The London School of Economics and Political Science

Essays in Productivity and Competition

John Benson Shannon

A thesis submitted to the Department of Economics
for the degree of Doctor of Philosophy

30 September 2024

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 14,922 words.

Statement of conjoint work

I confirm that Chapter 2 was jointly co-authored with Martin Pesendorfer and Julien Martin, and I contributed 30% of this work.

Acknowledgments

This thesis would not have been possible without the supervisory support of Martin Persendorfer, whose many meetings and detailed feedback greatly improved my work and shaped my scholarly thinking. I am also grateful to Alessandro Gavazza, Pasquale Schiraldi, Mark Schankerman, and Michael Gmeiner for helpful comments in seminars and engaging discussions. I am also indebted to Maitreesh Ghatak and Francesco Caselli for their selfless help and invaluable support with the thesis submission process. I extend warm thanks to my colleagues past and present who have sharpened my work, especially William Matcham, Michael Coen, Wolfgang Ridinger, Patrick Coen, and Peter Lambert. Finally, I would most like to thank my gracious wife Letitia for supporting me through every step of this journey, for never ceasing to encourage me, and most of all for having Patience with me.

Abstract

This thesis is composed of two chapters.

In the first chapter, I consider a model of production where the firm's labour input is the product of an endogenously organised hierarchy of workers of increasing skill. Production of the labour input is benefited by a kind of labour-augmenting technological shock, which I show can be identified using additional information about the composition of the firm's workforce and wages. I illustrate that the parameters of the model can be consistently estimated using Monte Carlo simulations.

In the second chapter, co-authored with Martin Pesendorfer and Julien Martin, we study the sale of oil and gas leases in New Mexico and assess the degree to which bidder behaviour is consistent with competitive equilibrium play. We use publicly available production and drilling cost data to reconstruct the value to bidders of each lease in our sample. We then test the implications of competitive bidding and reject the hypothesis that bidders cannot improve their expected payoffs by altering their bidding strategy. We find that an annual discount rate of over 17% is necessary to make submitted bids consistent with competitive bidding behaviour.

Contents

1 Quantifying Efficiency Gains from Reorganization **8**

1.1 Introduction 8

1.2 Model 10

1.2.1 The Labor Input 10

1.2.2 Organizational Efficiency 13

1.2.3 Efficiency vs. Productivity 15

1.3 Data 16

1.3.1 The Importance of Labor Composition 17

1.3.2 Is Organization Meaningful? 21

1.3.3 Variation Across Organization 23

1.3.4 Expansion Across Organizations 26

1.3.5 Expansion with Reorganization 26

1.4 Estimation 30

1.4.1 Controlling for Productivity 30

1.4.2 Monte Carlo Simulation 32

1.5 Conclusion 35

1.6 Appendix 39

1.6.1 ENIA Variables and Descriptions 39

1.6.2 Further Descriptive Results 40

2 Underbidding for Oil and Gas Tracts **42**

2.1 Introduction 42

2.2 Framework 44

2.3 Testable Implications 46

2.3.1 Independence 47

2.3.2 Best Response Test 48

2.3.3 Uniform Upward Deviation 49

2.4 Data 51

2.5 Conditional Independence Test 57

2.6 Underbidding 60

2.7 Uniform Upward Deviation 62

2.8 Conclusion 66

List of Figures

1.1 Relation between the number of managers and firm outcomes 19

1.2 Relation between the labor share of managers and firm outcomes 20

1.3 Relationship between managers and manager share 21

1.4 Joint density of value-added and labor, conditional on managers. 22

1.5 Distributions by Layer 25

1.6 Differences in Conditional Means 41

2.1 Histogram of Biprobit Correlation Coefficient Estimates 60

2.2 Simulated Profit Varying the Bid Scalar α. 62

2.3 Interest Rate and Cost Multiples Needed to Satisfy Uniform Upward Deviation Incentives: Top Five Bidders. 66

List of Tables

1.1 Summary Statistics 17

1.2 Firms Differ by Organization 24

1.3 Firms Expand Pyramidally 26

1.4 Firms that Reorganize Expand Differently 27

1.5 Regression of residual productivity on knowledge measures 29

1.6 Monte Carlo Results 33

1.7 Monte Carlo Results 34

2.1 Summary Statistics for Awarded Tracts 54

2.2 Bidding Returns for Top Bidders 56

2.3 Biprobit Correlation Coefficients. 59

2.4 Best Response Test: Optimal Bid Scalar $\hat{\alpha}^*$ 61

2.5 Testing Uniform Upward Deviation Incentives: $H_0 : D_T^S \leq 0$. 64

Chapter 1

Quantifying Efficiency Gains from Reorganization

1.1 Introduction

The question of how technical change impacts the productivity of different factors of production, as well as which factors matter for the production process, is of central economic importance. The simplest approach to this question is to aggregate observed and unobserved factors as much as possible, but this simplicity comes at the cost of weakened economic inference. Firms make many decisions in the production process, and overlooking the marginal conditions that govern these conditions will often lead to biased inference. Under most commonly assumed conditions, firm input decisions are endogenous, a problem which can only be overcome when it is possible to control for all variables affecting these decisions. Mismeasuring the inputs or failing to control for unobserved technical change will produce inconsistent estimates.

The latter problem has been addressed by an insight from the industrial organization literature first made by Olley and Pakes (1996) and developed in Levinsohn and Petrin (2003) and Ackerberg et al. (2015) – if firms make decisions based on what they know about productivity, then it is possible to invert the process and infer the firm's knowledge of productivity from its decisions. This approach, however, will only succeed if the channels by which unobserved productivity affects firm decisions have been properly specified. In the standard case of a single unobserved Hicks-neutral component to productivity, the inversion mapping is correctly specified if factor demands are strictly increasing in productivity. If, on the other hand, technical change has multiple dimensions, each affecting demands for a different subset of inputs, the standard approach will fail. This paper extends the approach of Olley and Pakes (1996) to a setting with multidimensional technical change.

A similar extension has been put forward by Doraszelski and Jaumandreu (2018), who allow for both a Hicks-neutral and a labor-augmenting dimension of productivity. They build on a similar insight as in Olley and Pakes (1996), using the firm's choice of an optimal input mix to identify labor-augmenting technical change. Based on ratio of the firm's marginal conditions for labor and material inputs, the ratio of these two inputs should not change when the ratio of their marginal costs does not, unless there is an unobserved term affecting their relative marginal products. Applying their method to a panel of Spanish manufacturing firms, they find that labor-augmenting productivity accounts for a significant share of productivity growth as well as the falling revenue share of labor. One challenge presented by the Spanish data, however, is the sharp distinction in the Spanish labor market between regular and temporary workers. Changes in the ratio of these two types of workers could affect the firm's choice of input mix between labor and material inputs. Consequently, the composition of the firm's labor input, as well as the way heterogeneous worker types are combined into an aggregate labor input, must be specified correctly.

This paper examines the importance of the composition of the labor input for firm choices using Chilean manufacturing data. Aside from covering a wide variety of industries over more than a decade, the Chilean data also disaggregate workers along multiple dimensions, allowing for a richer analysis of labor composition. I show that not only does the share of workers of different types affect a wide variety of firm outcomes, it also does so independently of the total number of workers in the firm. If multiple dimensions of worker heterogeneity are relevant to the firm's decisions, how can the firm's choice of an input mix be feasibly modelled?

A model for the firm's endogeneous organization of its labor force has been put forward by Garicano (2000) and subsequently developed in a macroeconomic context (Garicano and Rossi-Hansberg, 2004; Garicano and Rossi-Hansberg, 2006; Caliendo and Rossi-Hansberg, 2012). Workers vary along a single dimension of knowledge, defined as the probability of successfully producing output within a given span of time. Given the way that opportunities to use knowledge are shared within the firm, the firms choose an optimal mix of workers and knowledge levels, affecting not just the mix of different worker types but also their wages. Variation in the average wage, then, reveals something about the firm's chosen mix of workers and their skill levels. The implications of this model have been studied in the settings of particular industries as well as entire countries. Garicano and Hubbard (2016) find evidence of returns to hierarchical organization within law firms, and Garicano and Hubbard (2018) study how the technology governing that organization has changed over time. Caliendo et al. (2015) use French administrative data that links workers with hierarchically-coded occupations to firms and find evidence that the number of layers within a firm is statistically meaningful. Caliendo et al. (2020) repeat the exercise with similar data for the Portuguese manufacturing sector and incorporate a demand model to estimate the effects of reorganization due to a demand shock on a firm's revenue- and quantity-based productivity. This paper finds similar patterns in the Chilean plant data and uses the hierarchical model of firm organization to identify the labor-augmenting component of technical change.

The paper proceeds as follows. Section 1.2 introduces the organizational model and provides an illustration of its application to a broader model of the firm. Section 1.3 introduces the data and explores the importance of the composition of the firm's labor input. Section 1.4 describes the estimation strategy and illustrates its performance in a Monte Carlo simulation. Section 1.5 concludes.

1.2 Model

The theory of optimal firm organization was first derived by Garicano (2000) and subsequently developed by Caliendo and Rossi-Hansberg (2012). I follow Caliendo et al. (2020) in proposing a model where the firm's choice of optimal organization is nested within a standard profit maximization problem.

1.2.1 The Labor Input

Consider a firm with a Cobb-Douglas (value-added) production technology:

$$Y = AK^{\alpha}O^{\beta}$$

with productivity A, capital stock K, and labor input O. The labor input is denoted here by O rather than the traditional L to signal that it is not simply the total labor hours within the firm. Instead, the labor input is the outcome of the firm's organizational decision, in which workers with different skills and positions within the firm collaborate to use all the knowledge within the firm efficiently.

To create the labor input, workers in the production layer generate production problems at a constant rate, and each problem creates one unit of output if it can be solved by the firm. The problems are drawn from a distribution F with strictly decreasing density. A particular problem z can be solved if it falls within the production workers' knowledge set, denoted by the interval $[0, z_1]$. If the problem cannot be solved in the production layer (i.e. $z > z_1$), the problem is passed up to higher layers consisting of managers who listen to problems at the rate h. Managers in layer k learn to solve a range z_k of problems, so that workers in, e.g., the second layer learn to solve problems in the interval $[z_1, z_1 + z_2]$. Since managers in layer k spend all of their time listening to problems from lower layers, the employment at a layer k will be the listening rate h times the volume of unsolved production problems $n_1[1 - F(Z_{k-1})]$, where $Z_{k-1} = \sum_{l=1}^{k-1} z_l$. The share of production problems that can be solved by the firm is then $F(Z_L)$, where L is the number of layers, making the labor input $O = n_1 F(Z_L)$.

The cost of producing the labor input O is the sum of the wages of workers in each layer. Knowledge is costly to acquire, and wages compensate workers for the value of their knowledge as well as their time. Workers are paid a base wage of w per unit time and are compensated proportionally for their knowledge at rate c.[1] The wage for workers in layer k is then $w_k = w[1 + cz_k]$, leading to a total wage bill $W(n, z) = \sum_{k=1}^{L} n_k w[1 + cz_k]$.

The optimal organization minimizes the cost of producing a target labor input to the overall production technology, solving:

$$\min_{\{z_k, n_k\}_{k=1}^{L}} \sum_{k=1}^{L} n_k w[1 + cz_k]$$

subject to

$$n_1 F(Z_L) \geq O$$

$$n_k = hn_1 \left[1 - F(Z_{k-1})\right], \quad k = 2, \ldots, L$$

[1]This parametrization has been interpreted by Caliendo and Rossi-Hansberg (2012) as payments to teachers who earn a wage of w and can impart knowledge at the rate c per unit time. Alternatively, it can be interpreted as on-the-job learning, where workers are employed by the firm but spend time outside of their roles (producing or managing) acquiring knowledge. The only important assumption is that the firm is a price taker in the labor market and is thus rendered indifferent between training its own employees or hiring workers with pre-existing skillsets.

Given values of the organizational parameters (w, c, h) and the distribution of production problems F, this problem can be used to infer the wage bill given a value of the labor input. Since in practice we do not observe the labor input but do observe the firm's optimally chosen wage bill, it is more useful to consider the dual maximization problem:

$$\max_{\{n_k, z_k\}_{k=1}^{L}} \quad n_1 F(Z_L)$$

subject to

$$\sum_{k=1}^{L} n_k w \left[1 + c z_k\right] \leq W$$

$$n_k = h n_1 \left[1 - F(Z_{k-1})\right], \quad k = 2, \ldots, L$$

The solution to this problem provides a mapping from the total wage bill to the efficient labor input.[2] Imposing the additional restriction $N = \sum_{k=1}^{L} n_k$ (that is, the sum total of labor hous at each layer must be equal to an aggregate amount N), the firm's problem can be rewritten in per-capita terms by substituting out labor hours in each layer[3] and dividing by N:

$$\max_{\{z_k\}_{k=1}^{L}} \frac{F\left(Z_L\right)}{1 + h \sum_{k=2}^{L} \left[1 - F\left(Z_{k-1}\right)\right]} \tag{1.1}$$

subject to

$$\frac{w\left[1 + c z_1\right] + \sum_{k=2}^{L} \left(h\left[1 - F\left(Z_{k-1}\right)\right]\right) w \left[1 + c z_k\right]}{1 + h \sum_{k=2}^{L} \left[1 - F\left(Z_{k-1}\right)\right]} \leq \frac{W}{N} \tag{1.2}$$

Let $\Psi_L(W/N)$ denote the value function of this problem, i.e. the efficient per-capita labor input given the average wage W/N. It is then possible to infer the firm's optimal labor input (the result of the organizational problem) given the optimally-chosen wage bill and labor hours:

[2]The equivalence of cost minimization and output maximization is a textbook result, so I employ a simple proof by contradiction. Suppose that the chosen labor input $O = n_1 F(Z_L)$ maximizes profit but does not maximize the labor input given the expenditure W. Then an $O' > O$ exists that satisfies the expenditure constraint, which produces output $Y' > Y$ at the same cost. Since this strictly increases profits, the labor input O cannot be a profit-maximizing choice by the firm, a contradiction. QED.

[3]The total number of labor hours pins down the number of hours in each layer by the listening constraint:

$$\begin{aligned} N &= \sum_{k=1}^{L} n_k \\ &= n_1 \left(1 + \sum_{k=2}^{L} \frac{n_k}{n_1}\right) \\ &= n_1 \left(1 + h \sum_{k=2}^{L} \left[1 - F\left(Z_{k-1}\right)\right]\right) \end{aligned}$$

where the last equality uses the constraint $n_k = h n_1 \left[1 - F(Z_{k-1})\right]$ for $k = 2, ..., L$.

$$O = N \cdot \Psi_L \left(\frac{W}{N} \right) \tag{1.3}$$

Substituting this expression for O back into the firm's production technology yields this formulation of the firm's profit maximization problem:

$$\max_{K,N,W} AK^{\alpha} \left[N \cdot \Psi_L \left(\frac{W}{N} \right) \right]^{\beta} - rK - W \tag{1.4}$$

The organizational problem enters the profit maximization decision by endogenizing the wages paid to workers. This endogeneity does not reflect monopsony power – the firm is still a price taker in the labor market, but it now has the choice not only of *quantity* of labor but *quality* of labor. The effect of labor quality on output is captured by the organizational decision, entering the production process through the value function Ψ_L.

1.2.2 Organizational Efficiency

Thus far, the firm's organizational problem is completely deterministic: the wage will always vary proportionally with labor hours in the same way for all firms. To allow for variation in organizational efficiency, Caliendo et al. (2020) simply replace labor hours in the production function with the cost of choosing the optimal labor input: $C(O^*) = W$

$$Y = \widetilde{A} K^{\alpha} [C(O^*)]^{\beta}$$

At the optimal labor input, the cost function $C(O^*)$ divided by the input quantity O^* is equal to the average cost function $AC(O^*)$, implying that $O^* = C(O^*)/AC(O^*)$. Substituting this into the production function yields

$$\widetilde{A} = A[AC(O^*)]^{-\beta} = A \left[\frac{W/N}{\Psi_L \left(W/N \right)} \right]^{-\beta}$$

Clearly, the average cost only varies if the average wage varies. However, the production function in Equation 1.4 produces an average wage that does not depend on productivity A.[4] This implies that, given the organizational structure represented by Ψ_L, a one-to-one mapping between wages and labor hours exists. To infer the optimal total labor hours given a wage payment, we replace the prior constraint $\sum_{k=1}^{L} n_k = N$ with the constraint

[4]It is easy to show that ratio of the first order conditions for N and W is a function of the average wage W/N only. Since the average wage must always solve an equation that does not depend on A, variation in A will not cause variation in the average wage.

$n_L = 1$, following Caliendo et al. (2020). Denote the optimal aggregate labor hours under this constraint $N_L^*(W)$.[5] Then we can define the difference between the average cost implied by the model and the average cost chosen by the firm:

$$\frac{W/N}{\Psi_L(W/N)} = \frac{W/N_L^*(W)}{\Psi_L\left(W/N_L^*(W)\right)} \cdot \exp\left(-\xi^L\right)$$

The term on the lefthand side is the average cost implied by the organizational model when taking the firm's choice of labor hours N as given. The first term on the righthand side is the average cost implied by the theoretical restriction that there is only one worker in the top layer. The term ξ^L is the (log) difference between these two average costs. Positive values of ξ^L imply that the firm is able to acheive lower average costs than under the constrained model. Using lowercase letters to denote logs, and defining the functions $\psi_L(w) \equiv \ln \Psi_L(e^w)$ and $n_L^*(w) \equiv \ln N_L^*(e^w)$, we can solve for ξ^L:

$$\xi^L \equiv [n + \psi_L(w-n)] - [n_L^*(w) + \psi_L(w - n_L^*(w))] \tag{1.5}$$

This term is the analagous to the factor productivity term A in the firm's output production function. Higher values of A reflect a firm's ability to produce higher levels of output given the same level of inputs. Similarly, higher levels of ξ^L allow the firm to produce higher levels of organizational output (the labor input O) given the same levels of organizational inputs (wages W and labor hours N). Consequently, I will refer to ξ^L as the *organizational efficiency* (or simply *efficiency*) of the firm to distinguish it from the firm's *factor productivity* (or simply *productivity*).

