Coordination with Endogenous Contracts: A Structural Learning Model Approach

David J. Cooper *†
University of Iowa
University of East Anglia

Christos A. Ioannou ‡
Université Paris 1
Panthéon - Sorbonne

Shi Qi §
William and Mary

This draft: August 1, 2023

Abstract

Structural models lie at the heart of empirical economic analysis. In this study, we introduce a structural learning model with sophistication to examine the robustness of the results in Cooper, Ioannou, and Qi (2018) as well as investigate their broader implications. The model is fit to the full dataset using the method of simulated moments. *Ex ante* heterogeneity in the beliefs of unsophisticated learners as well as heterogeneity between unsophisticated and sophisticated learners play a central role in the model's ability to track the differences between treatments and features of the experimental design. Counterfactual exercises evaluate the robustness of the earlier results and investigate their broader implications.

*The usual disclaimer applies.

†Mailing Address: Department of Economics, University of Iowa, Iowa City, IA. Email: david-j-cooper@uiowa.edu

‡Mailing Address: Centre d'Économie de la Sorbonne, Université Paris 1 Panthéon - Sorbonne, Paris, France. Email: christos.ioannou@univ-paris1.fr

§Mailing Address: Department of Economics, William and Mary, Williamsburg, VA. Email: shiqi.econ@gmail.com

1 Introduction

If pay is based on group performance, the presence of strategic complementarities generated by a weak-link technology coupled with agents' pessimistic beliefs can lead to *productivity traps*; that is, all agents exert low effort in the (correct) expectation that any effort will be wasted given that no individual can unilaterally improve productivity. Coordinating on low effort levels is typically modeled as a Pareto dominated equilibrium in a weak-link game.

Experimental economists have studied a number of mechanisms for escaping productivity traps.[1] Increasing incentives to coordinate at high rather than low effort levels is a natural approach given existing evidence on the efficacy of performance pay (e.g. Lazear (2000); Eriksson and Villeval (2008); Dohmen and Falk (2011); Bandiera, Guiso, Prat, and Sadun (2015)). In spite of the presence of strategic uncertainty, increased incentives have been also shown to increase productivity in multiple studies of the weak-link game, where the sssignment of individuals to incentive contracts was random and *exogenous* (Brandts and Cooper (2006); Hamman, Rick, and Weber (2007); Brandts, Cooper, and Weber (2015b)). Yet, in field settings, workers often choose between jobs with differing incentive contracts, making the assignment of individuals to incentive systems *endogenous*. Existing studies, in non-interactive settings (e.g. Cadsby, Song, and Tapon (2007)), consistently find that much of the impact of increasing incentive pay can be attributed to selection caused by endogenous assignment. Selection works differently in interactive settings, like the weak-link game, but the same basic intuition applies. The effects of incentive pay may thus be systematically underestimated if assignment to incentive contracts is exogenous.

Cooper, Ioannou, and Qi (2018) (henceforth, referred to as CIQ for brevity) addressed this issue by studying how *endogenous* assignment of incentive contracts through a market mechanism affects the ability of groups to escape productivity traps. From a global point of view, CIQ showed that the overall effect of endogenous assignment to incentive contracts was neutral; a strong positive effect of *endogenous* assignment to high performance pay was offset by an equally large negative effect for groups that endogenously select low performance pay. CIQ decomposed the effect of endogenous assignment into selection and strategic anticipation effects. Selection effects occur because individuals who are relatively optimistic (pessimistic) about the

[1] See Cooper and Weber (2020) for a survey.

odds of successful coordination at high effort levels select into groups with low (high) *fixed pay* and high (low) *performance pay*. Outcomes in weak-link games are strongly correlated with initial beliefs, hence selection leads to higher (lower) effort with high (low) performance play. Strategic anticipation refers to the ability of some subjects to correctly anticipate the positive effects of selection. Strategic anticipation causes groups with high (low) performance pay to become even more optimistic (pessimistic) and achieve higher (lower) effort levels. CIQ reported that selection accounts for *at least 73%* of the positive effect of endogenous assignment to high performance pay contracts, but this is a lower bound and the effect could be larger.

CIQ left some important questions unanswered. Precisely how much of the effect of endogenous assignment to performance pay can be attributed to selection? How would the effects of endogenous assignment to performance pay change if the environment was different? For example, how would the results change if group sizes were larger or workers had differing ability levels? The primary purpose of this study is to address these questions by developing, fitting, and simulating a structural econometric model of learning in weak-link games.

Given the central role played by beliefs in weak-link games, we use a belief-based model of learning in the spirit of stochastic fictitious play (Fudenberg and Levine (1998)). This basic model cannot capture the effects of selection and strategic anticipation, so we modify it by adding heterogeneity along subjects' initial beliefs and subjects' sophistication in modelling the learning of others. Individuals are either unsophisticated or sophisticated types. Unsophisticated types are subdivided into optimists and pessimists. As the names imply, optimists (pessimists) initially expect other agents to choose high (low) effort levels. In each round, unsophisticated types best respond (with noise) to their beliefs which are updated from round to round based on the observed choices of others. Unsophisticated types do *not* explicitly model the decision rules of other agents, and therefore cannot anticipate how changed incentives or learning will affect the behavior of others. Similar to a level-k model, sophisticated types model all other individuals in their group as unsophisticated types, anticipating their learning and decision rules. Critically, because they explicitly model the decision making of other agents, sophisticated types anticipate the effects of selection.

We fit the model to the full dataset using the method of simulated moments. Simulations show that the fitted model captures the main qualitative features of the data. The structural model is then modified to estimate the true effect of selection. The intuition is simple: the 73% estimate reported by CIQ relied on a strong assumption that the experimenters could perfectly predict which incentive contract, high or low performance pay, an agent would be assigned to. In reality, these predictions were imperfect, biasing the estimated selection effect downwards. We modify the structural model to account for errors in predicting which incentive contract will be assigned to an agent. Doing this, we conclude that more than 90% of the effect of endogenous assignment to incentive contracts can be attributed to selection. In other words, virtually all of the effect of endogenous assignment to incentive pay is due to selection.

Including (i) heterogeneity in the initial beliefs of unsophisticated learners, and (ii) heterogeneity in the ability of agents to anticipate the decision rules of others improves the fit of the model to the data. This does not necessarily imply that adding heterogeneity along both dimensions improves the model's ability to track differences across treatments as treatment effects are *not* part of the objective function. We investigate the importance of heterogeneity by fitting alternative models that either eliminate one or both types of heterogeneity. Based on simulations, the model with no heterogeneity does poorly at tracking the data, largely missing the effect of endogenous selection. Adding only heterogeneous initial beliefs captures the major features of the data better than adding only heterogeneous sophistication, and adding sophisticated learners *in addition to* heterogeneous initial beliefs only marginally improves the model's ability to track the main qualitative features of the data. This is not to deny the presence of sophisticated learners in the population. Rather, the point is to distinguish between statistical and economic significance. Selection drives most of the effects of endogenous assignment to incentive contracts, and heterogeneous priors provide a parsimonious model of selection effects. Sophisticated types are necessary to capture strategic anticipation, but this only plays a minor role in generating the effects of endogenous assignment to incentive contracts.

Any experiment is necessarily a limited exercise, exploring a single environment and only considering a limited number of treatments. Having a structural model makes it possible to inexpensively explore counterfactuals; but an obvious concern with using simulations to make predictions beyond the confines of the original experimental design is the possibility that the model cannot predict behavior out of sample. The experimental design of CIQ makes it straight forward to address this issue. Three of their treatments (Random Assignment, Auction, and Sort) were used to identify the selection and strategic anticipation effects, and thus only data from these treatments are necessary to identify the parameters of the structural model. The fourth treatment (Reverse Sort) was used to illustrate the strength of selection. For our purposes, data from the Reverse Sort treatment serves as a good test of the model's ability to fit out of sample. We fit the data to the other three treatments and then use simulations to predict outcomes for the Reverse Sort treatment. The simulated results track the main features of the Reverse Sort treatment, giving us greater confidence in predictions based on the model.

Having established the model's validity, we use it to explore variations of the CIQ experimental design. In one variation, we allow for heterogeneous initial beliefs (pessimists vs. optimists) and heterogeneous ability. Heterogenous ability is introduced into the model through differing effort costs, where the higher the individual's ability the lower the effort cost. Social surplus is maximized when agents with low effort costs, regardless of their initial beliefs, select the high performance pay. Yet the simulations find that selection of the increased performance pay prioritizes optimism over low effort costs. This promotes efficiency through high effort choices, but fails to achieve full efficiency due to frequent selection of individuals with high effort costs. In another variation, we change the group size. We show that less of the impact of endogenous assignment can be attributed to the selection effect as the group size grows. Intuitively, efficient coordination becomes more difficult with larger group sizes. Selection into high performance pay always has a large positive effect, but with larger group sizes there remains more room for additional improvement due to strategic anticipation.

The contribution of our paper is not the learning model per se. The basic model is based on the well-known stochastic fictitious play model of Fudenberg and Levine (1998), and several existing models feature sophisticated learners, notably the sophisticated EWA model of Camerer, Ho, and Chong (2002). Nor are we the first to use a learning model to study behavior in weak-link games (e.g. Brandts and Cooper (2006); Brandts, Cooper, Fatas, and Qi (2015a)). Instead, our contribution lies in illustrating how a structural model of learning can be used to expand what can be learned from the experimental dataset. Some of what we do resembles what other papers have already explored. For example, Crawford and Broseta (1998) use a learning model to decompose the effects of buying the right to participate in a median game into optimistic subjects and forward induction effects closely akin to our selection and strategic anticipation effects. Their results are quite different, putting roughly even weight on the two factors. Both the games being played and the learning models differ quite a bit between their paper and ours. Other papers have identified the presence of sophisticated learners.

