiria and Mira (2010), Doraszelski and Pakes (2007), and Ackerberg et al. (2007) for references on methodological issues.

[5]Variants of such estimators have, for example, been suggested in Bajari et al. (2007), Pakes et al. (2007), Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008).

[6]Since the empirical dynamic games literature has largely focused on MPE, the *multiplicity issue* is typically referring multiplicity *within* the set of MPE. Abito (2015) discusses identification in repeated games without imposing a Markov structure.

[7]The same issue arises for making counterfactual predictions. If the counterfactual scenario allows for multiple equilibria one has to pick the equilibrium that agents would actually coordinate on in the counterfactual. An econometric suggestion of how to deal with this problem has, for example, been suggested in Aguirregabiria and Ho (2012).

counterfactual bias) if the econometrician ignores collusion.[8]

We find that the symmetric MPE organizes the comparative statics very well. This finding is consistent with results from the experimental literature on dynamic games. For example, Battaglini et al. (2015) and Vespa (2019) find that a large proportion of play can be rationalized by the MPE in other dynamic games.[9] Finally, experimental data has been used to evaluate structural estimates. The main focus of these studies is not to assess the accuracy of counterfactual predictions but of parameter recovery. See Bajari and Hortacsu (2005) and Ertaç et al. (2011) for the case of auctions, Brown et al. (2011) for the case of labor-market search, Fréchette et al. (2005) for bargaining, and Merlo and Palfrey (2013) for voter turnout models.[10]

2 Setup

We chose a parsimonious implementation of Ericson and Pakes (1995) with two firms, indexed by i. The time horizon is infinite, and agents discount the future by $\delta \in (0,1)$. In each period t, firms first face a *quantity stage* and then face a market *entry/exit stage*. A state variable tracks whether firm i is in the market in period t ($s_{it} = 1$) or not ($s_{it} = 0$). We assume that at $t = 0$ all firms start in the market. There are four possible values for the state of the game at time t: $\mathbf{s}_t = (s_{1t}, s_{2t}) \in S = \{(0,0), (0,1), (1,0), (1,1)\}$.

Model Quantity Stage

At the beginning of each period, the observable part of the state ($\mathbf{s}_t$) is common knowledge. If both firms are in the market ($\mathbf{s}_t = (1,1)$), firms simultaneously make a quantity choice, $q_{it} \in \{0,1\}$, with quantity stage profits given by:

$$\Pi_{it} = A \cdot (1 + q_{it}) - B \cdot q_{-it}. \tag{1}$$

[8]There is a literature that studies collusion in the laboratory. First, if the prisoner's dilemma is thought of as a reduced form oligopoly game, many insights on collusion -or cooperation- have been provided by the experimental literature on the repeated prisoner's dilemma (see Dal Bó and Fréchette, 2018 for a recent survey). There is also a literature that studies collusion in the context of market experiments (see Davis and Holt, 2008 for a survey).

[9]For other studies of dynamic games in the laboratory, see Battaglini et al. (2012), Kloosterman (2018), Saijo et al. (2014), and Vespa and Wilson (2017; 2018). There is also a literature that focuses on studying to what extent subjects can solve dynamic problems, but abstracts from aspects of strategic interaction that are the focus of this paper. See for example, Hey and Dardanoni (1988), Noussair and Matheny (2000), Lei and Noussair (2002), and Houser et al. (2004).

[10]There are also a limited number of studies that use field data to either evaluate structural estimates or aid structural estimation. See Arcidiacono et al. (2016), Conlon and Mortimer (2010), Conlon and Mortimer (2013), Keniston (2011), and Pathak and Shi (2017).

We require that $B > A$. A is a parameter that captures the effect of the own production decision on profits.[11] B measures the effect of competition: how firm i's profits are affected when the competitor increases production. The profit function is therefore a reduced form that captures the typical strategic tension inherent in a Cournot game, which is represented here with prisoner's dilemma payoffs. Selecting the higher quantity ($q_{it} = 1$) increases firm i's own market share but also imposes an externality on the other firm through the decrease in price. We will refer to the choice firms face when both are in the market as the quantity stage decision and to the choice of $q_{it} = 1$ ($q_{it} = 0$) as selecting the *high* (*low*) quantity. Once both firms have made their quantity choices, they learn the other firm's choice and the corresponding quantity-stage profits.

If at least one firm is out of the market ($\mathbf{s}_t = (s_{1t}, s_{2t}) \in \{(0,0), (0,1), (1,0)\}$), there is no quantity choice. The quantity stage profits of a firm that is out of the market are normalized to zero. If only firm i is in the market, its quantity stage profits are given by $2A$. This corresponds to the highest payoff in (1), as if $q_{it} = 1$ and $q_{-it} = 0$. Formally, the available action space ($\mathcal{Q}_i$) for the quantity stage depends on the state: $\mathcal{Q}_i(\mathrm{s} = (1,1)) = \{0,1\}$, and $\mathcal{Q}_i(\mathrm{s} \neq (1,1)) = \{\varnothing\}$.

Model Entry/Exit Stage

After the quantity stage, firms can decide whether they want to be in the market for the next period or not. This choice is captured by $a_{it} \in \{0,1\} = \mathcal{A}_i$, with $a_{it} = 1$, indicating that firm i chooses to be in the market in period $t+1$. If a firm that is currently in the market decides to exit, it collects a scrap value, which is iid, that is $\phi_{it} \sim \mathcal{U}[0,1]$.[12] At the beginning of the exit stage, a firm that is deciding on whether to exit or not is privately informed about the realized scrap value. Firms that stay in the market don't receive a scrap value.

A firm that is currently out of the market, but is deciding whether to enter or not, faces a similar situation. If the firm decides to enter, it must pay an entry fee. This entry fee is $C + \psi_{it}$. The fixed part C is common knowledge as well as the fact that the random part is iid and that $\psi_{it} \sim \mathcal{U}[0,1]$. The firm deciding whether to enter or not is privately informed of the realization of ψ_{it} before it makes its choice. Firms can re-enter the market if they are out, which means that exiting the market does not lead to an absorbing state.

Once firms make their entry/exit choices, period t is over, and period payoffs are realized.

[11] To guarantee that subjects make no negative payoffs in the laboratory we add a constant 0.60 to all payoffs

[12] While in empirical applications the error term is typically assumed to be distributed T1EV, we favored a uniform distribution because it is much easier to explain to subjects in the laboratory. Moreover, the bounded support rules out extremely large payoffs.

The dynamic entry/exit choice determines the evolution of the state from $\mathbf{s}_t$ to $\mathbf{s}_{t+1}$, and a new period starts.

Markov Perfection

In each period, total payoffs are pinned down by the state (s) and the random component of the entry/exit decision. Using these payoff-relevant variables, it is possible to compute the value function of the game at t, which for known market quantities (q_{it}, q_{-it}) is given by:

$$V_i(\mathbf{s}_t, \epsilon_{it}) = \max_{a_{it}, q_{it}} \Big\{ \mathbf{1}\{\mathbf{s_t} = (1,1)\} \cdot (A \cdot (1 + q_{it}) - B \cdot q_{-it}) + \mathbf{1}\{\mathbf{s_t} = (1,0)\} \cdot (2A) \\ + \epsilon_{it}(a_{it}, s_{it}) - \mathbf{1}\{a_{it} = 1, s_{it} = 0\} \cdot C + \delta \cdot \mathbb{E}_{\phi_{-i}, \psi_{-i}} \left[V_i(\mathbf{s}_{t+1}, \epsilon_{i(t+1)}) | \mathbf{s}_t, a_{it} \right] \Big\} \quad (2)$$

with

$$\epsilon_{it}(a_{it}, s_{it}) = \phi_{it} \cdot \mathbf{1}\{a_{it} = 0, s_{it} = 1\} - \psi_{it} \cdot \mathbf{1}\{a_{it} = 1, s_{it} = 0\},$$

where $\mathbf{1}\{\cdot\}$ is an indicator function. Following Maskin and Tirole (2001), a Markov strategy prescribes a choice for both stages of each period that depends only on payoff-relevant variables.[13]

Definition: A Markov strategy is a set of functions: i) prescribing a choice for each state in the quantity stage, $\rho_i : \mathrm{s} \to \mathcal{Q}_i$; and ii) an entry/exit choice for each value of the state and random component, $\alpha_i : (\mathbf{s}, \epsilon_i) \to \mathcal{A}_i$.

A Markov-perfect equilibrium (MPE) is a subgame-perfect equilibrium (SPE) of the game in which agents use Markov strategies. It is the typical solution concept for dynamic oligopoly games, as well as an essential assumption in the estimation procedures that we evaluate.

Definition: An MPE is given by Markov strategies $\big(\rho = [\rho_1, \rho_2], \alpha = [\alpha_1, \alpha_2]\big)$ and state transition probabilities $(F^{\alpha}(\mathbf{s}_{t+1}|\mathbf{s}_t))$ such that: i) $\rho(\mathrm{s} = (1,1)) = (1,1)$; ii) α maximizes the discounted sum of profits for each player, given $F^{\alpha}(\mathbf{s}_{t+1}|\mathbf{s}_t)$; and iii) α implies $F^{\alpha}(\mathbf{s}_{t+1}|\mathbf{s}_t)$.

[13] For simplicity in the text we refer to s_t as the state, but in the formulation of the value function payoff-relevant variables include the *endogenous* state $\mathbf{s}_t$ and the conditionally *exogenous* state ϵ_{it}. s_t is endogenous as it depends on the firm's choices, while ϵ_{it} is conditionally exogenous. That is, conditional on the firm being in the market or out of the market, ϵ_{it} is determined by an exogenous process. It is possible to expand the definition of the state so that it would also include part of the history of the game. In the limit all aspects of the history can be included, making the restriction to strategies that condition on the state irrelevant. The goal of this paper is to test how restrictive it is to focus on equilibria that ignore past play. For further discussion on why it is meaningful not to include elements of the history that are not payoff-relevant as part of the state see Mailath and Samuelson (2006).