Using this definition of efficiency, the labor input can be written as $O = N_L^*(W) \cdot \Psi_L\left(W/N_L^*(W)\right) \cdot \exp\left(\xi^L\right)$. This formulation makes it clear that organizational efficiency is ultimately about transforming wages into productive labor. Here we see another important parallel between organizational efficiency and factor productivity. When firms are not price takers in the output market, it is impossible to distinguish between unobserved factors that improve a firm's productivity (reducing its production costs) and those that

[5] As before, we can plug in the listening-time constraints to pin down N given n_L:

$$\begin{aligned} N &= \sum_{k=1}^{L} n_k \\ &= n_L \left(\sum_{k=1}^{L-1} \frac{n_k}{n_L} + 1 \right) \\ &= n_L \left(\frac{1}{h\left[1 - F\left(Z_{L-1}\right)\right]} + \sum_{k=2}^{L-1} \frac{1 - F\left(Z_{k-1}\right)}{1 - F\left(Z_{L-1}\right)} + 1 \right) \end{aligned}$$

The function $N_L^*(W)$ is equal to this expression when $n_L = 1$ and the knowledge levels z maximize $n_1 F(Z_L)$ subject to $n_k = n_1 h[1 - F(Z_{k-1})]$ for $k = 2, \ldots, L-1$, $n_L = 1$, and $\sum_{k=1}^{L} n_k w[1 - cz_k] \leq W$.

increase a firm's market power (raising its revenue) without imposing structure on output demand. Similarly, if firms are not price takers in the labor market, it is impossible to separate unobserved factors contributing to organizational efficiency (allowing a firm with the same knowledge stock to solve more production problems) and factors that contribute to monopsony power (allowing the firm to secure more knowledgeable workers for the same wages) without making additional assumptions about the structure of labor supply. The application of this definition of organizational efficiency to the study of monopsony power is an exciting avenue for future research.

1.2.3 Efficiency vs. Productivity

To further explore the distinction between efficiency and productivity, I follow the convention of Olley and Pakes (1996) and decompose (log) productivity into two terms: a component ω that is observed by the firm when it chooses its inputs and a component ε that is realized after input choices are made but before output is produced. Combining this decomposition with the efficiency residual ξ^L defined above yields the following production function (with all variables represented in logs):

$$y_{it} = \alpha_0 + \alpha k_{it} + \beta o_{it} + \omega_{it} + \varepsilon_{it} \tag{1.6}$$

where k_{it} is the firm's capital input and o_{it} the firm's labor input. The labor input (and consequently labor hours n_{it}) is determined by the firm's aggregate wage payment w_{it} and the efficiency shifter ξ^L_{it}:

$$o_{it} = n_{it} + \psi_L(w_{it} - n_{it}) = n^*_L(w_{it}) + \psi_L(w_{it} - n^*_L(w_{it})) + \xi^L_{it}$$

The efficiency shifter ξ^L_{it} enters the production function only through the choice of the labor input o_{it}. In this way, ξ^L_{it} acts as a kind of labor-augmenting technical change, increasing the amount of labor that enters the production function. The difference in interpretation lies in that it does not augment the labor stock directly but instead augments the production of the intermediate labor input from raw hours and wages. Increases in labor efficiency change the amount of labor time that contributes to production, rather than increasing the contribution for a constant amount of time. Evolution in efficiency would have the same effect on output as classic labor-augmenting change if the (perhaps nonlinear) cost of hiring workers were kept fixed, but it does not require this restriction. The efficiency term ξ^L_{it} more flexibly captures changes in the firm's labor costs, reflecting changes in the optimal mix of labor sub-inputs. If the optimal mix of different worker types varies endogenously, so too does the marginal cost of the labor input, which will in turn affect the optimality conditions governing the firm's demand for other inputs, rendering estimation techniques that do not control for the mixture of worker types inconsistent.

This raises the question of whether the mixture of worker types has an effect on firm output and demands for other inputs. To answer this question, I examine the dataset described in the following section.

1.3 Data

To estimate the model, I use data from the Encuesta Nacional Industrial Anual (ENIA) survey carried out by Chile's Instituto Nacional de Estadisticas (INE). The ENIA annually collects plant-level data from over 13,000 industrial manufacturing firms from 1995 to 2019, with between 45 and 60 industries represented in the sample each year.

In addition to recording plants' annual value-added and various production inputs, the data contain the number of workers (averaged throughout the year) and their wages disaggregated into various categories by contractual status and whether or not they are associated with the industrial process. Managers are a separate category, but aggregated together with contracted workers. All worker categories are disaggregated by sex, but wage payments to workers in each category are not.

Some variables are not available in all years, so I restrict the sample to the years 2000-2014, when all critical variables are observed (wages paid to managers are not recorded until 2000, and capital stock is not recorded after 2014). This results in a dataset covering 13,304 firms across 62 industries with 117,912 plant-level observations.

Table 1.1a shows summary statistics for all firms in the sample period, and Table 1.1b shows summary statistics for the number of workers and wages across categories (monetary values are in billions of real Chilean pesos).

Table 1.1a provides a sense of the scale and distribution of standard firm outcomes. The median firm produces 243 milllion pesos in value-added whil employing 27 workers who are paid wages of 156 million pesos. Capital stock is slightly higher than wages for the median firm, but exhibits much wider variation across the distribution, ranging from only 5 million at the 10th percentile to 4.8 billion at the 90th percentile. All distributions are

Table 1.1: Summary Statistics

Monetary variables are measured in billions of 2015 Chilean pesos.

(a) Aggregated variables

Variable	**Mean**	**St. Dev.**	**10th pctile**	**Median**	**90th pctile**
Value-added	3.810	36.374	0.028	0.243	3.965
Labor	84.755	239.969	8.647	27.112	189.728
Wages	1.082	16.266	0.030	0.156	1.632
Capital stock	6.318	107.039	0.005	0.182	4.801
Investment	1.314	29.238	0.000	0.007	0.603
Electricity	0.210	2.323	0.001	0.007	0.167
Fuel	0.189	2.923	0.000	0.007	0.123

(b) Disaggregated labor variables by category

	Workers		**Wage**	
Category	**Mean**	**St. Dev.**	**Mean**	**St. Dev.**
Managers	2.571	6.654	0.114	2.124
Contracted, industrial	58.175	181.252	0.673	11.080
Contracted, nonindustrial	13.888	55.017	0.205	4.619
Subcontracted, industrial	6.407	53.358	0.066	2.499
Subcontracted, nonindustrial	2.243	24.528	0.021	0.899
Homeworkers	0.233	3.058	0.003	0.305
All workers	84.755	239.969	1.082	16.266

heavily right-skewed, with the mean being several times larger than the median.

Table 1.1b shows the average use of different labor categories across firms. Contracted industrial workers are the largest subcategory, followed by contracted nonindustrial workers, with the latter earning slightly higher wages on average (after dividing the reported mean wage by the reported mean number of workers). Subcontracted workers form a smaller but non-negligible component of the work force, but homeworkers are quite rare.

1.3.1 The Importance of Labor Composition

How meaningful are these subcategories of labor to firm outcomes? If labor is homogenous, the share of any particular subcategory of labor should be essentially uncorrelated with other endogenous choices made by the firm, such as input demands or output. At the same time, increasing the amount of any subcategory of the labor input should on average increase both the firm's output directly and also its demands for other inputs due to complementarities between the labor input and the marginal productivity of other inputs. On the other hand, if the share of one labor subcategory within the aggregate labor stock is correlated with other endogenous outcomes, it implies that firms choose these shares optimally, which in turn implies that the share affects the firm's profit and labor should not be treated as homogenous.

To assess whether the composition of labor affects firm outcomes, I plot the variation of a variety of endogenous outcomes against both the level and the share of subcategories of labor. To control for unobservable factors that are constant across firms, I also detrend all variables using industry-specific time trends, so that a quantity x_{ijt} pertaining to firm i in industry j at period t is scaled by $\overline{x}_{\cdot jt}$, defined as the average across firms at period t within industry j of the quantity x:

$$\ddot{x}_{ijt} \equiv \frac{x_{ijt}}{\overline{x}_{\cdot jt}}$$

The resulting variable $\ddot{x}_{ijt}$ represents the quantity of interest x_{ijt} measured as a percentage of the average within industry j in period t. Firms with below-average quantities will have detrended quantities between zero and one, which will become negative when plotted on a log scale.

Here I focus on one particular subcategory: managers. Figure 1.1 depicts the variation between the number of managers and a number of firm choices, such as the firm's demand for capital and intermediate inputs, the wages paid to workes in the firm, and the firm's output. All relationships are positive and hold both in the raw and detrended data.[6] Firms with more managers tend to hire more workers, maintain higher capital stocks, produce more output, pay higher wages, etc. Results like this would follow from any model that treats managers as an input to the firm's production technology.

Figure 1.2 depicts the variation between the same outcomes with the share of managers within the firm (the number of managers divided by the total number of workers). This figure, in contrast to Figure 1.1, shows that the share of managers within the firm's workforce is negatively correlated with all of the same firm outcomes. As the share of managers in the workforce increases, the size of the workforce decreases, together with output, wages, capital, etc. Managers are responsible for more production workers, but both they and the workers they manage tend to be paid less and produce less output with lower levels of complementary inputs. Firms with a high share of managers also maintain less capital per worker on average, so the decrease in capital is not merely proportional to the decrease in labor.

Furthermore, these changes cannot be accounted for merely by a negative correlation

[6] Since logs are taken after averaging, the detrended variables will in general not have mean zero by Jensen's Inequality.

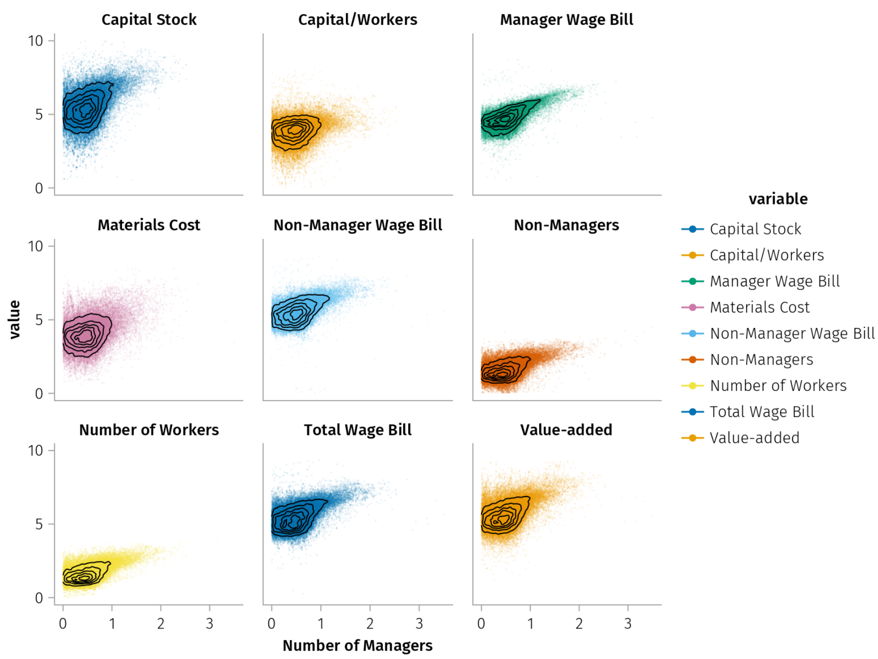

(a) Raw data

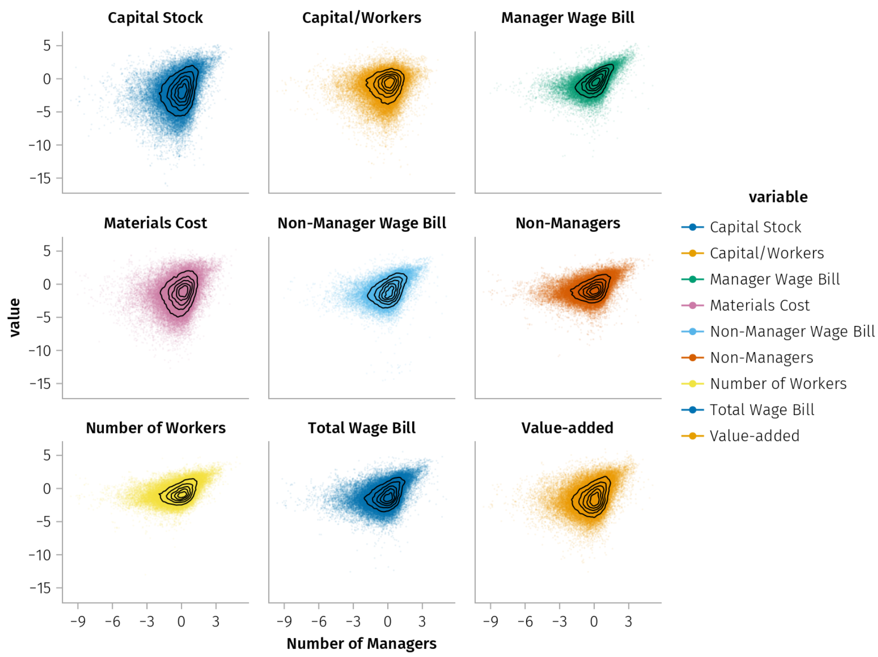

(b) Detrended by industry-year

Figure 1.1: Relation between the number of managers and firm outcomes

All variables plotted on $\log_{10}$ scale.

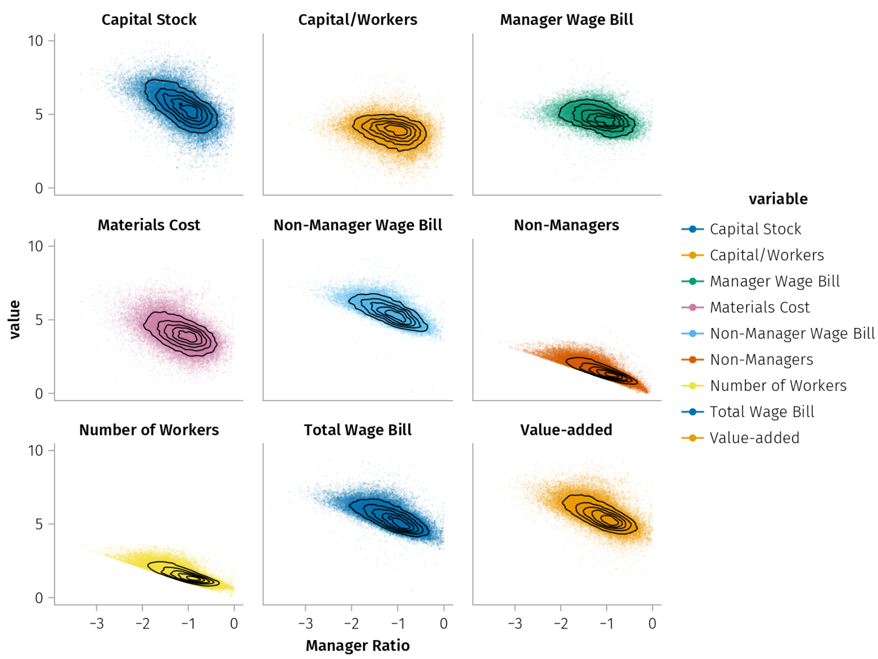

(a) Raw data

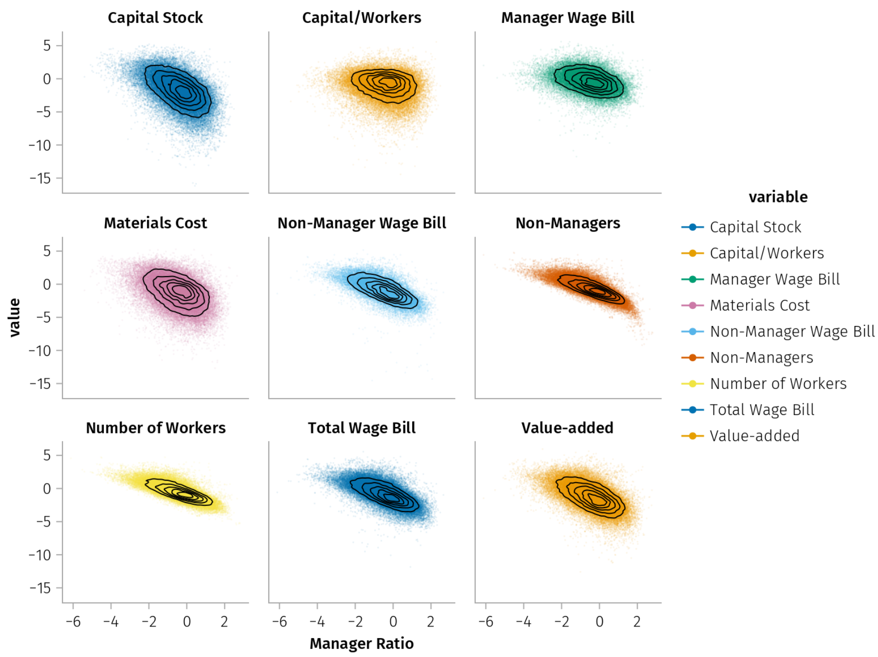

(b) Detrended by industry-year

Figure 1.2: Relation between the labor share of managers and firm outcomes

All variables plotted on $\log_{10}$ scale.