What differs in our paper is the focus on economic rather than statistical significance; allowing for sophisticated learners improves the fit of the model, but does little to improve the model's ability to track major features of the data. The greatest source of novelty in our work comes from our analysis of counterfactuals. For any good experiment, it is easy to think of informative new treatments. All experiments are limited by the choice of specific parameter values (e.g. group size, effort costs to name a few). In a world of finite budgets and subject pools, it is unrealistic that all possible treatments and constellations of parameters can be explored via new sessions. We argue that counterfactuals provide a useful tool for expanding the domain of experimental results. Obviously, this must be done with a degree of caution. We careful validate our model's ability to predict out of sample and recognize that predictions based on simulation exercises necessarily lose validity the further we get from the original experimental environment. Subject to these caveats, counterfactual exercises vastly expand our understanding of how endogenous assignment to incentive contracts affects performance.

The paper adheres to the following plan. We provide next a brief description of the experimental design in CIQ. Then, in Section 3, we present the structural model of learning. In Section 4, we provide the details of the estimation procedure while, in Section 5, we conduct counterfactual exercises. Finally, in Section 6, we conclude and offer suggestions for future research.

2 Experimental Design

The experimental design in Cooper, Ioannou, and Qi (2018) applies a variation of the Turnaround game, where game play consists of twenty rounds, split into two blocks of ten rounds each. The subjects play in groups of four. In the first block, the subjects are randomly assigned to their group to play the Coordination game with a low incentive pay. The first block is designed to induce coordination failure. In the second block, the assignment of subjects into groups as well as their assignment to low or high pay contracts varies by the treatment. The incentive system in the experimental design implements a linear revenue-sharing scheme as *only* the group output, which is equivalent to the minimum effort of the group, is observed. The payoff of a subject depends on a fixed pay, the cost of effort based on the subject's effort choice, and a bonus that depends on the minimum effort of the group. Two different types of contracts are used in the experiment. Relative to the contract with high incentive pay (i.e. Contract 2), the contract with low incentive pay (i.e. Contract 1) assumes a higher fixed pay, but a lower bonus factor. In summary, Contract 2 features higher incentives to coordinate at high effort levels than Contract 1.

CIQ implement the following four treatments. In the *Random* treatment, the subjects are randomly assigned to groups and contracts. In the *Auction* treatment, the subjects participate in a reverse English auction in which bids indicate willingness to take lower fixed pay in exchange for increased performance pay. The auction mimics, albeit imperfectly, the selection process in the market mechanism. In the *Sort* treatment, the subjects' characteristics and initial choices are used to predict which incentive contract the subjects would have been assigned to had the selection occurred via the market mechanism. Predicted contract assignments are exogenously implemented to (imperfectly) reproduce the market's outcome.[2] The *Reverse Sort* treatment, flips the contract assignments relative to the Sort treatment. Thus, low performance pay is assigned to subjects who would probably be assigned to high performance pay through the market mechanism (and vice versa). The interested reader should refer to the paper for the details.

3 Structural Learning Model

We next present the model's setup. We then indicate the explicit assumptions in the beliefs and decision rules of (sophisticated and unsophisticated) subjects in each treatment.

[2]Therefore, this treatment preserves the effect of selection as it inherits the market's tendency to assign optimists to high performance pay, but eliminates the effect of strategic anticipation by depriving subjects of any information allowing them to anticipate the selection process.

3.1 Setup

Let $e \in \{1,2,3,4,5\}$ correspond to the effort level chosen by a subject. If a subject chooses $e = 1$, then this maps to an effort level of 0; if a subject chooses $e = 2$, then this maps to an effort level of 10, and so on and so forth. At the highest effort level, a subject chooses $e = 5$, which maps to an effort level of 40. Let $m \in \{1,2,3,4,5\}$ correspond to the minimum effort level of the other 3 subjects in the group. Payoffs of subject i are determined by a contract consisting of a base wage, an effort cost, and a bonus rate. There are two contracts, henceforth denoted by superscript c. In Contract 1, $c = 1$, base wage is $W_1 = 300$ ECUs, and in Contract 2, $c = 2$, base wage is $W_2 = w$ ECUs, which varies across different experimental sessions. The payoff of subject with contract c, who chooses effort level e and faces other subjects' minimum effort level m, is

$$\begin{aligned}
\bar{\pi}^1(e,m) &= W_1 - 5 \cdot (e-1) \times 10 + B_1 \cdot (\min\{e,m\} - 1) \times 10, \\
\bar{\pi}^2(e,m|w) &= w - 5 \cdot (e-1) \times 10 + B_2 \cdot (\min\{e,m\} - 1) \times 10.
\end{aligned}$$

For the following exposition, we model the experiment for the duration of twenty rounds, split into two blocks of 10 rounds each. Subjects' payoffs differ by the contract they are assigned in the beginning of a block. A subject retains the same contract throughout the entire block. Without loss of generality, let $\pi_{it}(e,m)$ denote the payoff function of subject i in round t. All subjects are assigned to Contract 1 in Block 1. In addition, every subject is randomly assigned to a group of four subjects at the beginning of Block 1. In the beginning of Block 2, subjects' contracts and groups are re-assigned. Half of the subjects are assigned to Contract 1, while the other half is assigned to Contract 2. The re-assignment rules differ across treatments. We consider four treatments: Random, Auction, Sort, and Reverse Sort.

The model consists of two types of subjects: the sophisticated subjects, and the unsophisticated ones. We denote unsophisticated learners with superscript u, and denote sophisticated learners with s. Unsophisticated learners are randomly assigned to two subtypes that differ in their initial beliefs about others' behavior. *Optimists* initially believe that all other subjects will choose 40 (i.e. the highest possible effort level). *Pessimists* initially believe that all other subjects will choose 0 (i.e. the lowest possible effort level).[3] Sophisticated learners are also randomly assigned to two

[3]Using a specification with two subtypes rather than a continuum makes fitting the model com-

subtypes. One subtype has sophisticated learners with 'strategic anticipation' while the other has sophisticated learners without. The two subtypes differ in how they model belief formation during and after the auction in the respective treatment. The model generates dropout times as noisy best responses to beliefs about effort levels in Block 2 under each contract.[4] The distinguishing feature of sophisticated types with strategic anticipation is that they take into account how the auction changes the expected proportion of optimists assigned to each contract. Unsophisticated types and sophisticated types without strategic anticipation do not account for optimists being more likely to be assigned to Contract 2 in Block 2. Sophisticated types with strategic anticipation form beliefs by modeling the decision-making of unsophisticated types, both optimists and pessimists, in the auction. As the value of base wage changes, they update the expected composition of unsophisticated subtypes (optimists vs. pessimists) under each contract. Sophisticated types with strategic anticipation understand that pessimists tend to choose Contract 1 and optimists tend to choose Contract 2. This makes them drop out later in the auction and choose higher (lower) effort levels if they are assigned to Contract 2 (Contract 1).[5]

Next, we provide the detailed model specification of the two types of subjects. For notational convenience, the contract c superscript is suppressed in the following subsections.

3.1.1 Unsophisticated Subjects

An unsophisticated subject updates his/her beliefs based on the current-round experience. Let $N_{-it}(e)$ be the number of subjects other than subject i choosing effort level e. Let $W_{it}(e)$ be the subject's previous-period experience; thus, the updating rule is

$$W_{it+1}(e) = \delta \cdot W_{it}(e) + N_{-it}(e),$$

putationally feasible and eases interpretation of the results. Our qualitative results do not depend on this simplifying assumption.

[4]Given that we are, in essence, using a Vickrey auction, beliefs about others' strategies in the auction do not play an important role.

[5]The learning model only includes sorting based on beliefs, as opposed to sorting based on other personal characteristics. This is appropriate given that personal characteristics have little effect on dropout times beyond their effect on the probability of choosing 40 in the first round. This indicates that personal characteristics affect dropout times through their impact on initial beliefs.

where δ is the depreciation parameter to be estimated from the data. This updating rule is drawn from the standard stochastic fictitious play model. The initial period weights for the optimists are $[0, 0, 0, 0, \omega]$, and the weights for the pessimists are $[\omega, 0, 0, 0, 0]$. In other words, optimists initially believe that all other subjects will choose 40 and pessimists initially believe that all other subjects will choose 0. The initial value of ω is a parameter to be estimated.

The proportion of unsophisticated subject types in the population are parameters estimated from the data. The probability of being subtype l is μ_{pes}, and the probability of being subtype h is μ_{opt}. An unsophisticated subject believes that the probability of any other subject in the group choosing effort level e in the next round $t+1$ is

$$p^u_{it+1}(e) = \frac{W_{it+1}(e)}{\sum_{k=1}^{5} W_{it+1}(k)}.$$

With the belief $p^u_{it+1}(e)$, an unsophisticated subject can infer the probability of minimum effort being m in round $t+1$ is $q^u_{it+1}(m)$, where

$$q^u_{it+1}(m) = \left(\sum_{k=m}^{5} p^u_{it+1}(k)\right)^3 - \left(\sum_{k=m+1}^{5} p^u_{it+1}(k)\right)^3, \quad \forall m \in \{1, 2, 3, 4\} \tag{1}$$
$$q^u_{it+1}(5) = (p^u_{it+1}(5))^3.$$

Notice that in the initial period, optimists' values are $p^h_1(5) = 1$ and $q^h_1(5) = 1$. Conversely, pessimists' values are $p^l_1(1) = 1$ and $q^l_1(1) = 1$. Given belief $q^u_{it}(m)$, the expected payoff of a subtype u player choosing effort e is

$$E\pi^u_{it}(e) = \sum_{m=1}^{5} q^u_{it}(m) \cdot \pi_{it}(e, m) \quad \text{for } u = h, l \text{ and } \forall e.$$

We assume that the probability of a type u subject choosing effort level e follows a logit specification, where

$$\Phi^u_{it}(e) = \frac{\exp(\lambda \cdot E\pi^u_{it}(e))}{\sum_{k=1}^{5} \exp(\lambda \cdot E\pi^u_{it}(k))} \quad \text{for } u = h, l \text{ and } \forall e. \tag{2}$$

Each unsophisticated subject takes a random draw to determine his actual effort choice based on this probability Φ^u_{it} in each round t.