In an MPE, firms have no means of enforcing anything but the static Nash equilibrium in the quantity choice. Since $B > A$, firms will always choose the *high* quantity when both are in the market. If any firm were to select the *low* quantity and use a strategy that conditions only on the state, the other could systematically take advantage by selecting the *high* quantity. The quantity choices in an MPE lead to the following reduced form for the quantity stage payoffs: $\Pi_{it} = s_{it} \cdot (2A - B \cdot s_{-it})$. In other words, a firm earns the highest quantity stage profits ($2A$) when it is alone in the market and the defection payoff ($2A - B$) when both firms are in the market at the same time. With the quantity-stage profits set, an MPE specifies entry/exit probabilities for each value of the state. The existence of MPE equilibria is discussed in Doraszelski and Satterthwaite (2010), but our assumptions guarantee that there is a symmetric MPE consisting of a set of state-specific cutoff strategies. In the next section, we compute such equilibria for specific parametrizations that we implement in the laboratory.

3 Experimental Design

Standard Treatments

The underlying primitives of the model are A, B, C, δ, and the distribution of ϵ. The goal of the structural estimation procedure will be to recover estimates for A, B, and C using experimental data, assuming that the econometrician does know δ and the distribution of ϵ. We will generate data in the laboratory using three different values for A. For our purposes, it is useful to have A as a treatment variable, as it affects whether collusion in the quantity choice can be supported as an SPE or not. We will describe collusive equilibria later in this section, but to intuitively see why, notice –in Equation 1– that as A increases, the temptation of deviating from a *low* to a *high* quantity increases as well. Hence, for a given δ, it will be more difficult to support collusion when the own effect on profits (A) is larger.

In all our treatments, parameter B, which measures the effect on player's own profits if the other player increases production, is set to 0.60. Depending on whether the value of A is small (A_S), medium (A_M), or large (A_L), the coefficients are, respectively: $A_S = 0.05$, $A_M = 0.25$, or $A_L = 0.40$. The three payoff matrices in Table 1 display the quantity stage payoffs for the case when both subjects are in the market for each of the three values of A. Finally, we set $C = 0.15$. For presentational purposes in the laboratory, all payoffs are multiplied by 100.

Treatment variable A is used to perform the main exercise of the paper. For example, we

	0	1
0	65	5
1	70	10

$A_S = 0.05$

	0	1
0	85	25
1	110	50

$A_M = 0.25$

	0	1
0	100	40
1	140	80

$A_L = 0.4$

Table 1: Quantity-choice payoffs for the row player in the laboratory

will recover estimates from the baseline treatment, A_L, (where collusion is not an SPE) and then make predictions for the treatment with A_S (where collusion is supported as an SPE). If collusion is indeed present in the data for A_S, then the counterfactual prediction that assumes MPE play will entail a large prediction error.

Characterization of the Stationary MPE

We compute for each treatment (three values for A) the cutoff-strategies corresponding to the symmetric MPE and report them in italics in Table 2.[14] For the dynamic entry/exit choice, the equilibrium MPE strategy provides the probability of being in the market next period, conditional on the current state, $p(s)$, and we refer to the vector with such probabilities as $\mathbf{p}$.[15] Given the uniform distributions for the random entry fee and the random exit payment, we can interpret these probabilities as thresholds. Consider, for example, the A_S treatment when $s = (1, 0)$. In that case, the strategy prescribes for the agent in the market to exit if the random exit payoff is higher than the threshold 0.458.[16] In other words, when the firm is in the market ($s(1, \cdot)$), the threshold indicates the lowest scrap value at which the firm would exit, and we will refer to these as exit thresholds. Entry thresholds (when the state is $s(0, \cdot)$) include only the random part of the entry fee and indicate the highest random entry fee for which the firm would enter. For example, if in the current state both agents are out ($s = (0, 0)$), the MPE probability of being in the market next period for the A_S treatment is 0.161. This means that, to enter the market in that state, an agent is willing to pay a random entry fee of up to 0.161.

[14] In Appendix A we provide details behind these computations.

[15] The table reports the conditional probability of being in the market next period for the agent whose current state is the first component of s. For example, if $s = (1, 0)$, then $p(1, 0)$ reports the corresponding conditional probability for the firm that is currently in the market.

[16] In the laboratory random entry and exit payoffs are multiplied by 100 given the normalization. For predictions and when we report results in the text we omit the normalization and will thus refer to random exit and entry payoffs as numbers between 0 and 1.

Table 2: Cutoff-strategies for each treatment: MPE and CE

	A_S		A_M		A_L	
Conditional probability	MPE ($\mathbf{p}$)	CE ($\mathbf{p}_c$)	MPE ($\mathbf{p}$)	CE ($\mathbf{p}_c$)	MPE ($\mathbf{p}$)	CE ($\mathbf{p}_c$)
$p(1,0)$	*0.458*	0.519	*0.688*	0.823	*0.880*	0.925
$p(1,1)$	*0.360*	0.512	*0.596*	0.784	*0.781*	0.870
$p(0,0)$	*0.260*	0.368	*0.498*	0.663	*0.681*	0.757
$p(0,1)$	*0.161*	0.362	*0.408*	0.624	*0.583*	0.702
Is collusion in quantity choice an SPE:	YES		YES		NO	
Gains from collusion in quantity choice only:	450.8%		75.9%		32.1%	
Gains from collusion in quantity + dynamic choice:	481.1%		93.2%		52.0%	

Note: This table presents the conditional probability for the firm whose current state is the first component of s. $p(s)$ indicates the probability of being in the market next period conditional on being in state s in the current period. The probabilities are presented for each of the three values of A as indicated in the top row. Predictions are presented for the $MPE(\mathbf{p})$ as well as the case where players collude in the quantity choice, CE ($\mathbf{p}_c$). The bottom panel of the table indicates whether the collusive equilibrium can be supported as an SPE and how high the gains over the MPE would be. Full collusion refers to the joint monopoly case with computation in Appendix A where firms not only coordinate their static quantity production choice but also coordinate in the entry/exit choices.

The stationary MPE reported in Table 2 makes clear predictions within and across treatments. In Appendix B we explicitly formulate these comparative-statics hypotheses.

Collusive Equilibrium

An assumption underlying the symmetric MPE is that agents play a Nash equilibrium in the quantity stage. In principle, however, it is possible for agents to attain higher than MPE payoffs in equilibrium if they collude in their quantity choices. We now present a non-Markovian strategy that can support collusion in the quantity choice (details in Appendix A).

Assume that both agents select the *low* quantity whenever they are in the market. In such a collusive arrangement the value of being in the market is higher. The entry and exit probabilities under such stage game collusion are denoted as $\mathbf{p}_c$ (Table 2). When both agents are in the market, the outcome of the quantity choice is observed before the exit decision is made. This enables agents to use trigger strategies: as long as they have colluded in the past, they will make their entry/exit decisions according to $\mathbf{p}_c$. If any agent ever deviates to the *high* quantity, then all entry/exit decisions made from then onwards follow the MPE thresholds (**p**), and agents choose the stage-Nash quantities.

For the discount value that we implement in the laboratory ($\delta = 0.8$), the collusive strategy is an SPE for A_S and A_M. In fact, Table 2 shows that the gains from collusion relative to the MPE are large: 450.8% and 75.9% for A_S and A_M, respectively. For A_L, there are incentives to deviate from the collusive strategy, and it does not constitute an SPE.

From now on, we will refer to the SPE that uses the collusive strategy as the *collusive equilibrium* (CE), although this is simply one of possibly many collusive equilibria. We think of the characterized collusive equilibrium as a natural benchmark for collusive behavior.[17]

We now want to highlight some aspects of the comparison between MPE and CE probabilities. First, the ordering of the probabilities across the two is unchanged. In other words, the difference between the MPE and the CE probabilities is quantitative but not qualitative. Second, the CE probabilities predict a much smaller effect of competition. Because players always select the low quantity, the payoff from being alone in the market is closer to the one of being in the market together. Third, and perhaps most importantly, along the equilibrium path, the CE is consistent with a Markov strategy and looks like an MPE. The only difference for the computation of the probabilities is the assumption on quantity stage payoffs when both agents are in the market.

The CE also delivers a prediction for entry and exit thresholds following defection. Exit and entry thresholds should respond to market behavior according to the collusive strategy: the market is relatively less valuable after the other agent deviates from collusive behavior. As a consequence, agents would be willing to leave for lower scrap values and would be willing to pay less in order to re-enter the market.[18]

Optimization errors and No Quantity Choice Treatments

To explain why we introduce a second treatment variable, it is useful to discuss some features of the environment presented so far. An important reason why we use a binary action space for quantities is that it makes the trade-off very stark. Consider the alternative, where the quantity choice is made in a continuous or in a large discrete action space. The problem for

[17]This collusive equilibrium, however, does not support the most efficient outcome from the firms' perspective. To achieve efficiency, firms would also need to coordinate their entry/exit choices. In Appendix A we also provide the computation of entry/exit thresholds under joint maximization. However, given the private nature of scrap values and the random portion of the entry fee, coordination is difficult and we favor a collusive equilibrium benchmark that does not rely on such demanding conditions. More importantly, the data are not consistent with the predictions under joint maximization of profits.