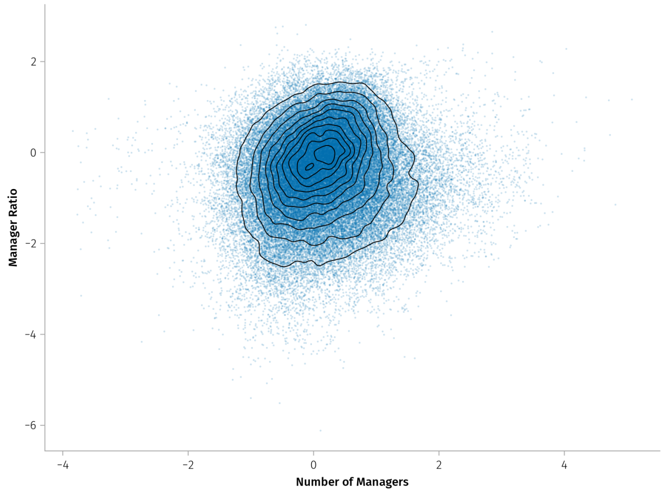

Figure 1.3: Relationship between managers and manager share

between the share of managers and the total number of managers, as Figure 1.3 that these are (weakly) positively correlated. The positive effects of increasing the *level* of managers must be offset by the negative effects of increasing the *share* of managers.

Finally, Figure 1.4 examines the joint density of total workers with output conditional on the total managers and share of managers in the firm. The same positive correlation between output and labor observed in Figure 1.1 can be observed in Figure 1.4a, and the same negative correlation as in Figure 1.2 appears in Figure 1.4b. The addition of a third dimension sheds light on how the two vary together. The joint density shifts around as the number or share of managers changes, but the conditional distribution of output (holding the other two variables fixed) shifts as well. Firms with a lower manager ratio have a distribution of output conditional on the total number of workers. This implies that the share of managers must have some effect on the production process and is not merely an irrelevant share of a homogeneous input. Similar effects can be shown for other labor subcategories (e.g. production workers vs. non-production workers, contracted workers vs. non-contracted workers). If various dimensions of the labor input are relevant for firm outcomes, what is the "correct" dimension to consider?

1.3.2 Is Organization Meaningful?

While the data contain information about the number of workers within different categories, these categories may not directly reflect the organizational structure as it affects productivity. Caliendo et al. (2015) and Caliendo et al. (2020) both assess the impact of firm organization on observable outcomes by first constructing a measure of the num-

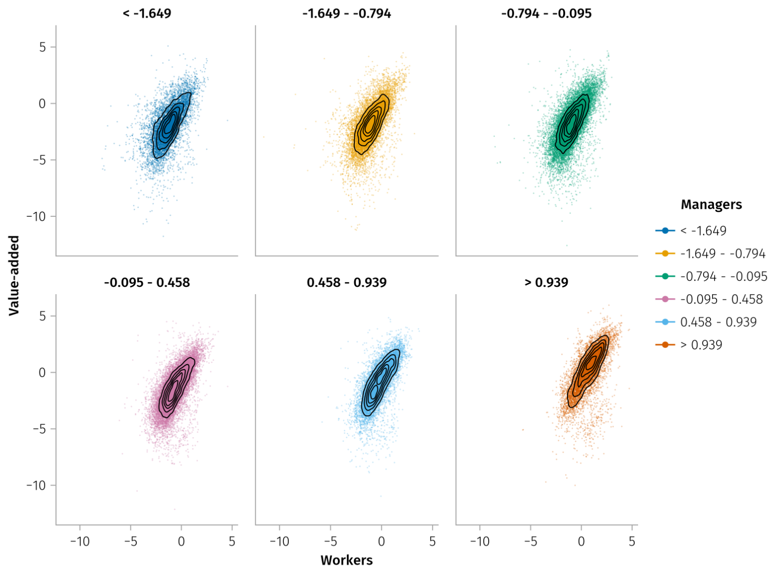

(a) Managers

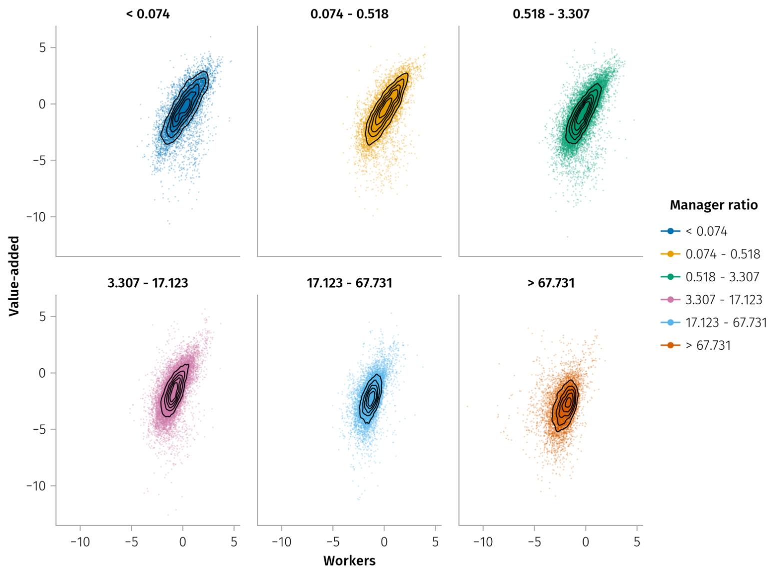

(b) Manager ratio

Figure 1.4: Joint density of value-added and labor, conditional on managers.

Value-added and labor are detrended by industry-year, and all variables except the manager ratio are reported in $\log_{10}$ scale.

ber of layers in the firm and then comparing the distributions of various outcomes across organizational groups. The goal of this analysis is to establish that the classification of employees into layers is economically meaningful rather than merely reflecting arbitrary classification schemes that do not vary systematically across firms.[7] Both papers present three key findings.

First, the literature finds that firms with different numbers of layers have different aggregate outcomes. Firms with more layers tend to produce more, hold more capital stock, employ more workers, and pay higher wages.

Second, firms with different numbers of layers choose to expand in the shape of a pyramid, with lower layers growing proportionally more than higher layers when the firm increases its output. Conversely, wages for workers in higher layers grow more quickly than for workers in lower layers.

Third, firms that restructure by adding or removing layers do not simply expand or contract; employment and wages move in opposite directions in layers that exist both before and after the change. Firms that add a layer of management choose to hire more workers, but while the number of workers in the pre-existing layers rises, the average wages within each of those layers falls. Similarly, firms that remove a layer of management also shrink the size of the layers that remain, but the workers in those layers are paid higher wages.

I replicate the patterns in these findings in order to establish that an economically meaning classification of firms into different organizational hierarchies exists in ENIA data. To identify firms with different organizations, I use the presence or absence of managers within the firm to form a classification scheme of two layers.[8] While the ENIA dataset does not record as many categories as datasets used previously in the literature, the three key patterns outlined above can still be clearly seen in the data.

1.3.3 Variation Across Organization

Do firms that are organizationally different have different characteristics? To examine whether the organizational measures condsidered above have any statistical relevance,

[7]See Caliendo et al. (2015), page 820, and Caliendo et al. (2020), Appendix B, page 10.

[8]While the number of managers in the data is subject to the addition of statistical noise, it is still possible to cleanly distinguish between cases where the number of managers must be zero and the number of managers must be strictly positive under the assumption that managers must be paid a positive wage. See Appendix B for details.

I test the effects of firm organization on the distribution of endogenous firm outcomes. Specifically, I consider capital, labor, average wages, and value-added by regressing the log of each variable on an organizational dummy, defined as the presence of at least one manager paid a positive wage, and a set of plant and industry-year fixed effects. Table 1.2 displays the resulting estimates.

Table 1.2: Firms Differ by Organization

Regression of log outcomes on an organizational dummy and plant and industry-year fixed effects

Outcome	**Estimate**	**Std. Error**	**p-value**	**Observations**
Capital	0.108	0.014	0.000	63,392
Labor	0.136	0.007	0.000	67,169
Wages	0.218	0.009	0.000	67,178
Average wages	0.084	0.007	0.000	67,141
Value-added	0.113	0.011	0.000	64,344

The estimated coefficient on the organizational measure is positive and statistically significant for all variables. The interpretation of these estimates is not causal; the claim is only that partitioning firms by their organizational status creates two subpopulations that are different from each other in a statistically meaningful way. Firms with at least one managerial layer have higher levels of input (both capital and labor stock), higher average wages, and higher production.[9] These findings are robust across different specifications of fixed effects, but other organizational measures fail to produce significantly different distributions of outcomes (see Figure 1.6 in Section 1.6.2 for details).

Figure 1.5 shows the estimated density for each outcome both before and after removing the plant and industry-year fixed effects.[10]

[9] Increases in the ratio of managers to workers, other than the change from zero by introducing managers, tends to reduce the firm's factor demands and output while increasing wages (table forthcoming). Firms with a higher concentration of managers will tend to be smaller because firms with multiple managerial layers can more efficiently use the managers' knowledge and time, resulting in fewer managers relative to production workers. If managers were merely a complementary worker category, the effects could be reversed.

[10] These fixed effects are removed by running the regression:

$$\ln y_{it} = \beta L_{it} + x_{it}'\gamma + u_{it}$$

where x_{it} is a vector of fixed effects. The estimated fixed effects $\widehat{\gamma}$ are then subtracted from $\ln y_{it}$ and replaced with the (log) median value m of the outcome for a firm without a managerial layer:

$$\ln \ddot{y}_{it} = \ln y_{it} - x_{it}'\widehat{\gamma} + m$$

The figure shows the estimated local polynomial densities forln $\ddot{y}_{it}$.

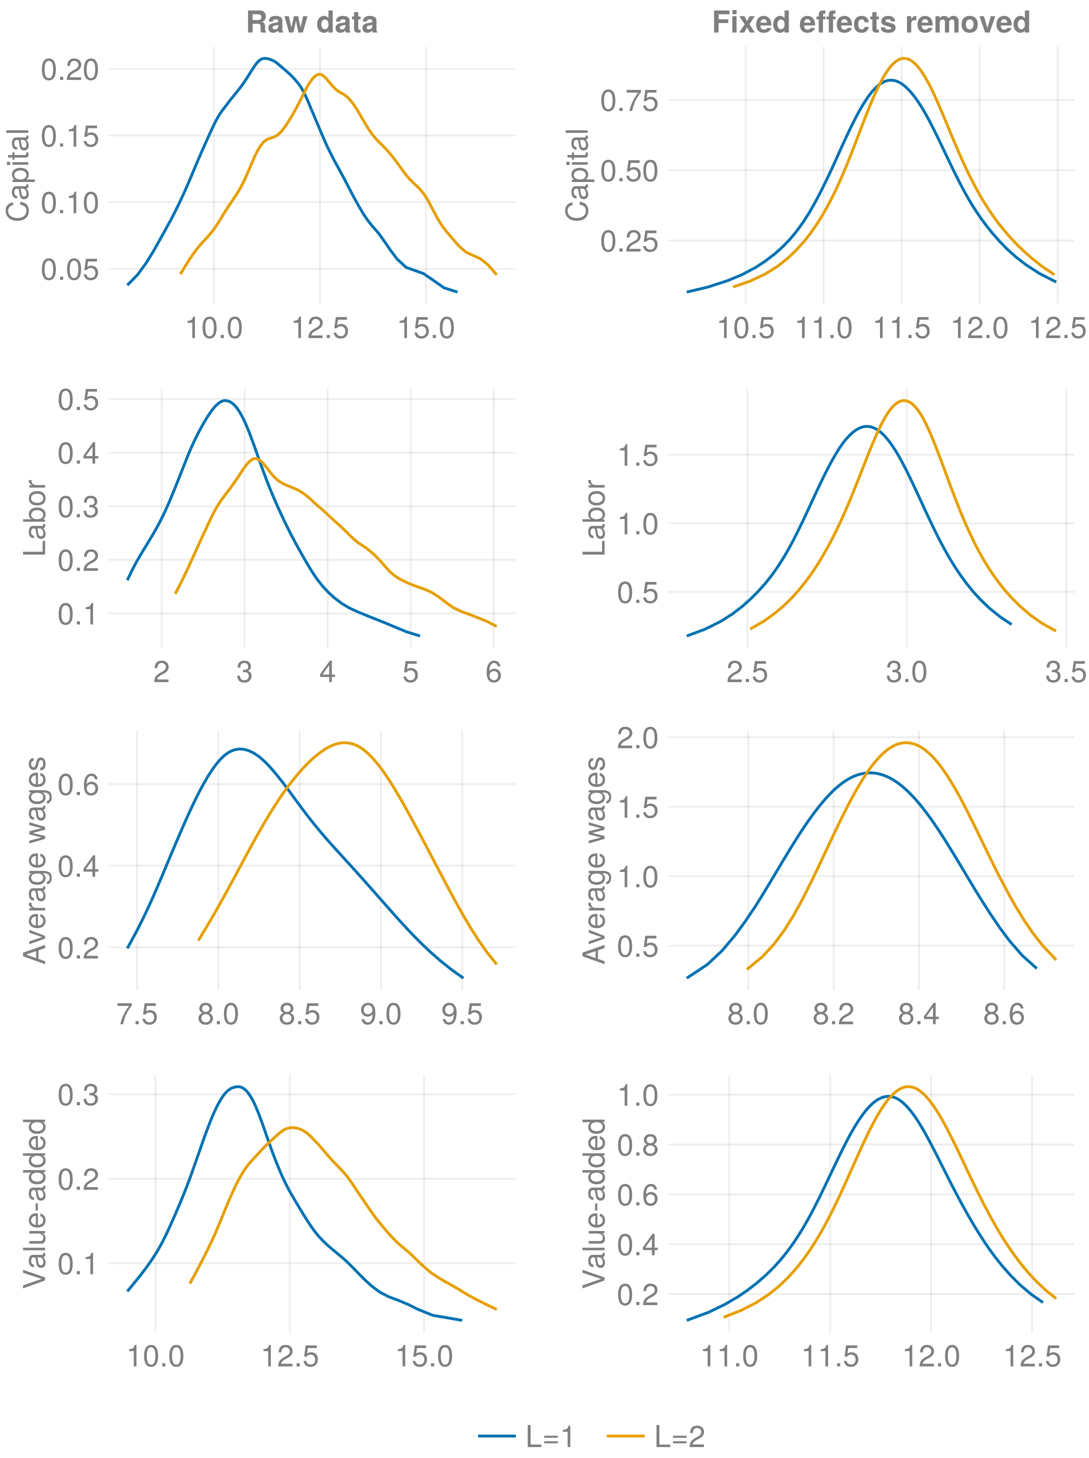

Figure 1.5: Distributions by Layer

Firms tend to maintain a hierarchical structure, with 77.8% maintaining a hierarchy in wages (higher wages paid to managers than non-managers) and 99.0% maintaining a hierarchy in the number of workers.

1.3.4 Expansion Across Organizations

Another feature of the organizational model is that as firms expand their production, they do so pyramidally: the lower a layer is, the more its numbers of workers increases and the less their wage increases.

Table 1.3: Firms Expand Pyramidally

Variable	**No. Layers**	**Layer**	**Estimate**	**Std. Error**	**p-value**	**Obs.**
Workers	1	1	0.263	0.011	0.000	15,837
Workers	2	1	0.298	0.008	0.000	31,406
Workers	2	2	0.259	0.009	0.000	31,438
Workers	2	2 - 1	-0.037	0.005	0.000	31,406
Wage	1	1	0.244	0.013	0.000	15,815
Wage	2	1	0.272	0.009	0.000	31,389
Wage	2	2	0.295	0.010	0.000	31,438
Wage	2	2 - 1	0.022	0.006	0.000	31,389

The average firm increases its production layer proportionally more than its management layer when it expands. Conversely, the wage of the management layer rises proportionally more than wage of the production layer. Not only are the estimated coefficients different, but their difference is statistically significant (as seen by looking at the row where the number of layer is "2 - 1", or the log change in the management layer minus the log change in the production layer).

1.3.5 Expansion with Reorganization

The organizational model predicts that reorganizing by adding or dropping layers is different than merely expanding or contracting. While a firm that expands by adding layers also expands in total size and average knowledge, both the size and knowledge levels of pre-existing layers shrink. To test this implication in the data, I examine whether the distribution of log changes in detrended firm outcomes changes significantly conditional on reorganization. I run regressions of the form:

$$\Delta \ln \widetilde{y}_{it} = \gamma_{-1} 1_{\{\Delta L_{it}=-1\}} + \gamma_0 1_{\{\Delta L_{it}=0\}} + \gamma_1 1_{\{\Delta L_{it}=1\}} + x_{it}{}' \beta + \varepsilon_{it}$$

where x_{it} is a vector of fixed effects or other controls and ΔL_{it} is the change number of layers.

Table 1.4: Firms that Reorganize Expand Differently

Regression of log detrended variables on reorganization and plant fixed effects.
$^{***}\, p < 0.001$, $^{**}\, p < 0.01$, $^{*}\, p < 0.05$

Variable	$\Delta L = -1$	$\Delta L = 0$	$\Delta L = 1$
Value-added	-0.041*	-0.000	0.048**
Total workers	-0.043***	0.000	0.036***
Production workers	0.064***	0.000	-0.073***
Normalized workers	0.598***	0.001	-0.627***
Average wage	-0.016	-0.001	0.041**
Average production wage	0.049***	-0.001	-0.033*
Observations	2569	50376	2540

Table 1.4 shows the results of the regression of log detrended variables on reorganization dummies and plant fixed effects.