3.1.2 Sophisticated Subjects

A sophisticated subject is denoted with superscript s. Our modeling of the sophisticated subjects and their learning process is similar to that of the sophisticated EWA model of Camerer, Ho, and Chong (2002). Sophisticated subjects assume all other subjects in the group are unsophisticated. A sophisticated subject knows how unsophisticated subjects' beliefs evolve and can therefore anticipate other subjects' behaviors. However, a sophisticated subject cannot identify other subjects in the group, therefore, cannot track individual players' behaviors over time. We assume a sophisticated subject i uses the average of other subjects' $N_{-jt}(e)$ to formulate the updating rules of the unsophisticated types. For example, if a sophisticated type observes effort levels $5, 3, 2, 1$ in the group, and knows his own effort choice, which is 2, then he can infer that the subject choosing 5 would have $N = [0, 0, 1, 1, 1]$, the subject choosing 3 would have $N = [1, 0, 0, 1, 1]$, and the subject choosing 1 would have $N = [1, 0, 1, 1, 0]$. So the average is $\bar{N} = [2/3, 0, 2/3, 1, 2/3]$. The same updating rule is used where

$$W_{it+1}(e) = \delta \cdot W_{it}(e) + \bar{N}_{-it}(e).$$

An unsophisticated person's effort choice probability Φ_{it}^{us} can be readily computed. Notice that Φ_{it}^{us} is the same for all other subjects in the group. This probability is different from the actual choice probability unsophisticated subjects used.

A sophisticated subject does not know the type of the other subjects. He believes the probability of another subject being type l is $\tilde{\theta}_i$ and the probability of another subject being type h is $1 - \tilde{\theta}_i$. The latter may or may not be consistent with the distribution of unsophisticated subject types in the population, where the probability of being subtype l is $\theta = \frac{\mu_{pes}}{\mu_{pes}+\mu_{opt}}$ and the probability of being subtype h is $1 - \theta$.

In Block 1, beliefs about θ are correct, so $\tilde{\theta}_i = \theta \; \forall i \in \{i : \text{Sophisticated}\}$. In Block 2, beliefs may be incorrect if sophisticated subjects fail to anticipate the effect of selection. The two subtypes of sophisticated subjects differ in their ability to anticipate the effects of selection caused by the auction (i.e. strategic anticipation). We denote the proportion of the subtype with the ability of strategic anticipation in the population μ_{sa}. Therefore, the probability of the subtype without strategic anticipation is $1 - \mu_{pes} - \mu_{opt} - \mu_{sa}$.

On one hand, a sophisticated subject without strategic anticipation still assumes $\tilde{\theta}_i = \theta$ after the auction has taken place. In other words, a sophisticated learner without strategic anticipation does not adjust his beliefs post-auction for the effect of selection. On the other hand, a sophisticated subject with strategic anticipation alters his post-auction beliefs $\tilde{\theta}_i$ based on his second-block contract assignment in the Auction treatment. No such adjustment takes place in the other treatments as the subject has no reason to anticipate selection effect, so that $\tilde{\theta}_i = \theta$ in all other treatments.

A sophisticated subject's belief of the distribution of effort choices by another subject $j \neq i$ is

$$p^s_{t+1}(e) = \tilde{\theta}_{it+1}\Phi^{u,h}_{it+1}(e) + (1 - \tilde{\theta}_{it+1})\Phi^{u,l}_{it+1}(e) \quad \forall e.$$

With the belief $p^s_{it+1}(e)$, a sophisticated subject can infer the probability of minimum effort being m in round $t+1$. We denote this probability as $q^s_{it+1}(m)$, which is similarly defined as in Equation (1). Therefore, the expected payoff of sophisticated subject i choosing effort level e in the next round is

$$E\pi^s_{it}(e) = \sum_{m=1}^{5} q^s_{it}(e)\pi_{it}(e, m) \quad \forall e.$$

We assume that the probability of a type s subject choosing effort level e follows a logit specification, where

$$\Phi^s_{it}(e) = \frac{\exp(\lambda E\pi^s_{it}(e))}{\sum_{k=1}^{5}\exp(\lambda E\pi^s_{it}(k))} \quad \forall e. \tag{3}$$

The actual effort choice based on this probability Φ^s_{it} in each round t is randomly drawn for each sophisticated subject.[6]

[6] Both Camerer, Ho, and Chong (2002) and Brandts, Cooper, Fatas, and Qi (2015a) allow for the possibility that some sophisticated types are forward looking. In other words, some sophisticated types can anticipate that their current actions will affect others' future choices and engage in strategic teaching to manipulate future outcomes. We have developed and fit a version of our model that includes forward looking types, but cannot reject the null hypothesis that the fraction of such types is zero. We have therefore dropped this feature from this model. None of our conclusions would be affected if forward looking types were included in the model.

3.1.3 Reset

At the beginning of Block 2 (or round 11), reassignment of groups and contracts occurs. We assume that the experience weights of subjects are reset at this point. The reset experience weight $\tilde{W}_{i,11}$, is a weighted average of realized experience at the end of Block 1, $W_{i,11}$, and a subject's initial beliefs weight $W_{i,0}$; that is,

$$\tilde{W}_{i,11} = \rho \cdot W_{i,0} + (1 - \rho) \cdot W_{i,11},$$

where the reset parameter ρ is fit from the data.

3.2 Treatments

We consider the four treatments in CIQ, which differ in the rules of assigning subjects into contracts and of forming new groups within each contract in the second block.

3.2.1 Random

Each individual subject draws a random shock ε_i from a common distribution. We rank ε_i's from high to low. Then, half of the subjects who have ε_i's above the median shock value are assigned to Contract 1 in Block 2. The other half of the subjects who have ε_i's below the median shock value are assigned to Contract 2 in Block 2. Subjects are then randomly assigned into groups. Sophisticated subjects do not change their beliefs of the composition of the unsophisticated subtypes, so $\tilde{\theta}_i = \theta$, $\forall i \in \{i : \text{Sophisticated}\}$.

3.2.2 Auction

In the Auction Treatment, Contract 2 base wage ranges from 400 to 0. Every 5 seconds, the base wage drops by 5 ECUs. In every instant in time, a subject needs to decide whether to press a 'Contract 1' button to drop out of the Auction process.

An unsophisticated subject, given reset experience weight $\tilde{W}_{i,11}$, can formulate beliefs about the distribution of minimum effort in a future group $q^u_{i,11}$, and the expectation on payoffs in round 11 under Contract 1 and under Contract 2 with any base wage; that is,

$$E\bar{\pi}_{i,11}^{u,c=1}(e) = \sum_{m=1}^{5} q_{i,11}^{u}(m) \cdot \bar{\pi}^{c=1}(e,m),$$
$$E\bar{\pi}_{i,11}^{u,c=2}(e|w) = \sum_{m=1}^{5} q_{i,11}^{u}(m) \cdot \bar{\pi}^{c=2}(e,m|w).$$

Using the expected payoffs, we can calculate the effort choice probability $\Phi_{it}^{u,c=1}(e)$ and $\Phi_{it}^{u,c=2}(e|w)$ for every possible contract. Therefore, the expected payoff of choosing Contract 1 is always $E\bar{\pi}_{i}^{u,c=1} = \sum_e \Phi_{it}^{u,c=1}(e) \cdot E\bar{\pi}_{i,11}^{u,c=1}(e)$, and expected payoff of having Contract 2 with base wage w is $E\bar{\pi}_{i}^{u,c=2}(w) = \sum_e \Phi_{it}^{u,c=2}(e|w) \cdot E\bar{\pi}_{i,11}^{u,c=2}(e|w)$. Assume subjects make noisy dropout decisions based on a logit decision rule, where the logit parameter is ζ. At any given w, the probability of subject i choosing Contract 1 for any given wage w_i is

$$d(w_i) = \frac{\exp(\zeta \cdot E\bar{\pi}_{i}^{u,c=1})}{\exp(\zeta \cdot E\bar{\pi}_{i}^{u,c=1}) + \exp(\zeta \cdot E\bar{\pi}_{i}^{u,c=2}(w))}.$$

The above probability, $d(w)$, represents the likelihood of an unsophisticated subject dropping out at base wage w, *conditional on* the subject not having chosen to drop out at any base wage $w' > w$. We can iteratively define the *unconditional* probability of the subject dropping out at the base wage w, $p^{u}_{dropout}(w)$, as follows:

$$p^{u}_{dropout}(w=400) = d(w=400); \qquad D^{u}(w=400) = p^{u}_{dropout}(w=400);$$
$$p^{u}_{dropout}(w) = (1 - D^{u}(w+5)) \cdot d(w); \quad D^{u}(w) = D^{u}(w+5) + p^{u}_{dropout}(w), \forall w \in \{395, 390, ..., 5, 0\},$$

where $D^{u}(w)$ is the probability of the subject dropping out before or at base wage w.