[18]For further details, see Appendix B where we explicitly state these hypotheses.

subjects who want to collude would be more challenging as they first have to coordinate on the collusive quantity and, in addition, on which quantity to use for punishments. With just two choices there are no such coordination problems; the difference in payoffs between collusive and stage-Nash outcomes is easier to see.

From an experimental design perspective, it would have been ideal to keep things as simple in the dynamic decision as well, which would have meant a small number of discretized entry cost and scrap values. However, both for equilibrium existence and estimation, the model requires a continuum of scrap values and entry fees - see section 4. One consequence of this choice is that entry/exit decisions are more demanding on subjects. For example, consider the case when one subject is in the market and the other is out. For the A_S parameter, the MPE in Table 2 indicates that the subject should exit if the scrap value is 0.458 or higher. For a scrap value realization of 0.95 most subjects will quickly realize that exiting is worthwhile, but if the realization were 0.48, the difference in payoffs from exiting and staying in the market is rather small. In other words, the incentives to evaluate if a threshold of 0.47 is preferred to 0.49 are weak, and it seems reasonable to expect - ex ante - that subjects' choices will not exactly meet the theoretical predictions in entry and exit thresholds. We will refer to such differences as *optimization errors*.

It is possible for optimization errors to create a bias in the structurally recovered parameter estimates,[19] but the challenges with optimizing entry/exit thresholds are present in both baseline and counterfactual. Since the main focus of our exercise is to evaluate whether the incentives to collude affect counterfactual predictions (and not on whether subjects can optimally solve a dynamic programing problem), we introduce a second treatment variable to control for such discrepancies from the theoretical benchmark.

For each value of A, we conduct an additional treatment where agents *do not* make a quantity choice. The quantity-stage payoffs when both are in the market, are those prescribed by the unique Nash equilibrium. We refer to these treatments as *No Quantity Choice* and to treatments that do involve a quantity choice as *Standard* treatments. There can be deviations from the optimal policy in the No Quantity Choice treatments, but such errors are, by definition, unrelated to collusion incentives. Therefore, the question is whether errors in counterfactual predictions are higher in the Standard treatments than in the No Quantity Choice treatments. This comparison holds the rest of the dynamic environment and therefore the propensity for

[19] In subsection 5.2 we discuss a specific mechanism that can create a bias.

other optimization errors fixed.[20]

To summarize, we implement a 3×2 between-subjects experimental design, where we explore three different levels of A along one dimension and whether or not players can choose stage-game quantities along the other. Before describing our data, we will outline the structural estimation procedure.

4 Estimation Procedure

The first part gives a brief introduction to the estimation procedure. We also present results of a Monte Carlo study (this is an actual Monte Carlo study and we are not referring to the experiment here). The Monte Carlo serves two purposes. First, it demonstrates that the estimator indeed recovers the underlying parameters consistently if the model is correctly specified. Second, it helps to quantify the bias that results if data is coming from collusive interactions and the model is mis-specified.

4.1 Two Stage Estimator

We follow the *two-stage approach* (Aguirregabiria and Mira, 2007, Bajari et al., 2007, Pakes et al., 2007 and Pesendorfer and Schmidt-Dengler, 2008), which is computationally tractable and, under certain assumptions, does not suffer from the problem of multiple MPEs.

The procedure works under the assumption that the observed behavior is the result of an MPE. Under this assumption, a dataset containing the commonly observed payoff relevant states as well as the choices of firms (enter/exit), allows the researcher to observe the corresponding equilibrium policy function in the data.[21] The first stage of the procedure directly estimates the policy function, using non-parametric techniques such as kernel-density estimation or sieves.[22] In this case, it simply consists of four conditional choice probabilities ($\hat{p}(a|s)$)

[20]The characterization of the collusive equilibrium presented earlier is useful for our purpose as it describes a rationale for why using the Markov assumption at the estimation stage may lead to counterfactual prediction errors. In principle, it is possible that subjects in the Standard treatment are not following the symmetric MPE, but as long as the (possibly asymmetric) equilibrium that they follow does not depend on colluding at the quantity stage, such equilibrium is also available in the No Quantity Choice treatment.

[21]The data under consideration involves several sets of two firms, with each set interacting in a separate market under the rules of our theoretical model. For each set the econometrician observes whether the firm is in or out of the market in each period of time. From the econometric perspective cross-sectional and inter-temporal variation are equivalent.

[22]In our case, the procedure assumes that the econometrician knows the distribution functions for the private information terms. Specifically, in line with the empirical literature on dynamic games we assume that the econometrician knows the discount factor, $\delta = 0.8$, and that $\phi \sim \mathcal{U}[0, 1]$ and $\psi \sim \mathcal{U}[0, 1]$.

for the dynamic choice, one for each state. The second stage uses the estimated policy function to recover the structural parameters. Using the first-stage choice probabilities, one can invert the value function, $\hat{V}$. The inverted value function together with a parameter guess, $\hat{\theta} = \{\hat{A}, \hat{B}, \hat{C}\}$, can be used to obtain a set of predicted choice probabilities: $\Psi(a|s; \hat{V}, \hat{\theta})$. These predicted choice probabilities are then used for a simple moment estimator where $\hat{\mathbf{p}}$ is the vector of all choice probabilities, and $\mathbf{\Psi}(\theta)$ is the vector of respective predicted choice probabilities:[23]

$$\min_{\theta}(\hat{\mathbf{p}} - \mathbf{\Psi}(\theta))'\mathbf{W}(\hat{\mathbf{p}} - \mathbf{\Psi}(\theta))$$

Notice that the procedure only uses entry and exit decisions and does not rely on observing choices in the quantity stage. However, entry and exit decisions are affected by choices in the quantity stage. For example, firms that are colluding in quantities are more inclined to stay in the market, which becomes more valuable under collusion.

Some models treat the stage game as a function of only the number of firms in the market, and others explicitly use stage-game data. Here, we assume that the researcher does not have access to quantity data. However, since we do collect quantity-stage data, we will also report how estimates would differ if this additional information was taken into account.

4.2 Monte Carlo

The main purpose of the Monte Carlo simulation is to verify that the three parameters of interest (A, B and C) are consistently estimated if we assume the correct data generating process (i.e. the symmetric MPE). We also explore what estimates (bias) would result if, instead, the data is generated according to the CE identified earlier.[24]

The data-generating process matches the features of the data we collect in the laboratory, and details are presented in Appendix A.[25] The main message from the exercise can be sum-

[23] We use the identity matrix as a weighting matrix ($\mathbf{W}$).

[24] Notice that to the econometrician the CE on the equilibrium path looks like an MPE. Both for the MPE and for the CE, conditional on the state, choices are made according to probabilities that do not change in time. The punishment trigger will not be executed and there will be no "structural break" in the conditional choice probabilities. Specifically, recall that the econometrician only has access to data on entry and exit, and to recover the unknown structural parameters, the procedure assumes that the data is being generated from a symmetric MPE. If the MPE assumption holds, then agents condition the choices at t only on the payoff-relevant state at t. If the data was generated by collusive play and at some point one of the firms defects, then a portion of the data would follow CE probabilities (until the defection), and another portion of the data would follow MPE probabilities. In this case, with enough data, the MPE assumption can be shown to fail: agents would be conditioning on the state *and* on past behavior. But if there are no defections, along the equilibrium path play looks as in a symmetric MPE.

[25] We also quantify the bias in counterfactual predictions under collusion. We provide those results in subsection 5.3

marized in Figure 1, which focuses on the A_M case, but the findings for other A values are qualitatively unchanged. Consider the first observation, where the collusion rate is zero. In this case, the data is generated under the assumption of MPE play (no collusion in the market). The vertical axis presents the estimation for each of the three parameters of interest. As can be clearly seen, the estimation is very accurate, indeed recovering the true parameters when there is no collusion.

The same graphs also shows what happens when the data is instead generated under varying extends of collusive play. Different collusion rates on the x-axis correspond to different fractions of firm pairs that collude under the true DGP. Estimation, however, is conducted under the maintained assumption of no collusion. As collusion in the data goes up, we see that the estimates of A and C are entirely unaffected. However, the interaction parameter B becomes more downwards biased as the collusion rate increases.[26]

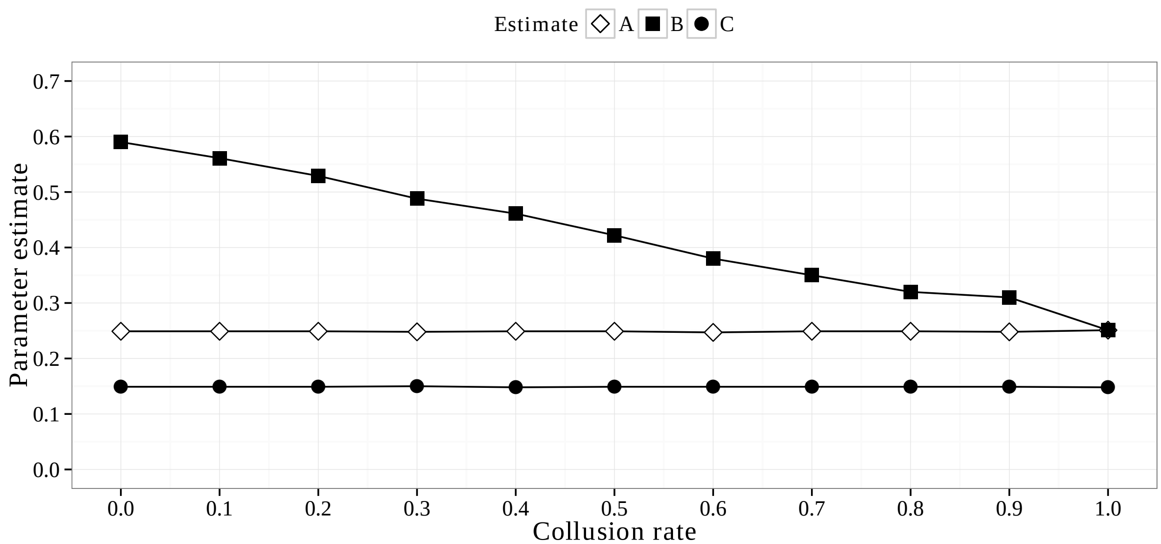

Figure 1: Parameter estimates under different collusion probabilities for the A_M treatment.