Adding (dropping) a layer tends to coincide with expansion (contraction) of the firm's production, workers, and average wages. In contrast, both the number and wages of production exhibit the opposite pattern, shrinking with the addition of a layer and expanding with the removal of a layer. This pattern is precisely what is predicted by the organizational model of the firm and is consistent with prior empirical findings by Caliendo et al. (2015) and Caliendo et al. (2020).[11]

Firms that reorganize expand everywhere *except* the wages paid to workers in preexisting layers.

[11] Both of these papers make use of data with more granularity in the organizational structure by matching worker-level data, including information about the skill level of the worker's occupation, to firm-level data.

Wages and Productivity

Another implication of the organizational model is a positive correlation between residual productivity and firm knowledge. This comes through two channels.

First, firms that pay higher wages per worker do so because they are compensating workers for their knowledge, which allows the firm to produce more of the labor input for a given number of labor hours. Since this additional productivity is not captured by the functional form of non-organizational production technologies, it gets added to the residual.

Second, firms organized with more layers make more efficient use of the firm's knowledge for any number of workers or wage level. If the firm's organization is not captured as an input to the production technology, its effect on output will be captured by the residual.

If productivity is decomposed into two scalar components, one known to the firm at the time of its production choices and the other unknown until production is realized, only the productivity known to the firm will be correlated with organizational decisions. To test if this correlation is present in the data, we need an estimate of this endogenous component of productivity.

I begin by estimating an ordinary Cobb-Douglas production technology, regressing log value-added y_{it} on log capital stock k_{it}, log labor stock n_{it}, and a set of industry-year fixed effects z_{it}:

$$y_{it} = \alpha_0 + \alpha k_{it} + \beta n_{it} + \gamma' z_{it} + \varepsilon_{it}$$

I take the resulting residual $\widehat{\varepsilon}_{it}$ and decompose it into a scalar component ω_{it} known to the firm in period t and a scalar component η_{it} that the firm does not learn until y_{it} is realized. I estimate ω_{it} using the inverse f of the firm's demand for intermediate material inputs m_{it}.[12] This is implemented here by approximating f with a cubic polynomial in the firm's electricity consumption.[13]

$$\widehat{\varepsilon}_{it} = \underbrace{f(m_{it})}_{\omega_{it}} + \eta_{it}$$

The fitted values from this regression $\widehat{\omega}_{it}$ are then regressed on an organizational outcome

[12] As established by Levinsohn and Petrin (2003), this inverse demand depends not only on material inputs but also on the choice of other endogenous inputs. This approach does not yield consistent estimates, but it is a serviceable approximation to illustrate a correlation in the data that motivates the more complete model.

[13] The constant term is omitted as it cannot be identified separately from the intercept in the previous regression.

Table 1.5: Regression of residual productivity on knowledge measures

(a) Aggregate across industries

Variable	Coefficient	Std. Error
Average wage	0.1141	(0.0019)
Managers	0.0398	(0.0024)
Managers paid a wage	0.0510	(0.0018)
Subcontracted workers	0.0341	(0.0017)
Industrial and nonindustrial	0.1030	(0.0026)

(b) Within industries (number of 95% significant coefficients)

Variable	Negative	Positive	Total estimates
Average wage	2	76	109
Managers	11	46	101
Managers paid a wage	7	52	103
Subcontracted workers	7	44	107
Industrial and nonindustrial	5	68	98

x_{it}, such as the log of the average wage in the firm or a dummy variable indicating the presence of a subcategory of workers:

$$\widehat{\omega}_{it} = \delta_0 + \delta x_{it} + u_{it}$$

Table 1.5 reports the resulting estimates $\widehat{\delta}$ at different levels of aggregation. Table 1.5a reports the estimated coefficients pooled across industries, all of which are positive and statistically significant. Table 1.5b reports the number of statistically significant coefficients from regressions run separately for each industry. Approximately half of the industry-specific coefficients are positive and significant, with positive coefficients considerably outnumbering negative coefficients. The effect of the average wage is particularly strong across industries, with 76 of 109 coefficients being positive and significant, compared to only 2 negative and significant coefficients.

Clearly, both the average wage and the composition of the labor input are related to the firm's beliefs about its productivity, as measured by its demand for intermediate inputs. If the variation in wages affected firms purely as cost shocks, firms would contract in the face of higher wages due to a decline in the marginal profit of labor; instead, higher wages are associated with greater levels of activity and consumption of intermediate inputs. Likewise, adding different types of workers is associated with increased production. Both the mix of workers and their wages affect productivity, in accordance with the organizational model.

1.4 Estimation

I now present my estimation strategy, based on Caliendo et al. (2020). First, I discuss the assumptions I make about the production process and the timing of firm choices and shock realizations. Next, I define the dynamic structure of the productivity shocks. Finally, I discuss my estimation strategy and the moment conditions I use to identify the effect of changes in firm structure on productivity. After outlining my estimation strategy, I verify its performance with a Monte Carlo simulation.

1.4.1 Controlling for Productivity

The inclusion of the firm's organizational problem introduces a new unobservable component to the estimation problem, but it also provides a model sufficient to identify the residual efficiency using wages and labor hours for any parameter values using Equation 1.2.2. After controlling for organizational efficiency, the only remaining source of endogeneity is the productivity term ω_{it}, as in existing approaches. All that remains is to make assumptions regarding the DGP sufficient to control for productivity in the moment conditions.

I make virtually identical assumptions as Ackerberg et al. (2015) regarding the DGP. I state the assumptions regarding the unobservables ω_{it} and ξ_{it} in terms of a vector $\zeta_{it} \equiv (\omega_{it}, \xi_{it})$ to allow for the possibility of correlation between the two.

Assumption 1 – Information Set

Firm i's information set in period t ($\mathcal{I}_{it}$) contains both current and previous productivity and efficiency shocks $\{\zeta_{i\tau}\}_{\tau=0}^{t}$. The idiosyncratic shocks ε_{it} satisfy $E[\varepsilon_{it}|\mathcal{I}_{it}] = 0$.

Assumption 2 – First Order Markov

Productivity evolves stochastically via a first order Markov process.

$$E[\zeta_{it+1}|\mathcal{I}_{it}] = E[\zeta_{it+1}|\zeta_{it}]$$

Assumption 3 – Timing of Input Choices

Capital and labor are chosen at in period t using the information set $\mathcal{I}_{it}$. The labor input is produced by a combination of wages and labor hours chosen conditional on

$\mathcal{I}_{it}$.

Assumption 4 – Scalar Unobservable

Material demand depends only on the scalar unobservable ω_{it} (after conditioning on other inputs):

$$m_{it} = h_t(k_{it}, o_{it}, \omega_{it})$$

Assumption 5 – Strict Monotonicity

Material demand $h_t(k_{it}, o_{it}, \omega_{it})$ is strictly increasing (and hence invertible) with respect to the unobservable term ω_{it}. There exists then a function $g(k_{it}, o_{it}, m_{it})$ such that

$$\omega_{it} = g(k_{it}, o_{it}, m_{it})$$

Using Assumption 5, we can replace ω_{it} in the production equation:

$$y_{it} = \alpha_0 + \alpha k_{it} + \beta o_{it} + g(k_{it}, o_{it}, m_{it}) + \varepsilon_{it}$$

By the assumption that the transitory shock ε_{it} is unobserved by the firm when it makes input choices, we have the following moment condition:

$$E[\varepsilon_{it}|k_{it}, o_{it}, m_{it}] = 0$$

This moment condition alone is insufficient to identify both the parameters (α, β) and the function g. In order to acheive identification, we need to control for productivity using instruments other than k_{it} and o_{it}. We acheive this by exploiting the stochastic (joint) evolution of productivity and efficiency outlined in Assumption 2. Productivity evolves by a first order Markov process, but we allow for the possibility that this process depends on past efficiency as well as past productivity:

$$E[\omega_{it+1}|\mathcal{I}_{it}] = E[\omega_{it+1}|\omega_{it}, \xi_{it}] \equiv f(\omega_{it}, \xi_{it})$$

This produces a second identifying equation:

$$y_{it} = \eta_0 + \alpha k_{it} + \beta o_{it} + f[g(k_{it-1}, o_{it-1}, m_{it-1})] + u_{it}$$

with the residual $u_{it} \equiv \varepsilon_{it} + \omega_{it} - E[\omega_{it}|\mathcal{I}_{it-1}]$. The moment condition used for identification is:

$$E[u_{it}|k_{it}, k_{it-1}, o_{it-1}, m_{it-1}] = 0$$

Finally, we require that the efficiency shocks be zero in expectation in order to identify the scale of the labor input cost function:

$$E\left[\xi_i\right] = 0$$

Nonlinear functions of current or lagged capital, inferred labor input, wages, and materials can be used as instruments to construct sufficient moment conditions to identify all the parameters of interest.

1.4.2 Monte Carlo Simulation

This section presents details and results of a Monte Carlo simulation to verify that the estimation technique can consistently recover the parameters of interest under a known data-generating process (DGP). Consider a simple DGP where i) labor and capital are chosen statically and ii) firms know their organizational efficiency but not their productivity.

Consider a cross-section of two-layer firms $i = 1, \ldots, N$. The sample $(Y_i, K_i, W_i, N_i)_{i=1}^{N}$ consists of output Y_i, capital stock K_i, wage bill W_i, and labor stock N_i for each firm i, with the convention that lowercase letters represent logs. Firms choose inputs to maximize expected profits:

$$\max_{K,O} E\left[K^{\alpha} O^{\beta} \exp\left(\alpha_0 + \varepsilon_i\right) - r_i K - C_2\left(O; w, c, h, \lambda\right) \exp\left(-\xi_i\right) |\mathcal{I}_i\right]$$

where the firm's information set $\mathcal{I}_i$ contains the efficiency shock ξ_i but not the productivity shock ε_i. The variables $(r_i, \varepsilon_i, \xi_i)$ drive variation in firms' input demands and output. Each random driver is drawn from a Beta$(10, 10)$ distribution shifted and scaled to a compact interval: r_i is distributed on $[0.4, 1.0]$, ε_i on $[-0.75, 0.75]$, and ξ_i on $[-0.25, 0.25]$.

Wages and labor stock are determined uniquely by the firm's choice of labor input. The solution to the organizational problem provides a one-to-one mapping between observed labor stock and the unobserved labor input $o_i = o(n_i; w, c, h, \lambda)$. Combining this with the observed (log) wage bill w_i permits recovery of the efficiency shock, so we can solve for both residuals as:

$$\varepsilon_i = y_i - [\alpha_o + \alpha k_i + \beta o(n_i; w, c, h, \lambda)]$$
$$\xi_i = \ln C_2(O_i; w, c, h, \lambda) - w_i$$

The identifying moment conditions are

$$E[\varepsilon_i | k_i, o_i, \xi_i] = 0$$
$$E[\xi_i] = 0$$

The first moment condition holds given the assumptions about firm behavior, and the second assumption is a normalization that identifies the scale w of labor input costs. The conditional moment restriction on ε_i is converted into unconditional moment restrictions involving nonlinear functions of k_i, o_i, n_i, and w_i.[14]

The estimator is constructed in three steps. First, the organizational parameters are used to infer a labor input $\widehat{o}_i$ and efficiency shock $\widehat{\xi}_i$ via the one-to-one mapping given the observed labor stock n_i and wage bill w_i. The mapping is estimated by numerically solving the firm's cost minimization problem. To avoid numerically computing this at every observation for each objective function evaluation, the problem is solved on a grid of values for labor input o and wage w given the organizational parameters, then a cubic spline interpolation over this grid is used to approximate values at each observation.

Second, the inferred values $\widehat{o}_i$ and $\widehat{\xi}_i$ are used to construct the unconditional sample moments:

$$\widehat{g} = N^{-1} \sum_{i=1}^{N} \left(\widehat{\varepsilon}_i, k_i\widehat{\varepsilon}_i, w_i\widehat{\varepsilon}_i, n_i\widehat{\varepsilon}_i, \widehat{o}_i\widehat{\varepsilon}_i, k_i\widehat{o}_i\widehat{\varepsilon}_i, k_i\widehat{o}_i^{\,-1}\widehat{\varepsilon}_i, k_i^2\widehat{\varepsilon}_i, \widehat{o}_i^2\widehat{\varepsilon}_i, w_in_i\widehat{\varepsilon}_i, w_in_i^{\,-1}\widehat{\varepsilon}_i, \widehat{\xi}_i, \widehat{\xi}_i\widehat{\varepsilon}_i, \right)'$$

Third, the vector of sample moments is used to evaluate the GMM objective function (with the identity weighting matrix) $Q(\theta) = \widehat{g}\,'\,\widehat{g}$.

Table 1.6: Monte Carlo Results

Mean parameter values for 1000 repetitions of the GMM estimator for different sample sizes. Standard deviations in parentheses are obtained from the distribution of estimates across repetitions.

Parameter	**Value**	**500**	**1500**	**5000**	**10000**	**15000**
α_0	1.25	1.2515	1.2514	1.25	1.2496	1.2497
		(0.0073)	(0.0058)	(0.0034)	(0.0035)	(0.0021)

[14] The instruments used for ε_i are 1, k_i, o_i, k_io_i, k_i/o_i, k_i^2, o_i^2, n_i, w_i, w_in_i, and w_i/n_i. Furthermore, ξ_i is instrumented with a constant, and ε_i and ξ_i are assumed to be uncorrelated.

Parameter	Value	500	1500	5000	10000	15000
α	0.3	0.2994	0.2995	0.3	0.3004	0.3003
		(0.0027)	(0.0017)	(0.0011)	(0.0011)	(0.0006)
β	0.5	0.5001	0.5	0.4999	0.4998	0.4999
		(0.0028)	(0.0017)	(0.0011)	(0.0008)	(0.0006)
w	0.607	0.6065	0.6075	0.6066	0.6067	0.6069
		(0.0046)	(0.0033)	(0.0024)	(0.0013)	(0.0015)
c	0.3	0.299	0.2993	0.3008	0.3006	0.3009
		(0.0037)	(0.0033)	(0.0018)	(0.0018)	(0.0014)
h	0.5	0.5006	0.5034	0.499	0.5009	0.4992
		(0.0211)	(0.0171)	(0.01)	(0.0069)	(0.0039)
λ	0.7	0.6998	0.7011	0.7019	0.702	0.7021
		(0.0094)	(0.0056)	(0.005)	(0.0034)	(0.0041)

Table 1.6 reports the resulting estimates for 1000 repetitions of the estimator at each sample size. Estimates are quite close to the true value for all parameters, and standard errors shrink with the sample size. The null that the estimate is equal to the true parameter value cannot be rejected at standard confidence levels for any parameter or sample size.

Table 1.7: Monte Carlo Results

Mean parameter values for 1000 repetitions of a simple OLS estimator for different sample sizes. Standard deviations in parentheses are obtained from the distribution of estimates across repetitions.

Parameter	Value	500	1500	5000	10000	15000
α_0	1.25	-0.3229	-0.2842	-0.2971	-0.3014	-0.2975
		(0.174)	(0.0997)	(0.0547)	(0.0378)	(0.0374)
α	0.3	0.2964	0.3028	0.3001	0.2988	0.2993
		(0.0469)	(0.0263)	(0.015)	(0.0106)	(0.0105)
β	0.5	0.6184	0.5942	0.6024	0.6055	0.603
		(0.1136)	(0.0644)	(0.0362)	(0.0249)	(0.0248)

Table 1.7, on the other hand, reports the results of an OLS regression of log output y_i on log capital k_i and log workers n_i using the same simulated data. Both the intercept and the labor coefficient are significantly biased under the misspecified model. Interestingly, the

labor coefficient is positively biased, overstating rather than understating the sensitivity of output to changes in aggregate labor hours. One might expect the opposite, since the true labor input is always strictly less than aggregate labor hours, requiring a greater change in aggregate hours to have the same effect on output as a given change in the labor input, which in turn implies a lower output elasticity for aggregate labor. The positive bias in the labor coefficient (and large negative bias in the intercept) are consequences of the nonlinear relationship between the two measures of labor, meaning that the influence of higher-order terms in the production function cannot be ignored. Because of this influence, it is not possible to determine *a priori* the sign of the bias stemming from the omission of the firm's true organizational technology.

1.5 Conclusion

The technology by which firms combine different types of workers has implications both for the optimal mix of worker types within the firm and for the interpretation technical change. This paper shows that labor composition affects firm choices across a wide variety of industries and develops an estimation approach to identify shocks to organizational efficiency using variation in the firm's choices of wages.

The identification of the organizational component of technical change is achieved by developing a model of endogenous hierarchical organization of a firm's labor force. The relevant labor input for a firm is not simply the sum total of workers in the firm but instead reflects the ability of skilled workers to successfully transform time spent in production into output. Firms increase the labor input not only by hiring new workers but by increasing worker skill, which drives up average wages. Firms differ in their organizational efficiency, the ability to organize workers into the labor input for a given cost, resulting in different optimal mixes of worker types across firms. Variation in the number of workers for a given wage identifies the organizational efficiency component of productivity separately from the Hicks-neutral component.

I show that the optimal mix of worker type affects a wide variety of endogenous firm outcomes by studying a panel of over 13,304 Chilean manufacturing firms from 2000-2014. I find that the level of workers of different types is positively correlated with firm outcomes such as output, capital stock, demand for material inputs, and wages; the share of workers, on the other hand, has varying effects. The share of managers is negatively correlated with all outcomes while being very nearly uncorrelated with the number of managers. Furthermore, the distribution of outcome conditional on total workers varies as the share of managers changes. This suggests that not only does the share of worker types matter, the organizational technology does not depend only on the share but also on the level of managers. Firms that choose to add managers to the mix pay their non-managers less than before and wages are positively correlated with residual productivity, another prediction of the hierarchy model of firm organization.

