MIT Open Access Articles

Estimating dynamic games of oligopolistic competition: an experimental investigation

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Salz, Tobias and Emanuel Vespa. "Estimating dynamic games of oligopolistic competition: an experimental investigation." RAND Journal of Economics 51, 2 (June 2020): 447-469. © 2020 The RAND Corporation.

As Published: <http://dx.doi.org/10.1111/1756-2171.12321>

Publisher: Wiley

Persistent URL: <https://hdl.handle.net/1721.1/130375>

Version: Original manuscript: author's manuscript prior to formal peer review

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike

Estimating Dynamic Games of Oligopolistic Competition: An Experimental Investigation*

Tobias Salz (MIT)

Emanuel Vespa (UCSB)

February 18, 2019

Abstract

We evaluate dynamic oligopoly estimators with laboratory data. Using a stylized entry/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 mis-predict 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.

*For helpful discussions of this project we would like to thank John Asker, Isabelle Brocas, Colin Camerer, Juan Carrillo, Allan Collard-Wexler, Guillaume Fréchette, Ali Hortaçsu, 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 Weizsäcker, Alistair Wilson, and Sevgi Yuksel. Vespa is grateful for financial support from the UCSB Academic Senate.

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 partic-

ular 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

¹To 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

²The 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.

³For 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 believe 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

⁴To 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 Akerberg et al. (2007) for references on methodological issues.

⁵Variants 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).

⁶Since 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.

⁷The 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 Ertac 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:

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

⁸There 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).

⁹For 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).

¹⁰There 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: $\mathcal{Q}_i(\mathbf{s} = (1, 1)) = \{0, 1\}$, and $\mathcal{Q}_i(\mathbf{s} \neq (1, 1)) = \{\emptyset\}$.

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.

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

¹²While 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 s_t to 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:

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

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, $\rho_i : s \rightarrow \mathcal{Q}_i$; and ii) an entry/exit choice for each value of the state and random component, $\alpha_i : (s, \epsilon_i) \rightarrow \mathcal{A}_i$.

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 ($\rho = [\rho_1, \rho_2]$, $\alpha = [\alpha_1, \alpha_2]$) and state transition probabilities ($F^\alpha(s_{t+1}|s_t)$) such that: i) $\rho(s = (1, 1)) = (1, 1)$; ii) α maximizes the discounted sum of profits for each player, given $F^\alpha(s_{t+1}|s_t)$; and iii) α implies $F^\alpha(s_{t+1}|s_t)$.

¹³For 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 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	1		0	1
0	65	5	0	85	25	0	100	40
1	70	10	1	110	50	1	140	80
	$A_S = 0.05$			$A_M = 0.25$			$A_L = 0.4$	

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.

¹⁴In Appendix A we provide details behind these computations.

¹⁵The 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.

¹⁶In 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

¹⁷This 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.

¹⁸For 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 programming 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

¹⁹In [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 ($\hat{p}(a|s)$)

²⁰The 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.

²¹The 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.

²²In 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 $\Psi(\theta)$ is the vector of respective predicted choice probabilities:²³

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

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-

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

²⁴Notice 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.

²⁵We 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 extents 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.

²⁶Simple 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).

²⁷Once 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.

²⁸The 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.

²⁹We 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

³¹This 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 they 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).

³²In [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 probabilities. 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.

³³The 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

³⁴The 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

³⁵Using 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.

³⁶The 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.

³⁷Regarding 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.

State (s_i, s_{-i})	Standard				No Quantity Choice			
	(1, 0)	(1, 1)	(0, 0)	(0, 1)	(1, 0)	(1, 1)	(0, 0)	(0, 1)
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	<i>A</i>	<i>B</i>	<i>C</i>		<i>A</i>	<i>B</i>	<i>C</i>	
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	
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	<i>A</i>	<i>B</i>	<i>C</i>		<i>A</i>	<i>B</i>	<i>C</i>	
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	
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	<i>A</i>	<i>B</i>	<i>C</i>		<i>A</i>	<i>B</i>	<i>C</i>	
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

³⁸We 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.

³⁹Our 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.

⁴⁰A 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.

⁴¹Figure 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(p)	MAPE(V)	MAE(p)	MAPE(V)	MAE(p)	MAPE(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(p) reports the mean absolute error in the prediction of probabilities. MAPE(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-

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

lute error in probabilities ($\text{MAE}(\mathbf{p})$) 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 ($\text{MAPE}(\mathbf{V})$; see Equation 6 in Appendix A for the expression used to compute the continuation values ($\mathbf{V}(\mathbf{s})$)). 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 $\text{MAPE}(\mathbf{V})$ 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 $\text{MAPE}(\mathbf{V})$, 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

⁴³The 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.

⁴⁴Being 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, $\hat{A}_M^{A_S} = 0.4 > \hat{A}_L = 0.14$.

⁴⁵For 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 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(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.

⁴⁶As 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.

⁴⁷Table 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, $\Delta\text{MAPE}(\mathbf{V})$, is the difference between Standard and No Quantity Choice $\text{MAPE}(\mathbf{V})$ computed using Table 6. The denominator is the Maximal $\text{MAPE}(\mathbf{V})$ 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

⁴⁸As 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.
- Akerberg, 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.

ONLINE APPENDICES

Appendix A

Solving for MPE equilibria

We express the equilibrium as a system of equations in terms of choice continuation values. Denote the per period profits by $\Pi^\alpha(s_t)$ and let $v^\alpha(1, s_t)$ and $v^\alpha(0, s_t)$ be the continuation value of entering the market given the state and exiting the market given the state respectively. In other words, the $v^\alpha(., s_t)$ are the non-random part of the value of choosing either of the two alternatives, given the current state. The α notation indicates that we are looking for pairs of values $v^\alpha(., s_t) \forall s_t \in \mathbf{S}$ and $F^\alpha(s_{t+1}|., s_t) \forall s_t \in \mathbf{S}$ such that these $v^\alpha(., s_t)$ imply the conditional distribution for the state transition process $F^\alpha(s_{t+1}|., s_t)$ and the state transition process induces the values $v^\alpha(., s_t)$. This gives us the following recursive expression for the value functions, which defines a system of equations:

$$v^\alpha(1, s_t) = \Pi^\alpha(s_t) + \delta \cdot \sum_{s_{t+1} \in \mathbf{S}} \int \max \{v^\alpha(0, s_{t+1}) + \epsilon'(0), v^\alpha(1, s_{t+1}) + \epsilon'(1)\} dG(\epsilon') F^\alpha(s_{t+1}|1, s_t) \quad \forall s_t \in \mathbf{S} \quad (3)$$

$$v^\alpha(0, s_t) = \Pi^\alpha(s_t) + \delta \cdot \sum_{s_{t+1} \in \mathbf{S}} \int \max \{v^\alpha(0, s_{t+1}) + \epsilon'(0), v^\alpha(1, s_{t+1}) + \epsilon'(1)\} dG(\epsilon') F^\alpha(s_{t+1}|0, s_t) \quad \forall s_t \in \mathbf{S} \quad (4)$$

We focus on symmetric Markov-perfect equilibria. Since we have four states $S = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$ we are solving a system of eight equations in eight unknowns, four of each $v^\alpha(1, s_t) \forall s_t \in \mathbf{S}$ and $v^\alpha(0, s_t) \forall s_t \in \mathbf{S}$. When the player is in the market $dG(\epsilon')$ is equal to one and $\epsilon(0)$ is drawn uniformly from $[0, 1]$, which is the support of the scrap value distribution. The entry cost $\epsilon(1)$ is always zero in this case. Likewise, if the player is outside of the market $dG(\epsilon')$ is equal to $\frac{1}{1+C}$ and $\epsilon(1)$ is drawn uniformly from $[C, 1 + C]$ whereas $\epsilon(0)$ is zero in these states. Due to the possibility of multiple equilibria we solve this system of equations from many different starting values. In practice we always found only one solution to this system of equations for each of our parameter constellations of interest. While this is a strong indication that in our simple case there exists a unique equilibrium, we cannot entirely rule out that there are other equilibria that our numerical solver did not find. Note that once the $v^\alpha(., s_t)$ are obtained they imply a set of cutoff values for the scrap value and the random part of the entry cost, which imply the choice equilibrium choice probabilities.

Collusive Equilibrium

Let $\mathbf{p}(a = 1|s_t, \delta)$ be the vector of MPE choice probabilities. For $\delta = 0.8$, these are the choice probabilities, which are identified in *italics* in Table 2. These probabilities maximize the value function under the assumption that when both agents are in the market they play the stage Nash equilibrium and receive a payoff of $2A - B$. Likewise, let $\mathbf{p}_c(a = 1|s_t, \delta)$ represent the probabilities that maximize the value function for the case that agents earn A , the collusive quantity choice outcome, whenever they are in the market.

The collusive quantity-stage outcome may be supported as an SPE if defection from the prescribed low quantity choice is punished. Strategies that support collusion have two phases: the collusion phase and the punishment phase. In the collusion phase both players select a low quantity in the first period and as long as both have always selected a low quantity in the past.⁴⁹ Under collusion they make their entry/exit decisions following the implied choice probabilities $\mathbf{p}_c(a = 1|s_t, \delta)$. We consider a punishment that is akin to grim-trigger: if one agent deviates from low production, then all entry/exit decisions are made according to $\mathbf{p}(a = 1|s_t, \delta)$; and whenever both agents are in the market, the choice is high quantity (stage-game Nash). To express punishments formally, let m_t be given by:

$$m_t = \begin{cases} 0 & \text{if } q_{i,r} = q_{j,r} = 0 \text{ for all } r \leq t \text{ when } s_t = (1,1) \\ 1 & \text{otherwise} \end{cases}$$

In each period, m_t is therefore updated according to the latest quantity decisions. Due to the timing assumptions agents know m_t before they make their entry/exit decisions. If both agents have selected a low quantity (whenever both have been in the market) up to and including period t , then m_t takes on a value of 0. If any agent selected high production in any period up to and including t , then m_t takes a value of 1. Let $\nu_{i,t}$ be the scrap value, or the random component of the entry cost, whichever corresponds to the state of the player. In our analysis we focus on grim-trigger strategies that specify an action for the quantity stage –in case both agents are in the market– and a decision rule for entry/exit choices. Hence, for each player i and time period t this class of strategies can be summarized by a pair (q_{it}, a_{it}) such that:

⁴⁹Note that both players start in the market.

$$q_{it} = \begin{cases} 0 & \text{if } m_{t-1} = 0 \text{ or } t = 1 \\ 1 & \text{if } m_{t-1} = 1 \end{cases}$$

and

$$a_{it}(s_t) = \begin{cases} 1 & \text{if } \nu_{it} \leq p_c(s_t, \delta) \text{ and } m_t = 0 \\ 0 & \text{if } \nu_{it} > p_c(s_t, \delta) \text{ and } m_t = 0 \\ 1 & \text{if } \nu_{it} \leq p(s_t, \delta) \text{ and } m_t = 1 \\ 0 & \text{if } \nu_{it} > p(s_t, \delta) \text{ and } m_t = 1 \end{cases}$$

Let $C(\delta)$ be the discounted value of collusion and let $D(\delta)$ be the discounted value of playing according to the symmetric MPE. Now, consider deviations. To establish that a collusive strategy can be supported as an SPE we check whether there is a profitable deviation from such a strategy. According to the one-shot deviation principle it is enough to consider strategies that deviate in period t but otherwise (for every following period) conform to the collusive one. Whenever both agents are in the market in period t , an agent can deviate from the production decision, from the exit decision, or from both. By construction of $\mathbf{p}_c$ there are no incentives to deviate only in the exit decision. If there were, then $\mathbf{p}_c$ would not have been computed correctly. An agent could deviate in the production decision. In that case, they would receive a quantity stage payoff of $2A$. If one agent deviates in the quantity stage of period t , then $m_t = 0$ and the exit decision in that period will be taken according to $\mathbf{p}$. From period $t+1$ onwards agents would be in the punishment phase for all future periods, but this involves playing according to the symmetric MPE, which is sub-game perfect. The payoff of a deviation is: $Def(\delta) = 2A + \mathbb{E}[\nu | \nu > p(s_{it} = 1, s_{it} = 1)] \cdot (1 - p(s_{it} = 1, s_{it} = 1)) + \delta \times D(\delta)$. The grim-trigger strategy is a sub-game perfect equilibrium for all δ such that: $Def(\delta) \leq C(\delta)$.

In order to determine whether the trigger strategies constitute an SPE we first compute $\mathbf{p}_c$ for each treatment, which are reported in the second column of [Table 2](#). We then use these probabilities to check in each case if $Def(0.8) < C(0.8)$. Our treatment parameters are chosen so that trigger strategies *can* support the quantity- stage collusive outcome for A_S and A_M but not for A_L . Finally, we provide a measure of how much higher gains can be under the trigger strategies in each treatment. This measure captures how high the collusion incentives are and is summarized by the percentage increase of collusion over the MPE payoffs: $GoC = 100 \cdot (C(0.8) - D(0.8))/D(0.8)$. The figures are reported in the last row of [Table 2](#).