[26]Simple algebra shows that the lower bound for the estimate of B must be A since the difference in earnings in the market when the other player is in versus out ($2A - A = A$) is entirely attributed to B.

5 Results

5.1 Experimental Sessions

We conducted three sessions for each of our 6 treatments with subjects from the population of students at UC Santa Barbara. Subjects participated in only one session, and each session consisted of 14 participants.[27] The average participant received approximately $19 and all sessions lasted close to 90 minutes.[28]

In a given session, subjects will play several repetitions of the supergame. Subjects are randomly rematched with another subject in the room each time a new supergame starts. Repetitions of the supergame allow subjects to gain experience with the environment. In total there are 16 supergames per session.[29] Additional details regarding the implementation of the experiment are given in Appendix F.

5.2 Overview

We first document some basic patterns in the data and then proceed to the structural estimation and counterfactual prediction. We find that subjects' behavior is qualitatively close to the predictions, which indicates that subjects are reacting to the main tensions in the environment. Moreover, we also find that subjects respond to the collusion incentives in the predicted manner so that, in principle, it is possible that the structural estimates – under the incorrect assumption of an MPE in the data – are biased and lead to large errors in counterfactual predictions. However, we also document that successful collusion is rare and that punishments after unsuccessful collusion attempts are consistent with a reversion to the MPE. The frequent breakdown of collusion is reflected in low additional bias in counterfactuals due to collusion (section 5.3).

[27]Once participants entered the laboratory instructions were read by the experimenter (see Appendix G with instructions) and the session started. Subjects only interacted with each other via computer terminals and the code was written using zTree (Fischbacher, 2007). At the end of the session payoffs for all periods were added, multiplied by the exchange rate of 0.0025$ per point, and paid to subjects in cash. One session of the treatment with A_S-No Quantity Choice treatment had 12 participants.

[28]The environment the theory is trying to capture involves much larger stakes. While we cannot infer whether behavior in this setting is sensitive to the size of the stakes, see Camerer (2003) for numerous examples where increasing the stakes did not lead to changes in reported behavior.

[29]We observe 315 dynamic games per treatment. Since the length of these games is random this translates into a random number of interactions, where an interaction is defined as a tuple of state and dynamic choice for each player. In A_L we observe 1933 repeated interactions in the treatment with quantity stage choice and 1954 in the treatment without; in A_M it is 2220 with quantity stage choice and 2052 without; and in A_S it is 2080 and 1928 respectively.

In Figure 2, we provide an overview of the entry/exit choices by focusing on the aggregate frequencies, which constitute the central input of the *first stage* in the estimation routine. For each treatment, the white diamonds display the estimated frequency of being in the market next period (vertical axis) for each possible current state (horizontal axis).[30] We also represent 95% confidence intervals around the estimate and the theoretical MPE probabilities of Table 2, which are shown as black circles.

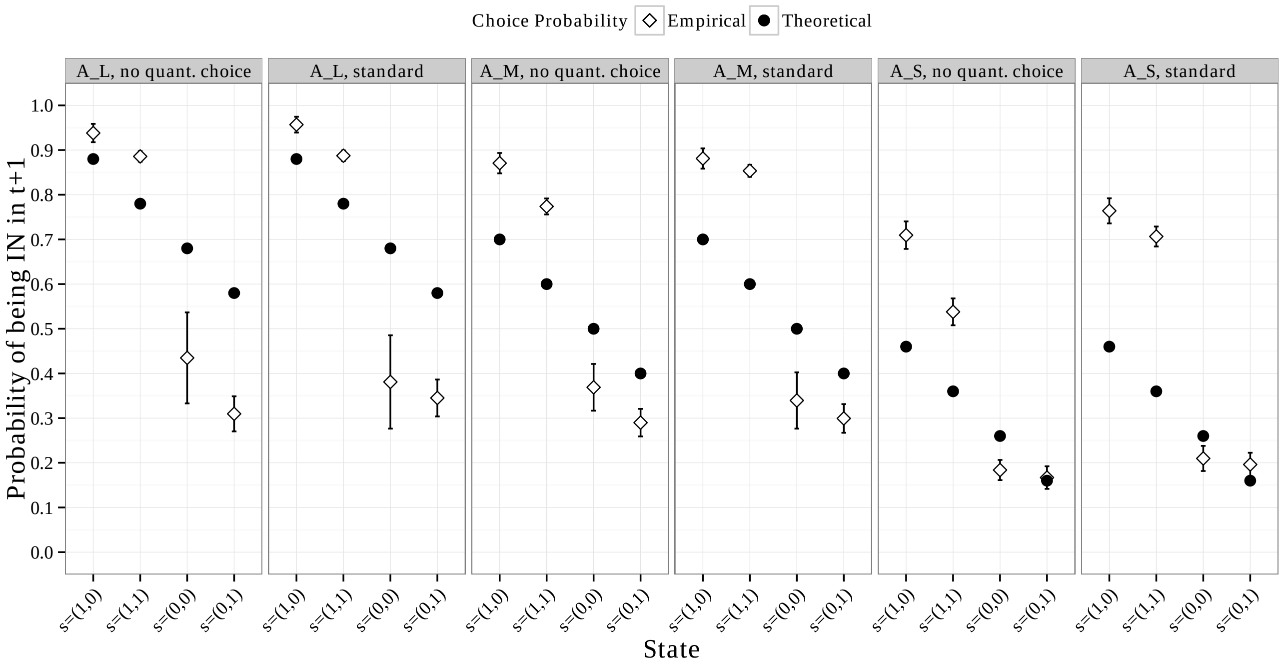

Figure 2: Probability of being in the market next period for each current state by treatment.

Support for Comparative Statics and Presence of Inertia

A first observation is that the data in all treatments can be organized fairly well by the MPE comparative statics. All 60 possible comparisons are in the predicted direction, and 50 are statistically significant at the 5% level or lower (2 more at the 10% level). The reader is referred to Appendix B and Appendix C for details on the comparisons. Evidence of behavior consistent with MPE comparative statics is not evidence of MPE play, because such comparative statics are not unique to the characterized MPE. Indeed, the CE shares several comparative statics. The observed comparative statics, therefore, do not allow us to determine which type

[30] Table 4 shows the first stage frequencies presented graphically in Figure 2.

of equilibrium better rationalizes choices. However, the evidence does indicate that subjects are responding to the economic incentives in a sensible manner.

While the data is in line with the MPE comparative statics, Figure 2 also shows that there is a quantitative deviation from the theoretic MPE probabilities. Subjects are more likely to stay in the market when they are already in (white diamonds are *above* the black circles) and less likely to enter if they are out (white diamonds are *below* the black circles). Relative to the prediction, subjects are demanding higher payoffs to leave and are willing to pay less to enter the market. This, in turn, means that subjects are more likely to remain in their current state than predicted by the MPE. We will refer to this phenomenon as subjects displaying *inertia* relative to the MPE.

Inertia can be a manifestation of what we earlier referred to as optimization errors. To see how, consider the case of a subject in the A_S treatment who is in the market while the other is out. According to the MPE, the subject should exit if the scrap value is higher than 0.458. Exiting the market for a high realization of the scrap value is not difficult to determine, but determining the lowest value at which to sell is more difficult. Given the challenge to compute the threshold, it is possible that subjects use exit thresholds that are more conservative than optimal, trying to avoid "selling" the company for a lower-than-optimal value. Likewise, subjects may use entry thresholds that are more conservative than the MPE, because they want to avoid paying a higher-than-optimal entry fee.[31] Under this rationale, the optimization errors are not centered around the MPE prediction but are systematically on one side of the predicted threshold.[32] It is possible that experience could reduce or eliminate inertia. However, we find that there is little to no change in inertia as the session evolves. Specifically, in Appendices B and C we present detail on choices as the sessions evolve, which allow to study if inertia changes with experience and more broadly possible learning effects. We document that there is little evidence of play moving in a systematic direction, suggesting that inertia might not simply disappear as subjects earn more experience with the environment.

Table 3 summarizes the absolute difference between theoretical MPE probabilities and empirical probabilities across treatments. We will refer to the absolute difference between MPE and empirical probabilities in the No Quantity Choice treatments as *inertia*. We distinguish

[31]This phenomenon is consistent with inertia described in experiments of choices under experience. The canonical example in such experiments is a decision problem in which subjects select between pressing button A or B, and know that each button will generate a payoff but are not told any details about the distributions generating payoffs. Instead, subjects can experiment and learn from experience the payoffs they receive when the click on each button. In this environment many subjects display inertia in the sense that they keep on pressing the same button even if recent observations suggest that switching may be preferable (see Erev and Haruvy, 2015 for a detailed exposition).

[32]In footnote 39 we argue that inertia cannot be generated by risk aversion.

those differences from the corresponding differences in Standard treatments, which might (in part) also be due to collusion. Average inertia (as measured in each treatment by the third column in Table 3) is comparable across No Quantity Choice treatments, ranging between 0.13 and 0.17. While the averages are similar, the presence of inertia in entry and exit thresholds changes with A. As A increases from A_S to A_L, inertia shifts from exit to entry thresholds.