To account for potential heterogeneity in the composition of labor, I develop an estimation approach incorporating the firm's decisions to produce the labor input. The labor input cost function is estimated parametrically, identifying both the unobserved labor input and the unobserved organizational efficiency shock. This technique is tested with a Monte Carlo simulation that recovers the both the organizational parameters and the parameters of the production function consistently when standard Cobb-Douglas estimates are biased. This implies that the organizational model can be used to augment existing production function estimation techniques to better control for the composition of the labor input.

References

Daniel A. Ackerberg, Kevin Caves, and Garth Frazer. Identification Properties of Recent Production Function Estimators. *Econometrica*, 83(6):2411–2451, 2015.

Lorenzo Caliendo and Esteban Rossi-Hansberg. The Impact of Trade on Organization and Productivity. *The Quarterly Journal of Economics*, 127(3):1393–1467, 2012.

Lorenzo Caliendo, Ferdinando Monte, and Esteban Rossi-Hansberg. The Anatomy of French Production Hierarchies. *Journal of Political Economy*, 123(4):809–852, 2015.

Lorenzo Caliendo, Giordano Mion, Luca David Opromolla, and Esteban Rossi-Hansberg. Productivity and Organization in Portuguese Firms. *Journal of Political Economy*, 128 (11):4211–4257, November 2020.

Ulrich Doraszelski and Jordi Jaumandreu. Measuring the Bias of Technological Change. *Journal of Political Economy*, 126(3):1027–1084, June 2018.

Luis Garicano. Hierarchies and the Organization of Knowledge in Production. *Journal of Political Economy*, 108(5):874–904, 2000.

Luis Garicano and Thomas N. Hubbard. The Returns to Knowledge Hierarchies. *The Journal of Law, Economics, and Organization*, 32(4):653–684, June 2016.

Luis Garicano and Thomas N Hubbard. Earnings Inequality and Coordination Costs: Evidence from US Law Firms. *The Journal of Law, Economics, and Organization*, 34 (2):196–229, April 2018.

Luis Garicano and Esteban Rossi-Hansberg. Inequality and the Organization of Knowledge. *The American Economic Review*, 94(2):197–202, 2004.

Luis Garicano and Esteban Rossi-Hansberg. Organization and Inequality in a Knowledge Economy. *The Quarterly Journal of Economics*, 121(4):1383–1435, 2006.

James Levinsohn and Amil Petrin. Estimating Production Functions Using Inputs to Control for Unobservables. *The Review of Economic Studies*, 70(2):317–341, 2003.

G. Steven Olley and Ariel Pakes. The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64(6):1263–1297, 1996.

1.6 Appendix

1.6.1 ENIA Variables and Descriptions

Variable	Description
B001	Average owners and executive male personnel
B002	Average owners and executive female personnel
B003	Average male workers with contracts associated with industrial process
B004	Average female workers with contracts associated with industrial process
B005	Average male workers with contracts not associated with industrial process
B006	Average female workers with contracts not associated with industrial process
B007	Total average workers with contracts - males
B008	Total average workers with contracts - females
B009	Total average workers with contracts
B010	Average male subcontracted workers associated with the production process
B011	Average female subcontracted workers associated with the production process
B012	Average male subcontracted workers not associated with the production process
B013	Average female subcontracted workers not associated with the production process
B014	Average male homeworkers
B015	Average female homeworkers
B016	Total average uncontracted workers - males
B017	Total average uncontracted workers - females
B018	Total average uncontracted workers
B019	Total average workers with and without contracts - males
B020	Total average workers with and without contracts - females
B021	Total average workers with and without contracts
B022	Compensation and other payments to owners and executive personnel
B023	Compensation, fee payments, and other payments to owners and executive personnel
B024	Compensation and other payments to workers with contracts associated with the industrial process
B025	Compensation, fee payments, and other payments to workers with contracts associated with the industrial process

Variable	Description
`B026`	Compensation and other payments to workers with contracts not associated with the industrial process
`B027`	Compensation, fee payments, and other payments to workers with contracts not associated with the industrial process
`B028`	Total compensation and other payments to workers with contracts
`B029`	Total compensation, fee payments, and other payments to workers with contracts
`B030`	Compensation and other payments to subcontracted workers associated with the production process
`B031`	Compensation and other payments to subcontracted workers not associated with the production process
`B032`	Payments to homeworkers
`B033`	Total compensation and other payments to uncontracted workers
`B034`	Withdrawals by owners, partners, and family members. Excludes compensation
`B035`	Dividends distributed during the period

1.6.2 Further Descriptive Results

First, I examine the effect on the conditional mean of these outcomes by running the regression:

$$y_{it} = \beta L_{it} + {X_{it}}'\gamma + \varepsilon_{it}$$

where L_{it} is equal to one if the firm belongs to a given organizational category and X_{it} is a set of fixed effects. Figure 1.6 reports the estimated values and significance levels of the coefficient β for each regression.

Organizational Measure	Definition
Contractual	Plant hires workers both with and without contracts
Industrial	Plant hires both industrial and nonindustrial workers
Managerial	Plant hires managers

Each organizational measure produces significant effects at least once (for example, all estimates for changes in labor are significant). However, only the managerial definition of organization produces consistent effects for all variables (and, remarkably, for all sets of fixed effects).

Figure 1.6: Differences in Conditional Means

The effect appears to hold across the conditional distribution for each outcome of interest.[15]

[15]This effect is confirmed by testing the difference in quantiles across the two distributions, using the bootstrap to compute the variance of the difference in quantiles. This test finds significant differences at all quantiles 0.05, 0.06, . . . , 0.95.

Chapter 2

Underbidding for Oil and Gas Tracts

2.1 Introduction

Oil and gas production in the state of New Mexico generates a tremendous amount of revenue. Home to one of the most productive oil basins in the US, New Mexico received $4.1 billion in oil and gas tax revenue in the fiscal year 2022.[1] Rights to drill on state lands are auctioned off each month by the New Mexico State Land Office (NMSLO), but the prices paid in these auctions are difficult to reconcile with the massive value provided to leaseholders. Though only one of eight leases are drilled, we find that the average profit from obtaining a lease is still well over eight times the price paid in the lease sale. Such a dramatic difference between price paid and value is cause to suspect collusive behavior on the part of bidders in these sales. This paper examines features of the New Mexico oil and gas lease sales that would facilitate collusion and proposes a series of statistical tests to rule out the possibility of BNE bidding under a wide range of equilibria.

Methods to detect collusion in first-price auctions (e.g., Porter and Zona (1993), Porter and Zona (1999), Bajari and Ye (2003), Chassang and Ortner (2019), Chassang et al. (2022), Kawai and Nakabayashi (2022), Kawai et al. (2023)) focus on detecting bidding anomalies inconsistent with BNE bidding. These methods provide antitrust authorities with a set of statistical screening devices which are aimed at settings where collusive bid patterns are not sufficiently sophisticated to disguise their intentions. We complement

[1] Reported by the New Mexico Taxation and Revenue Department:
https://www.tax.newmexico.gov/all-nm-taxes/oil-natural-gas-mineral-extraction-taxes/

the existing methods by incorporating ex post returns to construct direct tests for BNE bidding and minimum revenue. Information on ex post returns is available in many auction settings. Ex post returns are naturally available in oil and gas lease sales as lease values are measurable using publicly available production records, but can also be used more broadly in settings in which resale or secondary markets can be used to assess object value.

We consider the sale of oil and gas leases in first-price, sealed bid auctions by the NMSLO over a twenty-two-year period from 1994 to 2015. For each auction we observe each bid submitted, along with the identity of the bidder, as well as a land survey description of the tract. By connecting publicly available data on oil and gas production to geographic descriptions of leased tracts, we construct estimates of the profit generated by each lease. Examining the winners across the sample period, we find that the bidding market is highly concentrated, with the largest four bidders holding a market share of more than 50%. Since lease sales are held in person each month, these dominant bidders had ample experience with each other and many opportunities to interact.

First, we show that bidder participation decisions are correlated across bidder pairs, conditional on ex post tract value. We use a biprobit model to control for ex post value and bidder-sale fixed effects, and reject the null hypothesis of zero correlation at the 5% significance level for 21 of 48 bidder pairs. Since the conditional independence of bid distributions is an implication of the common values auction with independent conditional signal distributions, we reject the null hypothesis that the observed bids were generated by this BNE. We relax the assumption of independent signal distributions for our subsequent tests.

Second, we utilize the return data to test for the existence of profitable strategic deviations. For strategies to constitute an equilibrium they must maximize ex ante expected payoffs regardless of the information available to bidders at the time of the auction. We propose an underbidding test based on the Nash equilibrium condition that unilateral deviations cannot be profitable. Our test builds on the bid-scaling (winner's curse) test proposed in Hendricks et al. (1987) and is robust to the information structure available to bidders. We find that auction participants substantially underbid relative to the maximally profitable unilateral deviation. We find that when all bids of a bidder are multiplied by a factor of 3.2 holding rival bids constant, then the expected bidder payoff doubles, which is a violation of BNE bidding.

Finally, we examine uniform upward deviation incentives as considered in Feldman et al. (2017) and Bergemann et al. (2017). We find that winning bids are considerably below this incentive bound. Winning bids in a first-price auction are bounded away from zero; if rival bids are too low, upward deviation strategies will be profitable because the value of winning outweighs the cost of raising one's own bid. Bergemann et al. (2017) show that the lowest distribution of bids that can sustain a BNE corresponds to a particular "worst-case" equilibrium where the information structure is such that it minimizes auctioneer revenue subject to the uniform upward deviation constraint. We formulate a statistical test based on the non-profitability of uniform upward deviations using data on the distribution of ex post values and winning bids. We find that profitable upward deviations exist.

While the test results we obtain are suggestive of the presence of collusion, it is important to note that none of the tests we propose *prove* its existence. Section 2.8 discusses practical steps the auctioneer can take when facing bidders who seem likely to conspire to rig their bids.

Related Literature

The paper is organized as follows: Section 2.2 describes our framework. We describe the auction model and discuss the assumptions. Section 2.3 devises statistical testing procedures aimed at detecting collusion. Section 2.4 describes the market and highlights features that may facilitate collusion. Section 2.5 argues that bidders coordinate their bidding strategies. Section 2.6 shows that bidders underbid. Bids are too low to maximize ex ante bidder profit. Section 2.7 examines the uniform upward deviation incentive and shows that observed bids fail this bound. Section 2.8 concludes.

2.2 Framework

Our framework is the pure common values mineral rights model as described in Bergemann et al. (2017), which contains the classic mineral rights model proposed in Wilson (1977) as a special case.

A seller has one tract for sale. Bidders $i = 1, \ldots, N$ are risk-neutral and bid for the tract. The tract has a common value v contained in a compact interval $V = [\underline{v}, \overline{v}] \subset \mathbb{R}$. The value v is drawn from the cumulative distribution function (cdf) G with support V. The value distribution is common knowledge among bidders. Bidder i additionally receives private information about the value beyond knowing the prior distribution. This information comes from a signal $x_i \in [\underline{x}, \overline{x}] \subset \mathbb{R}$ that is correlated with the value v. We denote $\mathbf{X} = (X_1, \ldots, X_N)$ the random variables and $\mathbf{x} = (x_1, \ldots, x_N)$ the realizations. The joint distribution of signals and ex post tract value is $F(\mathbf{x}, v)$. The seller announces a minimum bid, or reserve price, $r \in \mathbb{R}_+$.

Denote the set of high bidders with $W(\mathbf{b}) = \{i \quad | b_i \geq b_j, \text{ for all } j = 1, .., N \text{ and } b_i \geq r\}$, where $\mathbf{b} = (b_1, \ldots, b_N) \in B^N = [0, \overline{v}]^N$ denotes the vector of bids. Let the probability that bidder i receives the good be $q_i(\mathbf{b}) = 1/|W(\mathbf{b})|$ if bidder i is among the high bidders, and $= 0$ otherwise.

A bidding strategy for player i is a mapping $\beta_i : [\underline{x}, \overline{x}] \longrightarrow B$ from signals to bids. Let Σ_i denote the set of strategies for bidder i and let $\beta \in \mathbf{\Sigma} = \times_{i=1}^{N} \Sigma_i$ denote a strategy profile.

Bidder i's ex ante payoff from the first-price auction is given by

$$U_i(\beta) = \int_{v \in V} \int_{\mathbf{x} \in [\underline{x}, \overline{x}]^N} [v - \beta_i(x_i)]\, q_i(\beta(\mathbf{x}))\, F(d\mathbf{x}, dv). \tag{2.1}$$

The profile β is a Bayesian Nash Equilibrium (BNE) if and only if $U_i(\beta) \geq U_i(\beta_i^{'}, \beta_{-i})$ for all $\beta_i^{'} \in \Sigma_i$.

Discussion of the Assumptions

Wilson (1977) and most of the subsequent empirical literature on common value auctions require stronger assumptions than stated above. Wilson assumes that the signal X_i is iid with continuous conditional cdf $F(.|v)$. The joint distribution of signals and ex post tract value is then $F(\mathbf{x}, v) = \prod_i F(x_i|v)G(v)$.[2]

Our data do not include information on bidders' signals. Proposition 4 in Laffont and Vuong (1996) establishes that the signal distribution $F(.|v)$ cannot be identified from bid data and ex post values alone. For example, monotone rescaling of signals results in observationally equivalent signal distributions. See also Somaini (2020) for identification results with interdependent signals. While signals are not identified, we can explore statistical

[2]Based on the iid signal assumption Wilson (1977) and Milgrom and Weber (1982) characterize the symmetric BNE. Let $Y_i = \max_{j \neq i} X_j$ and $F_{Y_i|X_i}(.|.)$ be the conditional distribution of Y_i given X_i. Let $u(x_i) = E[v|X_i = x_i, Y_i = x_i]$ be the expected value conditional on the own signal being x_i and the high rival signal being at most x_i. The equilibrium strategy satisfies the first-order differential equation $b'(x) = [u(x) - b(x)] \cdot \frac{f_{Y_i|X_i}(x|x)}{F_{Y_i|X_i}(x|x)}$ with boundary condition $b(\underline{x}) = u(\underline{x})$. The solution is

$$\beta_i(x_i) = u(x_i) - \int_{\underline{x}}^{x_i} L(y|x_i) du(y) \quad \text{where} \quad L(y|x_i) = \exp\left[-\int_{y}^{x_i} \frac{f(x|x)}{F(x|x)} dx\right].$$

properties of signals by examining bids instead. Bids are observed and are a strict monotone function of signals. If signals are independently distributed, then bids must be as well. We shall consider tests of this assumption using bids instead of signals in Section 2.5.

Prior empirical work on common value auctions typically adopts the Wilson BNE and imposes additional assumptions on the information structure to guarantee identification from observables. Bhattacharya et al. (2022) assume that bidder receive noisy signals of the quantity of oil in a tract. Hendricks et al. (2003) assume that the signal is an unbiased estimate of the ex post return conditional on winning. We shall consider tests of the null hypothesis that bids satisfy the conditional independence assumption in Section 2.5. A rejection of these tests may be indicative of bidder coordination prior to the auction, but could also suggest that the iid signal assumption of Wilson is not satisfied (if, e.g., signals are correlated with elements other than the ex post return). With this caveat in mind, our main statistical analysis will be based on the weaker set of assumptions described above which are robust to alternative information structures.

To summarize, our statistical analysis departs from the prior literature by using a weaker set of assumptions that is robust to all information structures, including the one proposed by Wilson. Our approach is robust to the specification of the signals and details of the Bayesian Nash equilibrium.

2.3 Testable Implications

This section describes testable implications of BNE bidding. We will formulate suitable statistical tests of these implications using the publicly available data on oil and gas lease sales from the New Mexico State Land Office (NMSLO).

The sealed-bid first-price auction has bidders submitting sealed bids and awards the item to the high bidder at his bid price. The identities of potential bidders are publicly known before every auction. On the day of the auction, the sealed bids are publicly revealed, and the high bidder wins. We let $\mathbf{b}^t$ denote the vector $(b_1^t, ..., b_N^t)$ and adopt the convention that $b_i^t = 0$ when potential bidder i refrained from bidding for lot t. We denote with z^t any information that is publicly available at time t, such as the oil and gas spot (and future) prices, that may affect bidders' signals $\mathbf{x}^t$ and thus bids $\mathbf{b}^t$. The variable v^t denotes the ex post return, which we calculate from the publicly observed drilling and production data.

We make the following assumption on the data generating process (DGP).

Assumption 1. *The observed data are* $(\mathbf{b}^t, v^t, z^t)_{t=1}^T$ *where* $(\mathbf{b}^t, v^t)$ are *identically and independently distributed conditional on exogenous covariates* $z^t \in Z$.

The assumption is commonly imposed in market games, see Tamer (2003).

We shall focus on three central implications of the mineral rights model, each requiring a decreasing amount of structure. First, we consider the classic Wilson model in which the submitted bids are independently distributed conditional on ex post returns and publicly available information at the time of the auction. Second, we relax the independence assumption and examine the null hypothesis that bidding strategies maximize ex ante expected returns, that is, that no bidder can systematically deviate from the equilibrium and receive strictly larger expected profits. Third, we examine winning bids and ex post returns data only, and consider the null hypothesis that bidders cannot find it attractive to uniformly bid upward. The last two properties hold regardless of bidders' information and the Bayesian Nash equilibrium. We shall describe these hypotheses in turn.