Joint Monopoly Entry/Exit Probabilities

In the fully collusive equilibrium firms not only collude in the quantity decision, they also coordinate their entry and exit choices. To implement such a collusive strategy firms have to solve a private monitoring problem since the random parts of the entry cost and the scrap values are only privately observed by the players. We do not solve this private monitoring problem. However, to still be able to obtain a benchmark on how much higher the gains from collusion would be if firms also collude in the dynamic decision, we solve for the case in which firms know the entry cost and scrap values of the other firm. Under this assumption the problem can be solved as a simple single agent dynamic programming problem. For each of the possible four states $s_t \in \mathbf{S}$ the combined firm has four possible actions $a_t \in \mathbf{A} = \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. We solve for the choice continuation values that summarize the non-random part of each of those four possible choices. In total there are sixteen equations:

$$v(a_t, s_t) = \Pi(s_t) + \delta \cdot \int \max_{a_t \in \mathbf{A}} \{v(a_t, s_{t+1}) + \epsilon'(a_t)\} dG(\epsilon') F(s_{t+1}|a_t, s_t) \quad \forall s_t \in \mathbf{S}, \forall a_t \in \mathbf{A} \quad (5)$$

Results are shown in [Table 8](#), which is similar to [Table 2](#) but includes the dynamic choice probabilities on the most collusive equilibrium. Interestingly, the probability to be in the market in each period is much lower compared to quantity-stage collusion for each of the four states. This result has a simple intuition. The combined market value is the same no matter whether there are two or only one firm in the market. Relative to the quantity-stage collusive equilibrium firms now coordinate on entry exit choices, which allows them to exploit the gains they can make from high scrap values and low entry cost. The prediction involves high turnover to exploit these gains.

Table 8: Cutoff-strategies for each treatment: MPE, CE and joint monopoly (MON)

Conditional probability	A_S			A_M			A_L		
	MPE (p)	CE (p _c)	MON	MPE (p)	CE (p _c)	MON	MPE (p)	CE (p _c)	MON
$p(1, 0)$	0.458	0.519	0.504	0.688	0.823	0.652	0.880	0.925	0.683
$p(1, 1)$	0.360	0.512	0.492	0.596	0.784	0.602	0.781	0.870	0.610
$p(0, 0)$	0.260	0.368	0.354	0.498	0.663	0.505	0.681	0.757	0.554
$p(0, 1)$	0.161	0.362	0.340	0.408	0.624	0.467	0.583	0.702	0.457
Is collusion in quantities an SPE?	YES			YES			NO		
Gains from collusion in quantities	450.8%			450.8%	481.1%	481.1%	75.9%	93.22%	93.22%

Note: This table presents the conditional choice probabilities for each of the four states as indicated in the left column. The choice probabilities are presented for each of the three market sizes, which we implement as treatments as indicated in the top row. Predictions are presented for the MPE(p) as well as the case where players collude in the marketstage, CE (p_c). In the bottom the table indicates whether the collusive equilibrium we highlight can be supported as an SPE and how high the gains over MPE would be.

Mapping Choice Probabilities to Value Functions

In this section we provide details on the estimation procedure. In a first stage we estimate the empirical dynamic choice probabilities P using simple frequency estimation, thus obtaining $\hat{P}$. This means we compute the fraction of time a player stays in the market or leaves the market respectively for each of the four states. Following the insight of Hotz and Miller (1993) the value function can be expressed in terms of these choice probabilities. Since all agents under the model assumption are planning with the equilibrium transitions that we estimate from the data we can directly solve for the value function in terms of these probabilities. We can write the value function as:

$$V = \sum_a P_a(\hat{P})[u_a + e_a(\hat{P}) + \delta \cdot F_a(\hat{P}) \cdot V] \Leftrightarrow$$

$$V = [I - \delta \cdot \sum_a P_a(\hat{P}) \cdot F_a(\hat{P})]^{-1} [\sum_a P_a(\hat{P})(u_a + e_a(\hat{P}))]$$

In our case with four states these objects take on a simple form:

$$V = [I - \delta \cdot [P_1(\hat{P}) \cdot F_1(\hat{P}) + P_0(\hat{P}) \cdot F_0(\hat{P})]]^{-1} [P_1(\hat{P}) \cdot (e_1(\hat{P}) + u_1) + P_0(\hat{P}) \cdot (e_0(\hat{P}) + u_0)] \quad (6)$$

with

$$u_1 = \begin{pmatrix} 0 \\ 0 \\ 2 \cdot A \\ 2 \cdot A - B \end{pmatrix} u_0 = \begin{pmatrix} -C \\ -C \\ 2 \cdot A \\ 2 \cdot A - B \end{pmatrix} e_1 = \begin{pmatrix} 0 \\ 0 \\ 1 - \frac{\hat{p}_3}{2} \\ 1 - \frac{\hat{p}_4}{2} \end{pmatrix} e_0 = \begin{pmatrix} -\frac{\hat{p}_1}{2} \\ -\frac{\hat{p}_2}{2} \\ 0 \\ 0 \end{pmatrix}$$

$$P_1 = \begin{bmatrix} \hat{p}_1 & 0 & 0 & 0 \\ 0 & \hat{p}_2 & 0 & 0 \\ 0 & 0 & \hat{p}_3 & 0 \\ 0 & 0 & 0 & \hat{p}_4 \end{bmatrix} P_0 = \begin{bmatrix} 1 - \hat{p}_1 & 0 & 0 & 0 \\ 0 & 1 - \hat{p}_2 & 0 & 0 \\ 0 & 0 & 1 - \hat{p}_3 & 0 \\ 0 & 0 & 0 & 1 - \hat{p}_4 \end{bmatrix}$$

$$F_1 = \begin{bmatrix} \hat{p}_1 & 1 - \hat{p}_1 & 0 & 0 \\ \hat{p}_3 & 1 - \hat{p}_3 & 0 & 0 \\ \hat{p}_2 & 1 - \hat{p}_2 & 0 & 0 \\ \hat{p}_4 & 1 - \hat{p}_4 & 0 & 0 \end{bmatrix} F_0 = \begin{bmatrix} 0 & 0 & \hat{p}_1 & 1 - \hat{p}_1 \\ 0 & 0 & \hat{p}_3 & 1 - \hat{p}_3 \\ 0 & 0 & \hat{p}_2 & 1 - \hat{p}_2 \\ 0 & 0 & \hat{p}_4 & 1 - \hat{p}_4 \end{bmatrix}$$

In the previous expressions, we use the following shortened notation: $p_1 = p(a = 1|s = (0, 0))$, $p_2 = p(a = 1|s = (0, 1))$, $p_3 = p(a = 1|s = (1, 0))$, $p_4 = p(a = 1|s = (1, 1))$. Note that Equation 6 expresses the value function only in terms of parameters and objects that are composed of observables. The following section shows briefly how the expected values $e_1(\hat{P})$ and $e_0(\hat{P})$ are obtained from choice probabilities as indicated above. For a given parameter guess and choice probabilities we therefore form a guess $\hat{V}$. Once we know $\hat{V}$, we can easily compute $\hat{v}(1, s_t) \forall s_t \in \mathbf{S}$ and $\hat{v}(0, s_t) \forall s_t \in \mathbf{S}$ from this.

Choice Probabilities and Expectations for the Uniform

Let $\mathbf{x}_{out}$ denote the states in which the player is out and $\mathbf{x}_{in}$ be the states in which he is in. For the entry-case under $\psi \sim U[0, 1]$ and $0 \leq (v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})) \leq 1$ we have:

$$E[\psi | \psi < v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})] = \int_0^{v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})} \psi \frac{1}{v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})} d\psi =$$

$$\frac{\psi^2}{2} \frac{1}{v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})} \Big|_0^{v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})} = \frac{(v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out}))}{2} = \frac{P(in)}{2}$$

For the exit case under $\phi \sim U[0, 1]$ and $0 \leq (v(1|\mathbf{x}_{out}) - v(0|\mathbf{x}_{out})) \leq 1$ we have:

$$E[\phi | v(1|\mathbf{x}_{in}) - v(0|\mathbf{x}_{in}) < \phi] = \int_{v(1|\mathbf{x}_{in}) - v(0|\mathbf{x}_{in})}^1 \phi \frac{1}{1 - (v(1|\mathbf{x}_{in}) - v(0|\mathbf{x}_{in}))} d\phi$$

$$= \frac{1 + (v(1|\mathbf{x}_{in}) - v(0|\mathbf{x}_{in}))}{2} = 1 - \frac{P(in)}{2}$$

Monte Carlo

Data Generating Process

We generate data either under the assumption of the symmetric MPE or the CE play for our parametrizations. We assume that there are 300 pairs of firms and each pair of firms is considered to be isolated from the rest. The interaction between each pair of firms ends after each period with probability 0.2, which corresponds to the discount implemented in the laboratory. For the Monte Carlo study we will estimate parameters using 100 such data sets and subsample each data set 30 times to obtain standard errors.

Table 9: Monte Carlo results

Parameter	A_S			A_M			A_L	
	Estimates		True value	Estimates		True value	Estimates	True value
	MPE	CE		MPE	CE			
A	0.052 (0.012)	0.044 (0.045)	0.05	0.249 (0.03)	0.246 (0.049)	0.25	0.390 (0.041)	0.4
B	0.620 (0.130)	0.032 (0.173)	0.6	0.611 (0.132)	0.24 (0.144)	0.6	0.575 (0.111)	0.6
C	0.142 (0.029)	0.148 (0.04)	0.15	0.146 (0.032)	0.148 (0.021)	0.15	0.153 (0.027)	0.15

Note: The table shows the results of the Monte Carlo estimation. For each of the three values of A it shows the estimates if the econometrician assumes the correct data generating process (MPE) as well as the estimates if the econometrician incorrectly assumes the MPE and the data is in fact coming from decisions on the equilibrium path of the highlighted collusive equilibrium (CE) for A_M and A_S . In the case of A_L , the CE is not an SPE. Estimates are averages over 100 datasets. Each dataset assumes 300 markets of an average length of five periods (market terminates randomly with probability 0.2). Standard errors are shown in parentheses below the estimates and are obtained by subsampling each data-set 30 times.

Monte Carlo Estimates

For each value of A , Table 9 presents the Monte Carlo estimates. In the “MPE” column we present estimates when firms’ play is generated according to the symmetric MPE, while in the “CE” column we display estimates when firms are assumed to follow the Collusive Equilibrium (and the econometrician wrongfully assumes MPE play). By comparing estimates in the “MPE” column with the True Value column treatment by treatment we verify that the parameters can be recovered with tight standard errors from a data set of modest size.⁵⁰

Comparing the CE estimates to the true value we notice that the bias from an incorrect assumption on the equilibrium shows up in B . The parameter that captures the competitive effect would be biased downwards. Intuitively, the estimator in recovering B compares the choices of firm i when firm $-i$ is in the market to those when firm $-i$ is not in the market. For

⁵⁰We also ran versions of the Monte Carlo in which players’ cost depends on individual specific cost shifters that are observable to the econometrician. Such cost shifters would help to considerably improve the standard error of the interaction term B and improved identification at the limits of the parameter space. Under the current specification, the only observable variable that shifts player i ’s action is the action of player $-i$. Because of this minimal structure of the model, parameter estimates become noisy for B values very close to zero, where the influence of the other player vanishes. However, in the experiment, other observable cost shifters (i.e. some variable x not determined endogenously) would have increased the number of necessary treatments and the complexity considerably, which is why we decided against such a setup.

example, consider market A_S and assume firm i is currently in the market. Because the CE data is generated according to the second column of Table 2, firm i will be in the market next period with probability 0.519 or 0.512, depending on whether firm $-i$ is in the market or not. This small difference in probabilities (a small effect of competition) will be rationalized with an estimate for B that is smaller than the true value. In other words, a lower estimate for B is consistent with the presence of collusion.

Monte Carlo Counterfactuals

For each baseline we perform an exercise in which there is full collusion in the Monte Carlo generated data, but the econometrician wrongly assumes that there is no collusion. Then the econometrician uses the estimates to predict behavior for other market sizes. To document the maximum possible error in counterfactual predictions we assume that in each counterfactual there is actually no collusion in the data.

The bias (as in Section 5.4 of the paper) is computed as the percentage deviation in the value of the firm from the true value of the firm under the correct data generating process for each state. The results are shown in Table 21. We then collapse these measures into a single measure assigning equal weight to each state, what we refer to in Section 5.4 as MAPE (V). No bias would result in a MAPE (V) of zero. These Monte Carlo results emphasize that proper collusion if not accounted for by the econometrician would severely bias counterfactual computations. Table 21 shows that the overestimate ranges from $\approx 36\%$ to $\approx 110\%$.

	A_L	A_M	A_S
Baseline/Counterfactual	MAPE(V)	MAPE(V)	MAPE(V)
A_L	-	35.4	56.5
A_M	60.9	-	107.6
A_S	107.3	113.5	-

Table 10: Maximal counterfactual error from collusion (in %)

This table (based on Monte Carlo results and NOT data) shows the maximal resulting bias if the econometrician wrongly assumes that firms do not collude in the baseline and predicts the outcomes in a market (of the same market size) under the correct assumption that firms in the counterfactual do not collude.