Compared to an unsophisticated subject, a sophisticated subject has the following information when making contract decisions. First, if $2n$ subjects participate in an auction, a sophisticated subject considers the other $2n-1$ subjects' contract choices. Second, a sophisticated subject does not know the type of other subjects, and assumes everyone else is unsophisticated. Furthermore, the subject knows the population distribution of subtypes, and there is a θ chance that the other subject is an unsophisticated pessimist. Third, a sophisticated subject forms beliefs on unsophisticated subjects' experience weight $\tilde{W}_{11}^{u}(m)$. The subject assumes that all unsophisticated subjects of the same subtype have the same beliefs. So using the same updating rules, the sophisticated subject can predict $D^u(w|\tilde{q}^u_{11})$ for each subtype of subject for any given base wage w, where D^u is the probability that a subtype u subject will choose to drop out of the auction at or before wage w.

A sophisticated type without strategic anticipation believes that the probability of having an unsophisticated pessimist within his own contract in Block 2 is still θ, which is the same as in the population. With the auction taking place, a sophisticated type with strategic anticipation has the ability to anticipate that optimists are more likely to be assigned to Contract 2, and pessimists are more likely to be assigned to Contract 1, so that the relative proportions of pessimists to optimists are different across contracts. Therefore, the probability of meeting a pessimist $\tilde{\theta}$ can vary by individual subject's past experience in Block 1 and depends on which contract a sophisticated subject is ultimately assigned to. We model strategic anticipation and how beliefs of $\tilde{\theta}$ are formed as follows.

Without loss of generality, assume there are n_h type h subjects and n_l type l subjects, where $n_l = (2n-1) - n_h$. From a sophisticated subject's perspective, the probability of facing n_h type h subjects and n_l type l subjects is $\binom{2n-1}{n_l}(\theta)^{n_l}(1-\theta)^{n_h}$. Given any n_l and n, the probability of exactly $\bar{n} \in (0, 2n-1]$ individuals dropping out at w is

$$
\begin{aligned}
&Prob(\bar{n}|n_l, n) \\
&= \sum_x \binom{n_l}{x}(D^l(w))^x(1-D^l(w))^{n_l-x} \cdot \binom{2n-1-n_l}{\bar{n}-x}(D^h(w))^{\bar{n}-x}(1-D^h(w))^{(2n-1-n_l)-(\bar{n}-x)},
\end{aligned}
$$

when $x \in [0, n_l]$ if $n_l < \bar{n}$, and $x \in [\bar{n} - n_h, \bar{n}]$ if $n_l \geq \bar{n}$.

Based on Bayes' rule, given any n_l, the probability that exactly $\bar{n} \in (0, 2n-1]$ individuals, and exactly $x \in [0, \min\{n_l, \bar{n}\}]$ of type l, have dropped out at or before any base wage w is

$$
\begin{aligned}
&P^{dropout}(x, w, \bar{n}|n_l) \\
&= \frac{\binom{n_l}{x}(D^l(w))^x(1-D^l(w))^{n_l-x} \cdot \binom{2n-1-n_l}{\bar{n}-x}(D^h(w))^{\bar{n}-x}(1-D^h(w))^{(2n-1-n_l)-(\bar{n}-x)}}{\sum_{k=0}^{\min\{n_l,\bar{n}\}} \binom{n_l}{k}(D^l(w))^k(1-D^l(w))^{n_l-k} \cdot \binom{2n-1-n_l}{\bar{n}-k}(D^h(w))^{\bar{n}-k}(1-D^h(w))^{(2n-1-n_l)-(\bar{n}-k)}}.
\end{aligned}
$$

Knowing the probability of other subjects' dropout wages w, a sophisticated learner decides upon his own dropout wage w'. In particular, a sophisticated subject understands that for any w to become the base wage of Contract 2 in the second block, two possibilities exist:

Case 1: exactly n unsophisticated subjects dropped out at or before w, and the sophisticated subject chose a dropout wage $w' < w$; or

Case 2: exactly $n-1$ unsophisticated subjects dropped out at or before w, and the sophisticated subject chose a dropout wage $w' \geq w$.

In case 1, the sophisticated individual gets Contract 2. If exactly number x of type l individuals are in Contract 1, the sophisticated subject knows that exactly $n_l - x$ individuals with Contract 2 are of subtype l, and exactly $n-1-(n_l-x)$ are of subtype h. So the sophisticated subject with strategic anticipation expects the probability of facing a subtype l in the Block 2 group to be $\tilde{\theta} = \frac{n_l - x}{n-1}$. A sophisticated subject's expected payoff conditional on getting Contract 2 is $E\bar{\Pi}^{s,c=2}_{i,11}(w, x, n_l)$. Similarly, in case 2, the sophisticated individual gets Contract 1. If exactly number x individuals with Contract 1 are of subtype l, then exactly $n-1-x$ of them are subtype h subjects. So the sophisticated subject with strategic anticipation expects the probability of facing a subtype l subject in the Block 2 group to be $\tilde{\theta} = \frac{x}{n-1}$. A sophisticated subject's expected payoff conditional on getting Contract 1 is $E\bar{\Pi}^{s,c=1}_{i,11}(x, n_l)$. Therefore, taking into account the probability of getting exactly x individuals of subtype l in either contract, the expected payoff of a sophisticated individual with strategic anticipation is:

$$\begin{aligned}
&\text{Case 1:} \quad \binom{2n-1}{n_l}(\theta)^{n_l}(1-\theta)^{2n-1-n_l}P^{dropout}(x, w, n|n_l)E\bar{\Pi}^{s,c=2}_{i,11}(w, x, n_l),\\
&\text{Case 2:} \quad \binom{2n-1}{n_l}(\theta)^{n_l}(1-\theta)^{2n-1-n_l}P^{dropout}(x, w, n-1|n_l)E\bar{\Pi}^{s,c=1}_{i,11}(x, n_l).
\end{aligned}$$

Sophisticated subjects of both subtypes can make a mistake in choosing contracts. The potential for mistakes is represented by an additional randomly drawn logit error term ϵ^c, which is i.i.d. across different contracts. Similar to unsophisticated subjects, we can derive the dropout wage w_i and the probability of dropping out for sophisticated types.[7]

For all subjects, we simulate their choices of dropout wages w_i's. We then rank dropout wage w_i's from high to low. Then, the half of the subjects who have w_i's

[7]There is no closed-form probability of base wage choices for the sophisticated individuals. To

above the median dropout wage are assigned to Contract 1 in Block 2. The other half of the subjects who have w_i's below the median dropout wage are assigned to Contract 2 in Block 2. The group assignment of subjects within a contract is random as in the actual experimental setting.

3.2.3 Sort and Reverse Sort

We use the same contract assignments from the simulation of the Auction treatment in the Sort treatment. Sophisticated subjects do not change their beliefs of the composition of the unsophisticated subtypes, so $\tilde{\theta}_i = \theta$, $\forall i \in \{i : \text{Sophisticated}\}$. However, two additional steps are added to simulate the group assignment of subjects.

First, the model allows for the possibility that the Sort and Reverse Sort treatments imperfectly replicate the selection mechanism from the Auction treatment by adding an additional noise term to the dropout time used to assign a subject to a group and contract in the respective treatments. The sorting of subjects is based upon $w_i + \varepsilon_i$, where $\varepsilon \sim N(0, \varsigma)$ is the sorting error. The standard deviation ς is a parameter estimated from the data. The standard deviation of this noise term determines how well the Sort and Reverse Sort treatments replicate the selection mechanism from the Auction treatment. A standard deviation of zero implies perfect replication. As the standard deviation of the noise term increases, more errors are made in terms of assigning individuals to the wrong contract. Then, the half of the subjects who have $w_i + \varepsilon_i$'s above the median dropout wage are assigned to Contract 1 in Block 2. The other half of the subjects who have $w_i + \varepsilon_i$'s below the median dropout wage are assigned to Contract 2 in Block 2.

In addition, subjects are sorted within a contract when assigned to groups as they were in the actual experimental setting. This means, for example, the four subjects with the four lowest predicted dropout wages $w_i + \varepsilon_i$ are sorted into the same group

simplify computation, we approximate $\tilde{\theta}_i^{s,c}$ by

$$\begin{aligned}
\tilde{\theta}_i^{s,c=1}(w) &= \sum_{n_l=0}^{2n-1} \sum_{x=0}^{\min\{n_l,n\}} \binom{2n-1}{n_l} (\theta)^{n_l} (1-\theta)^{2n-1-n_l} P^{dropout}(x, w, n-1|n_l) \frac{x}{n-1}, \\
\tilde{\theta}_i^{s,c=2}(w) &= \sum_{n_l=0}^{2n-1} \sum_{x=0}^{\min\{n_l,n\}} \binom{2n-1}{n_l} (\theta)^{n_l} (1-\theta)^{2n-1-n_l} P^{dropout}(x, w, n|n_l) \frac{n_l - x}{n-1}.
\end{aligned}$$

Monte Carlo tests indicate that simulated subjects' behaviors do not change too much under such an approximation.

in Contract 2, while the four subjects with the four highest predicted dropout wages are sorted into the same group in Contract 1.

As for the Reverse Sort treatment, we rank the predicted dropout wage $w_i + \varepsilon_i$'s from low to high. Then, the half of the subjects who have $w_i + \varepsilon_i$'s below the median dropout wage are assigned to Contract 1 in Block 2. The other half of the subjects who have $w_i + \varepsilon_i$'s above the median dropout wage are assigned to Contract 2 in Block 2. We also sort subjects into groups based on their dropout wages. Subjects are also sorted within a contract when assigned to groups. Sorting errors in the Sort and the Reverse Sort treatments are assumed to have the same standard deviation ς.