2.3.1 Independence

The data include detailed information on ex post drilling outcomes which allow us to calculate ex post returns for bidders, which we use as a control variable. Evidence of correlation in bids conditional on ex post returns is indicative of pre-play communication, which would be a violation of the Bayesian Nash equilibrium condition in Wilson's mineral rights model.

Implication 1. *Consider the assumption of Wilson (1977). The bids (signals)* B_i *and* B_j *are independently distributed for all* $i, j \in N$ *conditional on the ex post value realization* v.

The null hypothesis of conditional independence is

$$H_0^{B|X} : B_i \perp B_j | v \text{ for all } i, j \in N, \tag{2.2}$$

with the alternative its negation. The null can be tested for individual bidders or for bidder pairs. A violation of the null hypothesis suggests that the data were not generated from the BNE in the Wilson model. This result could indicate that the data were not generated from a BNE (e.g. because of collusion among a subset of bidders) or that the data were generated by a BNE in a game with a different information structure (e.g. if value signals are truly correlated even after conditioning on ex post values).

2.3.2 Best Response Test

Next, we relax the assumption on the information structure and go beyond the Wilson model. A simple yet powerful test examines the BNE condition ex ante using equation (2.1).

Bidder i's payoff realization from an auction is given by

$$U_i(b, v) = [v - b_i]\, q_i(b_i, b_{-i}).$$

Consider a unilateral deviation $\phi : B \longrightarrow B$ which results in the modified payoff realization

$$U_i(b, v|\phi) = [v - \phi(b_i)]\, q_i(\phi(b_i), b_{-i}). \tag{2.3}$$

The BNE condition requires that a unilateral deviation cannot be profitable on average. Consider the ex ante payoff under the deviation strategy:

$$U_i(\beta|\phi) = \int_{v\in V} \int_{\mathbf{x}\in[\underline{x},\overline{x}]^N} [v - \phi\left(\beta_i(x_i)\right)]\, q_i(\phi\left(\beta_i(x_i)\right), \beta_{-\mathbf{i}}(\mathbf{x}))\, F(d\mathbf{x}, dv). \tag{2.4}$$

This leads us to the following implication of BNE bidding.

Implication 2. *In any BNE under any information structure and for any bidder i, the function ϕ that maximizes the ex ante expected payoff in equation (2.4) must be the identity mapping.*

This property can be used to detect any deviation from BNE bidding. Our goal is to detect whether bidders systematically underbid. We follow Hendricks et al. (1987) in considering linear deviations, that is, deviation strategies that multiply all bids by a scalar $\alpha > 0$ such that $\phi(b) = \alpha \cdot b$. See also Chassang et al. (2022) who introduce multiple tests to assess whether bidders could gain by deviating from observed bidding histories for scenarios that include nonstationary settings as well as situations in which values are not observed.

Letting $\alpha^* = \arg\max_\alpha U_i(\beta, \alpha)$, Implication 2 of BNE bidding leads to the following null hypothesis:

$$H_0 : \alpha^* = 1.$$

Originally, this type of best response test was aimed at testing for the presence of the winner's curse (systematic overbidding, or $\alpha^* < 1$). We expand this test to check optimality of bidding more generally.

2.3.3 Uniform Upward Deviation

Our third test of BNE bidding focuses on the profitability of uniform upward deviations. These deviations can be examined empirically solely based on information from winning bids, and represent a special case of deviations described under Implication 2. Importantly, no information on losing bids is required for this test.

A uniform upward deviation to b is defined as a deviation $\widetilde{\beta}_i$ from strategy β_i where submitted bids are equal to $\widetilde{\beta}_i(x) = \max b, \beta_i(x)$ for any $x \in [\underline{x}, \overline{x}]$. The profitability of such a deviation is assessed by comparing the increase in payoff from winning additional auctions with the higher bid to the decrease in payoff from paying the higher bid in auctions that would have been won in equilibrium. This class of deviations has been studied before by Feldman et al. (2017), who use it to establish a lower bound on winning bids in correlated equilibria of private-value auctions, and Bergemann et al. (2017), who use the constraints imposed on the distribution of winning bids to derive a tight lower bound on expected revenue that is robust to the information structure. We use the observed distribution of winning bids and valuations to test the hypothesis that uniform upward deviations are not profitable.

Examining uniform upward deviations proves a valuable tool for empirical analysis, as it solely relies on realized ex post returns data and prices paid. It does not require information on losing bids, which may not be readily accessible. Prior empirical work specifies the information distribution of bidders or the details of the Bayesian equilibrium or information on the bidding histories. Our approach departs from the prior literature and is robust to the specification of the information structure, signals, and other details of the Bayesian equilibrium. We test whether observed returns and prices paid lead to profitable uniform upward deviations.

Let $H_i(b|v)$ denote the probability that bidder i wins with a winning bid less than or equal to b when the value is v. Note that this probability is well defined for any information structure and equilibrium strategies. For the sake of notational simplicity, we omit the explicit dependence on the information structure and equilibrium strategies. Consequently, $H(b|v) = \sum_{i=1}^{N} H_i(b|v)$ denotes the total probability that the winning bid is less than or equal to b when the value is v. In a BNE, a bidder receives the ex ante expected rent

$$\int_{\underline{v}}^{\overline{v}} \int_{0}^{\overline{b}} [v-p] H_i(dp|v) G(dv).$$

Suppose bidder i uniformly deviates up to b. For auctions where the winning bid is greater than or equal to b, outcomes and therefore payoffs are unchanged. For auctions where the winning bid is less than b, the deviator now wins the auction and receives a payoff of $v-b$. The ex ante expected payoff to the deviator is therefore given by:

$$\int_{\underline{v}}^{\overline{v}} [v-b] H(b|v) dG + \int_{\underline{v}}^{\overline{v}} \int_{b}^{\overline{b}} [v-p] H_i(dp|v) G(dv)$$

The attractiveness of the upward deviation depends on whether the gains outweigh the losses incurred from deviating:

$$\int_{\underline{v}}^{\overline{v}} [v-b] H(b|v) G(dv) - \int_{\underline{v}}^{\overline{v}} \int_{0}^{b} [v-p] H_i(dp|v) G(dv) \leq 0 \tag{2.5}$$

Let $S \subseteq \{1, \ldots, N\}$ denote a subset of bidders. By summing equation (2.5) across bidders in S, we obtain

$$|S| \int_{\underline{v}}^{\overline{v}} [v-b] H(b|v) G(dv) - \int_{\underline{v}}^{\overline{v}} \int_{0}^{b} [v-p] \sum_{i \in S} H_i(dp|v) G(dv) \leq 0. \tag{2.6}$$

Implication 3. *In any BNE under any information structure, inequality (2.6) holds for all b and all $S \subseteq \{1, \ldots, N\}$.*

We refer to Bergemann et al. (2017) Lemma 1 for a rigorous proof argument. Under the null of competitive bidding, we assume that the observed bids are generated by a profile of BNE bidding strategies that adhere to the uniform upward deviation constraint. The null hypothesis necessitates that the expected net deviation payoff is non-positive. We test Implication 3 using the null hypothesis:

$$H_0 : \frac{|S| \int_{\underline{v}}^{\overline{v}} [v-b] H(b|v) G(dv) - \int_{\underline{v}}^{\overline{v}} \int_{0}^{b} [v-p] \sum_{i \in S} H_i(dp|v) G(dv)}{\int_{\underline{v}}^{\overline{v}} \int_{0}^{\overline{b}} [v-p] \sum_{i \in S} H_i(dp|v) G(dv)} \leq 0, \tag{2.7}$$

which normalizes deviation gains by the ex ante expected rent for bidders in S.

The test for uniform upward deviation incentives is robust to details of the Bayesian game being played and the information available to bidders. The incentive constraint remains valid if bidders are asymmetrically informed, completely uninformed, or completely informed about the tract values. The bound also applies when there is unobserved auction heterogeneity, that is, when bidders publicly observe a part of the common value that is not recorded in the data. The bound is also robust to risk aversion – if bidders were risk averse, we would expect more aggressive bidding, pushing bids even higher above the bound.

Next, we shall describe the market and data.

2.4 Data

We study oil and gas lease sales held monthly by the New Mexico State Land Office (NMSLO) on the third Tuesday of each month between 1994 until 2015.[3] Every month the NMSLO distributes a list of leases which are sold at auction next month. Bidders can nominate an area for auction. Leases typically cover a land area of 320 acres.

Lease Sale Procedure

It is a legal requirement that the NMSLO awards leases by one of two types of auction formats[4]: (i) sealed-bid first-price auction and (ii) open-outcry English auction. For most of the sample period both formats were used in each monthly sale.[5] The assignment of auction format is mostly random, although conversations with the leading auctioneer during the sample period suggest that, in the case of split tracts (a tract larger than 320 acres that is split into two separate lots for the sale), the first price auction is used for the larger of the two halves or for the part closest to an existing lease owned by the bidder that nominated the tract. The monthly lease sale proceeds in two stages: First, the sealed bids for the first set of leases are unsealed and announced publicly. Every lease is awarded to the high bidder. Second, the set of English auction leases are awarded in sequence by means of an open outcry auction where all eligible bidders are present in the room. Every month the same set of bidders interact. We do not include auctions taking place after 2015 in our sample as the post-award production record is incomplete.

The lease duration is five years, during which time the winning bidder can drill a well. If oil or gas is found, then the lease can be extended until all minerals have been extracted.

[3]Data from NMSLO have been studied in the prior literature. Kong, Kong (2020, 2021) studies the relationship between first-price and English auctions. Bhattacharya et al. (2022) study the effect of post auction drilling decisions on the optimal auction design in terms of the royalty rate.

[4]The legal setting is described in New Mexico Statues, Chapter 19, Article 10, Section 19-10-17: https://law.justia.com/codes/new-mexico/2019/chapter-19/article-10/section-19-10-17/

[5]During the summer of 2016 the auctions moved online. Starting in 2019 most of the sales were conducted in first-price format only.

Any revenues arising from the well are subject to various revenue taxes and royalty payments at a rate depending on the type of lease. Additionally, leaseholders are charged a negligible rental rate, typically \$0.50 or \$1.00 per acre. The winning bidder pays corporate taxes on profits, which is a proportional deduction and thus neutral to the bidding calculus. The winning bidder pays the bid bonus immediately after the auction and receives the net return on future and uncertain minerals extracted.[6]

Our data contain detailed information on ex post drilling outcomes and production until May 2023 which enables us to calculate the future return following the prior literature. After the auction, the winner of the tract obtains the tract lease and the subsequent oil and gas extraction is publicly observed. Bhattacharya et al. (2022) and Hendricks et al. (2003) construct the value measure using the realized value of minerals extracted from ex post drilling activity. We follow this definition of the common value and construct our value estimates by matching publicly available production data to leases. We obtain a list of all oil and gas wells in New Mexico from the New Mexico Oil Conservation Division (OCD) describing the location of each well. We then match these well locations to the geographic descriptions of leased tracts provided by the NMSLO in the letting announcements. For each lease we aggregate the monthly oil and gas production of each well (collected from monthly production reports submitted to the OCD) and weigh them with deflated crude oil and gas prices. To account for production delays, we discount all returns to the date of the auction using a ten percent annual interest rate. We account for the royalty rate, which varies by the type of lease, and deduct a revenue tax of 7.1%[7]. Our final gross revenue measure equals the realization of the discounted ex post value of the oil field net of royalties and taxes. Our gross revenue measure is a lower bound on revenues, as it is based on observed production and does not include potential future production beyond 2023[8]. We define the common value v as the net return, which equals the gross revenue

[6]If during future production it is found that the well can be used to extract minerals for multiple leases, then bidders are by law required to enter a pooling agreement. While the share of leases with pooling agreements is relatively small in our data (less than four percent), it forces bidders to engage and cooperate with each other on those leases.

[7]According to Chapter 2 of the "Decision-Makers Field Guide (2002)," (available online at https://geoinfo.nmt.edu/publications/guides/decisionmakers/2002/dmfg2002_complete.pdf) there are six taxes imposed directly on oil and gas extraction and processing: (i) severance tax which amounts to 1.875% during the first 5-7 years of production and then increases to 3.75%; (ii) conservation tax of about 0.19%; (iii) emergency school tax of 3.15% for oil and 4% for gas minus drilling credit which is given some times; (iv) ad valorem production tax of about 0.39%; (v) natural gas processors tax of 0.45% and (vi) ad valorem equipment tax of 0.07%. Totaling these taxes amounts to about 7.1%. The rate of 7.1% is also reported as tax revenues obtained in the year 2000 on the value of oil and gas reported by the Taxation and Revenue Department, see Figure 5 in Decision-Makers Field Guide (2002). Following Ordin (2019) we observe that corporate profit taxes do not affect the bidding calculus as the tax rate applies proportionally.

[8]An alternative measure of gross revenues uses the oil and gas future prices at the time of the auction as weights to discounted quantities instead of the realization of the oil and gas prices at the time when production takes place. The resulting gross revenue measure is very similar in magnitude but on average slightly larger than the measure obtained using the realized prices.

measure minus the well cost.

Well costs are measured following the formula described in Kellogg (2014), which is based on drilling rig rental costs predicted by oil future prices. Kellogg studies oil extraction in the Texas region of the Permian Basin, while most of the wells we study fall in the New Mexico region of the Permian Basin. Kellogg notes a significant positive correlation between 18-month-ahead oil prices and rig dayrates. He regresses daily drilling rig rental rates on 18-month-ahead oil prices and obtains an R-squared of 0.64. This regression is used to predict the daily rig rental rate, which is then multiplied by the expected number of drilling days to get the expected rental cost, which is in turn multiplied by three to produce total expected drilling costs.[9]

We update Kellogg's previous estimate of 19.2 days to drill a well by incorporating two significant changes. First, we use a more conservative average of 34 days, calculated from a sample of 332 New Mexico drilling cost records from publicly available reports on pooling agreements[10]. Second, we take into account well cost heterogeneity by calculating separate day rate estimates based on the observed annual well cost depth quintile. This approach enables us to factor in variations in well costs associated with different well depths. Our revised Kellogg formula accounts for variation in well costs based on the observed well depth and is about 75% higher than Kellogg's.

Our well cost estimates do not account for heterogeneity in well costs across bidders or potential challenges bidders may encounter in accessing capital. While cost differences may exist across bidders, we follow Hendricks et al. (2003) and Bhattacharya et al. (2022) in assuming that these differences are small in magnitude or not known at the time of the auction. However, our counterfactual exercises explore scenarios where well costs become prohibitively costly for all but a small subset of bidders. In these exercises, we artificially reduce the number of potential competitors from thirteen bidders to seven and four. This artificial scenario allows us to analyze the impact of such constraints on the bidding process.

Table 2.1: Summary Statistics for Awarded Tracts

	All	First-Price	English
Number of Auctions	9,717	4,535	5,182
Gross Revenue (minus royalty and tax payment)	423.9	477.0	377.5
	(2,493)	(2,842)	(2,141)
Well Cost	171.1	173.2	169.1
	(658.8)	(700.5)	(620.0)
Net Revenue v	252.8	303.8	208.4
	(1,987)	(2,268)	(1,703)
Winning Bid	52.82	58.68	47.70
	(113.5)	(138.8)	(85.0)
Reserve Price	4.71	4.53	4.87
	(4.398)	(4.302)	(4.474)
Fraction Drilled	0.125	0.119	0.131
	(0.331)	(0.323)	(0.337)
Well Cost of Drilled Tracts	1,367	1,457	1,295
	(1,354)	(1,502)	(1,219)
Fraction Productive	0.111	0.107	0.114
	(0.313)	(0.308)	(0.317)
Disc Revenue of Productive Tracts	3,818	4,442	3,304
	(6,563)	(7,597)	(5,524)

The data consist of all awarded auctions between 1994 and 2015. Dollar figures are measured in thousand of 2000 US dollars. Standard deviations are in parentheses.

Descriptive Summary Statistics

Table 2.1 shows that 4,535 sales were held using first-price auctions while 5,182 using an English auction format. Table 2.1 considers only pre-2016 sales as the initial drilling decision can occur at the end of the lease term, as shown in Bhattacharya et al. (2022). All dollar magnitudes are deflated using 2000 dollars.

Strikingly, the bonus bid is very small relative to tract value. The bonus bid equals \$53,000 on average, which amounts to 21 percent of the average tract value. In comparison, the offshore lease sale literature has shown bids being much closer to the value of the tract. Hendricks et al. (1987) report winning bids in offshore sales equal 76 percent of tract value for wildcat sales and 49 percent for drainage sales.

A second surprising element is that only a small fraction of awarded tracts are drilled: about 12.5 percent.[11] Most drilled wells are productive (i.e. the well produced oil or gas).

[9] The scaling factor of three emerged from conversations with industry members who estimated that rig rental costs constitute on average one third of total drilling costs.

[10] Under pooling agreements, parties are legally obligated to submit an Authorization for Expenditure (AFE) to New Mexico's Oil Conservation Division, providing details on the anticipated cost of a proposed well. By parsing AFEs filed between 2000 and 2014, we extracted information on reported well depth, the number of days drilled and the total well cost. The relevant data can be found on the OCD website at https://wwwapps.emnrd.nm.gov/OCD/OCDPermitting/Data/Hearings/Cases.aspx.

[11] The drilling rate estimate is almost identical to the onshore drilling rate reported in the prior literature, see Bhattacharya et al. (2022).

The drilling rate and productivity rate are both low in comparison to that of offshore tracts, which Porter (1995) finds to be 78 percent and 35 percent, respectively. We hypothesize that the low initial drilling rate is a result of lease hoarding, which appears common for onshore leases[12]. We shall provide further evidence on the number of undrilled leases hoarded by individual bidders below.