Appendix B

In this appendix we study entry/exit choices in more detail. We first provide detailed hypotheses on comparative statics, which are afterwards tested.

Comparative Static Hypotheses

We now explicitly formulate MPE comparative statics in terms of entry and exit thresholds using Table 2 as a reference. We can then contrast these predictions under MPE play with what we actually observe in the data.

Comparative Statics 1 (CS1): Exit vs. Entry thresholds (within treatment). *Exit thresholds are predicted to be higher than entry thresholds.*

In addition, there are predictions on how entry and exit thresholds should vary as a response to the state of the other player (i.e. compare thresholds for $s(\cdot, 1)$ to $s(\cdot, 0)$). When the other is in the market there is competition, and quantity-stage payoffs correspond to the static Nash equilibrium; these payoffs are lower than the (static monopoly) payoffs the agent gets when the other is out of the market. In the equilibrium this is captured with a difference between the two exit thresholds and a difference between the two entry thresholds. We refer to these predictions as the “effect of competition” on thresholds.

Comparative Statics 2 (CS2): Effect of competition in thresholds (within treatment). *Fix the agent’s own current state. Thresholds are higher when the other player is currently out than when the other player is currently in the market.*

The MPE also provides comparative statics across treatments. As the value of A increases, the relative attractiveness of the market also increases and all corresponding thresholds are higher: agents demand higher exit payments to leave the market and are willing to pay higher entry fees to go in. This prediction is summarized below:

Comparative Statics 3 (CS3): Between treatments: *All thresholds increase monotonically with A .*

The Collusive Equilibrium also provides hypotheses on comparative statics. For example, 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.⁵¹ We now state this as a hypothesis.

⁵¹As mentioned earlier, the characterized collusive equilibrium is one of possibly many equilibria that support collusion in the quantity

Collusion Hypothesis 1 (CH1): Effect of Defection on Thresholds. *According to the Collusive Equilibrium, entry and exit thresholds are lower in all periods after defection in the quantity stage.*

Finally, we can also use [Table 2](#) and compare thresholds between treatments that allow and that do not allow for a quantity choice.

Collusion Hypothesis 2 (CH2): Standard vs. No Quantity Choice treatments. *Fix the value of A and compare Standard treatments to treatments with No Quantity Choice. There is evidence consistent with the presence of collusion if: 1) the effect of competition is lower in the Standard treatments; and 2) if thresholds for all states are higher in the Standard treatments.*

States

There are large differences across treatments in terms of states' frequencies, which are presented in [Table 11](#). Considering the unit of observation as a pair of subjects in each period of each supergame there are three possible states the pair can be at: both are out of the market, one out and the other in or both are in the market. The table shows the proportion of periods in which a pair was in either of these three states by treatment. Period 1 of every match is omitted as by definition all subjects start out, so that the table only shows the results of endogenous decisions.

There are clear patterns in the table. First, being out of the market is more likely when A is lower. The state when both are out reaches the highest share for A_S and its occurrence diminishes when A is higher. In fact, in only very few occasions do we observe pairs of subjects in this state for A_L . The opposite situation is observed for the state when both are in, which reaches the highest share for A_L . The lower likelihood of being out in the large market size treatment will have consequences on the accurateness of the average entry threshold estimate. This will be reflected as relatively larger confidence intervals as can be seen in [Figure 2](#).

The second pattern is that the likelihood of being out is higher when there is No Quantity Choice as long as A is not at the highest level. Consider, for instance, A_S . The proportion of times that both agents are out of the market is clearly larger when there is No Quantity Choice. This is consistent with the fact that quantity choices present the alternative to obtain higher payoffs and thus have the potential of making staying in the market more attractive. The effect

stage. However, any punishment phase that achieves collusion in the quantity stage is expected to involve lower thresholds. If collusion in the quantity stage cannot be supported, then being in the market is less valuable, which is reflected by lower thresholds.

Treatment	State		
	Both Out	One Out-One In	Both In
A_S , No Quantity Choice	34.8	51.2	14.0
A_S , Standard	22.9	49.7	27.4
A_M , No Quantity Choice	9.5	48.1	42.5
A_M , Standard	5.7	41.1	53.2
A_L , No Quantity Choice	2.8	32.5	64.7
A_L , Standard	2.6	31.5	65.9

Table 11: States by treatment after period 1 (in percentages)

is smaller for A_M and almost indistinguishable for A_L , when there is a negligible number of cases in either treatment where both agents are out of the market.

Within Treatment Hypotheses on Thresholds

To test for the main within treatment hypotheses (CS1, CS2), we will use panel data analysis. We conduct one regression per treatment, which are reported in [Table 12](#). The left-hand-side variable in all cases is the threshold selected by each subject. If in the corresponding period the subject is deciding to exit (enter) the market, the threshold variable captures their report for the exit (entry) threshold. On the right-hand side there are four variables. We exclude the state when the subject is out and the other is in the market ($s = (0, 1)$), which in theory corresponds to the case where the subject is least likely to enter the market next. Naturally, the excluded state will be captured by the constant, and we add a dummy for each of the three other possible states. Depending on the state, the corresponding dummy will report the increment to the baseline threshold defined by the constant.

CS1 predicts that exit thresholds are higher than entry thresholds. This translates into four comparisons by treatment, and in all 24 cases the differences are significant at the 1% level in the direction predicted by the theory. In other words, there is strong support for this hypothesis, which indicates that subjects do respond to one of the most basic incentives of the

game.

There is also evidence in favor of CS2. In this case, the hypothesis implies two comparisons by treatment: fixing s_i and testing whether there is a difference depending on the state of the other. For the case when the subject is out of the market, the outcome for the comparison is readily available in the estimates for coefficient ($s = (0, 0)$). In all cases the estimate is positive: subjects are willing to pay more to enter the market if the other is out. However, the coefficient is significant at the 5% level for A_M treatments and at the 10% level in two other cases. Moreover, while the theory predicts a 10-point difference (see Table 2), the estimate is quantitatively smaller in all cases. The effect of competition, however, is more evident when the subject is in the market. In this case, the difference between the coefficients ($s = (1, 1)$ and $s = (1, 0)$) is always as predicted by the theory and significant in all treatments. In a few words, all comparisons are in line with the prediction and all but two are significant at least at the 10% level.

Table 12: Panel regressions: Within treatment hypotheses

Variable	A_S		A_M		A_L	
	Standard	No Quantity Choice	Standard	No Quantity Choice	Standard	No Quantity Choice
Intercept	20.377*** (2.412)	19.273*** (2.494)	36.365*** (3.001)	33.879*** (3.666)	45.488*** (5.503)	37.391*** (3.471)
$s = (0, 0)$	2.310* (1.375)	1.150 (2.431)	5.828*** (1.077)	5.387*** (1.306)	2.302 (5.575)	1.658* (0.915)
$s = (1, 1)$	47.587*** (2.709)	34.768*** (5.136)	46.913*** (6.465)	42.930*** (6.102)	41.026*** (8.394)	49.589*** (2.272)
$s = (1, 0)$	54.119*** (3.306)	50.047*** (6.494)	51.202*** (5.700)	49.650*** (5.084)	46.221*** (6.678)	53.658*** (1.032)

Note: This table provides reduced form analysis of the within treatment comparative statics. The dependent variable is the selected threshold (entry or exit) that corresponds to the state. The independent variables are dummies for each state where the state ($s = (0, 1)$) is the excluded category. Standard Errors reported between parentheses, Significance levels: 1%(***), 5%(**), 10%(*), Standard Errors are clustered at the session level

The analysis so far has focused on centrality measures, which are key for computing structural estimates. To provide a broader perspective of our data, Figure 4 displays the cumulative distributions of thresholds by state for all treatments. Some patterns are present across treatments. First, the distributions are largely ordered as predicted by the symmetric MPE. Entry thresholds display lower values than exit thresholds, and within each case subjects largely

select higher values when the other is out. Second, most distributions suggest that values are centered around the mean. For example, in the case of the A_S -Standard treatment, entry thresholds display a significant mass between the relatively lower values (20 and 40), while most of the mass in the case of exit thresholds is between 70 and 80.

Between Treatment Hypotheses on Thresholds

We now test statements that involve comparisons that depend on the value of A or on whether there is a quantity stage or not. With this aim we conduct two panel regressions, one for exit and one for entry thresholds. More specifically, the “Entry Threshold” regression only considers periods when subjects had to select an entry threshold. The selected entry threshold constitutes the left-hand side variable. On the right-hand side there are two sets of dummies. The excluded group corresponds to the A_S -Standard case. The first set of dummies will capture the differential effect corresponding to the other five treatments. The second set of dummies interacts the treatment dummy with the state of the other player. That is, for each treatment there is a dummy that takes value 1 if the other is out of the market. The second regression uses the same controls, but considers exit thresholds on the left-hand side instead. The results are reported in [Table 13](#).

CS3 states that thresholds increase with market size. There are 24 comparative statics: for each of the four states there are three comparisons, and such comparisons can be made for treatments with and without a quantity stage. Not all comparative statics are statistically significant, but all differences are in the direction predicted by the hypothesis. In all treatments the difference in entry and exit thresholds is significant at the 1% level when comparing the A_S treatment to either of the other market sizes. When comparing the A_M to A_L , the difference between thresholds is statistically significant at the 5% level only for exit thresholds when there is No Quantity Choice. In other cases differences are not statistically significant. Overall this means that in 18 out of 24 comparisons differences are statistically significant.

Section 3 also presented hypotheses that would be consistent with the presence of collusion. Part 1 of CH2 claims that there is evidence consistent with the presence of collusion if the effect of competition is lower when there is a quantity choice. There would be evidence supporting the claim if fixing the market size, the interaction dummy is significantly higher when there is a quantity stage. Again, in all six comparisons the differences are in the direction predicted by the hypothesis. For entry thresholds the differences are significant at the 5% and

Figure 4: Cumulative distributions of thresholds across treatments

10% level for A_M and A_L , respectively. For exit thresholds, the differences are significant at the 1% level only for A_S .

Fixing the market size collusion is consistent with higher thresholds when there is a quantity choice, which constitutes part 2 of CH2. This hypothesis involves 12 comparisons using the estimates presented in [Table 13](#). For example, consider the entry threshold regression and A_M . The hypothesis claims two comparisons, depending on whether the other is in or not: i) the coefficient for A_M -S is higher than for A_M -NQ, and ii) adding the coefficients for A_M -S and the interaction A_M -S $\times$ Other Out is lower than the addition of the same coefficients but when there is no quantity choice. The direction of the differences is in line with the prediction in all cases, but differences are not significant with the exception of the exit threshold for A_S size when the other is in the market.

Table 13: Panel regressions: Between treatment hypotheses

Variable	Entry Threshold	Exit Threshold
Intercept	22.001*** (1.955)	67.131*** (1.951)
A_S -NQ	-3.110 (2.988)	-14.547*** (3.337)
A_M -S	10.046*** (3.566)	14.652*** (4.055)
A_M -NQ	10.858*** (2.681)	9.000*** (3.177)
A_L -S	13.498*** (2.578)	18.087*** (4.258)
A_L -NQ	12.320*** (4.207)	19.482*** (2.585)
A_S -S $\times$ Other Out	1.406* (0.830)	7.490*** (1.380)
A_S -NQ $\times$ Other Out	1.968 (2.595)	14.398*** (1.310)
A_M -S $\times$ Other Out	2.204*** (0.450)	4.479*** (0.635)
A_M -NQ $\times$ Other Out	3.945*** (0.348)	6.476*** (1.725)
A_L -S $\times$ Other Out	1.139*** (0.229)	3.521*** (1.219)
A_L -NQ $\times$ Other Out	4.297** (1.836)	3.802*** (1.390)

Note: Standard Errors reported between parentheses, Significance levels: 1%***), 5%(**), 10%(*), Standard Errors are clustered at the session level. S: indicates Standard treatment; NQ indicates No Quantity Choice treatment.

Effect of Quantity-Stage Choices on Thresholds

CH1 claims that entry and exit thresholds are lower after defection. We present the results of a random-effects probit regression where the left-hand side is the exit (or entry) threshold and on the right-hand side there is a set of dummy variables that capture the outcome for the last time subjects were in the market.⁵² Table 14 displays the results of these regressions for each treatment.

Several patterns emerge. First, consider exit thresholds. In A_S and A_M treatments subjects are more responsive to last period's outcome. In these cases, subjects are significantly more likely to select a higher exit threshold, while if the other defected in the previous market interaction they are more likely to select a lower threshold. This last effect is also present for A_L . When, instead, we look at entry thresholds, the pattern is less clear. It appears that in most cases subjects are less responsive to recent market behavior when they are out of the market.

