Dynamic incentivization in relational contracts: Evidence from Indian manufacturing

Aroon Narayanan

November 13, 2024

Abstract

This paper studies a setting in which a buyer uses relationships with multiple suppliers in order to screen and incentivize them for good performance, influencing outcomes via relationship management rather than monetary tools. This leads to starkly different dynamics than standard contracting, in particular on how relationships evolve, due to extended history dependance. I show that there is an additional *competition value* motive for exploration in such a setting, above and beyong the standard option value motive. After developing a novel structural method for Principal-Agent dynamic games, I estimate it using outcome-level data on all relationships between a buyer and their suppliers, engaged in small-scale manufacturing in India. I find that the buyer's long term pool of supplier is on average 28% better than the general pool of supplier, and the buyer's incentivization scheme improves outcomes by 7% on average.

1 Introduction

The making, breaking, and remaking of relationships is an integral part of the contracting process across the world, especially in developing countries. The relational nature of these environments often necessitates incentive provision via promises of closer future engagement and threats of relationship terminations, in order to ensure that participants in the relationship engage in surplus-enhancing activities when it is in effect.

This paper advances our understanding of relational contracts by examining a novel setting in the Indian manufacturing sector. I focus on a buyer that incentivizes small manufacturing firms to produce high quality output by promising them more orders in the future if they do so. The setting provides fertile ground for exploring how firms deploy dynamic tools to navigate the challenges of contract enforcement and quality assurance when traditional mechanisms are unreliable or prohibitively costly. Theoretically, this leads to a situation where the principal (i.e. the buyer) may need to offer fully history-dependent contracts at the optimum. I show that converting this problem into the continuation value space allows for a recursive formulation, utilizing the fact that effort today by the suppliers depends on the variation in future continuation values promised by the buyer. This allows me to build a structural model that can be taken to the data for estimation; as far as I know, this is one of the first instances of structural estimation of a Principal-Agent dynamic game. Empirically, the strength of the data lies in the fact that unlike standard datasets used in IO and organizational economics, outcomes of each order are observed for the entire length of each relationship. This allows me to track exactly what happens over the course of the relationship, and how the parties respond to what happens. Estimating the model speaks to some crucial issues in the development context - why multinational companies multi-source in developing countries, how much value is added by intermediation in such markets, and how much value is lost due to weak institutions. In an early preview of the results, I find that the buyer improves outcomes by 28% over the long term via selection of better types, and by 7% via incentivization. Even though the buyer selects better types over time, she keeps some intermediate types in the pool, partly in order to motivate the higher types to continue performing well.

While the existence of relational contracting has been consistently reaffirmed by many papers in recent years, the contours of how they are implemented are still being explored. In textbook relational contracting models (such as Levin (2003)), relationships are long and feature identical incentive each each period. This project rationalizes the fact that in practice most relationships end quite early, and values grow substantially over time in those that survive. Yet another novel contribution of the project is to show that in relational contracts, learning about the counterparty has a *competition value* in addition to the traditional option value. Having a larger pool of good suppliers increases a firm's outside option, allowing it to credibly threaten to replace the suppliers more often on bad performance, and hence ensuring higher quality today. This provides an additional impetus for exploration, which again reinforces why firms are willing to shop around for suppliers in developing countries.

The findings of this paper have important implications for both economic theory and policy. The continuation value based structural estimation framework I develop has wide applications beyond this specific setting. By demonstrating how relational contracts evolve over time and how firms use dynamic incentives to manage relationships, this paper contributes to a more nuanced understanding of firm behavior in settings with weak formal institutions. These insights can inform policies aimed at improving the functioning of supply chains in developing countries and highlight the role of intermediaries in facilitating market transactions. Moreover, shedding light on the mechanisms through which firms manage quality and reliability in their supplier base can help guide strategies for fostering industrial development and improving the integration of small manufacturers into global value chains.

This paper builds on and contributes to several strands of literature. First, it adds to the growing body of work on relational contracts, particularly their empirical applications. Early theoretical work by Bull (1987) and MacLeod and Malcomson (1989) laid the foundation for understanding self-enforcing agreements in repeated interactions. Levin (2003) provided a tractable framework for analyzing optimal stationary relational contracts, which has been extended in various directions. Empirical work in this area includes Macchiavello and Morjaria (2015)'s paper on the the temptation to renege in relational contracts in the Kenyan rose export market and Macchiavello and Morjaria (2021)'s study of how competition can degrade relational contracts. This paper contributes to this literature by introducing a structural approach to estimating relational contracts, allowing for a more direct application of theoretical tools.

Second, this paper relates to the literature on structural estimation of dynamic games. Following the seminal work of Rust (1987) and Hotz and Miller (1993) on single agent dynamic estimation, Bajari et al. (2007) and Pesendorfer and Schmidt-Dengler (2008) developed methods for estimating dynamic games. Our approach builds on these methods, extending them to a setting with a principal-agent structure and relational incentives. This combination of relational contracts and structural estimation of dynamic games is a novel contribution to the literature.