Table 3: Differences between MPE and empirical probabilities

	A_L			A_M			A_S		
	Exit	Entry	Average	Exit	Entry	Average	Exit	Entry	Average
Standard	0.097	0.269	0.183	0.216	0.134	0.175	0.326	0.043	0.184
No Quantity Choice	0.082	0.260	0.171	0.171	0.123	0.147	0.216	0.041	0.128

Note: This table presents a summary of the differences between MPE and empirical probablilities. In each state we first compute the absolute difference between equilibrium MPE probabilities and empirical probabilities reported in Table 4. 'Exit' columns present the average in $p(1,0)$ and $p(1,1)$ (related to exit thresholds). 'Entry' columns present the average in $p(0,1)$ and $p(0,0)$ (related to entry thresholds). The average over all probabilities is presented in the 'Average' columns.

The presence of inertia will bias the structural estimates, but it is important to highlight that inertia is present in both Standard and No Quantity Choice treatments. Hence, it is not a phenomenon that is due to collusion.

Evidence of Short-lived Collusion

We can directly observe evidence for collusion by inspecting quantity-stage choices in Standard treatments. In the first period, all subjects start in the market, and their quantity-stage choices can be used as a measure of collusion attempts. The rate of *attempted collusion* refers to the proportion of subjects who selected the low quantity in the first period, while the rate of *successful collusion* captures the proportion of subjects who selected the low quantity and have a partner who also selected the low quantity in the first period.[33] The dark-shaded bars in the left panel display the average attempted and successful collusion rates by treatment (we also represent 95% confidence intervals around the estimate). Attempts to collude are highest for the treatments where collusion can be supported in equilibrium (A_S and A_M) and lowest in the treatment where collusion is not sub-game perfect (A_L). This indicates that, in the aggregate, subjects are responding to the incentives to collude.

[33] The measure of attempted collusion is often referred to in the experimental repeated-games literature as the first-period cooperation rate. The cooperation rate in later periods is endogenous, as it is affected by earlier choices within the supergame.

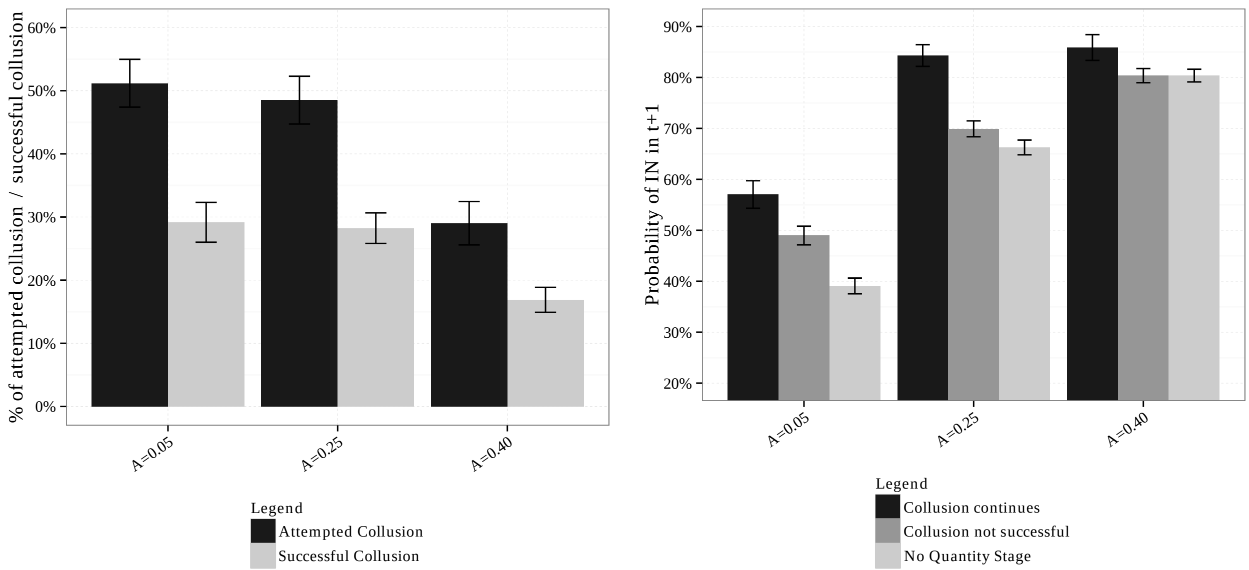

Figure 3: Collusion: Intentions, successes, and failures

For collusion to have an effect on structural estimates, it is necessary that the patterns of entry and exit are affected by quantity-stage choices. For each treatment, the left-most bar (black) of the right panel of Figure 3 shows the proportion of subjects who select to be in the market next period, conditional on both subjects colluding in the current period. The bar in the center (dark gray) shows the proportion of subjects who select to be in the market next period if at least one subject did not select the collusive quantity.[34] These differences, which are present in all treatments, are consistent with the CE: it shows that when collusion is successful it can have an effect on entry-exit choices, hence introducing a bias in the structural estimates.

However, to understand how much this can bias the estimates, it is crucial to determine how often collusive attempts are successful. A first observation from the left panel of Figure 3 is that even in the treatments where collusion attempts are highest, approximately 50% of subjects make choices that are consistent with the stage-Nash equilibrium and hence are not trying to collude. If, however, in the first period one of the two subjects does not collude in the

[34] The proportion of subjects who select to be in the market next period is a measure for the probability of being in the market next period. The figure for the A_S treatment is lower than for other treatments because in the A_S treatment the outside option is relatively more attractive than in other treatments. Notice, in addition, that the bar in the center (at least one of the subjects did not collude in the quantity stage) is of comparable magnitude to the right-most bar that represents the corresponding no-quantity-choice treatment, where by definition no subject can collude. The figure shows a difference only in the case of the A_S treatment.

quantity stage the belief that collusion will take place in the future is likely lower.[35] The rate of *successful collusion* is presented in the left panel of Figure 3 in light gray. The figure shows that slightly more than a quarter of subjects succeed in colluding in the first period. In other words, in treatments where collusion can be supported in equilibrium, about three-quarters of subjects do not experience a first-period outcome that would foster a belief of collusion for future periods. Finally, comparing the two right-most bars in the right panel of Figure 3, we observe that the dynamic choice after unsuccessful collusion in standard treatments is close to the dynamic choice in the no-quantity choice treatments.

The previous observations are representative of broader patterns in supergame choices. A formal study of quantity-stage choices throughout the supergame is presented in Appendix C, where we conclude that successful collusion represents approximately 12% of all choices. In cases where collusion does not succeed, the analysis shows that a large proportion of choices are consistent with the punishments of the CE (using the stage-Nash equilibrium). Overall, the analysis indicates that while a large proportion of subjects intends to collude, and while successful collusion attempts can have an impact on entry-exit decisions, there are relatively few successful cases. In addition, while there is evidence that quantity-stage collusion can affect thresholds, most differences are small and not statistically significant.

5.3 Structural Results

Parameter Estimates

The structural parameter estimates are reported in Table 4.[36] We consider the estimates of A first. The estimates are below the true parameters in all treatments. However, for both the Standard and the No Quantity Choice treatments we find that $\hat{A}_L > \hat{A}_M > \hat{A}_S$. Moreover, for a fixed true value of A, the estimates of the Standard and the No Quantity Choice treatments are relatively close to each other. For instance, for A_M the estimates are $\hat{A} = 0.14$ and $\hat{A} = 0.17$ for the Standard and No Quantity Choice treatments, respectively.[37]

[35] Using the Monte Carlo exercise reported in Figure 1 as a reference, the attempted collusion rates indicate that the bias in parameter B would be considerably below the maximum possible bias. Still, if 50% of subjects *successfully* collude, the bias can be substantial. But their expectations will likely be lower than 50% as the rate of *successful collusion* is lower.

[36] The estimates in Table 4 use data from all supergames in each session. In Appendix D we present estimates constraining the number of supergames included and show that the estimates are robust to changes in the sample. Only in the A_L-Standard treatment do we observe an increase in A once we restrict data to the last eight supergames.

[37] Regarding the precision of the estimates, Table 4 shows that A is estimated with small standard errors in all treatments except for the A_L-No Quantity Choice treatment.

Table 4: Estimated and theoretical entry/exit probabilities for each treatment along with parameter estimates.

	Standard				No Quantity Choice			
State (s_i, s_{-i})	$(1,0)$	$(1,1)$	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$	$(0,0)$	$(0,1)$
$\mathbf{A = 0.4, B = 0.6, C = 0.15}$								
MPE	0.880	0.781	0.681	0.583				
CE	0.925	0.870	0.757	0.702				
Empirical	0.967	0.887	0.381	0.345	0.938	0.886	0.435	0.310
Parameter	A	B	C		A	B	C	
Estimates	0.18	0.11	0.54		0.22	0.22	0.56	
Std. Err	0.01	0.03	0.03		0.05	0.13	0.02	
$\mathbf{A = 0.25, B = 0.6, C = 0.15}$								
MPE	0.701	0.602	0.502	0.403				
CE	0.768	0.737	0.612	0.580				
Empirical	0.881	0.854	0.339	0.299	0.871	0.774	0.369	0.290
Parameter	A	B	C		A	B	C	
Estimates	0.14	0.05	0.55		0.17	0.19	0.47	
Std. Err	0.01	0.03	0.02		0.01	0.04	0.02	
$\mathbf{A = 0.05, B = 0.6, C = 0.15}$								
MPE	0.459	0.360	0.260	0.161				
CE	0.519	0.512	0.368	0.362				
Empirical	0.764	0.707	0.210	0.196	0.710	0.540	0.184	0.166
Parameter	A	B	C		A	B	C	
Estimates	0.10	0.07	0.53		0.08	0.20	0.43	
Std. Err	0.01	0.03	0.02		0.01	0.04	0.02	

Note: The table provides an overview over all estimates (standard errors) as well as first stage probabilities (Empirical), the latter of which can be compared to the theoretical MPE and Collusive probabilities. Theoretical probabilities coincide for Standard and No Quantity Choice treatments. Results are shown for each of the six treatments with different market sizes (A) and conditional on whether there is a quantity choice or not.