Outcome Last Quantity Stage	A_S				A_M				A_L			
	Exit Threshold		Entry Threshold		Exit Threshold		Entry Threshold		Exit Threshold		Entry Threshold	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
(Collude, Collude)	5.551***	0.956	0.997	0.947	5.688***	0.622	4.478**	1.916	0.108	0.618	.704	3.479
(Collude, Defect)	-6.867***	1.172	-1.848*	1.01	-1.965***	0.772	-1.013	1.657	-1.940***	0.719	4.361	2.761
(Defect, Collude)	-1.319	-1.17	-0.306	0.931	-0.377	0.773	0.538	1.484	-0.876	.126	4.885	3.073
Constant	69.484***	1.868	37.84***	2.639	81.81***	2.246	47.08***	2.982	85.88***	.124	49.7***	3.396

Notes: The dependent variable is the threshold selected by subject i . The right-hand side variables include possible outcomes from last quantity stage, where the first action corresponds to the choice of subject i . Significant at: *** 1%, **5%, *10%

Table 14: Effect of past market choices on thresholds

Choices As the Session Evolves

In principle it is possible that choices in the aggregate change as the session evolves. Figures 5 and 6 display average exit and entry thresholds for each supergame for each possible state a subject may be at, for Standard and No Quantity Choice treatments, respectively. Visual inspection suggests that in most cases a trend is not evident, which we indeed confirm with statistical analysis.⁵³ If we add a dummy for each supergame to the regressions of Tables

⁵²In cases where subjects are deciding on an exit threshold the last period for which there is an outcome is the current period. The reference is to the last period in which there was a market choice in the case of entry thresholds.

⁵³The case of entry thresholds for A_L when the other is not in the market does display volatility, but this is due to the relatively low number of observations (see Table 11).

12 and 13 the message is similar. In a few cases there is a significant effect of a particular supergame; when such effect is present it happens in the earlier supergames of the session and is quantitatively very small.

It is also possible that subjects change the thresholds within a supergame. This may be because they are following a strategy that conditions the threshold on past play (i.e. as in the CE) or because they follow a strategy that conditions on a particular period. We know that some subjects may be conditioning their thresholds on past play given that some aggregate choices are consistent with CH1. In order to test if there is a strong pattern in aggregate thresholds depending on the period, we include the regressions in Table 12: a) a set of period dummies and b) interactions of each period dummy with the dummy that takes value 1 if the other is not in the market. Results show that there is no clear pattern that indicates a period effect at the aggregate level.⁵⁴

Summary

We now summarize the main findings in this appendix:

- There is broad support in the data for the comparative statics predicted by the symmetric MPE. Out of 60 comparisons implied by CS1, CS2 and CS3 all differences are in the predicted direction and 50 (52) are significant at least at the 5% (10%) level.
- There is evidence that is consistent with market collusion having an effect on threshold choices. The effect of market collusion on threshold choices appears to be higher for A_S and A_M , and for exit thresholds.
- The evidence does not indicate substantial changes in aggregate behavior as the session evolves.

⁵⁴Given the large set of controls we do not report these regressions, but they are available upon request.

Figure 5: Evolution of thresholds: No Quantity Choice treatments

Figure 6: Evolution of thresholds: Standard treatments

Appendix C

This appendix provides additional analysis on the quantity-stage choice. Figure 7 displays the cooperation rate taking all periods of a supergame into consideration and basically reproduces the same broad patterns presented in Figure 8, which only takes into account cooperation in the first period of each supergame. Cooperation rates after period 1 will be endogenously affected by behavior within the supergame. In this section we use two approaches to better understand the determinants of quantity-stage choices. First, we take a non-structural approach and use panel regression analysis to study how cooperation in period t is affected by behavior in previous periods and supergames. Second, we use a structural method to study which strategies better capture subjects' choices. Finally, we study further the connections between quantity-stage and entry/exit-stage choices.

Figure 7: Cooperation rates (all)

Cooperative Behavior: A Non-Structural Approach

To further study decisions in the quantity stage we run random effects probit regressions where the market action is the variable on the left-hand side (1: cooperation) and the right-hand side includes a series of usual controls. We control for the subjects' last choice (Own past action) and their partner's past action (Other's past action) last time they were in the market, and we include period 1 decisions to control for dynamic unobserved effects. There are also three dummies to capture the state in the last period, where the state in which both subjects are

Figure 8: Cooperation rates (first period)

out is omitted. Match dummies and period dummies are also included, but for space reasons omitted in [Table 15](#).

Several patterns are consistent across treatments. First of all, the likelihood of cooperation is higher when the subject or the other colluded last time they were in the market. In fact, the probability of cooperating is higher when the other colluded previously. Second, the dummy for the likelihood of cooperation if the state last period was (1, 0) is the most negative in all treatments. This indicates that subjects are least likely to cooperate coming from a situation when they were in the market, but the other was out. This also suggests that some subjects may choose to be less cooperative in the market in order to incentivize the other to leave the market.

Recovering Strategies Using SFEM

Cooperation rates provide one measure of collusion, but there are techniques –the Strategy Frequency Estimation Method (SFEM) of Dal Bó and Fréchette (2011)– that recover which strategies best rationalize the data. This would tell us to what extent choices in treatments where collusion can be supported as an SPE are consistent with the CE.⁵⁵

⁵⁵In principle, another alternative consists of regressing choices on past play. However, such an exercise can be misleading. For example, if only a small proportion of subjects are consistently conditioning their choices on some aspect of the history, the associated coefficient can

Variable	A_S		A_M		A_L	
	Coeff.	Std. Err.	Coeff.	Std. Err.	Coeff.	Std. Err.
Own past action	0.369***	0.124	.784***	0.095	.836***	0.105
Other's past action	0.857***	0.121	.966***	0.103	1.259***	0.119
Own action in period 1	2.191***	0.150	2.039***	0.124	1.336***	.126
Other's action in period 1	0.659***	0.120	0.809***	0.107	0.738***	.124
$s_{t-1} = (0, 1)$	-0.262	0.191	-0.943***	0.186	-0.215	.198
$s_{t-1} = (1, 0)$	-1.256***	0.223	-1.264***	0.186	-0.929***	.246
$s_{t-1} = (1, 1)$	-0.913***	0.119	-1.049***	0.105	-0.710***	.116

Note: Significant at: *** 1%, **5%, *10%

Table 15: Random effects probit results

In order to outline how the SFEM works, consider an infinitely repeated prisoners' dilemma. The game involves just a static decision, where in every period the agent faces a binary choice (cooperate or defect), as if both subjects were in the market. In that simpler environment there is a large set of possible strategies σ an agent can follow. Strategies may depend on past behavior, and it is possible to compute for each $\sigma \in \Sigma$ what choices the subject would have made had she been exactly following strategy σ . On the other hand we have the subject's actual choices. The unit of observation is a *history*: the set of choices a subject made within a supergame. The SFEM procedure works as a signal detection method and estimates via maximum likelihood how close the actual choices are from the prescriptions of each strategy. The output is the frequency for each strategy in the population sample.

To describe the method in further detail, assume that the experimental data has been generated for an infinitely repeated prisoners' dilemma and define ch_{igr} as the choice of subject i in period p of supergame g , $ch_{igr} \in \{Cooperate, Defect\}$. Consider a set of K strategies that specify what to do in round 1 and in later rounds depending on past history. Thus, for each history h , the decision prescribed by strategy k for subject i in period p of supergame g can be computed: $h_{igr}(h^k)$. A choice is a perfect fit for a history if $ch_{igr} = h_{igr}(h^k)$ for all rounds of the history. The procedure allows for mistakes and models the probability that the choice

turn out to be significant. We could mistakenly conclude that past play does have an effect in the population, while it is actually driven by a small share. The technique we use, on the other hand, allows us to estimate the proportion of choices that can be better rationalized as conditioning on past play. We would, thus, observe if such proportion is a small share or not. More generally, using standard regression analysis to evaluate whether behavior is Markovian or not can be very misleading; see Vespa (2019) for more details.

corresponds to a strategy k as:

$$Pr(ch_{igp} = h_{igp}(h^k)) = \frac{1}{1 + \exp\left(\frac{-1}{\gamma}\right)} = \beta. \quad (7)$$

In (7) $\gamma > 0$ is a parameter to be estimated. As $\gamma \rightarrow 0$, then $Pr(ch_{igp} = h_{igp}(h^k)) \rightarrow 1$ and the fit is perfect. Define y_{igp} as a dummy variable that takes value one if the subject's choice matches the decision prescribed by the strategy, $y_{igp} = 1 \{ch_{igp} = h_{igp}(h^k)\}$. If (7) specifies the probability that a choice in a specific round corresponds to strategy k , then the likelihood of observing strategy k for subject i is given by:

$$p_i(s^k) = \prod_g \prod_p \left(\frac{1}{1 + \exp\left(\frac{-1}{\gamma}\right)} \right)^{y_{igp}} \left(\frac{1}{1 + \exp\left(\frac{1}{\gamma}\right)} \right)^{1-y_{igp}} \quad (8)$$

Aggregating over subjects: $\sum_i \ln(\sum_k \phi_k p_i(s^k))$, where ϕ_k represents the parameter of interest, the proportion of the data which is attributed to strategy s^k . The procedure recovers an estimate for γ and the corresponding value of β can be calculated using (7). The estimate of β can be used to interpret how noisy the estimation is. For example, with only two actions a random draw would be consistent with $\beta = 0.5$.

Our environment is more complex than an infinitely repeated prisoners' dilemma; it involves a dynamic (continuous) and a static (discrete) choice. While the SFEM procedure is designed to study discrete choices, we can still use it to learn about the strategies that rationalize our subjects' quantity choices. A necessary condition for CE is that subjects follow a grim-trigger strategy whenever both are in the market. Likewise, a necessary condition for the symmetric MPE is that both subjects always defect from cooperation in the static choice. We use the SFEM procedure to study if subjects' behavior in the quantity stage is consistent with these necessary conditions.

We proceed in the following manner. First, for each *history* in our dataset we only keep the static choices. All subjects make a quantity choice in period 1, but it is possible for example that the next period with a market choice is period 4. In our constrained dataset we would only keep the quantity stage choices for rounds 1 and 4 and would interpret them as the first and the second quantity choices. In this way we obtain a dataset that resembles the dataset coming from an infinitely repeated prisoners' dilemma. Second, we define a set of strategies $K \subset \Sigma$ following the literature (see for example Dal Bó and Fréchette (2011) or Fudenberg et al. (2012)). We include in K five strategies that have been shown to capture most behavior in

infinitely repeated prisoners' dilemma: 1) Always Defect (AD), 2) Always Cooperate (AC), 3) Grim-Trigger (Grim), 4) Tit-for-Tat, and 5) Suspicious-Tit-for-Tat.⁵⁶

	All data			Last 8 Supergames		
	A_S	A_M	A_L	A_S	A_M	A_L
AD	0.394*** (0.077)	0.403*** (0.096)	0.549*** (0.109)	0.463*** (0.114)	0.404*** (0.121)	0.604*** (0.154)
AC	0.064 (0.052)	0.094* (0.050)	0.024 (0.031)	0.123** (0.057)	0.124** (0.054)	0.026 (0.038)
Grim	0.205*** (0.088)	0.339*** (0.125)	0.047 (0.047)	0.172* (0.106)	0.335** (0.151)	0.000 (0.043)
Tit-for-Tat	0.285*** (0.088)	0.111 (0.083)	0.188 (0.126)	0.206* (0.128)	0.087 (0.103)	0.211* (0.128)
Susp.-Tit-for-Tat	0.052	0.052	0.192	0.036	0.048	0.158
γ	0.488*** (0.044)	0.419*** (0.054)	0.384*** (.033)	0.404*** (0.039)	0.357*** (0.051)	0.333*** (0.047)
β	0.885	0.916	0.931	0.922	0.943	0.953

Note: Significant at: *** 1%, **5%, *10%. See [Appendix C](#) for the definition of γ . $\beta = \frac{1}{e(-1/\gamma)}$

Table 16: Strategy frequency estimation method results

[Table 16](#) presents the results of the estimation for the three Standard treatments using all data and using the last eight super games.⁵⁷ The estimates uncover clear patterns in subjects' choices. Consider first always defect (AD). In all cases this is the strategy with the highest frequency, around 40% for A_S and A_M and close to 60% for A_L . Second, comparing across treatments we observe that Grim displays the opposite pattern of AD: while clearly non-existent for A_L , there is a large and significant mass in other cases.⁵⁸

⁵⁶In the cases of Tit-for-Tat and Suspicious-Tit-for-Tat from the second choice onwards the subject would simply select what the other chose the previous time, but these strategies differ in the period 1 choice. Tit-for-Tat starts by cooperating, while Suspicious-tit-for-tat starts with defection.

⁵⁷We compute the standard deviations for the estimates bootstrapping 1000 repetitions. The procedure leaves unidentified the standard error for the K -th strategy. The estimate of β can be used to interpret how noisy the estimation is. For example, with only two actions a random draw would be consistent with $\beta = 0.5$. Notice that in all cases the estimate of β is relatively high, indicating that the set of strategies used for the estimation can accurately accommodate the data.