4 Estimation Procedure

Due to the structural complexity and non-linearity of the learning model, there is no closed-form solution to the estimator. We therefore rely upon a moment-based simulation method (Method of Simulated Moments) to estimate the model. This is a mixture model, so we do not estimate the type or subtype of any specific individual, but estimate instead the distribution of types and subtypes within the population. To ensure a reasonable efficiency of our estimator, we generate 100 simulated datasets for each session of each treatment. In each of the 100 simulations, for each simulated subject, we redraw the noise term of effort choice in each round, the noise term of auction (in the Auction treatment only), and the type of the individual (while maintaining the same proportions in the population).

The moments are the average efforts of each round in Block 1 (10 moments), the average efforts of each round in Block 2 by contract and by treatment (10 rounds $\times$ 2 contracts $\times$ 4 treatments). So we have a total of 90 moments. Let the set of moments be denoted by Γ. For any particular set of parameters Θ, we take a set of the aforementioned random draws for each subject and simulate subjects' behavior in an experimental session. For each session, we simulate the model $J = 100$ times using different sets of random draws. Each subject receives a random draw in the beginning of time to determine his/her own type and subtype (whether he/she is a sophisticated subject with/without strategic anticipation, a subtype l subject or a subtype h subject). The type and subtype of a subject are fixed for the entire simulated experimental session. The unconditional probabilities of subject types are estimated parameters. In particular, the fraction of pessimistic unsophisticated types is μ_{pes}, the fraction of optimistic unsophisticated types is μ_{opt}, and the fraction of sophisticated types with strategic anticipation is μ_{sa}. The fraction of sophisticated types without strategic anticipation is $1 - \mu_{pes} - \mu_{opt} - \mu_{sa}$. The set of noise terms and type draws are fixed for all iterations of the estimation routine. To be consistent, we also use the same noise/type draws for all of the estimation and counter-factual exercises in the paper.

For each $j \in \{1, 2, ..., J\}$, we determine the simulated moments $\widehat{\Gamma}_m(\Theta)$. Then, the optimal estimator is defined as $\Theta^* = \arg\min_\Theta \left[\Gamma - \frac{1}{J}\sum_{m=1}^{J} \widehat{\Gamma}_j\right]' W \left[\Gamma - \frac{1}{J}\sum_{j=1}^{J} \widehat{\Gamma}_j\right]$, where W is the optimal weighting matrix. Standard errors are obtained by bootstrapping. The estimation procedure takes into account that subjects' behaviors are potentially correlated across periods at the individual level. In other words, draws of the error term (see Equation 2 and 3) are independently distributed across different simulated subjects, but can be correlated for the same individual across different rounds. The estimated parameter values from fitting the full model to the data are presented in Table 1.

4.1 Alternative Model Specifications

The full model contains four different types of players, yielding a relatively complex model. It is reasonable to ask whether this complexity serves any useful purpose, either in terms of improving our ability to fit the data or, more importantly, our ability to track major features of the data and to predict out of sample. We fit four models, the full model and three models that simplify the full model by reducing the number of types. Ordered from least to most complex, the four models are as follows.

Model 1: Rather than having optimists and pessimists, all unsophisticated learners initially believe that all other subjects will choose effort 40 with probability θ^m and effort 0 with probability $1 - \theta^m$, where the parameter θ^m is estimated from the data. There are no sophisticated learners ($\mu_{pes} + \mu_{opt} = 1$). Model 1 has no *ex ante* heterogeneity and no sophisticated learners.

Model 2: There are two subtypes of unsophisticated learners, optimists and pessimists, modeled as in the full model. There are no sophisticated learners ($\mu_{pes}+\mu_{opt} = 1$). Model 2 allows for *ex ante* heterogeneity in beliefs but contains no sophisticated

Table 1: ESTIMATED LEARNING MODEL PARAMETERS

Parameter	Brief Description	Estimated Values
μ_{pes}	Fraction of pessimistic unsophisticated types	0.185***
		(0.027)
μ_{opt}	Fraction of optimistic unsophisticated types	0.706***
		(0.034)
μ_{sa}	Fraction of sophisticated types with strategic anticipation	0.031***
		(0.003)
ω	Initial experience weight	86.135***
		(29.176)
δ	Experience weight depreciation factor	0.859***
		(0.029)
λ	Effort choice logit parameter	0.458***
		(0.025)
ρ	Reset of belief in round 11	0.998***
		(0.025)
ζ	Auction dropout choice logit parameter	0.173***
		(0.065)
ς	Sort error standard deviation	85.900*
		(44.670)
SSE		1,202.20

Notes: The standard errors are provided in the parentheses. SSE is an acronym for the sum of squared errors. *** and * indicate statistical significance at the 1%, and 10%, respectively. The estimated proportion of sophisticated types without strategic anticipation in the population is $1 - \mu_{pes} - \mu_{opt} - \mu_{sa}$.

learners.

Model 3: This model adds a different type of heterogeneity than in Model 2. Similar to Model 1, all unsophisticated learners initially believe that all other subjects will choose effort level 40 with probability θ^m and effort level 0 with probability $1-\theta^m$, where the parameter θ^m is estimated from the data. Model 3 includes sophisticated learners. Since there is only one type of unsophisticated learner, there can be no selection of different types of unsophisticated learners into contracts. It follows that Model 3 does not include sophisticated learners with strategic anticipation as there is nothing for them to anticipate. We estimate the fraction of unsophisticated learners

μ_{unsoph} while the fraction of sophisticated learners is $1 - \mu_{unsoph}$. Model 3 does not allow for *ex ante* heterogeneity in the beliefs of unsophisticated learners, but does allow for heterogeneity between unsophisticated and sophisticated learners.

Model 4: This is the full model presented in the manuscript. Model 4 allows for *ex ante* heterogeneity in the beliefs of unsophisticated learners and heterogeneity between unsophisticated and sophisticated learners.

Table 2: Estimated Parameters of Alternative Models

Parameter	Model 1	Model 2	Model 3	Model 4
μ_{pes}				0.185***
				(0.027)
μ_{opt}		0.715***		0.706***
		(0.041)		(0.034)
μ_{sa}				0.031***
				(0.003)
μ_{unsoph}			0.654***	
			(0.075)	
ω	43.787***	131.556**	79.6547**	86.135***
	(17.892)	(64.128)	(36.900)	(29.176)
δ	0.990***	0.870***	0.990***	0.859***
	(0.026)	(0.032)	(0.036)	(0.029)
λ	0.020***	0.400***	0.021***	0.458***
	(0.002)	(0.073)	(0.004)	(0.025)
ρ	0.938***	0.999***	0.999***	0.998***
	(0.026)	(0.027)	(0.024)	(0.025)
ζ	2.189	0.169***	0.585***	0.173***
	(2.827)	(0.075)	(0.185)	(0.065)
ς	0.060	59.987	0.163*	85.900*
	(0.084)	(43.701)	(0.090)	(44.670)
θ^m	0.998***		0.999***	
	(0.025)		(0.028)	
SSE	6,041.98	1,934.62	3,600.39	1202.20

Notes: The standard errors are provided in the parentheses. SSE is an acronym for the sum of squared errors. ***, ** and * indicate statistical significance at the 1%, 5% and 10%, respectively.

We fit all four models to the full dataset; the parameter estimates are presented in Table 2.[8] After fitting the model, we generate 100 simulated datasets for each of the fitted models. Table 3 compares the average effort in Block 2, broken down by treatment and contract, for the experimental data and the simulated data.

Table 3: Average Effort Across Alternative Models

	Random		Auction		Sort		Reverse Sort	
	Contr. 1	Contr.2	Contr.1	Contr.2	Contr.1	Contr.2	Contr.1	Contr.2
Data	21.17	31.34	9.64	39.81	13.45	37.48	30.03	16.06
Model 1	17.63	31.42	17.58	31.48	17.64	31.41	17.68	31.40
Model 2	19.02	28.13	8.17	39.90	8.51	39.30	38.18	16.36
Model 3	15.21	29.03	11.17	33.15	11.15	33.18	19.84	25.06
Model 4	16.75	28.62	7.23	39.95	8.53	37.32	31.43	19.67

Notes: The Table displays a comparison of average efforts in the second block by contracts and by treatments across alternative models. The actual averages are provided in the first row.

Fitting a model to the data does not guarantee either the ability to track existing data or to make predictions. The former point can easily be seen by comparing the experimental data and simulated data for Model 1 in Table 3. At no point are the simulated data from Model 1 particularly close to the experimental data. The failure to track the data is especially vivid in the Reverse Sort treatment, where even the flipped relationship between the two contracts is incorrect in the simulated data. It is not surprising that Model 1 does poorly at tracking the data, as it contains no mechanism for capturing either the selection effect or the effect of strategic anticipation. Models 2 and 3 do better than Model 1 at tracking the data, but still have some failures (notably in the Reverse Sort treatment).