Interestingly, the auctioneer's revenues are higher for first-price auctions than for English auctions, both in terms of royalty payments and bonus bids, see also Kong (2020). For split tracts, in which both auction formats were used and the assignment (according to conversations with the lead auctioneer) is essentially random, the first-price auction generates 23% higher cash bonus bid and also a 33% higher royalty return per acre. We can reject the null of identical bonus bids across the two formats at the one percent significance level. Yet, the null of identical royalty returns cannot be rejected. During the year 2019 the NMSLO began awarding leases exclusively in the more favorable first-price format.

The revenue ranking is surprising in the light of the classic theoretical work on symmetric BNE bidding equilibria in standard auction formats. Milgrom and Weber (1982) derive that English auctions generate more revenues than first-price auctions on average. One explanation for the revenue superiority of the first-price auction format is that some coordination or collusion arises in English auctions. Avery (1998) shows that bidders may use initial jump bids to signal their intention to rivals, which gives rise to multiple equilibria in English auctions, some of which may have drastically reduced revenues. Indeed, we observe jump bidding in English auctions conducted online, where the timestamp of each bid is recorded.[13]

Suspiciously low English auction bids arise also in split tract sales, where bidders should arguably have the same value estimate for both halves. There are 335 occasions where a bidder failed to win the English auction although the bidder submitted a bid in the first-price auction that was (substantially) higher in per-acre terms than the selling price in the English auction. On 194 of these occasions, the bidder failed to win both the first-price auction and the English auction. On average the bidder's losing first-price auction

[12] Accoding to a Wilderness Society's article from December 15, 2015, hoarding is common in the oild and gas industry. For instance, suspension of federal leases has affected 3.25 million acres in April 2015. See https://www.wilderness.org/articles/blog/land-hoarders-oil-and-gas-companies-are-stockpiling-your-public-lands.

[13] The NMSLO online English format is similiar to eBay sales where the current standing winning price is revealed to rival bidders and not the submitted bid. In the online English lease sale in January 2019, Slash Exploration LP started the bidding with a bid substantially above the reserve, and two attempts by rival bidders during the next sixteen hours to outbid Slash Exploration failed, resulting in Slash Exploration winning the lease.

Table 2.2: Bidding Returns for Top Bidders

Bidder	No of Bids	No of Wins	Return v	Bid b	ROI
			on average across auctions won		
YATES PETROLEUM CORP	5,810	4087	186.90	30.67	10.81
DANIEL E GONZALES	828	592	506.52	65.57	5.00
DOUG J SCHUTZ	838	587	195.40	64.21	-0.11
CHASE OIL CORPORATION	348	249	83.76	39.96	4.48
FEDERAL ABSTRACT COMP.	381	205	265.35	12.87	1.96
SLASH EXPLORATION LP	683	164	12.36	6.33	3.00
FEATHERSTONE DEV. C.	376	149	224.42	22.40	5.19
MARBOB ENERGY CORP	278	130	636.99	78.40	6.05
BAR CANE INC	220	129	338.71	60.95	3.86
RONALD MILES	369	122	812.46	91.09	5.17
THE BLANCO COMP.	617	103	378.37	14.39	57.25
FRINGE	N/A	4280	249.82	61.44	5.28

The data consist of all awarded auctions between 1994 and 2021. Dollar figures are measured in thousand 2000 US dollars.

bid was 120 percentage points higher than the final English auction price. While these bid patterns seem odd, they can in fact arise as a BNE when bidders have beliefs that they will be outbid in the English auction. Since BNE bidding in English auctions may resemble coordination or collusion, our subsequent analysis focuses on first-price auction sales.

Factors Facilitating Collusion

There are several factors in the lease sale market that may facilitate coordination or collusion. We think of collusion as an implicit or explicit arrangement to limit competition among market participants and to increase profits.

The market we study is concentrated, with three firms winning half of all leases sold at auction. In such a setting, a small set of firms coordinating their actions can have a big impact on market price. Table 2.2 reports summary statistics for bidders who won more than 100 leases between 1994 and 2021 and together account for two thirds of all bids submitted and 60% of auction awards. The table also includes a "fringe" bidder accounting for all remaining bids. We report dollar measures on average across all auctions won by the bidder. Yates Petroleum Corp has a market share of about 38% in the number of leases with more than 1,000 active (undrilled) leases (320,000 acres) held during any calendar year between 2000 and 2015. It was acquired by EOG Resources in 2016 for $2.5 billion. Yates Petroleum operated beyond the New Mexico region and held about 1.5 million acres in at least seven US states at the time of acquisition. The next largest bidders are Daniel E Gonzales and Doug J Schutz, both of whom have a 5% market share each, which amounts to more than 500 leases. These two bidders held on average more than 130 active (undrilled) leases during the period 2000 to 2015.

The rate of return from winning an auction (ROI), measured by the profits (the tract value minus the bonus bid) divided by the bonus bid, is substantial, equaling more than 899 percentage points on average across bidders. The high percentage arises as the lease acquisition cost is very low relative to the return. We shall examine the null that the bonus bid is too low to be consistent with competitive bidding in Sections 6 and 7.

Leases are homogenous products that can be resold in the future. Competition is only in terms of price, so a cartel need only coordinate in the price dimension to collude. Sales occur regularly at monthly intervals, with bidders gathering in person for each sale. Consequently, bidders know the identity of other potential bidders before they bid. Bidders may have formed relationships with each other at prior sales or as a result of pooling agreements they are required by law to enter into when a well spans multiple leases. Bidders participating in the NMSLO's auctions have faced allegations of illicit conduct in other states.[14] Multiple leases are awarded at each sale date, allowing bidders to divide the market without using side payments. Additionally, leases can be resold at subsequent periods, providing a mechanism for bidders to implement a suitable market division. The frequency of sales makes it costly for firms to deviate from any agreement. To summarize, the market exhibits characteristics that facilitate collusion. It is a concentrated market, a homogeneous product is sold, multiple leases are sold at every sale date, and sales occur at regular monthly intervals.

Next, we shall conduct statistical tests to examine whether we can reject the null of competitive bidding.

2.5 Conditional Independence Test

Competitive behavior requires that bidders submit their bids independently of each other conditional on the information available to them. Bidding strategies cannot be coordinated

[14] On March 15, 2012, the US Department of Justice filed a law suit alleging bid-rigging in Colerado, see https://www.justice.gov/atr/case/us-v-sg-interests-i-ltd-et-al. Reuters reported on June 25, 2012, that email exchanges between Chesapeake Energy Corp and a competitor apparently intended to avoid bidding against each other in Michigan, see https://www.reuters.com/article/us-chesapeake-land-deals-idUSBRE85O0EI20120625. On March 1, 2015, the US Department of Justice indicted the CEO of Chesapeake Energy Corp for bid rigging in Oklahoma, see https://www.bloomberg.com/news/articles/2016-03-01/chesapeake-co-founder-mcclendon-indicted-over-lease-bid-rigging.

or correlated; a player's strategy should be a function of their signal only. Coordination among competing bidders or information sharing is not legal at auctions. In contrast, when bidders coordinate or communicate prior to the auction, then we may expect bids to be correlated beyond the information available to bidders individually. This section considers tests aimed at distinguishing these two hypotheses based on the assumptions of Wilson (1977).

It could be argued that such patterns of coordination could also arise due to exogenous variations over time. For example, bidder pairs may be more active in certain seasons, or respond in the same way to variation in oil prices or any other exogenous shock. For example, a pattern of bid rotation would emerge when bidders are less likely to bid if a sizable number of leases have been won in the preceding sale.

We exploit the timing of individual auctions to account for these alternative explanations. Our data have multiple auctions taking place on the same date. A sale occurs once a month, and all auctions within a sale have an identical bid submission deadline. In total our data have 259 sales dates between 1994 and 2015 with an average 18 first price auctions taking place per sale.

Wilson's mineral rights model assumes the signal x_i is drawn independently from a conditional cdf $F(x_i|v,z)$, where v denotes the common value of the oil field and z are auction date variables. We do not observe the signal realizations $\mathbf{x}$ but do observe the bids $\mathbf{b}$ and the ex post outcome v and control for z using auction date fixed effects. In Wilson's BNE a bid is a strict monotone function of a bidder's signal, $b_i = b(x_i)$. Since signals are drawn independently from $F(x_i|v,z)$, and bids are a function of one signal only, bids will be distributed independently after conditioning on ex post returns. These assumptions lead to the following null hypothesis for competitive bidding:

$$H_0^{B|X} : b_i \perp b_j | v, z \text{ for all } i, j \in N, \tag{2.8}$$

with the alternative hypothesis:

$$H_1^{B|X} : b_i \not\perp b_j | v, z \text{ for some } i, j \in N.$$

As most bids in our data are equal to zero (the convention we adopt to represent that a bid was not submitted), our test statistic aggregates bids into the binary bid submission variable $s_i = \{1 \text{ if } b_i > 0; \ = 0 \text{ otherwise } \}$.

We test the null, $s_i \perp s_j | v, z$ for any pair $i, j \in N$ with a parametric bivariate probit framework. The biprobit controls for bidder-specific date effects and ex post return values (linear and quadratic). The correlation coefficient in the biprobit measures the bidder-pair correlation not accounted for by ex post return and date fixed effects z^t.

Wilson additionally assumes that the conditional cdf F is identical for all bidders. Asymmetries in the distribution of bids could arise if, e.g., the variance in the signals differs across bidders. We are not concerned with testing for the presence of such asymmetries; here we only assume that bids are independently (not necessarily identically) distributed.

Table 2.3: Biprobit Correlation Coefficients.

Pairs	H_0 of Zero Correlation rejected at			Sign of Correlation Coefficient	
	10%	5%	1%	Positive	Negative
48	22	21	12	33	15

Test results are reported for the null of a zero correlation coefficient in the bivariate probit for bidder pairs. The data consist of all pairs in which both bidders are active on sales involving at least 200 auctions. Explanatory variables include ex post return, ex post return squared and bidder specific sale-date fixed effects.

Table 2.3 indicates that a quarter to one half of bidder pairs have correlation coefficients that are significantly different from zero, depending on which significance level is considered. The majority of coefficients are positive. Consequently, we can reject the null of independence in the mineral rights model. The evidence suggests that if the game played by bidders is the same as in the mineral rights model, bid submission decisions are coordinated. This coordination arises among a quarter to one half of bidder pairs mentioned in Table 2.2.

The evidence so far has been inconsistent with competitive bidding in the mineral rights model, but could be produced by a different information structure. In the presence of unobserved heterogeneity – characteristics of the tract that are (i) informative of its value, (ii) observed by bidders, and (iii) not recorded in the data – some bidder pairs would exhibit correlation in participation decisions driven by the unobserved heterogeneity. Because unobserved heterogeneity affects the value signal in the same direction (either positively or negatively) for all bidders, any correlation induced by it should be positive, see Krasnokutskaya (2011). To examine whether the correlation we observe could be explained by unobserved heterogeneity, we examine the distribution of participation correlation coefficients across bidder pairs.

Figure 2.1 plots the histogram of correlation coefficient estimates. Some bidder pairs appear to refrain from bidding against each other, while the majority of bidder pairs

Figure 2.1: Histogram of Biprobit Correlation Coefficient Estimates

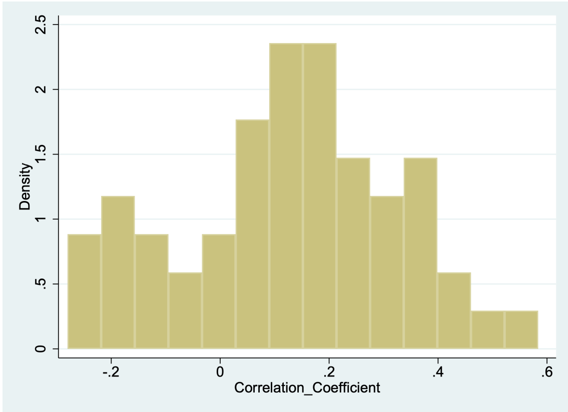

complement each other in bid submission. There are both sizable negative and positive correlation coefficients, with surprisingly little mass at zero. The histogram differs from that of a normal distribution in that it has a hump at -0.2 and another hump at 0.2. Such a bimodal distribution is inconsistent with the strictly positive correlation that would arise under unobserved heterogeneity.

The empirical evidence considered so far cannot be reconciled with competitive bidding under the mineral rights model. Having rejected the null of independent bid distributions, we are left with two possibilities. First, the bidding strategies that generated the data could be coordinated rather than competitive, e.g. some bidders refrain from bidding or submit "phony" bids. Alternatively, the correlation in bids could be explained by the underlying information structure mediating ex post returns and bids. If signals are positively correlated across some pairs and negatively correlated across others, the observed patterns of positive and negative bid correlation could arise in a competitive equilibrium. The next two sections consider statistical tests of bidding in the common values BNE that are robust to the underlying information structure.

2.6 Underbidding

This section examines whether bidders systematically under- or overbid.

If bidders coordinate in order to suppress bid payments, then such behavior will be detectable by finding the existence of a profitable deviation in the bidder's choice problem. Since we observe all bids, as well as the ex post return, we can measure the observed average payoff and use it to test the null that a systematic deviation cannot be profitable.

We develop a test procedure that is applicable regardless of the underlying information structure. We examine deviations from observed bidding in which all bids of a bidder are multiplied by a positive scalar, holding rival bids fixed. Of course, richer deviation strategies can be permitted and the test augmented. Nevertheless, in our case, even a scalar deviation results in a substantial profit increase.

Recall the null hypothesis

$$H_0 : \alpha^* \equiv \arg\max_{\alpha} \int_{v \in V} \int_{\mathbf{x} \in [\underline{x}, \overline{x}]^N} [v - \alpha \cdot \beta_i(x_i)]\, q_i(\alpha \cdot \beta_i(x_i), \beta_{-\mathbf{i}}(\mathbf{x}))\; F(d\mathbf{x}, dv) = 1.$$

We estimate the expectation using the sample average to compute the following test statistic:

$$\widehat{\alpha}^* = \arg\max_{\alpha} \frac{1}{T} \sum_t \left[v^t - \alpha \cdot b_i^t\right] q_i(\alpha \cdot b_i^t, b_{-i}^t),$$

where the objective function is the average payoff realization when bid b_i^t in all auctions is multiplied with the scalar parameter α.

A slope estimate $\widehat{\alpha}^* < 1$ suggests the potential presence of the winner's curse or risk aversion, see Matthews (1983) and Maskin and Riley (1984). On the other hand, $\widehat{\alpha}^* > 1$ suggests that bidder i underbid. Hendricks et al. (1987) propose a precursor bid scaling test and (weakly) reject the winner's curse in offshore sales. Kong (2020) develops an alternative approach that does not rely on ex post returns and finds evidence of risk aversion. Chassang et al. (2022) propose bid scaling tests designed for scenarios in which the econometrician may not observe values.

We obtain the sampling distribution for the test statistic using a block bootstrap to account for correlation due to the exogenous variables z^t by resampling blocks of auctions where an individual block consists of a six-month sequence of auctions. Kunsch (1989) shows the approach is consistent when the exogenous variable z^t is stationary.

Table 2.4: Best Response Test: Optimal Bid Scalar $\hat{\alpha}^*$

	Overall	Bidder	
		Top-5	Non-Top-5
Bid Scalar Estimate $\hat{\alpha}^*$	3.23	3.19	2.34
5th and 95th Quantile	[2.00,3.72]	[2.07,3.19]	[2.00,3.82]

The confidence region is estimated using 199 bootstrap samples by resampling from the set of all auctions using the block bootstrap.

Table 2.4 reports the scalar estimate overall and for two subgroups of bidders: the five bidders in Table 2.2 who submit the most bids and all other bidders. Bidders in all groups substantially underbid. That is, holding rival bids constant, an individual bidder would optimize their expected payoff across auctions by increasing their own bid by a factor of more than three. Clearly, we can reject the null of BNE bidding in first-price auctions, suggesting that bidders coordinate to keep prices substantially below market value. Note that the test does not establish that tripling bids would result in BNE bidding as rival bids were held constant when computing α^*; if one bidder were to triple her bid, a profitable deviation may still exist for rivals.

Figure 2.2: Simulated Profit Varying the Bid Scalar α.

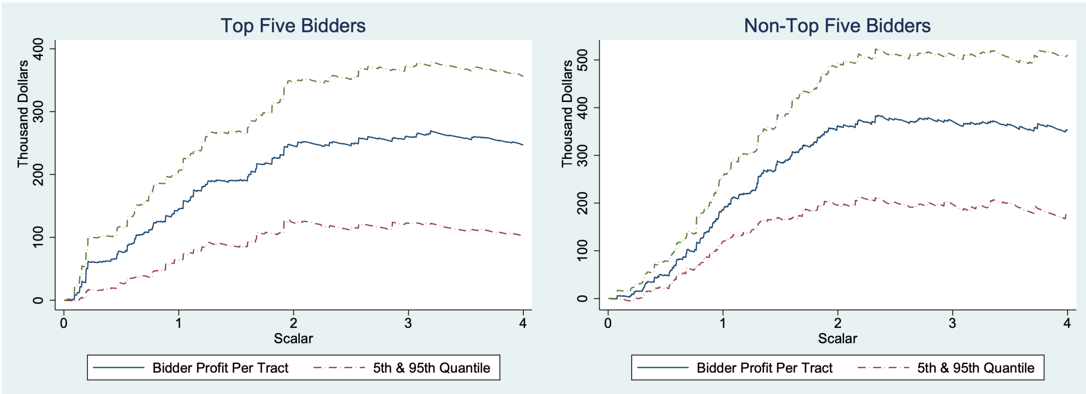

Empirical auction work typically imposes assumptions on the information structure available to bidders and the details of the BNE. A key advantage of the underbidding test is that these details do not have to be specified. The test is robust to informational assumptions and the details of the BNE because it utilizes observable ex post oil extraction returns and bid data. Ex post returns are readily available in our setting and are available more broadly in settings where the lots can be resold or there is an active secondary market for the good.