⁵⁸The proportion corresponding to Grim is relatively higher for A_M than for A_S . This is consistent with cooperation being more attractive for A_S . It may be that attracted by the gains of cooperation subjects are more willing to forgive and start a new cooperative phase, which is feasible using Tit-for-Tat.

More importantly, notice that strategy AC displays a frequency estimate of approximately 12% that is significant for A_M and A_S . This is the frequency of successful cooperation. In other words, AC captures the mass that may be particularly influential in determining how strong the Markov assumption for structural estimation is. A strategy such as Grim or Tit-for-Tat can only be identified if subjects deviate from cooperation: along the cooperative phase both strategies are identical. But once subjects enter a punishment phase market behavior is closer to the stage Nash, and hence, discrepancies with respect to the MPE assumption for the quantity choice are only present for the periods prior to defection.

Appendix D

Robustness of Structural Estimates

Tables 17, 18, and 19 provide estimates of A , B , and C as the session evolves. The first row in each table shows the estimates when we use all the sample (as reported in the text) and each row reports the estimation as we exclude earlier matches. The last row uses the last five matches of the session.

Overall the estimates for the medium and small market display relatively minor changes as the sample is restricted. We do notice some changes in the estimates of A and B in large market treatments. For the Standard treatment we notice a relatively large change when the sample is restricted to matches 9-16 and onwards. In the No Quantity Choice treatments we notice changes mainly in the estimate of B starting when the sample is restricted to matches 4-16. These changes in the estimates are consistent with the fact that in the large market treatments there are relatively few observations when both subjects are out of the market. As a session evolves the estimates rely on even fewer observations in this state.

Table 17: Sample used and estimates of A

Matches included	A_L		A_M		A_S	
	Standard	No Quantity Choice	Standard	No Quantity Choice	Standard	No Quantity Choice
1-16	0.18	0.22	0.14	0.17	0.10	0.08
2-16	0.18	0.22	0.14	0.17	0.10	0.08
3-16	0.18	0.21	0.15	0.18	0.10	0.08
4-16	0.18	0.17	0.15	0.17	0.10	0.08
5-16	0.18	0.16	0.15	0.18	0.09	0.08
6-16	0.18	0.16	0.15	0.17	0.10	0.08
7-16	0.22	0.16	0.14	0.17	0.10	0.09
8-16	0.23	0.16	0.14	0.17	0.10	0.08
9-16	0.42	0.16	0.14	0.18	0.10	0.09
10-16	0.46	0.16	0.14	0.18	0.09	0.09
11-16	0.68	0.15	0.14	0.21	0.09	0.10

Table 18: Sample used and estimates of B

Matches included	A_L		A_M		A_S	
	Standard	No Quantity Choice	Standard	No Quantity Choice	Standard	No Quantity Choice
1-16	0.11	0.22	0.05	0.19	0.07	0.20
2-16	0.10	0.20	0.05	0.19	0.07	0.20
3-16	0.10	0.19	0.06	0.19	0.07	0.20
4-16	0.10	0.08	0.07	0.19	0.07	0.19
5-16	0.09	0.07	0.05	0.19	0.07	0.17
6-16	0.09	0.07	0.06	0.18	0.09	0.17
7-16	0.20	0.07	0.04	0.18	0.09	0.19
8-16	0.22	0.06	0.05	0.19	0.09	0.19
9-16	0.71	0.05	0.06	0.20	0.09	0.21
10-16	0.78	0.07	0.05	0.23	0.06	0.22
11-16	1.33	0.04	0.04	0.28	0.04	0.25

Table 19: Sample used and Estimates of C

Matches included	A_L		A_M		A_S	
	Standard	No Quantity Choice	Standard	No Quantity Choice	Standard	No Quantity Choice
1-16	0.54	0.56	0.55	0.47	0.53	0.43
2-16	0.55	0.56	0.55	0.47	0.53	0.43
3-16	0.55	0.57	0.56	0.46	0.52	0.44
4-16	0.56	0.56	0.56	0.47	0.52	0.43
5-16	0.55	0.57	0.55	0.48	0.52	0.44
6-16	0.56	0.58	0.57	0.48	0.52	0.45
7-16	0.56	0.58	0.58	0.48	0.52	0.45
8-16	0.57	0.57	0.58	0.47	0.52	0.45
9-16	0.54	0.57	0.57	0.47	0.52	0.46
10-16	0.55	0.57	0.57	0.46	0.53	0.46
11-16	0.52	0.58	0.57	0.43	0.53	0.46

Robustness of Counterfactual Exercises

In this section we consider an alternative counterfactual experiment to measure the prediction bias due to collusion. We use the parameters we estimate for the Standard treatments and predict behavior in the No Quantity Choice treatments. We then contrast this prediction to actual behavior in No Quantity Choice treatments. The results are shown on [Table 20](#), where for a fixed value of A the baseline corresponds to the treatment with no quantity choice, and the counterfactual to the Standard treatment.

Baseline/Counterfactual	A_S	A_M	A_L
Standard/No Quantity Choice			
A_S	24.9 %	-	-
A_M	-	7.2 %	-
A_L	-	-	11.6 %

Table 20: Counterfactual Error From Collusion (MAPE(V)), Standard predicts No Quantity Choice

This table shows the resulting bias in counterfactual computations if the econometrician wrongly assumes that firms do not collude in the baseline (Standard treatment) and predicts the outcomes in a market (of the same market size) where firms do not collude. Then we contrast the prediction to actual behavior in the No Quantity Choice treatments and compute prediction errors.

The counterfactual bias varies from 7% to 25%. By definition, the highest incentives to collude are in the A_S market. Therefore, it is natural that the largest bias takes place in this case, where the magnitude means that averaging across states the difference between predicted continuation value and actual continuation value is 25%. On the one hand, a twenty-five percentage points bias in continuation values is relatively large. On the other hand, how large are these magnitudes relative to a case where firms fully collude in the baseline? To answer this question we perform a Monte Carlo exercise where in the baseline we generate data under full collusion and we estimate parameters under the assumption of no collusion. We then compare the prediction to counterfactual data where there is no collusion. Those results are shown in [Table 21](#)

These results show that the counterfactual bias due to collusion that we observe in the data is smaller than what would be observed if firms were successful at colluding. In the case of the small and medium market it is 11% to 12% of the maximal possible bias and in the large market it is 38% of the maximal possible bias. This means that especially in the scenario where

Baseline/Counterfactual	A_S (Standard)	A_M (Standard)	A_L (Standard)
A_S (No Quantity Choice)	197.1 %	-	-
A_M (No Quantity Choice)	-	62.9 %	-
A_L (No Quantity Choice)	-	-	30.2 %

Table 21: Maximal Counterfactual Error From Collusion, Standard predicts No Quantity Choice (**Monte Carlo Results**)

This table (based on Monte Carlo results and NOT data) shows the maximal resulting bias (measured by $\text{MAPE}(\mathbf{V})$) if the econometrician wrongly assumes that firms do not collude in the baseline and predicts the outcomes in a market (of the same market size) where firms do not collude.

the incentives to collude are high and where they should be quantitatively important we only observe a small fraction of the possible bias.

Appendix E

One of the salient deviations from the theory in our data is that subjects do not enter and exit the market as often as the MPE predicts in this setting. [Table 3](#) summarizes the absolute difference between theoretical MPE probabilities and empirical probabilities across treatments. We refer to the absolute difference between MPE and empirical probabilities in the No Quantity Choice treatments as inertia. 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.

In this appendix we describe three approaches to shed more light on inertia. In the first part of the appendix, we first study three alternative ways to explicitly account for inertia at the estimation stage. The goal is to obtain estimates closer to the true values. We will show that the three procedures actually lead to improvements. The first one directly corrects entry thresholds and uses the fact that we observe the actual thresholds that subjects pick. We then propose two different extended models that explicitly account for inertia and estimate an additional parameter (which do not require to observe the thresholds and could therefore be applied in a standard setting). The latter of the two is the most successful, and for this procedure we also report the full exercise on counterfactuals. We do document that counterfactual prediction errors are lower relative to the counterfactuals we report in the main text. However, the qualitative comparison between counterfactual prediction errors in Standard and No Quantity Choice treatment is qualitatively similar.

In the second part of the appendix we take a different approach and use Monte Carlo simulations to study how inertia affects parameter estimates. We show that inertia, for all reasonable ranges, biases the estimates of C (the entry cost) upwards and B (the competition parameter) downwards. However, the effect on A can be ambiguous. This is consistent with what we observe in the data. A is biased downwards in A_L and A_M but upwards in A_S . We illustrate with Monte Carlo simulations that the upwards bias in $\hat{A}$ can result when inertia is mostly present in exit thresholds as in A_S .⁵⁹

⁵⁹Notice that inspection of [Table 3](#) confirms this is the case in our data. That is, 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.

E.1. Three Procedures to Account for Inertia at the Estimation Stage

Subject-Specific Effects

We first explore to what extent we can recover estimates closer to the true parameter values by accounting for a subject-specific (“fixed-effect”) deviation from the theoretical threshold. Creating a dataset that controls for the average deviation from the MPE threshold of each subject would allow us to reduce the amount of inertia in the data and should result in estimates closer to the true values.

Recall that our experimental design asks subjects directly for their entry and exit thresholds (such thresholds, however, are typically not observed in a standard empirical application). The resulting binary entry and exit choices (which the estimation method uses) are computed based on those thresholds. To generate a dataset that controls for average deviations from theoretical predictions we proceed in three steps. We first compute a subject-specific mean deviation from the theoretically predicted ones. Second, we adjust each threshold by this amount. Finally, we compute a new set of entry-exit choices using the same random draws from the experimental sessions and the adjusted thresholds.⁶⁰ With this alternative dataset, based on “adjusted” entry/exit thresholds, we re-estimate the parameters. For comparability, the top panel of [Table 22](#) reproduces the main estimates in the paper, and the second panel shows the estimation results correcting for subject-specific effects.

What these results show is that this correction broadly moves all estimates in the right direction. Specifically, the entry cost C is substantially reduced, and the A estimates are moving in slightly closer to the true parameters. The estimate of B is also moving closer to the true value of 0.6 in almost all treatments, though it still is far from such value. Of course, correcting the data in this way is not feasible in a standard empirical application because (i) the thresholds are not directly observed and even if they were (ii) the theoretical thresholds are unknown.

Inertia-Augmented Model

We now present our first augmented model that allows to partially correct for inertia without a priori knowledge of thresholds. Since our experiment was not designed to estimate such a

⁶⁰Of course, because the choices implied by adjusted thresholds and the thresholds selected by the subject may not coincide, this alternative data is no longer dynamically consistent. For example, using the adjusted thresholds we may compute that the firm exits in period t , but would be actually observed next period in the market because this is the state that actually occurred in the laboratory.

Treatment	Standard			No Quantity Choice		
	A_S	A_M	A_L	A_S	A_M	A_L
Original (Reproduced from Table 4)						
<i>A</i>	0.1	0.14	0.18	0.08	0.17	0.22
<i>B</i>	0.07	0.05	0.11	0.2	0.19	0.22
<i>C</i>	0.53	0.55	0.54	0.43	0.47	0.56
Thresholds Corrected for Subject-Specific Effects						
<i>A</i>	0.02	0.13	0.3	0.02	0.19	0.22
<i>B</i>	0.24	0.12	0.32	0.3	0.39	0.22
<i>C</i>	0.2	0.19	0.18	0.16	0.18	0.56
Inertia-Augmented Model						
<i>A</i>	0.08	0.12	0.15	0.07	0.15	0.29
<i>B</i>	0.07	0.05	0.1	0.52	0.19	0.43
<i>C</i>	0.18	0.18	0.19	0.25	0.16	0.18
<i>D</i>	0.01	0.01	0.01	0.01	0.1	0.01
Myopic Inertia-Augmented Model						
<i>A</i>	0.04	0.19	0.25	0.04	0.21	0.23
<i>B</i>	0.47	0.69	0.89	0.47	0.83	0.72
<i>C</i>	0.23	0.04	0.06	0.23	0.06	0.09
<i>D</i>	0.15	0.14	0.12	0.15	0.08	0.09

Table 22: Inertia: Parameter estimates

The top part of this table reproduces the parameter estimates of Table 4 and the bottom shows the estimates including inertia parameter D , which captures a potential time inconsistent behavior where the firm assumed that there is an entry or exit cost that has to be paid for changing the market status today but not at any point in the future. True values of parameters are: $A_S = 0.05$, $A_M = 0.25$, $A_L = 0.4$, $B = 0.6$, $C = 0.15$.

model, we see the model we introduce here and in the third approach as more speculative. An intuitive explanation for the observed behavior is that subjects perceive some general switching cost to change the market state. This would call for a straightforward extension of the model with one extra parameter, the perceived switching cost, that needs to be estimated. Unlike the entry cost it would have to be paid both for entering and exiting the market (this is, intuitively, also how it would be separately identified). We now present our first augmented model and start by re-writing the expressions that solve for the MPE recursively. For $x \in \{0, 1\}$, the following function captures the continuation value.

$$\Gamma(x, \mathbf{s}_t) = \sum_{\mathbf{s}_{t+1} \in \mathbf{S}} \int \max\{v^\alpha(0, \mathbf{s}_{t+1}) + \epsilon'(0), v^\alpha(1, \mathbf{s}_{t+1}) + \epsilon'(1)\} dG(\epsilon') F^\alpha(\mathbf{s}_{t+1}|x, \mathbf{s}_t) \forall \mathbf{s}_t \quad (9)$$

A first aspect to notice is that inertia can be thought of as working asymmetrically for a fixed agent i 's state. If agent i is in the market, which happens when the state is in set $\mathbf{s}_t^1 = \{(1, 0), (1, 1)\}$, then 'inertia' can be thought of as providing the agent an extra payment (D^1) if the agent stays in the market, but not if the agent decides to leave. Likewise, if the agent is not in the market, that is if the state is in the set $\mathbf{s}_t^0 = \{(0, 0), (0, 1)\}$, then inertia would provide an extra payment (D^0) if the agent stays out, a payment the agent would not receive if she chooses to enter. That is, the extra payments, D^0 and D^1 act by making it more expensive to leave the state in which the agent is located.