The entry cost C is on average estimated to be 3.4 times higher than the true value, and this bias is present in all treatments. Both Standard and No Quantity Choice treatments show higher estimates for C of a comparable magnitude.

Regarding the estimates of B, we make three main observations. First, the estimates range from 0.05 to 0.22 and are well below the true value of 0.6 in all treatments. Second, the estimates are lower in the Standard treatments than in the No Quantity Choice treatments.[38] For a fixed value of A, the estimate for the No Quantity Choice treatment at least doubles that of the Standard treatment. Third, the estimates in all treatments with No Quantity Choice are quite close to each other, but in the Standard treatments the estimates are lower for A_S and A_M.

In all treatments we therefore report structural estimates that are quantitatively far from the true values. We find two main sources for the differences. First, the presence of inertia in entry/exit thresholds can, in principle, introduce a bias in all coefficients. The second source is collusion: the estimates of B are further downwards biased in Standard relative to No Quantity Choice treatments.[39]

Inertia

In Appendix E, we document the effects of inertia in detail. Specifically, in Section E.1., we present three procedures that modify the structural model to explicitly account for inertia. We prefer the third procedure, which we call myopic inertia-augmented model, because it recovers estimates that are much closer to their values and in many cases significantly reduces counterfactual bias. The model involves estimating a fourth parameter, which can be interpreted as a perceived cost for either entering or exiting the market. Inertia in this model is myopic in the sense that the agent does not expect to pay this cost for entering and exiting in future periods.[40] The myopic inertia-augmented model recovers a parameter of B closer to the true value and a lower estimate for C. These findings are consistent with Monte Carlo

[38] We have implemented a non-parametric test based on a bootstrap procedure to determine the significance of those differences. In the A_L treatment those differences are not significant, but for the other two market sizes one can reject at the 5% level that the competition parameters are coming from the same distribution.

[39] Our estimations and equilibrium computations are based on the assumption of risk neutrality. However, allowing for risk aversion cannot rationalize the data. We numerically computed equilibria under risk aversion and it moves the probability that a firm wants to be in the market next period up in all four states. The intuition is the following. There are two ways to earn money in the market. The first one is by staying in the market and collecting rents and the second is to earn scrap value by entering the market for low cost and exiting for high resale values. But the latter way to earn profits is more "risky" due to the randomness of entry cost and scrap values. Players with higher risk aversion therefore want to stay in the market more often.

[40] A possible justification behind this extension is that the agent considers a realized state differently than a potential state that can happen in the future. That is, if the agent is in the market, the agent has an "extra" valuation for the current state because that is what she has now. But when she considers the future, she doesn't attach an "extra" valuation to being in the market tomorrow because that state has not yet materialized.

simulations that we report in Section E.2. Using simulations, we show that inertia always biases the estimates of C upwards, and that the estimates of A and B can be biased upwards or downwards. The pattern of entry and exit that would bias the estimate of B upwards is not present in our data, and consistently, we do not observe an upwards bias in any estimate of B reported in Table 4. However, A is biased downwards in A_L and A_M but upwards in A_S. The Monte Carlo simulations illustrate that the upwards bias in $\hat{A}$ can result when inertia is mostly present in exit thresholds, as we observe in A_S.

While inertia is present in all treatments, the second source of bias (collusion) affects Standard and No Quantity Choice treatments differently. A question then is to what extent the counterfactual bias from inertia will interact with the one from collusion. Monte Carlo simulations with inertia *and* collusion (Section E.2 of Appendix E) show that those two biases do not interact.[41] In other words, the simulations indicate that, even with inertia, the bias from collusion will still only show up in B. This insight will be important for discussing the results of the counterfactual bias due to collusion.

Correcting for Collusion by using Empirical Market Quantity Choice

The estimation so far assumed that the econometrician has no information on quantity choices. As was pointed out, deviations from Nash stage-game behavior can only occur as part of an equilibrium with dynamic punishment strategies. As the empirical analysis of quantity choices has revealed, collusion often breaks down, which implies a distribution of quantity choices that is in between the fully collusive outcome and static Nash. Under the assumption that subjects hold correct beliefs about implied quantity choices, we can adjust for actual market behavior by imputing the observed frequencies. Our Monte Carlo simulation shows that such an adjustment will only affect the estimate of B. For A_L, this adjustment increases B from 0.11 to 0.15, for A_M from 0.05 to 0.09, and for A_S from 0.07 to 0.13. Using the available quantity choice data, one can therefore move the estimates of B closer to the true value. Consistent with the equilibrium prediction, these adjustment are in percentage terms much larger in the A_M and A_S treatments. However, even after adjustment, the estimates of B are far from their true value, which mirrors the findings from the No Quantity Choice treatments and foreshadows our discussion of the effect of collusion on counterfactual predictions.

[41] Figure 9 reproduces the analysis of Figure 1 for the A_M treatment adding inertia, and we document the same pattern. As the collusion rate increases, the estimate of B is biased downwards, but the estimates of A and C are unaffected.

Counterfactual Calculations

We now come to the main exercise, which is to use the recovered parameters to predict behavior in another treatment, and then compare predicted behavior with actual laboratory behavior in those treatments. Unlike in typical applications with observational data, we observe each treatment and therefore each counterfactual scenario. This means that we can estimate the parameters for each treatment and also run the counterfactual for the respective remaining treatments. In total, this amounts to six different counterfactuals for the Standard case and six different counterfactuals for the No Quantity Choice case.

We use the term *baseline treatment* for the treatment that provides the estimated parameters to predict behavior elsewhere. To obtain the counterfactual parameters, we scale the recovered market size parameter by the factor that would make the true value of A in the baseline equal to the counterfactual true value of A. For example, for the A_L-No Quantity Choice treatment, we recovered $\hat{A} = 0.22$. If we want to predict behavior in the case of A_M, we scale the estimated parameter by $5/8$ (0.25/0.4). The other two parameters, $\hat{B}$ and $\hat{C}$, are kept constant. The results are presented in Table 5.[42]

Table 5: Counterfactual predictions

Baseline	Prediction A_L		Prediction A_M		Prediction A_S	
	MAE($\mathbf{p}$)	MAPE($\mathbf{V}$)	MAE($\mathbf{p}$)	MAPE($\mathbf{V}$)	MAE($\mathbf{p}$)	MAPE($\mathbf{V}$)
A_L						
Standard	-	-	0.11	37.8%	0.20	68.5%
No Quantity Choice	-	-	0.08	40.8%	0.13	64.9%
A_M						
Standard	0.13	82.3%	-	-	0.18	65.1%
No Quantity Choice	0.12	84.3%	-	-	0.10	50.9%
A_S						
Standard	0.36	949.1%	0.41	621.8%	-	-
No Quantity Choice	0.36	716.2%	0.32	326.4%	-	-

Note: The first column indicates the baseline treatment and subsequent columns the counterfactual. MAE($\mathbf{p}$) reports the mean absolute error in the prediction of probabilities. MAPE($\mathbf{V}$) reports the mean absolute percentage error in the prediction of continuation values.

For each of the six counterfactuals, Table 5 reports two measures. The first measure captures errors in the predicted probabilities of being in the market next period. The mean abso-

[42]More detailed results, conditional on states, are shown in Appendix D.

lute error in probabilities (MAE($\mathbf{p}$)) is the absolute difference between the actual and predicted probabilities, averaged across states.[43] Our preferred measure captures the mean absolute percentage error in firm values (MAPE($\mathbf{V}$); see Equation 6 in Appendix A for the expression used to compute the continuation values (V(s))). It is the absolute percentage difference between actual and predicted firm values, averaged across states.

Overall, we find that collusion does not increase counterfactual prediction errors substantially. When the baseline is A_L and there is No Quantity Choice, the MAPE($\mathbf{V}$) for the A_M counterfactual is 40.8%, which is close to 37.8% in the Standard treatment. If the prediction error in the No Quantity Choice treatment had been substantially smaller than in the Standard treatment, it would have meant that collusion is a driver of prediction errors. Instead, we observe a prediction error of similar magnitude. When the counterfactual treatment is A_S, there is an increase in the MAPE($\mathbf{V}$), but again there are only small differences between treatments with and without quantity choice. The same pattern holds for counterfactuals where A_M is the baseline.

Lastly, Table 5 shows that when the baseline is A_S, prediction errors are substantially larger, and it appears as if the bias shrinks (by about 25% and 48%) if we don't allow for collusion. In this case, however, our counterfactual exercise is misleading, due to extreme prediction that the model makes. When we use the $\hat{A}_S$ estimate to obtain a counterfactual value of A, the resulting value is above the actual estimate. For example, we compute the counterfactual for the A_L-Standard treatment as: $\hat{A}_L^{A_S} = \frac{0.4}{0.05}\hat{A}_S = 0.64 > \hat{A}_L = 0.18$. In other words, the prediction for the A_L parameter, using A_S as the baseline (0.64), is 3.5 times the best estimate that we have from behavior in the A_L-Standard treatment (0.18). The market value in the counterfactual (0.64) is in fact so large, that being in the market becomes an absorbing state. This is shown in the last set of graphs in Figure 10 (Appendix E).[44] When A_L or A_M are used as baseline treatments, it is also the case that the counterfactual predictions of A are far from the estimated values. However, such distortions do not predict an absorbing state.[45]

To make it easier to judge whether the counterfactual biases due to collusion are large

[43]The table presents the simple average across states. It is also possible to weigh states depending on how frequently they are visited. Since the frequency of visits to a state depends on the treatment, there are two possible weights: using the baseline or the counterfactual weights. Qualitatively using either weights would not change the findings we report.