Finally, our theoretical approach draws on the literature on solving dynamic games, particularly in the context of repeated moral hazard problems. The seminal work of Spear and Srivastava (1987) introduced recursive methods for analyzing repeated agency problems. Fudenberg et al. (1994) provided a general framework for analyzing repeated games with imperfect public monitoring. Our theoretical contribution lies in adapting these methods to a setting with multiple agents and a principal, incorporating learning and exploration motives into the dynamic contracting problem.

2 Model

A buyer, Z (she), sources orders from downstream customers, and selects a supplier $s \in S$ (he) to fulfil each order.

Orders arrive in discrete time $t = 1, 2, 3, ... \in T$ according to an order arrival process $O : T \to \mathbb{N}$, where O_t is the number of orders that need to be assigned at time t. The arrival process is iid across time, i.e. the realization of O_t is independent of the past realizations of O and identical across time, so the arrival process is identified by the parameter $P(O_t = n)$, which is the probability that there are n orders to assign in the current period. An order o has observable characteristics X_o , such as downstream customer identity, delivery date, metric tonnage, payment to supplier etc, and must be assigned in the same period that they arrive.

The stage game at each time t has the following sequential form. First, the buyer assigns a supplier to each order that has arrived at time t. Then, the assigned suppliers select an effort level and engage in production. This leads to an outcome realization for each order that the two parties commonly observe.

The outcome of each order is denoted $y \in \{0, 1\}$. The probability of y = 1 is given by $P(y = 1) = p(\theta, e)$, where θ is the type of the supplier and e is the effort put in by the supplier, which has a private cost $c(e, \xi)^1$. c is assumed to be convex in e, with c''(0) bounded away from 0.

The suppliers get flow benefit b from being assigned an order, and 0 otherwise. Z gets flow benefit $\frac{\beta}{1-\beta}b^2$ from an outcome of y = 1 and $\cot -\phi$ from y = 0. There is no incentivization via monetary transfers in this setting, either via monetary bonuses in the same period or via variation in future transfers³.

Z induces effort via the choice of an order assignment policy $A : H_t \times O_t \to S$, where H_t is the set of all public histories at time t. Here $A(h_t, o_t)$ specifies how orders are assigned at time t given the history of the game so far and the number of orders to be assigned today. We assume Z lacks commitment power, and hence A must be credible in the dynamic sense. Note that any policy can be rewritten as a policy that tracks a preference ordering of suppliers,

¹We assume that c() is such that the first best effort always leads to an interior probability of the high outcome.

²Here β is the buyer's bargaining weight in its relationship with the suppliers

 $^{^{3}}$ This is a simplification of the setting based on our observation of the contracts. A justification for the use of allocation incentives vs monetary ones based on relative credibility is in the works, but will not change this basic feature.

which is just $B: H_t \to \succeq (S)$.

3 Theory

3.1 Intractability in the time domain

We begin with the observation that the space of policies $B(h_t)$ is infinite dimensional, with a decision to be made at every possible history for infinite time. Moreover, the optimal policy will in general not be stationary, since it would depend in complicated ways on the outcome history in order to provide the best possible incentives for effort.

At the same time, the supplier's optimal response at a given time t to a policy has infinite dependence on all future decisions according to the policy and her own responses to them i.e. the supplier's effort today depends on the effort she will put in at all future points in time. The reason for this is quite intuitive - supplier effort today is determined by how attractive she views doing well today, which itself is determined by whether she gets some respite from working as a result of doing well today.

These factors make working in time space forbiddingly intractable, which motivates a transformation to the continuation value space, along the lines of Spear and Srivastava (1987).

3.2 Moving to the continuation value space

For the purposes of illustration, consider the case of a single supplier. We will relax the assumption that orders always need to be assigned, so that the supplier can be dropped or punished by withholding orders.

Suppose at some time t, Z promises some level of discounted lifetime value to the supplier starting from time t. Then, Z's problem can be recast as that of "optimally" picking future continuation values for the supplier conditional on outcome today, in a way that stays true to this promised value. The key step here that allows this to take on a recursive formulation is the fact that effort today by the supplier is pinned down by the difference in continuation values that she is promised on the high and low outcomes respectively. So for any promised continuation value today, Z effectively just needs to decide the optimal variation in next period's promised continuation values, which pins down effort, and then pick the optimal levels of future continuation values that are consistent with the chosen variation and the promised value today. This is how we build a Markovian recursive problem that Z must solve - in the continuation value space rather than in time.