The recursive expressions that solve for MPE are given by:

$$v^\alpha(1, \mathbf{s}_t^1) = \Pi^\alpha(\mathbf{s}_t^1) + D^1 + \delta \cdot \Gamma(1, \mathbf{s}_t^1) \forall \mathbf{s}_t^1 \quad (10)$$

$$v^\alpha(1, \mathbf{s}_t^0) = \Pi^\alpha(\mathbf{s}_t^0) + \delta \cdot \Gamma(1, \mathbf{s}_t^0) \forall \mathbf{s}_t^0 \quad (11)$$

$$v^\alpha(0, \mathbf{s}_t^1) = \Pi^\alpha(\mathbf{s}_t^1) + \delta \cdot \Gamma(0, \mathbf{s}_t^1) \forall \mathbf{s}_t^1 \quad (12)$$

$$v^\alpha(0, \mathbf{s}_t^0) = \Pi^\alpha(\mathbf{s}_t^0) + D^0 + \delta \cdot \Gamma(0, \mathbf{s}_t^0) \forall \mathbf{s}_t^0 \quad (13)$$

Unfortunately, we cannot identify D^0 separate from D^1 , so we impose that $D^0 = D^1 = D$.⁶¹

⁶¹As mentioned earlier, inertia is mostly present in exit thresholds when A is low (A_S), but shifts towards entry thresholds as A increases to A_L (see Table 3). Thus, if it were possible to estimate D^0 separately from D^1 , the fit would be better. Our design was not planned to study inertia and our ability to structurally characterize inertia is thus limited.

Hence, the inertia-augmented model involves estimating one extra parameter (D) and the estimates are presented in the third panel of [Table 22](#).

The estimate of parameter D is relatively small in all treatments. With respect to the parameters in the original model, we note a substantial drop of parameter C , which is closer to the true value in all treatments (relative to the original estimates). There are only small changes in the estimates of A and the estimates of B , while in most cases higher than in the original estimation, remain far from the true value. Even though such a model seems an intuitive alternative to capture inertia, it does not lead to a marked improvements of the estimates.

Myopic Inertia-Augmented Model

The findings with the inertia-augmented model have lead us to explore another alternative, in which inertia is present but in a time-inconsistent manner. 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.

Since in this model inertia only affects the valuation in the present, we can re-express the continuation values as follows:

$$v^\alpha(x, \mathbf{s}_t) = \Pi^\alpha(\mathbf{s}_t) + \delta \cdot \Gamma(x, \mathbf{s}_t) \quad \forall \mathbf{s}_t \quad (14)$$

for $x \in \{0, 1\}$.

However, the current valuations are affected by inertia. We express the value functions of a problem with myopic inertia as follows.

$$\tilde{v}^\alpha(1, \mathbf{s}_t^1) = \Pi^\alpha(\mathbf{s}_t^1) + D^1 + \delta \cdot \Gamma(1, \mathbf{s}_t^1) \quad \forall \mathbf{s}_t^1 \quad (15)$$

$$\tilde{v}^\alpha(1, \mathbf{s}_t^0) = \Pi^\alpha(\mathbf{s}_t^0) + \delta \cdot \Gamma(1, \mathbf{s}_t^0) \quad \forall \mathbf{s}_t^0 \quad (16)$$

$$\tilde{v}^\alpha(0, \mathbf{s}_t^1) = \Pi^\alpha(\mathbf{s}_t^1) + \delta \cdot \Gamma(0, \mathbf{s}_t^1) \quad \forall \mathbf{s}_t^1 \quad (17)$$

$$\tilde{v}^\alpha(0, \mathbf{s}_t^0) = \Pi^\alpha(\mathbf{s}_t^0) + D^0 + \delta \cdot \Gamma(0, \mathbf{s}_t^0) \quad \forall \mathbf{s}_t^0 \quad (18)$$

We cannot identify D^0 separate from D^1 . Since we know that inertia is mostly present in exit threshold in A_S (see Table 3), we estimate $D^1 = D$, $D^0 = 0$ in A_S . Inertia shifts towards entry thresholds in A_M and A_L , so we estimate $D^0 = D$, $D^1 = D$ in these two treatments.⁶²

We estimate D together with the other parameters using a slight modification to the estimation procedure described on Appendix A. We start with some initial values for all parameters and compute the system of equations in (9)-(14). In addition, we use (9)-(14) and the initial guesses to compute (15)-(18). We then use (9)-(14) to compute predicted choice probabilities that we contrast to observed choice probabilities. The procedure is repeated until the parameter estimates provides predicted choice probabilities that are sufficiently close to the observed choice probabilities.

The bottom panel of Table 22 provides the estimates. The estimate of parameter D varies from 0.05 to 0.15, depending on the treatment. With respect to the parameters in the original model, we also note a substantial drop of parameter C , consistent with all corrections for inertia (and the Monte Carlo exercises we report in the next section). As with other corrections, the changes in estimates of A are more nuanced, but there is a large upward adjustment in parameter B , which is closer to the true value of 0.6 in all treatments. Without controlling for inertia the parameter is substantially underestimated, with values between 0.05 and 0.22 and none of the previous corrections is very successful in bringing the estimate close to the true value.⁶³

Given that the myopic augmented-inertia model is relatively more successful in moving all estimates closer to the true values, we inspect it further and use it to report counterfactual predictions. The two top panels of Table 22 reproduce, as a reference, the predictions for counterfactuals that we report in the text. The two bottom panels provide the same computations, but using the myopic augmented-inertia model.

A first observation is that in 10 of 12 predictions, the myopic inertia-augmented model involves a lower prediction error than the original model reported in the text.⁶⁴ This suggests that controlling for inertia can indeed reduce the prediction error. A second observation is that while the prediction error is reduced in the majority of cases, there is no indication of higher prediction error in the Standard treatments relative to the No Market Choice treatments. This

⁶²An empirical researcher would not have information related to which thresholds are mostly affected by inertia, but the purpose of this exercise is to provide an estimation that explicitly controls for inertia as well as possible.

⁶³The Monte Carlo exercises in the next section will show that inertia biases the parameter B downwards.

⁶⁴In the two cases where the prediction is higher in the myopic inertia-augmented model, the increase in the error is relatively small. From 0.38 to 0.39 when the A_L Standard is the baseline and A_M Standard the predicted counterfactual, and from 0.52 to 0.67 when the A_M No Market Choice is the baseline and A_S No Market Choice the predicted counterfactual.

suggests that the possibility of collusion is not leading to an increase in the prediction error when we control for inertia. In a few words, the main comparative static that we report in the paper for counterfactual exercises holds when we use the myopic inertia-augmented model.⁶⁵

Baseline/Counterfactual	A_S	A_M	A_L
Standard			
A_S	-	621.8	949.1
A_M	65.1	-	82.3
A_L	68.5	37.8	-
No Market Choice			
A_S	-	326.4	716.2
A_M	51.6	-	84.3
A_L	64.9	40.8	-
Standard (Myopic Inertia)			
A_S	-	96.9	98.3
A_M	65.4	-	43.9
A_L	66.1	35.3	-
No Market Choice (Myopic Inertia)			
A_S	-	29.3	30.7
A_M	42.7	-	70.2
A_L	50.3	39.2	-

Table 23: Counterfactual prediction error

This table shows the counterfactual predictions using the $MAPE(V)$ measure if we account for inertia by using the myopic inertia-augmented model.

E.2. Monte Carlo Study on Inertia

In this section we use a series of Monte Carlo exercises to study how differences between observed probabilities and MPE probabilities can affect the estimates. We first provide an

⁶⁵The counterfactual exercise we report here does not fully remove inertia from counterfactual prediction errors. We conduct the counterfactual exercise scaling the A estimate and keeping the estimates of B , C and D from the baseline. As explained earlier, at the estimation stage we assume that the researcher knows that inertia operates mostly through exit thresholds in A_S and through both thresholds in other treatments. But for the counterfactual exercise we keep the estimate of D from the baseline. Since inertia can operate differently in baseline and counterfactual, the difference in counterfactual prediction errors can still capture some effect from inertia.

informal discussion on the identification of each coefficient that we illustrate with simulations. In a second step we outline how inertia affects the estimates in each of our treatments. What we document is consistent with our findings in the previous section. In particular, inertia will bias the estimate of C upwards and B downwards.

Table 24 reports several Monte Carlo simulations that connect entry-exit probabilities and coefficient estimates. The baseline reference is simulation 0. The data for the simulation is generated according to probabilities in the first four columns.⁶⁶ These probabilities correspond to the MPE equilibrium in the A_M treatment, and as shown in the last three columns of the Monte Carlo simulation, recovers the true parameters of this treatment. We now illustrate how small optimization errors can generate a bias in each coefficient.

We start with the estimate of A . Notice first that the structural procedure assigns a quantity-stage contemporaneous payoff that includes A whenever the subject is in the market.⁶⁷ This suggests that the estimate of A will be large if subjects display a high propensity to be in the market next period regardless of the state. Consider simulations 1 and 1' of Table 24. In the case of simulation 1 (1') the data is generated with probabilities equal to those of simulation 0 minus (plus) 0.05. In other words, being in the market next period is less (more) likely for all states in simulation 1 (1'). In line with the intuition, the estimate for A in simulation 1 (1') is below (above) the estimate in simulation 0.

Parameter B captures the effect of competition. Intuitively, $\hat{B}$ is identified from comparing the subject's choices depending on whether the other is in the market or not; that is, $p(1, 0) - p(1, 1)$ and $p(0, 0) - p(0, 1)$. In principle, the larger these differences the more the subject reacts to the presence of the other in the market, which means that the effect of competition is larger and the estimate of B will be correspondingly higher. This intuition is confirmed by simulations 2 and 2'. Notice that in the baseline simulation 0 the difference in probabilities is 0.1 in both cases.⁶⁸ In simulation 2 (2') we reduce (increase) the difference in probabilities to 0.05 (0.15) and accordingly observe a reduction (increase) in $\hat{B}$ relative to the baseline.

Finally, the estimate of C will be high when the probabilities indicate that subjects do not want to pay the fixed fee to enter the market, which is consistent with subjects being very likely to remain in the market if they are already in the market and to stay out if they are already out.

⁶⁶For each simulation we proceed as described in section 4. We assume that there are 300 pairs of firms, and generate 100 data sets (according to the probabilities described in the table) that we use to estimate parameters.

⁶⁷Recall that the quantity-stage contemporaneous payoff when the subject is in the market is $100 \times (2A + 0.60)$ if the other is out and $100 \times (2A - B + 0.60)$ if the other is also in. Meanwhile the quantity-stage payoff when the subject is out of the market is 100×0.6 .

⁶⁸ $p(1,0)-p(1,1)=0.70-0.60=0.1$; $p(0,0)-p(0,1)=0.50-0.40=0.1$.

In other words, the estimate $\hat{C}$ is large if there is a large difference between probabilities when the subject is in the market ($p(1, \cdot)$) and probabilities when the subject is out of the market ($p(0, \cdot)$). In simulation 3 (3') we subtract (add) 0.05 to $p(1, \cdot)$ simulation 0 probabilities and add (subtract) 0.05 to $p(0, \cdot)$ simulation 0 probabilities. Correspondingly, we find a lower (higher) estimate of C in simulation 3 (3') relative to the baseline.

Table 24: Monte Carlo simulations and optimization errors

Simulation	Probabilities				Estimates		
	$p(1, 0)$	$p(1, 1)$	$p(0, 0)$	$p(0, 1)$	$\hat{A}$	$\hat{B}$	$\hat{C}$
0	0.70	0.60	0.50	0.40	0.25	0.60	0.15
1	0.65	0.55	0.45	0.35	0.21	0.60	0.15
1'	0.75	0.65	0.55	0.45	0.41	0.50	0.16
2	0.65	0.60	0.45	0.40	0.16	0.28	0.19
2'	0.75	0.60	0.55	0.40	0.34	0.91	0.08
3	0.65	0.55	0.55	0.45	0.43	1.17	0.01
3'	0.75	0.65	0.45	0.35	0.20	0.40	0.26
4	0.95	0.85	0.44	0.40	0.33	0.01	1.47
5	0.80	0.60	0.50	0.30	0.27	0.77	0.17

Note: $p(s)$ indicates the probability of being in the market next period conditional on being in state s in the current period.