Table 4 explores the ability of the various models (the full model and the three simplified models) to replicate one of the most important features of the experimental

[8] The inclusion of both subtypes of sophisticated learners improves the fit of the full model. On one hand, compared to the full model, if the model only allows for sophisticated learners with strategic anticipation, the sum of squared errors increases from 1202.2 to 1483.8. On the other hand, if the model only allows for sophisticated learners without strategic anticipation, the sum of squared errors increases from 1202.2 to 1432.7. The estimated fraction of sophisticated learners with strategic anticipation is small (i.e. only 3.1%). Both strategic anticipation and sorting errors (captured by ς) contribute to differences in average effort between the simulated Auction and Sort treatments. If the model is fitted to the data without sorting error ($\varsigma = 0$), the fraction of sophisticated learners with strategic anticipation increases to 9.8%.

data. CIQ calculate that 73% of the effect of endogenous assignment to high performance pay can be attributed to selection. One of our main uses of the learning model is to explore how this figure would change if the Sort treatment perfectly replicated the selection mechanism from the Auction treatment. This exercise has little value if the model does a poor job of approximating the 73% figure with imperfect replication. The first column of Table 4 reports, for the experimental data and simulated data for each of the four models, the percentage of the effect of endogenous assignment to high performance pay that can be attributed to selection.[9] The final column of Table 4 examines one of the most striking features of our data; that is, the reversal in average effort between the two contracts in the Reverse Sort treatment. Specifically, for the experimental data and the simulated data, the last column reports the difference in average effort in Block 2 between Contract 2 and Contract 1 in the Reverse Sort treatment.

Table 4: Comparison of Selection Across Alternative Models

	% Due to Selection Contract 2	Reverse Sort Contract 2 - Contract 1
Data	73%	−13.97
Model 1	−20%	13.72
Model 2	95%	−21.83
Model 3	101%	5.22
Model 4	77%	−11.76

Notes: The first column provides the effect of endogenous assignment to high performance pay that can be attributed to selection across the alternative models. The last column reports the difference in average effort in the second block between Contract 2 and Contract 1 in the Reverse Sort treatment.

The full model (Model 4) does well along both dimensions, but it is the most complex of the four models. Is that complexity necessary to track major features of the experimental data, or is it sufficient to only incorporate heterogeneous initial beliefs (but no sophisticated learners) as in Model 2, or to only add sophisticated learners (but no heterogeneity of initial beliefs for unsophisticated learners) as in Model 3?

[9] This is the difference between average effort in Block 2 in the Sort and Random Assignment treatments divided by the difference between average effort in Block 2 in the Auction and Random Assignment treatments.

Table 4 indicates that the complexity of the full model is necessary. Because Model 2 has no mechanism for incorporating strategic anticipation, it overestimates the effect of selection (95% in the simulated data vs. 73% in the data) and, as a side effect, overestimates the difference between Contracts 1 and 2 in the Reverse Sort treatment (-21.83 in the simulated data vs. -13.97 in the data). While we don't expect any model to be perfect, these shortcomings are far larger than those for Model 4. Model 3 is even worse, missing the flipped relationship between Contracts 1 and 2 in the Reverse Sort treatment. Model 3 also allows for selection, but similar to Model 2 cannot capture the effects of strategic anticipation. Strategic anticipation relies on sophisticated learners recognizing that the auction will sort different types of unsophisticated learners into different contracts. When there is no heterogeneity of unsophisticated learners, this mechanism breaks down. The added complexity of Model 4 improves the model's ability to fit the data in the standard sense, meaning that the added parameters are statistically significant (see Table 1) but, more importantly, the added features are necessary to capture the most interesting features of our experimental data.

4.2 Out-of-Sample Prediction

As shown in the previous subsection, subjects' heterogeneity adds complexity to the model but improves the fit and is necessary for the fitted model to track important features of the experimental data. Fitting a model to the data does not guarantee that the fitted model will do a good job of tracking existing data. The differences between treatments are not part of the objective function per se, and fitting one treatment well may come at the expense of fitting another poorly or missing qualitative or quantitative differences between treatments.[10]

To validate our model's ability to predict out-of-sample, we estimate the model using data from the Random, Auction and Reverse Sort treatments, and simulate the model using the resulting parameters to predict outcomes in the Sort Treatment. In CIQ, the Sort treatment is designed to separate the effect of selection from the effect

[10] To see this point, consider an incredibly simple-minded model of individual decision-making in our experiments; thus, in all rounds of all treatments under all contracts, subjects pick effort level $i \in \{0, 10, 20, 30, 40\}$ with probability p_i. There is nothing that prevents us fitting this simple-minded model to our data, but it is obvious that, even using fitted values for p_0, p_{10}, etc., it would fail to track the dynamics in the data as well as the differences between the various treatments. It would also fail to predict data from new treatments.

of strategic anticipation. Therefore, we use the model's ability to predict the Sort treatment to validate its ability to make out-of-sample predictions.

A comparison of the estimated parameters for the model fitted from the full dataset and the model fitted from three of the four treatments is presented in Table 5. The parameters change very little. As a result, the differences in goodness of fit, as represented by the sum of squared errors, are negligible. This means, even when the model is fitted to the data of only three treatments, its predictions of the fourth (Sort) treatment are almost as good as the fully fitted version. In Block 2 of the Sort treatment under Contract 2, the predicted average effort of 37.32 based on fitting the model to only three treatments is almost identical to the average effort of 37.30 when fitting all four treatments, and the experimental data of 37.48. As a point of comparison, if we fit Models 1, 2, or 3 to the data from the Random, Auction and Reverse Sort treatments and simulate data for the Sort treatment, simulated average effort in Block 2 for the Sort treatment differs substantially for all three models from the observed data (30.33, 39.30, and 33.48 for Models 1, 2, and 3, respectively).

Based on the same out-of-sample prediction simulations, Table 6 expands on the preceding points. We carry out a similar exercise to what was reported in the first column in Table 4. The full model does well at predicting average effort in the Sort treatment, and hence does a good job of predicting the percentage of the effect of endogenous assignment to high performance pay that can be attributed to selection. The other models once again do poorly.

One point to note here is that the full model might not be the best possible model of learning from the vast universe of possible models either in terms of tracking the major features of our experimental data or predicting out-of-sample within the framework of our experiment. Identifying the best model is well beyond the scope of this paper. Rather, our goal is to show that our model does a reasonably good job of tracking major features of the experimental data as well as predicting out-of-sample while a simplified version of our model would do far worse. This gives us some confidence in the results of the counterfactual exercises based on the full model that come next.

Table 5: Out-of-Sample Predictions of Estimated Learning Parameters

Parameter	Brief Description	Fit 3, Predict Sort	Fit 4
μ_{pes}	Fraction of pessimistic unsophisticated types	0.186***	0.185***
		(0.019)	(0.027)
μ_{opt}	Fraction of optimistic unsophisticated types	0.708***	0.706***
		(0.016)	(0.034)
μ_{sa}	Fraction of sophisticated types with strategic anticipation	0.030***	0.031***
		(0.003)	(0.003)
ω	Initial experience weight	84.253***	86.135***
		(25.669)	(29.176)
δ	Experience weight depreciation factor	0.863***	0.859***
		(0.031)	(0.029)
λ	Effort choice logit parameter	0.457***	0.458***
		(0.040)	(0.025)
ρ	Reset of belief in round 11	0.990***	0.998***
		(0.014)	(0.025)
ζ	Auction dropout choice logit parameter	0.172***	0.173***
		(0.067)	(0.065)
ς	Sort error standard deviation	86.033**	85.900*
		(35.094)	(44.670)
SSE		1,206.27	1,202.20

Notes: In the Fit 3, Predict Sort column, we estimate the model using data from the Random, Auction and Reverse Sort treatments, and simulate the model using the resulting parameters to predict outcomes in the Sort treatment. The Fit 4 column, is reproduced from Table 1 for comparison purposes. The standard errors are provided in the parentheses. SSE is an acronym for the sum of squared errors. ***, ** and * indicate statistical significance at the 1%, 5% and 10%, respectively. The estimated proportion of sophisticated types without strategic anticipation in the population is $1 - \mu_{pes} - \mu_{opt} - \mu_{sa}$.

5 Counterfactual Simulations

In the following Subsection, we highlight the usefulness of structural models to evaluate the robustness of the earlier results. In the subsequent Subsection, we depict the vastness of structural models in investigating the broader implications of the results.

Table 6: Comparison of Selection Across Alternative Models When Fit 3 to Predict Sort

	% Due to Selection Contract 2	Reverse Sort Contract 2 - Contract 1
Data	73%	-13.97
Model 1	96%	12.19
Model 2	95%	-22.51
Model 3	101%	5.17
Model 4	77%	-11.92

Notes: The first column provides the effect of endogenous assignment to high performance pay that can be attributed to selection across the alternative models. The last column reports the difference in average effort in the second block between Contract 2 and Contract 1 in the Reverse Sort treatment. In both columns, we estimate the model using data from the Random, Auction and Reverse Sort treatments, and simulate the model using the resulting parameters to predict outcomes in the Sort treatment.

5.1 Evaluating Robustness

We consider next an important caveat to the main results in CIQ. The study finds that around 73% of the effect of making high performance pay endogenous is due to selection. But we know that this understates the effect of selection because the Sort treatment replicates *imperfectly* the selection mechanism from the Auction treatment. To estimate the true effect of selection, we compare the results from two simulation exercises. The first simulates the full model including the noise term for predicted dropout times. We use the estimated parameters (Fit 4) of the full model (Model 4). The second set of simulations is identical to the first except we do not include an error term for the predicted dropout times in the Sort and Reverse Sort treatments. In other words, we set the standard deviation parameter $\varsigma = 0$. This allows us to estimate what would happen in the counterfactual case where we perfectly replicated the selection mechanism from the Auction treatment.