[44]Being in the market is an absorbing state if $p(1,0) = p(1,1) = 1$. With these probabilities if the subject is ever predicted to be in the market, she will not leave. The same qualitative outcome happens when A_S is the baseline and A_M is the counterfactual treatment. For the Standard treatments, for example, $\hat{A}_M^{A_S} = 0.4 > \hat{A}_L = 0.14$.

[45]For example, consider the Standard treatments when A_L is the baseline and A_S the counterfactual. In this case, $\hat{A}_S^{A_L} = \frac{0.05}{0.4}\hat{A}_L = 0.023 < 0.10 = \hat{A}_S$. The difference between the predicted value (0.023) and the actual estimate (0.10) is large and it does introduce prediction errors, but there is no absorbing state in the counterfactual as the first row of Figure 10 shows.

or small we now provide a benchmark. The benchmark makes again use of Monte Carlo results, where we now compute the maximal bias in counterfactuals due to collusion.[46] This benchmark is computed under the assumption that there is no other confounding deviation from the Markov perfect equilibrium besides collusion and we refer to this benchmark as *Maximal*.

To allow for a direct comparison with the counterfactual bias resulting from our estimates we therefore take the following approach. We conduct a second set of counterfactual predictions where we assume the econometrician knows the counterfactual entry cost $\hat{C}$ and the market size parameter $\hat{A}$ but uses the estimate of B from the baseline treatment. To understand our approach for the comparison it is important to keep in mind that collusion only biases B and that other parameters are still unbiased, which is is true even under inertia (Figure 9 in Appendix E). In other words, the exercise "controls" for confounding variation in the data that leads to biases in parameters other than B. While in the the No Quantity Choice treatments the bias in $\hat{B}$ can not stem from collusion the $\hat{B}$ in Standard treatments might be biased both because of collusion and other confounding factors. In particular, inertia will also lead to downwards bias in B. Therefore, the comparison of counterfactual bias in Standard and No Quantity Choice treatments isolates the portion of the bias due to collusion. The computations are summarized in Table 6.[47]

The main message from this second exercise is consistent with the findings we documented earlier, as the prediction errors due to collusion are rather small. The two highest prediction errors reported in Table 6 (in terms of MAPE($\mathbf{V}$)) take place when A_M-Standard is the baseline to predict A_L (24.0%) and when A_L-Standard is the baseline to predict A_M (18.1%). Net of the counterfactual errors in the No Quantity Choice cases (9.3 and 6.5%, respectively), the proxy for the prediction errors due to collusion are 14.7 and 11.6%. We then compute the ratio between this difference and *Maximal*. We report those ratios in Table 7 and range from 0 to 33%. That is, in the worst case, the observed bias due to collusion is a third of the bias in the *Maximal* simulations. In most cases the ratio is substantially lower.

[46] As a reminder, this maximal bias would result if the econometrician recovers data from a baseline without collusion and then predicts behavior for a counterfactual with full collusion but under the maintained assumption of markov play, which means the model in the counterfactual is mis-specified.

[47] Table 26 in Appendix E presents the predicted probabilities. In Appendix E we also report other counterfactual predictions, such as the predicted probability of observing a monopoly and a duopoly; see Table 26.

Table 6: Counterfactual predictions based only on bias in $\hat{B}$

Baseline	Prediction A_L		Prediction A_M		Prediction A_S	
	$MAE(\mathbf{p})$	MAPE($\mathbf{V}$)	$MAE(\mathbf{p})$	MAPE($\mathbf{V}$)	$MAE(\mathbf{p})$	MAPE($\mathbf{V}$)
A_L						
Standard	-	-	0.04	18.1%	0.02	11.6%
No Quantity Choice	-	-	0.01	6.5%	0.04	4.7%
A_M						
Standard	0.05	24.0%	-	-	0.01	6.5%
No Quantity Choice	0.03	9.3%	-	-	0.04	0.8%
A_S						
Standard	0.04	5.8%	0.02	6.5%	-	-
No Quantity Choice	0.03	1.2%	0.00	6.1%	-	-

Note: The first column indicates the baseline treatment and subsequent columns the counterfactual. $MAE(\mathbf{p})$ reports the mean absolute error in the prediction of probabilities. MAPE($\mathbf{V}$) reports the mean absolute percentage error in the prediction of continuation values.

To sum up, we have documented that in most cases the counterfactual bias in the Standard treatment and the No Quantity Choice treatment are close. The cases where they are not close entail an extreme prediction that the firm never wants to leave the market. Compared to the *Maximal* benchmark, the bias due to collusion is small. After correcting for inertia, one never mis-predicts the value of the firm by more than 24%. This mis-prediction amounts to, at most, 33% of the possible bias and is often well below this.

In Appendix D we report on a different robustness exercise to test for the impact of collusion on counterfactual prediction errors. We use the parameters we estimate for the Standard treatments to predict behavior in the No Quantity Choice treatments. This can be interpreted as predicting behavior in a new market that has the same primitives but might have a different conduct. We then contrast this prediction to actual behavior in No Quantity Choice treatments. If collusion were large, the prediction error would be large as well. With this exercise we essentially reproduce our main findings. The observed counterfactual bias due to collusion represents a relatively small portion of the maximal counterfactual prediction error due to collusion. Finally, in a related exercise, we use a model where inertia is captured directly in the estimation. With this *inertia-augmented* model (Appendix E) we reproduce the counterfactual prediction exercise. We find that the levels of the counterfactual prediction errors are lower once we control for inertia and that the qualitative conclusions from comparing Standard and No Quantity Choice treatments are similar.

Baseline/Counterfactual	A_L	A_M	A_S
	$\frac{\Delta\text{MAPE}(\mathbf{V})}{\text{Maximal MAPE}(\mathbf{V})}$		
A_L	-	33%	12%
A_M	24%	-	5%
A_S	4%	0%	-

Table 7: Observed relative to maximal counterfactual error from collusion

For each comparison, this table reports the ratio in which the numerator, ΔMAPE($\mathbf{V}$), is the difference between Standard and No Quantity Choice MAPE($\mathbf{V}$) computed using Table 6. The denominator is the Maximal MAPE($\mathbf{V}$) that can result from collusion, which is reported in Table 21 of Appendix A.

6 Conclusion

The assumption of Markov play for dynamic oligopoly games is commonplace in applications, but it is often not possible to test its validity based on only observational data. The laboratory provides an environment where we can systematically evaluate conditions under which the restriction to Markov perfect equilibria (MPE) may introduce biases in estimation and counterfactual predictions. One possible source of these biases is collusion if agents condition their behavior on past play and reach sub-game perfect equilibria that are not MPEs. In this paper, we explore the resulting bias in counterfactual predictions if collusion is ignored. We take a simplified version of Ericson and Pakes (1995) to the laboratory and construct a series of treatments for which we characterize an MPE but where it is also possible for a collusive equilibrium (CE) to emerge. The estimator, however, assumes that the data is coming from a Markovian equilibrium, which rules out collusion. We use the lab to test how strong this restriction is.

Our experimental exercise provides several insights. First, the MPE prediction for the quantity stage is often wrong. We find that a large proportion of subjects intend to collude, particularly when the incentives are higher. If cooperation were successful, there would be large biases in the estimators and large prediction errors in counterfactuals due to the assumption of Markov play. However, we also document that cooperation often breaks down and that successful cooperative attempts are relatively rare.[48] As a result, we find that the structural parameter affected by collusion (B) is more biased in treatments where the incentives to collude

[48] As highlighted in the Introduction, the evidence that tacit collusion breaks down is not exclusive to the laboratory but consistent with evidence from the field.

are higher. The central question, however, is whether the extra bias leads to large prediction errors. Our results suggest that this is not the case. The prediction errors that we can attribute to collusion are far from the potential bias that Monte Carlo simulations suggest.

In studying the possible prediction errors introduced by collusion, we uncover a different, yet systematic, deviation from theoretical predictions in all treatments. We refer to such deviation as inertia and show that it arises even when collusion is not possible. The bias that can be attributed to collusion pales in comparison to what can be attributed to inertia. Our design is not equipped to uncover the mechanisms that lead to this behavior. However, we explore a number of alternatives to rationalize inertia, and we found that a model in which players myopically perceive an additional switching cost best explains the data. Given its prominent presence in the data, it is a feature that needs further exploration.

To the best of our knowledge, this is the first paper that uses the laboratory to study the relevance of the Markov restriction for the estimation of dynamic games and counterfactual computations. Our paper illustrates how laboratory methods can be used to substantiate behavioral assumptions that are required for structural estimation. Further experimental research can help to better understand if the quantitative deviations from the MPE that we document are a feature of our environment or if they are present in other settings as well. Experimental methods may be especially attractive to tackle the problem of equilibrium multiplicity in counterfactuals. In an experiment, the researcher has control over model specification and can observe not only the parameters, but also true counterfactual behavior as implemented by experimental treatments.

References

Abito, Jose Miguel, "How much can we identify from repeated games?," 2015.

Ackerberg, Daniel, C Lanier Benkard, Steven Berry, and Ariel Pakes, "Econometric tools for analyzing market outcomes," *Handbook of econometrics*, 2007, *6*, 4171–4276.