Now, we use Monte Carlo simulations to evaluate how inertia can affect the estimates. Suppose that MPE probabilities are given by $\mathbf{p}$. We say that $\mathbf{p}'$ exhibits inertia relative to $\mathbf{p}$ if $p'(1, \cdot) > p(1, \cdot)$ and $p'(0, \cdot) < p(1, \cdot)$. Notice that the exercise described in simulation 3' precisely involves the deviation from simulation 0 probabilities introduced by inertia. Inertia makes transitions between states more rare and is always rationalized with a higher estimate of $\hat{C}$.

With respect to the estimates of A and B inertia can introduce an upwards or downwards bias. The case presented in simulation 3' shows a downwards bias in $\hat{A}$ relative to simulation 0 (0.20 vs. 0.40). This downwards bias in $\hat{A}$ is consistent with the estimates for A_L and A_M

reported in [Table 4](#). In these two treatments there is evidence of inertia in probabilities related to both entry and exit thresholds. In the A_S treatment inertia is present almost exclusively in probabilities related to exit thresholds ($p(1, \cdot)$). Simulation 4 reproduces some of the inertia conditions present in the A_S treatment. Relative to the simulation 0 baseline, we add 0.25 to both probabilities related to exit thresholds, and subtract 0.06 from $p(0, 0)$.⁶⁹ In this case, there is an upwards bias in $\hat{A}$ (0.33 vs. 0.25). Clearly, the equilibrium MPE probabilities in A_S differ from those in the simulation 0 baseline, but we observe a similar effect in the estimates, in particular, with $\hat{A}$ biased upwards.

In all $\hat{B}$ reported in [Table 4](#) there is a downwards bias relative to the true B parameter. This is also consistent with the reports in simulations 3' and 4 that also involve inertia. Although we do not observe it in our data, it is possible to introduce inertia in a way that would bias B upwards. Simulation 5 involves a change in probabilities relative to simulation 0 that is consistent with inertia, but the 'inertia' in $p(1, 0)$ and $p(0, 1)$ is larger than the inertia in $p(1, 1)$ and $p(0, 0)$. In other words, inertia increases how one subject responds to the presence of the other in the market. In this case, we observe an upwards bias in $\hat{B}$ (0.77 vs. 0.60).

We conclude this section with another set of Monte Carlo simulations that focus on how the presence of inertia can bias the coefficients for different levels of collusion. As in [section 4](#), the data are generated from a model in which a proportion $x \in [0, 1]$ of the pairs of firms collude. We use the setup of the A_M treatment, but add (subtract) 0.14 points to exit (entry) threshold probabilities. In other words, the exercise reported in [Figure 9](#) is comparable to [Figure 1](#) except that it includes inertia in the data. The main observations remain unchanged: the estimates of A and C are basically unaffected by the presence of collusion, which does affect the estimate of B , biasing it downwards as collusion increases.

Counterfactuals

[Table 25](#) and [Figure 10](#) provide further details of the counterfactual predictions presented in [subsection 5.3](#). [Table 25](#) presents the predictions for $p(s)$ for the counterfactual exercise reported in [Table 5](#) (predicted 1) and the second exercise reported in [Table 6](#) (predicted 2). The table also displays actual observed probabilities in each case. [Figure 10](#) shows the actual probabilities and the predictions for the first counterfactual exercise (see [Table 5](#)).

⁶⁹As a consequence there is only a small difference between probabilities related to entry thresholds, which is consistent with what we find for A_S (see [Appendix B](#)).

<i>Baseline A_L</i>							
	state	A_S			A_M		
		predicted 1	predicted 2	actual	predicted 1	predicted 2	actual
<i>Standard</i>	$p(1, 0)$	0.56	0.74	0.76	0.79	0.87	0.88
	$p(1, 1)$	0.50	0.68	0.71	0.72	0.80	0.85
	$p(0, 0)$	0.02	0.21	0.21	0.24	0.31	0.34
	$p(0, 1)$	0.00	0.15	0.20	0.18	0.24	0.30
<i>No Quantity</i>	$p(1, 0)$	0.58	0.66	0.71	0.87	0.87	0.87
	$p(1, 1)$	0.47	0.59	0.54	0.72	0.76	0.77
	$p(0, 0)$	0.01	0.21	0.18	0.27	0.37	0.37
	$p(0, 1)$	0.00	0.12	0.17	0.12	0.26	0.29
<i>Baseline A_M</i>							
	state	A_S			A_L		
		predicted 1	predicted 2	actual	predicted 1	predicted 2	actual
<i>Standard</i>	$p(1, 0)$	0.58	0.76	0.76	1.00	0.98	0.95
	$p(1, 1)$	0.55	0.73	0.71	1.00	0.95	0.89
	$p(0, 0)$	0.03	0.23	0.21	0.55	0.44	0.38
	$p(0, 1)$	0.00	0.20	0.20	0.53	0.41	0.34
<i>No Quantity</i>	$p(1, 0)$	0.59	0.66	0.71	1.00	1.00	0.94
	$p(1, 1)$	0.49	0.58	0.54	0.97	0.89	0.89
	$p(0, 0)$	0.09	0.22	0.18	0.57	0.45	0.43
	$p(0, 1)$	0.01	0.14	0.17	0.50	0.33	0.31
<i>Baseline A_S</i>							
	state	A_M			A_L		
		predicted 1	predicted 2	actual	predicted 1	predicted 2	actual
<i>Standard</i>	$p(1, 0)$	1.00	0.88	0.88	1.00	0.97	0.95
	$p(1, 1)$	1.00	0.83	0.85	1.00	0.93	0.89
	$p(0, 0)$	1.00	0.34	0.34	1.00	0.43	0.38
	$p(0, 1)$	1.00	0.28	0.30	1.00	0.39	0.34
<i>No Quantity</i>	$p(1, 0)$	1.00	0.87	0.87	1.00	1.00	0.94
	$p(1, 1)$	1.00	0.77	0.77	1.00	0.91	0.89
	$p(0, 0)$	0.81	0.37	0.36	1.00	0.45	0.43
	$p(0, 1)$	0.78	0.28	0.29	1.00	0.33	0.31

Notes: Predicted 1 (Predicted2) presents probabilities related to the counterfactual exercise reported in [Table 5](#) ([Table 6](#)).

Table 25: Counterfactual calculations

Figure 9: Parameter estimates under different collusion probabilities for the A_M treatment with inertia.

Table 26 provides information on counterfactual calculations of the probability of observing a monopoly or a duopoly using the counterfactual exercise reported in Table 5.

Baseline treatments: Row 1- A_L , Row 2- A_M , Row 3- A_S

Figure 10: Counterfactual predictions

<i>Baseline A_L</i>							
	probability	A_L		A_S		A_M	
		predicted	actual	predicted	actual	predicted	actual
<i>Standard</i>	p_2	0.72	0.74	0.00	0.24	0.26	0.61
	p_1	0.24	0.23	0.04	0.35	0.38	0.29
<i>No Quantity</i>	p_2	0.74	0.71	0.00	0.13	0.25	0.50
	p_1	0.25	0.25	0.03	0.32	0.47	0.35
<i>Baseline A_M</i>							
	probability	A_M		A_S		A_L	
		predicted	actual	predicted	actual	predicted	actual
<i>Standard</i>	p_2	0.57	0.61	0.00	0.24	1.00	0.74
	p_1	0.30	0.29	0.03	0.35	0.00	0.23
<i>No Quantity</i>	p_2	0.49	0.50	0.01	0.13	0.94	0.71
	p_1	0.36	0.35	0.18	0.32	0.06	0.25
<i>Baseline A_S</i>							
	probability	A_S		A_M		A_L	
		predicted	actual	predicted	actual	predicted	actual
<i>Standard</i>	p_2	0.24	0.24	1.00	0.61	1.00	0.74
	p_1	0.35	0.35	0.00	0.29	0.00	0.23
<i>No Quantity</i>	p_2	0.12	0.13	0.04	0.50	1.00	0.71
	p_1	0.32	0.32	0.23	0.35	0.00	0.25

Table 26: Counterfactual calculations for the probability of observing a monopoly (p_1) or a duopoly (p_2)

Appendix F

This appendix describes several additional details of the experimental implementation. We implement the infinite time horizon as an uncertain time horizon (Roth and Murnighan, 1978). After each period of play, there is one more period with probability $\delta = 0.8$. We implement the uncertain time horizon using a modified block design (Fréchette and Yuksel, 2013). Subjects play the first five periods without being told after each period whether the supergame has ended or not. Once period five ends they are informed whether the game ended in any of the first five periods. Only periods prior to ending count for payoff (including the period when the game ended). From the sixth period onwards subjects are told period by period whether the game ended or not, and cumulative payoffs are computed for all periods until the game ends. This procedure allows us to collect information for several periods without affecting the theoretical incentives.⁷⁰

Our sessions are divided into two parts. The difference between parts is in how subjects report their dynamic choice to the interface. In Part 1, in the exit (entry) stage subjects are informed of the randomly selected exit payment (entry fee) and then decide whether to exit or not (enter or not). In Part 2 subjects first specify an exit threshold (entry threshold), that is a number between $[0, 100]$, with the understanding that if the exit payment is higher than the threshold (entry fee is lower than the threshold) they will exit the market (enter the market).⁷¹ Part 1 consists of 1 supergame and Part 2 consists of the remaining 15 supergames.⁷²

⁷⁰See Fréchette and Yuksel (2013) for a comparison between this and other alternatives to implement infinite time horizons in the laboratory.

⁷¹[Appendix G](#) presents the instructions, screenshots of the interface and describes how subjects made their choices. In the case of the entry fee subjects specify a threshold in $[15, 115]$, which includes the fixed portion of the entry fee. For the purpose of analysis in the paper we will always present entry thresholds net of the fixed entry fee.

⁷²The structural estimation procedure uses only information on whether subjects are in the market or not for estimation. Our implementation in part 2 provides us with additional information: we know the threshold of their decision. We use the additional information to evaluate the aggregate information content of only using the binary information for being in the market or not. We find the binary information to be consistent with thresholds if aggregate estimates on frequencies per state are of a comparable magnitude. If this were not the case, then using only binary information may already introduce a bias in the estimation. However, in the data we do not find that using only whether firms are in the market or not would lead to a bias.

Appendix G

This appendix provides the instructions for the Standard treatment for A_L . The instructions consist of two parts. The first part presents the environment for the first cycle, and part 2 introduces the thresholds for entry/exit decisions. Instructions for No Quantity Choice treatments are identical except that we do not present a table for the quantity choice decision, and instead subjects are told they would receive Nash payoff when they are both in the market. After the instructions there is a set of figures with screen shots of the interface.

INSTRUCTIONS

Welcome

You are about to participate in an experiment on decision-making. What you earn depends partly on your decisions, partly on the decisions of others, and partly on chance. Please turn off cell phones and similar devices now. Please do not talk or in any way try to communicate with other participants. We will start with a brief instruction period. During the instruction period you will be given a description of the main features of the experiment. If you have any questions during this period, raise your hand and your question will be answered so everyone can hear.

General Instructions: Part 1

1. This experiment is divided into 16 cycles. In each cycle you will be matched with a randomly selected person in the room. In each cycle, you will be asked to make decisions over a sequence of rounds.
2. The number of rounds in a cycle is randomly determined as follows:
 - After each round, there is an 80% probability that the cycle will continue for at least another round of payment.
 - At the end of each round the computer rolls a 100-sided die.
 - If the number is equal to or smaller than 80, there will be one more round that will count for your payments.
 - If the number is larger than 80, then subsequent rounds stop counting toward your payment.

- For example, if you are in round 2, the probability that the third round will count is 80%. If you are in round 9, the probability round 10 also counts is 80%. In other words, at any point in a cycle, the probability that the payment in the cycle continues is 80%.
3. You interact with the same person in all rounds of a cycle. After a cycle is finished, you will be randomly matched with a participant for a new cycle. In each round, your payoff depends on your choices and those of the person you are paired with. In each round there is a market stage and an entry/exit stage. In the entry/exit stage you and the other will decide whether to enter or exit the market. We first explain the market stage and later we explain the entry/exit stage.
 4. At the beginning of each cycle (in Round 1) you and the other start in the market. You and the other will first make the market stage choices and then decide whether you want to stay in the market or exit.

Market Stage

5. When you and the other are both in the market, your payoff depends on your choice and the choice of the other:
 - If you select 1, and the other selects 1, your payoff is 100, and the other's is 100.
 - If you select 1, and the other selects 2, your payoff is 40, and the other's is 140.
 - If you select 2, and the other selects 1, your payoff is 140, and the other's is 40.
 - If you select 2, and the other selects 2, your payoff is 80, and the other's is 80.