The results of the two simulation exercises are shown in Figure 1 with the Sort treatment on the left and the Reverse Sort on the right. Based on the experimental data, CIQ calculated that 73% of the effect of making high performance pay endogenous is due to selection. Doing the same calculation using simulated data from the full

Figure 1: Simulations with Tobit Error

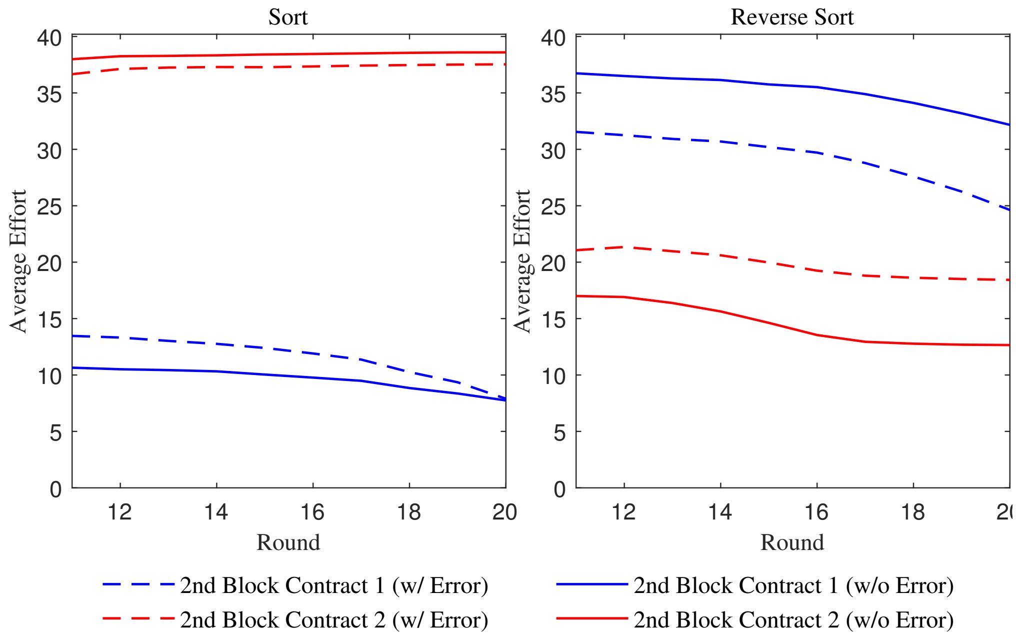

Notes: The Figure displays the two simulation exercises. The first exercise simulates the full model including the noise term for predicted dropout times, where the estimated parameters of the full model are used. The second exercise is identical to the first except we do not include an error term for the predicted dropout times in the Sort and Reverse Sort treatments.

model, which includes errors in assigning subjects to contracts for the Sort treatment, yields a figure of 77%. There is little difference between the selection effect calculated using the experimental and simulated data. For the simulations without errors in the assignment of subjects to contracts, we find that, in the Sort treatment, 93% of the effect of making high performance pay endogenous is due to selection. The difference between the two sets of simulations is of more interest than the levels. If the effect due to selection is underestimated by about the same amount in the experimental data and the simulated data with errors, the true effect of selection is roughly 90% of the effect of making high performance pay endogenous.

5.2 Broader Implications

The first counterfactual exercise considers an additional dimension of heterogeneity: ability. Heterogeneous ability is introduced into the model through differing effort costs. Using simulations, we investigate how multidimensional selection affects the efficiency properties of endogenous selection into incentive contracts.

In the Baseline simulations, all subjects had the same effort cost of 5 ECUs per unit of effort expended as in the original experimental design. As a counterfactual, we allow effort costs to vary within the subject population. In particular, half the population has an effort cost of 3, and the other half has an effort cost of 7 (henceforth, referred to as 'Three-Seven' simulations). We subdivide the Three-Seven simulations into two variants. The first (Three-Seven/Hetero) uses the same distribution of types (pessimists, optimists) as the Baseline simulations. Three-Seven/Hetero differentiates simulated subjects along two dimensions: initial beliefs and ability, thus allowing for multidimensional sorting. The second variant, Three-Seven/Homo, imposes homogeneous beliefs. Specifically, simulated subjects in Three-Seven/Homo initially believe that the other three members of their group will all choose 40 with probability 0.805 and will otherwise all choose 0. This value of 'optimism' was chosen to equalize the initial effort levels for the Three-Seven/Hetero and Three-Seven/Homo simulations.

In both settings, the average effort cost is five. However, half of the simulated subjects in the Three-Seven simulations face an effort cost higher than the Contract 1 bonus rate of 6. Hence, simulated subjects with an effort cost of 7 should only cooperate if they are assigned to Contract 2 with a bonus rate of 10. As a result, the Three-Seven simulations have lower average efforts in either Block 1 or, under Contract 1, in Block 2 of the Random and the Auction treatments than in the Baseline simulations.[11] The three sets of simulations have comparable average effort levels under Contract 2 for Block 2. In the Random Treatment, the Baseline has an average effort of 28.6 under Contract 2, compared to 27.8 in the Three-Seven/Hetero simulations; both of these values are substantially higher than the average of 15.72 in the Three-Seven/Homo simulations. In the Auction Treatment, almost all groups coordinate efficiently under Contract 2 for all three sets of simulations – the average

[11] The average effort in Block 1 is 16.9 in the Baseline as compared to 10.1 in Three-Seven/Hetero and 2.21 in Three-Seven/Homo. In Block 2 under Contract 1, in the Baseline simulations, average efforts are 16.8 and 7.2 in the Random and Auction treatments, respectively. The parallel values are 11.4 and 0.8 in Three-Seven/Hetero, and 3.23 and 0.00 in Three-Seven/Homo.

effort is essentially 40 in all three cases.

Figure 2: Simulations of Selection with Asymmetric Costs and Beliefs

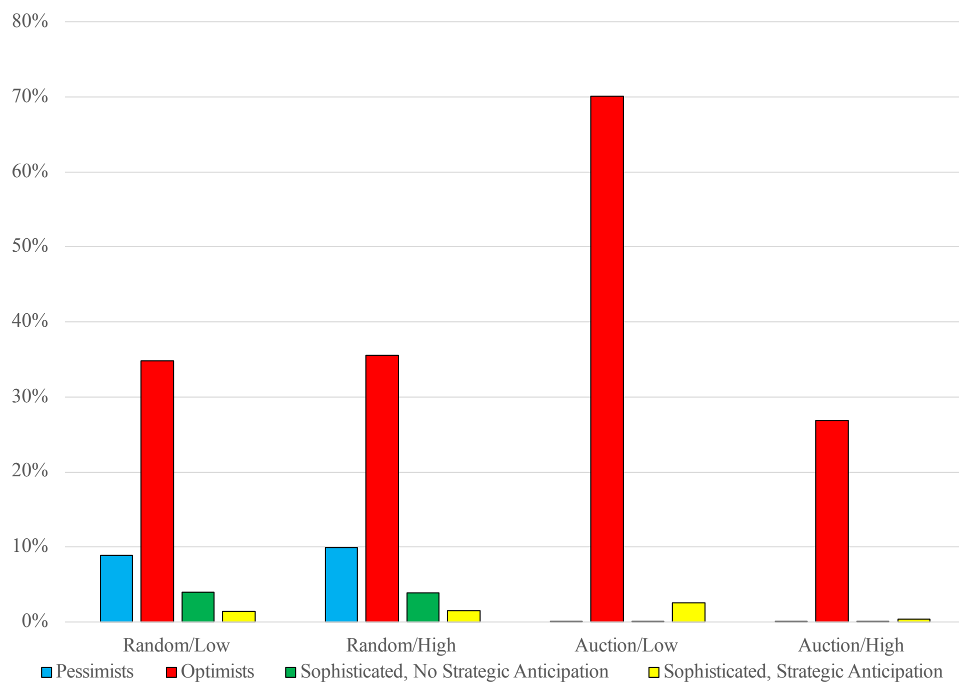

Notes: The Figure shows the distribution of subjects assigned to Contract 2 by their beliefs and effort costs in both the Random and Auction treatments of the Three-Seven/Hetero experiment. Note that a different cluster of bars is shown for each cost level in each treatment.

In the Baseline simulations, efficient coordination is sufficient to maximize surplus. This is not true with asymmetric costs. If Contract 2 yields efficient coordination while Contract 1 leads to coordination at very low effort levels, surplus is maximized by assigning low-cost types to Contract 2. This raises an interesting question: how does selection into Contract 2 work for Three-Seven/Hetero where multidimensional sorting is possible? It is obvious that all low-cost optimists would be endogenously assigned to Contract 2, and all high-cost pessimists would be assigned to Contract 1. What is less obvious is where the high-cost optimists and low-cost pessimists will be assigned. The key insight comes from Figure 2, which shows the distribution of subjects assigned to Contract 2 by their beliefs and effort costs in both the Random and Auction treatments of the Three-Seven/Hetero experiment. Note that a different cluster of bars is shown for each cost level in each treatment. For example, in the

Auction treatment, 70% of the subjects assigned to Contract 2 are low-cost optimists and 27% are high-cost optimists.

Figure 2 shows that virtually everyone assigned to Contract 2 in the Auction treatment will be an optimist. To the extent that there are not enough low-cost optimists to fill all the openings in Contract 2, the remaining slots are taken by high-cost optimists. Virtually all of the low-cost pessimists are assigned to Contract 1. Underlying this, even high-cost optimists find Contract 2 more attractive than Contract 1 because they anticipate a large bonus. Pessimists never think Contract 2 is attractive because they anticipate no bonuses. Low-cost pessimists would earn more under Contract 2 if efficient coordination occurred, but they view this as extremely unlikely. If effort levels are expected to be close to zero in equilibrium, it no longer matters what effort costs are.