Aguirregabiria, Victor and Arvind Magesan, "Identification and estimation of dynamic games when players' beliefs are not in equilibrium," 2016.

_ **and Chun-Yu Ho**, "A dynamic oligopoly game of the US airline industry: Estimation and policy experiments," *Journal of Econometrics*, 2012, *168* (1), 156–173.

_ **and Pedro Mira**, "Sequential estimation of dynamic discrete games," *Econometrica*, 2007, *75* (1), 1–53.

_ **and** _ , "Dynamic discrete choice structural models: A survey," *Journal of Econometrics*, 2010, *156* (1), 38–67.

Arcidiacono, Peter, Karthik Muralidharan, Eun young Shim, and John D Singleton, "Valuing School Choice: Using a Randomized Experiment to Validate Welfare Evaluation of Private School Vouchers," 2016.

Bajari, Patrick and Ali Hortacsu, "Are Structural Estimates of Auction Models Reasonable? Evidence from Experimental Data.," *Journal of Political Economy*, 2005, *113* (4), 703.

_ **, C Lanier Benkard, and Jonathan Levin**, "Estimating dynamic models of imperfect competition," *Econometrica*, 2007, *75* (5), 1331–1370.

Battaglini, Marco, Salvatore Nunnari, and Thomas R Palfrey, "Legislative bargaining and the dynamics of public investment," *American Political Science Review*, 2012, *106* (02), 407–429.

_ **,** _ **, and** _ , "The Dynamic Free Rider Problem: A Laboratory Study," *mimeo*, 2015.

Blonigen, Bruce A, Christopher R Knittel, and Anson Soderbery, "Keeping it fresh: Strategic product redesigns and welfare," *National Bureau of Economic Research Working Paper*, 2013.

Borkovsky, Ron N, Paul B Ellickson, Brett R Gordon, Victor Aguirregabiria, Pedro Gardete, Paul Grieco, Todd Gureckis, Teck-Hua Ho, Laurent Mathevet, and Andrew Sweeting, "Multiplicity of equilibria and information structures in empirical games: challenges and prospects," *Marketing Letters*, 2014, pp. 1–11.

Bresnahan, Timothy F and Peter C Reiss, "Entry in monopoly market," *The Review of Economic Studies*, 1990, *57* (4), 531–553.

Brown, Meta, Christopher J Flinn, and Andrew Schotter, "Real-time search in the laboratory and the market," *The American Economic Review*, 2011, *101* (2), 948–974.

Camerer, Colin, *Behavioral game theory: Experiments in strategic interaction*, Princeton University Press, 2003.

Collard-Wexler, Allan, "Demand Fluctuations in the Ready-Mix Concrete Industry," *Econometrica*, 2013, *81* (3), 1003–1037.

Conlon, Christopher T and Julie Holland Mortimer, "Effects of product availability: Experimental evidence," Technical Report, National Bureau of Economic Research 2010.

_ **and** _ , "All Units Discount: Experimental Evidence from the Vending Industry," 2013.

Dal Bó, Pedro and Guillaume R Fréchette, "The evolution of cooperation in infinitely repeated games: Experimental evidence," *The American Economic Review*, 2011, *101* (1), 411–429.

_ **and** _ , "On the determinants of cooperation in infinitely repeated games: A survey," *Journal of Economic Literature*, 2018, *56* (1), 60–114.

Davis, Douglas D and Charles A Holt, "The Effects of Collusion in Laboratory Experiments," *in Handbook of Experimental Economics Results, North Holland Press*, 2008.

De Paula, Aureo, "Econometric analysis of games with multiple equilibria," *Annu. Rev. Econ.*, 2013, *5* (1), 107–131.

Doraszelski, Ulrich and Ariel Pakes, "A framework for applied dynamic analysis in IO," *Handbook of industrial organization*, 2007, *3*, 1887–1966.

— **and Mark Satterthwaite**, “Computable Markov-perfect industry dynamics,” *The RAND Journal of Economics*, 2010, *41* (2), 215–243.

Dutta, Prajit K, “A folk theorem for stochastic games,” *Journal of Economic Theory*, 1995, *66* (1), 1–32.

Erev, Ido and Ernan Haruvy, “Learning and the economics of small decisions,” *The handbook of experimental economics*, 2015, *2.*

Ericson, Richard and Ariel Pakes, “Markov-perfect industry dynamics: A framework for empirical work,” *The Review of Economic Studies*, 1995, *62* (1), 53–82.

Ertaç, Seda, Ali Hortaçsu, and James W Roberts, “Entry into auctions: An experimental analysis,” *International Journal of Industrial Organization*, 2011, *29* (2), 168–178.

Fershtman, Chaim and Ariel Pakes, “A dynamic oligopoly with collusion and price wars,” *The RAND Journal of Economics*, 2000, pp. 207–236.

Fischbacher, Urs, “z-Tree: Zurich toolbox for ready-made economic experiments,” *Experimental economics*, 2007, *10* (2), 171–178.

Fréchette, Guillaume R and Sevgi Yuksel, “Infinitely Repeated Games in the Laboratory: Four Perspectives on Discounting and Random Termination,” *mimeo*, 2013.

— **, John H Kagel, and Massimo Morelli**, “Behavioral identification in coalitional bargaining: An experimental analysis of demand bargaining and alternating offers,” *Econometrica*, 2005, *73* (6), 1893–1937.

Fudenberg, Drew, David G Rand, and Anna Dreber, “Slow to anger and fast to forgive: cooperation in an uncertain world,” *The American Economic Review*, 2012, *102* (2), 720–749.

Goettler, Ronald L and Brett R Gordon, “Does AMD spur Intel to innovate more?,” *Journal of Political Economy*, 2011, *119* (6), 1141–1200.

Harrington, Joseph E, “Detecting cartels,” *Handbook of antitrust economics*, 2008, *213*, 245.

— **and Andrzej Skrzypacz**, “Private monitoring and communication in cartels: Explaining recent collusive practices,” *The American Economic Review*, 2011, *101* (6), 2425–2449.

Hey, J.D. and V. Dardanoni, "Optimal consumption under uncertainty: An experimental investigation," *The Economic Journal*, 1988, *98* (390), 105–116.

Hotz, V Joseph and Robert A Miller, "Conditional choice probabilities and the estimation of dynamic models," *The Review of Economic Studies*, 1993, *60* (3), 497–529.

Houser, Daniel, Michael Keane, and Kevin McCabe, "Behavior in a dynamic decision problem: An analysis of experimental evidence using a Bayesian type classification algorithm," *Econometrica*, 2004, *72* (3), 781–822.

Keniston, D, "Experimental vs. Structural Estimates of the Return to Capital in Microenterprises," *Unpublished manuscript, Yale University*, 2011.

Kloosterman, A., "An Experimental Study of Public Information in Markov Games," *mimeo*, 2018.

Lei, V. and C.N. Noussair, "An experimental test of an optimal growth model," *American Economic Review*, 2002, pp. 549–570.

Mailath, George J and Larry Samuelson, *Repeated games and reputations*, Vol. 2, Oxford university press Oxford, 2006.

Marshall, Robert and Leslie Marx, "The Economics of Collusion," 2012.

Marshall, Robert C, Leslie M Marx, and Matthew E Raiff, "Cartel price announcements: The vitamins industry," *International Journal of Industrial Organization*, 2008, *26* (3), 762–802.

Maskin, Eric and Jean Tirole, "Markov perfect equilibrium: I. Observable actions," *Journal of Economic Theory*, 2001, *100* (2), 191–219.

Merlo, Antonio and Thomas R Palfrey, "External validation of voter turnout models by concealed parameter recovery," *mimeo*, 2013.

Noussair, C. and K. Matheny, "An experimental study of decisions in dynamic optimization problems," *Economic Theory*, 2000, *15* (2), 389–419.

Pakes, Ariel, "Methodological Issues in Analyzing Market Dynamics," *mimeo*, 2015.

— , **Michael Ostrovsky, and Steven Berry**, "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," *The RAND Journal of Economics*, 2007, *38* (2), 373–399.

Pathak, Parag A and Peng Shi, "How Well Do Structural Demand Models Work? Counterfactual Predictions in School Choice," Technical Report, National Bureau of Economic Research 2017.

Pesendorfer, Martin and Philipp Schmidt-Dengler, "Asymptotic least squares estimators for dynamic games," *The Review of Economic Studies*, 2008, *75* (3), 901–928.

Porter, Robert H, "A study of cartel stability: The Joint Executive Committee, 1880-1886," *The Bell Journal of Economics*, 1983, pp. 301–314.

Roth, Alvin E and J Keith Murnighan, "Equilibrium behavior and repeated play of the prisoner's dilemma," *Journal of Mathematical psychology*, 1978, *17* (2), 189–198.

Ryan, Stephen P, "The costs of environmental regulation in a concentrated industry," *Econometrica*, 2012, *80* (3), 1019–1061.

Saijo, T., K. Sherstyuk, N. Tarui, and M. Ravago, "Games with Dynamic Externalities and Climate Change Experiments," *mimeo*, 2014.

Schmidt-Dengler, Philipp, "The Timing of New Technology Adoption: The Case of MRI," *mimeo*, 2006.

Sweeting, Andrew, "Dynamic product positioning in differentiated product markets: The effect of fees for musical performance rights on the commercial radio industry," *Econometrica*, 2013, *81* (5), 1763–1803.

Vespa, Emanuel, "An Experimental Investigation of Strategies in the Dynamic Common Pool Game," *mimeo*, 2019.

— **and Alistair J Wilson**, "Experimenting with Equilibrium Selection in Dynamic Games," *mimeo*, 2017.

— **and** — , "Experimenting with the Transition Rule in Dynamic Games," *mimeo*, 2018.