While the evidence of underbidding rejects the hypothesis of BNE bidding, it does require knowledge of losing bids. The next section provides a test that solely requires knowledge of winning bids.

2.7 Uniform Upward Deviation

This section examines whether a uniform upward deviation is profitable. This test is applicable regardless of the underlying information structure available to bidders and regardless of the BNE played. While the test does not provide a comprehensive examination of rationality, it addresses a crucial concern when facing potentially colluding bidders. Are the observed winning bids too low to be in equilibrium?

The uniform upward deviation constraint, which leads to a tight lower bound on revenues in first-price auctions as shown in Bergemann et al. (2017), is attractive for empirical purposes. It requires minimal data, solely requiring knowledge of winning bids and ex post values. It does not involve knowledge of losing bids, the underlying information structure, or the equilibrium bid strategies.

The uniform upward deviations test is operationalized based on realizations of ex post returns and winning bids. Replacing expectations in (2.7) with their sample analogues yields

$$\frac{|S|\sum_{t=1}^{T}[v^t - b\cdot a^t]\cdot 1(b\cdot a^t \geq p^t) - \sum_{i\in S}\sum_{t=1}^{T}[v^t - p^t]\cdot 1(p_i^t = p^t)\cdot 1(b\cdot a^t \geq p^t)}{\sum_{i\in S}\sum_{t=1}^{T}[v^t - p^t]\cdot 1(p_i^t = p^t)} \leq 0, \tag{2.9}$$

where v^t is the observed value, p^t the winning price, p_i^t the winning price when bidder i won the auction, and a^t the acreage of the tract. We specify deviation bids as linear in acreage, with $b \in B = \left\{\left(\frac{p^t}{a^t}\right)\right\}_{t=1}^{T}$.

As a test statistic, we examine the maximal deviation gains normalized by profits,

$$D_T^S = \max_{b\in B}\frac{|S|\sum_{t=1}^{T}[v^t - b\cdot a^t]\cdot 1(b\cdot a^t \geq p^t) - \sum_{i\in S}\sum_{t=1}^{T}[v^t - p^t]\cdot 1(p_i^t = p^t)\cdot 1(b\cdot a^t \geq p^t)}{\sum_{i\in S}\sum_{t=1}^{T}[v^t - p^t]\cdot 1(p_i^t = p^t)}. \tag{2.10}$$

We consider the null hypothesis

$$D_T^S \leq 0.$$

To obtain the sampling distribution for the test statistic, we use a block bootstrap method to account for correlation arising from the exogenous variables z^t by resampling auction blocks consisting of six-months of sales. The use of the block bootstrap enables us to obtain a consistent measure of the sampling distribution for D_T^S when the exogenous variables z^t are stationary, as shown in Kunsch (1989). We use 200 bootstrap samples for the calculations.

Table 2.5 presents the results of the test. The column labeled "All, when $|\hat{S}|$ equals 13" considers the set of all potential bidders, whose cardinality we estimate as the maximum number of bidders across all auctions, which is a consistent and superefficient estimator. While this is a commonly used approach in the empirical auction literature, we examine the robustness of the results to alternative definitions. The column labeled "All, when $|\hat{S}|$ equals 7" artificially reduces the set of potential number of bidders to seven. We also include a specification where the potential number of bidders is artificially reduced from thirteen to four.

The column labeled "Top Five" considers the five bidders in Table 2.2 who submit the

Table 2.5: Testing Uniform Upward Deviation Incentives: $H_0 : D_T^S \leq 0$.

Set of Bidders when $\lvert\hat{S}\rvert$ equals		All 13	All 7	Top-Five 5	Top-Three 3
Baseline Estimates	D_T^S	4.87	2.335	4.11	3.47
		(1.13)	(0.56)	(1.02)	(1.30)
	10-% level	3.45	1.66	3.32	2.19
	5-% level	2.95	1.40	2.79	1.80
	1-% level	1.93	0.83	1.97	1.30
Robustness Checks					
(i) Using Future Prices	D_T^S	6.09	2.96	4.64	3.85
		(0.97)	(0.49)	(1.19)	(1.74)
	10-% level	4.98	2.38	3.70	2.46
	5-% level	4.54	2.17	3.28	1.92
	1-% level	3.26	1.57	2.56	1.41
(ii) Artificially setting $\lvert\hat{S}\rvert = 4$	D_T^S	1.05	1.05	2.14	2.08
($\lvert\hat{S}\rvert = 3$ in column "Top-Five" and		(0.29)	(0.29)	(0.64)	(0.93)
$\lvert\hat{S}\rvert = 2$ in column "Top-Three")	10-% level	0.73	0.73	1.57	1.16
	5-% level	0.64	0.64	1.36	1.01
	1-% level	0.32	0.32	1.02	0.63
(iii) Multiplying Well Costs by 2	D_T^S	3.55	1.71	2.69	1.91
		(3.12)	(1.68)	(0.89)	(1.13)
	10-% level	1.45	0.70	1.68	0.62
	5-% level	0.83	0.34	1.13	0.12
	1-% level	0.00	0.00	0.15	0.00
(iv) Real Interest Rate of 15%	D_T^S	2.28	1.09	3.09	2.38
		(1.14)	(0.60)	(1.02)	(1.33)
	10-% level	1.28	0.52	1.76	0.93
	5-% level	0.87	0.36	1.29	0.69
	1-% level	0.10	0.11	0.72	0.16

Standard deviations are reported in parenthesis. The standard deviations of variables and the confidence levels of the null hypothesis are estimated using 200 bootstrap samples by resampling using the block bootstrap.

most bids. To ensure the active participation of these bidders, we select the sub-period from February 2007 until October 2013, during which all five bidders regularly submitted bids, participating in 78% of the sales.[15] Additionally, we examine a specification where the potential number of top five bidders is artificially reduced to three.

The column labeled "Top Three" considers the three bidders in Table 2.2 who submit the most bids.

According to Table 2.5, the optimal upward deviation gain equals 4.87 times profit for the average bidder. It quintuples profit for "Top Five" bidders, quadruples profit for "Top Three" bidders and triples profit for all bidders "when $|\hat{S}|$ equals 7". We can reject the null hypothesis that the deviation is not profitable at the 1% confidence level for all groups of bidders. Therefore, we reject the null of BNE bidding.

Robustness Checks

How robust is the test result? Table 2.5 reports robustness checks which relax one or more assumptions used to calculate the test statistic. We shall comment on the column "Top Five" bidders and the test results for "All" bidders "when $|\hat{S}|$ equals 7" and "Top Three" bidders are described in Table 2.5.

First, we examine the scenario where bidders did not form correct expectations about future prices correctly and instead evaluated future prices with the currently expected future price, see Hendricks et al. (1987). In this case, the test statistic increases to 4.64 for "Top Five" bidder, and we reject the null of no deviation gains at all confidence levels.

Second, we consider artificially reducing the potential number of bidders by replacing the consistent estimator $|\hat{S}| = \max N^t = 13$ with $|\hat{S}| = 4$. This implies the deviation benefit in equation (2.9) involves four bidders, while the cost include all bidders. For the "Top Five Bidders" scenario we artificially replace the deviation benefit number with an even smaller number $|\hat{S}| = 3$. For "Top Five" bidders the test statistic D_T^S becomes 2.14. We rejects the null of no deviation gains for "Top Five Bidders" at all significance levels. We can reject BNE bidding.

Third, we artificially doubled well cost estimates. This adjustment reduces the test statistic D_T^S for "Top Five" bidders to 2.69. We can reject the null of no deviation gains at all confidence levels for top "Top Five" bidders and at the 5 percent level for other

[15] Doug J Schutz and Slash Exploration LP commenced bidding in February 2007. Daniel E Gonzales stopped bidding in October 2013.

Figure 2.3: Interest Rate and Cost Multiples Needed to Satisfy Uniform Upward Deviation Incentives: Top Five Bidders.

(a) Realized Prices (b) Future Prices

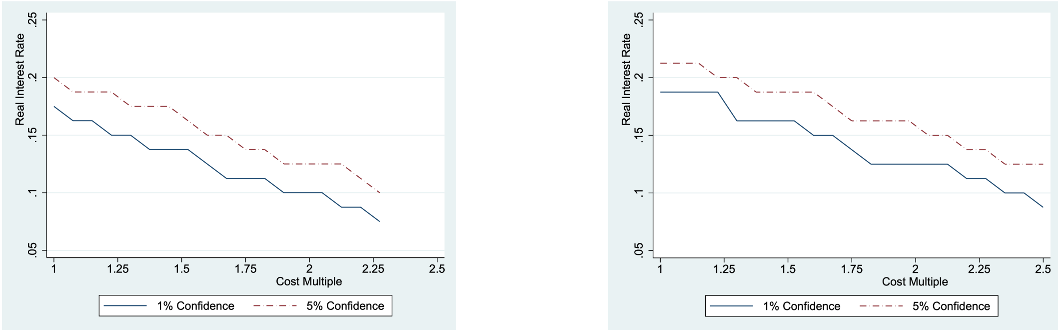

groups.

Fourth, we artificially increased the real interest rate to 15%. This modification reduces the test statistic D_T^S for the "Top Five" bidders scenario to 3.09. We can reject the null of no deviation gains for "Top Five Bidders" at all significance levels and therefore reject BNE bidding.

Finally, we illustrate graphically the joint levels of interest rates and cost multipliers required to justify the observed bids as a BNE. Figure 2.3 displays the test results indicating the rejection region of the null that $D_T^S \leq 0$ for "Top Five" bidders. We use the block-bootstrap method to obtain the sampling distribution for the test statistics. The use 1000 bootstrap samples for the calculations. Points situated to the south-west of the depicted line reject the null $D_T^S \leq 0$ at the reported significance level, while points located to the north-east cannot reject BNE. The Figure shows that a real interest rate in excess of 17.5%, or a cost increase in excess of 100% is necessary to rationalize the data as a BNE.[16] The right figure, using future prices instead of realized prices, shows that the rejection regions expands outward, requiring even higher interest rates or cost multiples. Consequently, we conclude that significantly elevated interest rates or cost multiples are required to rationalize the observed bids as a BNE.

2.8 Conclusion

This paper documents evidence of systematic underbidding in oil and gas lease sales in New Mexico. Features of this market are favorable towards bidder collusion. Leases cover

[16]The figure is created by using a discrete grid of well costs and interest rates. At each point in this grid, we employ block bootstrap sampling of the test statistic to determine the rejection region. This procedure can yield a local up-down pattern due to inherent randomness of the sampling process.

small homogeneous units of land and are awarded at regular time intervals at in-person auctions. The buyer's side is highly concentrated, with half of all leases won by only four bidders who know each other well and interact regularly.

Using the ex post value of leased tracts, we test for the presence of non-competitive bidding in three ways. First, we test whether bidder participation decisions are uncorrelated conditional on ex post returns and find statistically significant evidence of both positive and negative pairwise correlation. Second, we test whether bidders maximize their expected profit (holding rival strategies constant) and find that bidders could approximately double their expected profit by more than tripling each submitted bid. Finally, we test whether bidders can increase payoffs by uniform upward deviations. We find that bidders could increase their expected profit by more than fourfould which is inconsistent with BNE bidding.

There are several steps NMSLO can take, some of which it has already taken, to combat low auction revenues and to move toward the best-case outcome. First, the NMSLO could raise the reserve price, which has occurred in recent years. Using information from prior production outcomes of neighboring tracts, the reserve price could be raised much further to a level close to the predicted lease value. Second, information about lease values, from geological studies and historic production data on neighboring tracts can be made available to bidders along with the lease sale announcement, which would reduce informational asymmetries between bidders and encourage competition. Third, NMSLO has made changes in regulation that make it more difficult for firms to acquire leases and renew them without drilling for oil. This makes the practice of hoarding land to protect any information rent a bidding ring may have more expensive, as it necessitates the drilling of wells. Fourth, barriers to entry could be reduced by attracting new bidders, which was encouraged with the shift to online auctions in 2016. Fifth, the identities of bidders could be concealed, making it more difficult to detect deviations from the collusive agreement. NMSLO introduced this practice when it moved to online auctions. Sixth, packages of tract can be offered at auctions instead of individual tracts and lease sales could take place at less frequent time intervals. Doing so will increase the benefits from deviation from a collusive agreement making collusion more difficult to sustain. Taken together, these steps may limit the potential for collusion which is a primary concern for the NMSLO if it is to meet its objective of "optimizing revenues while protecting [New Mexico's] heritage and [its] future."

While it is premature to fully assess the impact of these measures on oil and gas exploration, their influence is already reflected in auction revenues. When comparing the four-year period before and after 2016, auction revenues per acre nearly tripled.

References

Christopher Avery. Strategic Jump Bidding in English Auctions. *The Review of Economic Studies*, 65(2):185–210, 1998.

Patrick Bajari and Lixin Ye. Deciding between Competition and Collusion. *The Review of Economics and Statistics*, 85(4):971–989, 2003.

Dirk Bergemann, Benjamin Brooks, and Stephen Morris. First-Price Auctions with General Information Structures: Implications for Bidding and Revenue. *Econometrica*, 85 (1):107–143, 2017.

Vivek Bhattacharya, Andrey Ordin, and James W. Roberts. Bidding and Drilling under Uncertainty: An Empirical Analysis of Contingent Payment Auctions. *Journal of Political Economy*, 130(5):1319–1363, May 2022.

Sylvain Chassang and Juan Ortner. Collusion in auctions with constrained bids: Theory and evidence from public procurement. *Journal of Political Economy*, 127(5):2269–2300, 2019.

Sylvain Chassang, Kei Kawai, Jun Nakabayashi, and Juan Ortner. Robust Screens for Noncompetitive Bidding in Procurement Auctions. *Econometrica*, 90(1):315–346, 2022.

Michal Feldman, Brendan Lucier, and Noam Nisan. Correlated and Coarse Equilibria of Single-Item Auctions. 2017.

Kenneth Hendricks, Robert H. Porter, and Bryan Boudreau. Information, Returns, and Bidding Behavior in OCS Auctions: 1954-1969. *The Journal of Industrial Economics*, 35(4):517–542, 1987.

Kenneth Hendricks, Joris Pinkse, and Robert H. Porter. Empirical Implications of Equilibrium Bidding in First-Price, Symmetric, Common Value Auctions. *The Review of Economic Studies*, 70(1):115–145, 2003.

Kei Kawai and Jun Nakabayashi. Detecting Large-Scale Collusion in Procurement Auctions. *Journal of Political Economy*, pages 000–000, January 2022.

Kei Kawai, Jun Nakabayashi, Juan Ortner, and Sylvain Chassang. Using Bid Rotation and Incumbency to Detect Collusion: A Regression Discontinuity Approach. *Review of Economic Studies*, 90(1):376–403, January 2023.

Ryan Kellogg. The Effect of Uncertainty on Investment: Evidence from Texas Oil Drilling. *The American Economic Review*, 104(6):1698–1734, 2014.

Yunmi Kong. Not Knowing the Competition: Evidence and Implications for Auction Design. *RAND Journal of Economics (Wiley-Blackwell)*, 51(3):840–867, September 2020.

Yunmi Kong. Sequential Auctions with Synergy and Affiliation across Auctions. *Journal of Political Economy*, 129(1):148–181, January 2021.

Elena Krasnokutskaya. Identification and Estimation of Auction Models with Unobserved Heterogeneity. *The Review of Economic Studies*, 78(1):293–327, 2011.

Hans R. Kunsch. The Jackknife and the Bootstrap for General Stationary Observations. *The Annals of Statistics*, 17(3):1217–1241, 1989.

Jean-Jacques Laffont and Quang Vuong. Structural Analysis of Auction Data. *The American Economic Review*, 86(2):414, May 1996.

Eric Maskin and John Riley. Optimal Auctions With Risk Averse Buyers. *Econometrica (pre-1986)*, 52(6):1473, November 1984.

Steven A Matthews. Selling to Risk Averse Buyers With Unobservable Tastes. *Journal of Economic Theory*, 30(2):370–400, August 1983.

Paul R. Milgrom and Robert J. Weber. A Theory of Auctions and Competitive Bidding. *Econometrica*, 50(5):1089–1122, 1982.

Andrey Ordin. Investment and Taxation: The Case of Oil and Gas in the Permian Basin. Working Paper, 2019.

Robert H Porter. The Role of Information in U.S. Offshore Oil and Gas Lease Auctions. *Econometrica*, 63(1):1, January 1995.

Robert H. Porter and J. Douglas Zona. Detection of Bid Rigging in Procurement Auctions. *Journal of Political Economy*, 101(3):518–538, 1993.

Robert H Porter and J Douglas Zona. Ohio School Milk Markets: An Analysis of Bidding. *The Rand Journal of Economics*, 30(2):263–288, 1999.

Paulo Somaini. Identification in Auction Models with Interdependent Costs. *Journal of Political Economy*, 128(10):3820–3871, October 2020.

Elie Tamer. Incomplete Simultaneous Discrete Response Model with Multiple Equilibria. *The Review of Economic Studies*, 70(1):147–165, 2003.

Robert Wilson. A Bidding Model of Perfect Competition. *The Review of Economic Studies*, 44(3):511–518, 1977.