From an efficiency point of view, effort costs only affect overall surplus when subjects actually exert efforts. It is surplus maximizing to have both low-cost optimists and low-cost pessimists assigned to Contract 2. This is what happens in Three-Seven/Homo; all simulated subjects have the same beliefs, hence sorting can only take place on the basis of effort costs. Not surprisingly, the Auction Treatment almost entirely assigns the low-cost types to Contract 2. The effort levels are basically identical to those in Three-Seven/Hetero (0 for Contract 1 and 40 for Contract 2), but the average effort cost is 36% higher in the heterogenous beliefs' case. Therefore, the overall efficiency is lower in the heterogeneous beliefs' case. When selection is possible along multiple dimensions, selection may not take place along the dimension that is *most* important for efficiency. Having either optimists or low-cost types assigned to Contract 2 (where efficient coordination occurs) is desirable but, ideally, you would get enough optimists to generate efficient coordination while getting all of the low-cost types in Contract 2. Selection favors optimism over low costs because costs only matter when high effort is expected. Thus, while selection still has an enormously positive effect, it need not resolve all inefficiencies. This is not an insight we could have derived from our original experiment or from other existing experiments with selection along a single dimension. Using our structural model to do counterfactual exercises makes it possible to gain new insight from existing data at relatively low cost.

In the second exercise, we hold all of the model's parameters fixed and enlarge the simulated sessions to allow for 20 groups while varying the number of groups receiving Contract 2 from 0 to 20. For this exercise, we simulate the average effort for each contract under the Auction treatment. Figure 3 graphs the average effort in Round 20 as a function of the percentage of groups assigned to Contract 2. Average effort is broken down by contract with the dotted line showing the weighted average across contracts. For any proportion of Contract 2 groups, average effort is far higher in Contract 2 than in Contract 1. But as the proportion of groups using Contract 2 gets large, the average effort with Contract 2 drops. Once all groups use Contract 2, selection cannot affect performance with Contract 2 since both the optimists and the pessimists are assigned to Contract 2. Thus, strategic anticipation does not play a role. With no selection, there is nothing to anticipate. The average effort, weighted across contracts, rises gradually as the proportion of groups using Contract 2 rises. Moving from a regime with only Contract 1 to a regime with only Contract 2 has a modest positive effect on productivity. This effect must equal the direct incentives' effect as measured by the difference between Contracts 1 and 2 in the Random treatment. With no selection and, by extension, no strategic anticipation, only the direct effect of higher incentives to coordinate at high effort levels remains.

As a final counterfactual exercise, we increase the group size from four to six. In the experimental data of the CIQ study, effort levels were quite high for Contract 2 in the Sort treatment. This leaves little room for improvement in the Auction treatment (limiting the measured effect of strategic anticipation). Using groups of six makes it harder to coordinate at high effort levels, implying that the ceiling at 40 plays a smaller role. Simulating the game with groups of four, 78% of the improvement with endogenous assignment to Contract 2 is due to selection (slightly more than the 73% figure in the real data). The proportion of the improvement with endogenous assignment attributed to selection drops to 66% when we simulate the game with groups of six. Moving to the larger groups decreases the importance of selection and, by extension, increases the importance of strategic anticipation. The mixture of subjects' types and subtypes is held fixed across the two simulation exercises, hence the increased importance of strategic anticipation must be due to a less binding ceiling allowing more room for improvement in the Auction treatment. There are two points to be taken from this counterfactual exercise. First, the measured effect of strategic anticipation depends both on the distribution of types and subtypes within the population and the strategic environment. Second, while the quantitative effect of selection is reduced, the qualitative conclusions change little. Most of the effect of

Figure 3: SIMULATIONS OF AUCTION TREATMENT WITH 20 GROUPS

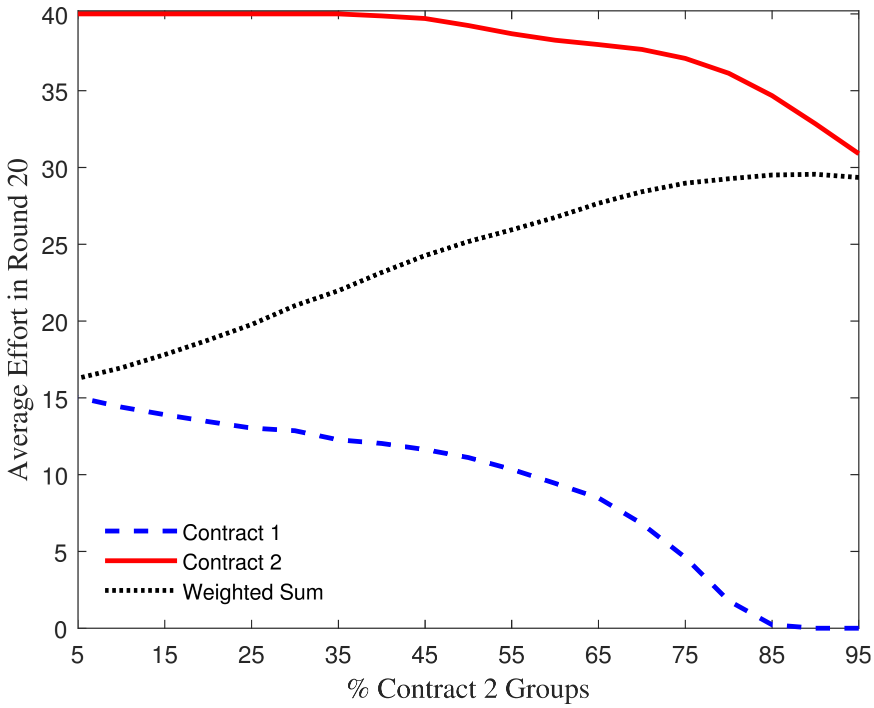

Notes: The Figure graphs the average effort in Round 20 as a function of the percentage of groups assigned to Contract 2. Average effort is broken down by contract with the dotted line showing the weighted average across contracts.

endogenous assignment to incentive contracts is due to selection.

6 Concluding Remarks

Structural models lie at the heart of empirical economic analysis. Their breadth stretches from offering an organizing principle to understand the data to simulating counterfactuals. This study uses a structural learning model that combines stochastic fictitious play with sophisticated learning. Here, the model is used to evaluate the robustness of the results in the study of Cooper, Ioannou, and Qi (2018) and to investigate their broader implications. The model's parameters are estimated from the data. *Ex ante* heterogeneity in the beliefs of unsophisticated learners as well as heterogeneity between unsophisticated and sophisticated learners play a central role in our model's ability to track the data.

Ultimately, structural models can widen and further the scope of the original experiments. First, an estimated structural model identifies unobservable subjects' characteristics, such as pessimistic and optimistic beliefs, strategic anticipation, and experience weights. Second, using these identified characteristics, researchers can conduct counterfactual experiments without conducting additional experiments. Counterfactual experiments help researchers break the limitation of the original experimental designs, which can often be limited by extraneous factors, such as financial support, time constraints and the size and composition of the experimental subject pool. Finally, structural model estimations can make *ex ante* predictions that inform the design of future experiments. For example, the model here finds that the importance of strategic anticipation depends on the strategic environment. Future experimental designs can test this prediction by manipulating contract assignment mechanisms and the levels of subjects' information.

References

Bandiera, Orianna, Luigi Guiso, Andrea Prat, and Rafaella Sadun. "Matching Firms, Managers, and Incentives." *Journal of Labor Economics* 33, 3: (2015) 623–81.

Brandts, Jordi, and David J. Cooper. "A Change Would Do You Good: An Experimental Study on How to Overcome Coordination Failure in Organizations." *American Economic Review* 96: (2006) 669–93.

Brandts, Jordi, David J. Cooper, Enrique Fatas, and Shi Qi. "Stand by Me - Experiments on Help and Commitment in Coordination Games." *Management Science* 62, 10: (2015a) 2916–36.

Brandts, Jordi, David J. Cooper, and Roberto Weber. "Legitimacy, Communication and Leadership in the Turnaround Game." *Management Science* 61, 11: (2015b) 2627–45.

Cadsby, Bram C., Fei Song, and Francis Tapon. "Sorting and Incentive Effects of Pay for Performance: An Experimental Investigation." *The Academy of Management Journal* 50, 2: (2007) 387–405.

Camerer, Colin F., Teck-Hua Ho, and Juin-Kuan Chong. "Sophisticated EWA Learning and Strategic Teaching in Repeated Games." *Journal of Economic Theory* 104: (2002) 137–88.

Cooper, David J., Christos A. Ioannou, and Shi Qi. "Endogenous Incentive Contracts and Efficient Coordination." *Games and Economic Behavior* 112: (2018) 78–97.

Cooper, David J., and Roberto Weber. "Recent Advances in Experimental Coordination Games." In *The Handbook of Experimental Game Theory*, edited by M. Rigdon M. Capra, R. Croson, and T. Rosenblatt, Cheltenham: Edward Elgar Publishing, 2020.

Crawford, Vincent, and Bruno Broseta. "What Price Coordination? The Efficiency-Enhancing Effect of Auctioning the Right to Play." *American Economic Review* 88: (1998) 198–225.

Dohmen, Thomas, and Armin Falk. "Performance Pay and Multidimensional Sorting: Productivity, Preferences, and Gender." *American Economic Review* 101: (2011) 556–90.

Eriksson, Tor, and Marie-Claire Villeval. "Performance-Pay, Sorting and Social Motivation." *Journal of Economic Behavior and Organization* 47, 3: (2008) 530–48.

Fudenberg, Drew, and David K. Levine. *The Theory of Learning in Games*. The MIT Press, 1998.

Hamman, John R., Scott Rick, and Roberto Weber. "Solving Coordination Failure with All-or-None Group Level Incentives." *Experimental Economics* 10, 3: (2007) 285–303.

Lazear, Edward P. "Performance Pay and Productivity." *American Economic Review* 90, 5: (2000) 1346–61.