The table below summarizes all the possible outcomes:

		Other's Choice	
		1	2
Your Choice	1	100,100	40,140
	2	140,40	80,80

In this table, the rows indicate your choices and the columns the choices of the person you are paired with. The first number of each cell represents your payoff, and the second number (in italics) is the payoff of the person you are paired with.

6. If in any round you are in the market and the other is out, your payoff will be equal to 140.
7. If in any round you are out of the market you make a payoff of 60.
8. Once the market stage is over, you will start the entry /exit stage.

Entry/Exit Stage

9. Exit decision. In each round when you are in the market, you will have to decide whether you want to exit the market or not. If you exit the market you will receive an exit payment. The exit payment is a random number between 0 and 100. All numbers are equally likely. The randomly selected exit payment will be presented to you on the screen. You will have to indicate whether you want to take the exit payment and exit the market or not take the payment and stay in the market.
10. The exit payment is selected separately for each participant. That means that you will have one exit payment and when the other is selecting whether to exit or not, they will have another randomly selected exit payment. The exit payment is selected randomly in each round. This means that exit payments in different rounds will likely be different.
11. Entry decision. If in any round you are out of the market, you have to choose whether you want to enter the market or not. To enter the market you have to pay an entry fee. The entry fee is a random number between 15 and 115. All numbers are equally likely. The randomly selected entry fee will be presented to you on the screen. You will have to indicate whether you want to pay the entry fee and enter the market or not pay the fee and stay out of the market.
12. The entry fee is selected separately for each participant. That means that you will have one entry fee and when the other is selecting whether to enter or not, they will have another randomly selected entry fee. The entry fee is selected randomly in each round. This means that entry fees in different rounds will likely be different.
13. In each round after round 1 you first face the market stage and then the entry /exit stage. If you are in the market in that round you will have to decide whether to exit or not. If you are out of the market you will have to decide whether to enter or not.

Payoffs

14. In each cycle you start with 30 points, and you will make choices for the first 5 rounds without knowing whether or not the cycle payment has stopped. At the end of the fifth round the interface will display on the screen the results of the 100-sided die roll for each of the first 5 rounds.
15. If the roll of the 100-sided die was higher than 80 for any of the first five rounds, the cycle will end, and the last round for payment is the first where the 100-sided die roll is higher than 80.
 - The interface subtracts entry fees that you pay, adds exit payments, and adds all points that you make in the market stages of all rounds that count for payment within a cycle.
 - For example, assume that the 100-sided die in the first five rounds results in: 40, 84, 3, 95, 65. Because 84 is higher than 80, payments will stop after the second round. The interface will add your market and entry/exit payoffs for rounds 1 and 2.
16. If the 100-sided die rolls were lower than or equal to 80 for the first five rounds, there will be a sixth round. From the sixth round onwards the interface will display the 100-sided die roll round by round. The cycle will end in the first round where the 100-sided die roll is higher than 80.
 - The interface subtracts entry fees that you pay, adds exit payments, and adds all points that you make in the market stages of all rounds that count for payment within a cycle.
 - For example, assume that the 100-sided die in the first five rounds results in: 51, 24, 13, 80, 55. Because all numbers are equal to or lower than 80 there will be another round, so the cycle continues to round 6. After you make your choices for round 6 you are shown that the 100-sided die for that round is 52, which is lower than 80 so there will be a seventh round. After round 7 you are shown that the 100-sided die for that round is 91. Because 91 is higher than 80 the cycle is over and the interface will add your payoffs for all rounds 1 through 7.
17. If at any point in the cycle your total payoff for the cycle is less than 0, the cycle is over.

18. Your total payoffs for the session are computed by adding the total payoffs of all 16 cycles. These payoffs will be converted to dollars at the rate of 0.0025\$ for every point earned.

Are there any questions?

Summary

Before we start, let me remind you that:

- The length of a cycle is randomly determined. After every round there is an 80% probability that the payment cycle will continue for another round.
- In Round 1 of each cycle you and the other start in the market.
- Each Round has a market stage and an entry/exit stage.

1. Market Stage Payoffs

		Other	
		In the Market	Out of the Market
You	In the Market	See Payoff Table	140
	Out of the Market	60	60

2. Exit/Entry Decision

- If you are out and decide to enter, you will pay the entry fee. If you stay out, you do not have to pay any fee.
- If you are in and decide to leave, you will be paid the exit payment. If you decide to stay in, you will not receive an extra payment.
- You interact with the same person in all rounds of a cycle. After a cycle is finished, you will be randomly matched with a participant for a new cycle.
- Part 1 consists of 1 cycle. Once the first cycle is over we will give you brief instructions for Part 2 that will consist of 15 cycles. The only difference between Part 1 and Part 2 will be on how you report your choices to the interface. Other than that Part 1 and Part 2 are identical.

General Instructions: Part 2

1. The only difference in Part 2 is on how you report to the interface your entry/exit decisions.

Exit Decision

2. Instead of deciding if you want to Exit or Stay In the market for a particular Exit Payment, you will report an Exit Threshold.
3. You will report your Exit Threshold before you learn the Exit Payment that was randomly selected.
4. The Exit Threshold specifies the minimum Exit Payment you would take to exit the market. If the Exit Payment were to be higher than your choice for the Exit Threshold, then you would exit the market and receive the Exit Payment. If the Exit Payment were to be equal to or lower than your choice for the Exit Threshold, then you will Stay In the market and not receive the Exit Payment.
5. After you submit your choice for the Exit Threshold the interface will show you the randomly selected Exit Payment and will implement a choice for your Exit Threshold.

Entry Decision

6. Instead of deciding if you want to Enter or Stay Out of the market for a particular Entry Fee, you will report an Entry Threshold.
7. You will report your Entry Threshold before you learn the Entry Fee that was randomly selected.
8. The Entry Threshold specifies the maximum Entry Fee below which you are willing to pay to enter the market. If the Entry Fee were to be higher than or equal to your choice for the Entry Threshold, then you would not enter the market and not pay the Entry Fee. If the Entry Fee were to be lower than your choice for the Entry Threshold, then you would enter the market and pay the Entry Fee.
9. After you submit your choice for the Entry Threshold the interface will show you the randomly selected Entry Fee and will implement a choice for your Entry Threshold.

Are there any questions?

Screenshots of the Interface

Figure 11 displays the first screen that subjects see when the experiment starts in the case of a Standard treatment with $A = 0.4$. At the top left subjects are reminded of general information: the cycle and rounds within the cycle. The blank part on the left side of the screen will be populated with past decisions as the session evolves. In round 1 of every cycle both start in the market, which they are reminded of at the top right. The quantity stage table is presented below and subjects are asked to select a row. In the laboratory we refer to the quantity stage as the market stage.

Figure 12 shows a case where the first row has been selected. As soon as a row is selected a 'submit' button appears. Subjects can change their choice as long as they haven't clicked on the 'submit' button. Figure 13 shows an example of the feedback subjects get in a case where the subject selected row 2 and the other selected row 1.

After a quantity stage subjects face the entry/exit stage. Figure 14 shows an example of an exit stage in cycle 1. Subjects are presented with a randomly selected scrap value and they simply indicate if they exit or stay in the market. An entry stage is similar, except that subjects decide between 'enter' or 'stay out.'

Figure 15 shows an example when there is no quantity decision in the quantity stage. Subjects in this case are simply informed of their quantity stage payoff. This screen is qualitatively similar to what subjects in the No Quantity Choice treatment see if they are both in the market. This screenshot, which corresponds to a case in round 5, also shows on the left side the table with past decisions in the current cycle. As the session evolves subjects also have access to choices for previous cycles. They simply enter the number for the cycle for which they wish to see their feedback in the box after 'History for Cycle' and click on 'Show.'

Cycle: 1 Round: 1	Market Stage: You In, Other In Please select a row <table border="1" style="margin: auto; border-collapse: collapse;"> <thead> <tr> <th colspan="3" style="text-align: center; padding: 2px;">The Other's Choice</th> </tr> <tr> <th style="width: 10%;"></th> <th style="width: 20%;"></th> <th style="width: 20%; text-align: center;">1</th> <th style="width: 20%; text-align: center;">2</th> </tr> </thead> <tbody> <tr> <td rowspan="2" style="width: 10%; text-align: center; vertical-align: middle;">Your Choice</td> <td style="width: 20%; text-align: center;">1</td> <td style="text-align: center;">100, 100</td> <td style="text-align: center;">40, 140</td> </tr> <tr> <td style="text-align: center;">2</td> <td style="text-align: center;">140, 40</td> <td style="text-align: center;">80, 80</td> </tr> </tbody> </table>	The Other's Choice					1	2	Your Choice	1	100, 100	40, 140	2	140, 40	80, 80
The Other's Choice															
		1	2												
Your Choice	1	100, 100	40, 140												
	2	140, 40	80, 80												

Figure 11: Quantity Stage of Standard Treatment with $A = 0.4$.

Cycle: 1 Round: 1	Market Stage: You In, Other In Please select a row <table border="1" style="margin: auto; border-collapse: collapse;"> <thead> <tr> <th colspan="3" style="text-align: center; padding: 2px;">The Other's Choice</th> </tr> <tr> <th style="width: 10%;"></th> <th style="width: 20%;"></th> <th style="width: 20%; text-align: center;">1</th> <th style="width: 20%; text-align: center;">2</th> </tr> </thead> <tbody> <tr> <td rowspan="2" style="width: 10%; text-align: center; vertical-align: middle;">Your Choice</td> <td style="width: 20%; text-align: center;">1</td> <td style="text-align: center;">100, 100</td> <td style="text-align: center;">40, 140</td> </tr> <tr> <td style="text-align: center;">2</td> <td style="text-align: center;">140, 40</td> <td style="text-align: center;">80, 80</td> </tr> </tbody> </table> <div style="margin-top: 5px;"> <input style="width: 40px; height: 15px; border: 1px solid black; background-color: red; color: white; font-size: 8px;" type="button" value="Submit"/> </div>	The Other's Choice					1	2	Your Choice	1	100, 100	40, 140	2	140, 40	80, 80
The Other's Choice															
		1	2												
Your Choice	1	100, 100	40, 140												
	2	140, 40	80, 80												

Figure 12: Example where the subject selects row 1.

Cycle: 1		Round: 1			
		Market Stage: You In, Other In			
		Outcome in This Round			
		The Other's Choice			
		Your Choice		1	2
			1	100, 100	40, 140
2	140, 40	80, 80			
		In this round your payoff will be: 140			
		<input type="button" value="MoveOn"/>			

Figure 13: Example of Feedback when the subject selects 2 and the other selects 1.

Cycle: 1		Round: 1	
		In the Current Round: You In, Other In	
		Market Stage: You selected 1	
		Market Stage: Other selected 2	
		EXIT STAGE	
		If you exit the market, you receive an exit payment of: 35	
		Please click on "Exit" or "Stay In"	
		<input type="button" value="Exit"/> <input type="button" value="Stay In"/>	

Figure 14: Example of an exit stage decision in cycle 1.

Cycle: 1
Round: 5

History for Cycle: 1

Show

Round	Market Stage: You	Market Stage: Other	Market Payoff	Entry Fee	Exit Payment
1	In	In	80	-	-
2	In	In	80	-	97
3	Out	Out	60	-	-
4	Out	Out	60	23	-

Market Stage: You In, Other Out

The other is not in the Market. In this round your payoff will be: 140

MoveOn

Figure 15: Example of quantity stage where there is No Quantity Choice.

Figure 16 presents an example of the exit decision for part 2 of the session. Subjects can select a threshold by clicking anywhere on the black line. Once they click on the horizontal black line a red vertical line appears with a red number indicating the choice. In the example the subject selects a threshold of 51. Once a choice is made the interface indicates with arrows the values of the scrap value for which the subject would exit or stay in the market. Subjects can change their choice by clicking anywhere else on the black line. They can also adjust their choice by clicking on the plus/minus buttons at the bottom. Each click in the plus (minus) button adds (subtracts) one unit to the current threshold. Once they click on 'submit' their choice is final.

Finally, Figure 17 shows an example of the feedback that subjects receive after they make an exit decision. First they are informed of the randomly selected scrap value (exit payment), then they are reminded of the threshold they finally submitted. Given these two values they are informed of the final decision.

Cycle: 1	Round: 1	In the Current Round: You In, Other In
		Market Stage: You selected 1 Market Stage: Other selected 2
		EXIT STAGE Make Your choice for the EXIT PAYMENT threshold
		For an Exit Payment higher than 51 exit
		For an Exit Payment lower or equal than 51 stay in the market <div style="text-align: center;"> - + </div> <div style="text-align: center;"> <input type="button" value="Submit"/> </div>

Figure 16: Example of an exit stage decision in cycles 2-16

Cycle: 1	Round: 1	In the Current Round: You In, Other In
		The Exit Payment this round is: 60 The Exit Threshold you selected was: 51 Because the Exit Payment is higher than your Exit Threshold you will Exit the Market and will be paid the Exit Payment.
		<input type="button" value="Move On"/>

Figure 17: Feedback after exit stage decision in cycles 2-16