Formally, suppose Z promises a continuation value of w to the supplier. Then, consider the problem of choosing the optimal $w^{f}(1)$ and $w^{f}(0)$ i.e. the continuation values to promise conditional on high outcome and low outcome respectively. This problem can be written as:

$$V(w) = \max_{w^f} \left\{ \sum_{(y) \in \{0,1\}} p(y)u(y) + \delta V(w^f(y)) \right\}$$
(1)

where the probability of high output $p(y) = yp(e^*, \theta) + (1 - y)(1 - p(e^*, \theta))$, e^* is supplier effort, and V is the value to Z from at a given continuation value. This is subject to the condition that value promised today must be consistent with what is promised tomorrow

$$w = b\mathbb{1}_{a=1} + \delta \sum_{(y) \in \{0,1\}} P(y) w^f(y)$$

where $\mathbb{1}_{a=1}$ captures whether a = 1 i.e. the supplier is assigned an order this period. Note that a is implicitly specified when w and $w^f(y)$ are specified, since only one value of a will satisfy the above equation.

We still need to specify the optimal supplier effort level e^* in the above problem. First, note that given $w^f(1)$ and $w^f(0)$, e can be deduced from the following problem that the supplier solves:

$$e^* = \max_{e} \left\{ p(e)w^f(1) + [1 - p(e)]w^f(0) - c(e) \right\}$$

Clearly e^* is a function only of $w^f(1)$ and $w^f(0)$ - in fact only of $\Delta w := w^f(1) - w^f(0)$. Moreover, we have the following system of equations for $w^f(1)$ and $w^f(0)$ in w and Δw :

$$\Delta w = w^{f}(y = 1) - w^{f}(y = 0)$$
$$w = b\mathbb{1}_{a=1} + \delta \sum_{(y) \in \{0,1\}} p(y) w^{f}(y)$$

The solutions to this system of equations, which represent the future continuation values that can be induced, will be labelled as $F(w, \Delta w)$. An important consequence of recasting the problem this way is that the optimal effort level e^* is fixed once Δw is fixed, and hence the probability of the high outcome p(y) is fixed. Thus for any given current continuation value, fixing the difference in future continuation values fixes supplier effort and provides us with a set of future continuation values that the buyer can induce. This allows us to rewrite Equation 1 into a proper constrained Bellman:

$$V(w) = \max_{\Delta w} \left\{ \max_{(w^f)\in F(w,\Delta w)} \sum_{(y)\in\{0,1\}} p(y|\Delta w)u(y) + \delta V(w^f(y)) \right\}$$
(2)

Some loose ends remain to be addressed. The set of w_i , say \mathcal{W}' , that can be induced by some supplier assignment policy need not be an interval of the real line - in fact it need not even be a dense subset of the real line. As a corollary, the set of Δw that can be chosen given any w_i need not be an interval either. This can mean that the value function V can be hard to evaluate numerically, with implications for tractability. This motivates a further analytical innovation - in the theoretical analysis, we allow Z to randomize between supplier assignment policies. As a result, the set of w that can be induced will lie in $\mathcal{W} = [\inf \mathcal{W}', \sup \mathcal{W}']$. Denoting the objective in Equation 2 as $U(w_i, \Delta w)$, the value function becomes:

$$V(w_i) = \max\left\{\max_{\alpha \in [0,1]: \alpha w' + (1-\alpha)w'' = w_i} \left[\alpha V(w') + (1-\alpha)V(w'')\right], \max_{\Delta w} U(w_i, \Delta w)\right\}$$
(3)

Now, we can state our first result.

Proposition 1. The value function V is concave. Moreover, it has at least one interior maximizer.

Proof. To see concavity, note that a linear combination of two points on the value function would be one of the candidates in the inner maximization in Equation 3. As a result, the linear combination of values of the two points cannot be greater than the value at the combination. A concave function on a bounded interval must attain a maximum. The fact that at least one is interior can be seen from the fact that the maximizer cannot be at the extreme points. At the lower extreme, the supplier gets her lowest promised continuation value, which means she never gets employed. At the other extreme, she gets her highest promised continuation value, which means she is always employed. Z can improve on the first policy by employing her sometimes, and on the second by firing her sometimes.

The following proposition hints at some sort of tenure system being optimal in this setting.

Proposition 2. The optimally chosen $w_f(1|w)$ is increasing in w, while Δw is decreasing in w.

Moreover, given the value function V and the optimal control distribution \tilde{w} , the associated allocation decision can be constructed inductively as follows:

- 1. Calculate $w_0 = \arg \max_w V(w)$
- 2. $B(h_0) = 1$ iff $w^* = b + \delta \mathbb{E}[\tilde{w}(w_0)]$
- 3. For each $w_t \in \operatorname{supp}[\tilde{w}(w_{t-1})]$: $B(h_t; w_t) = 1$ iff $w^t = b + \delta \mathbb{E}[\tilde{w}(w_t)]$

3.3 Multiple suppliers

The case of multiple suppliers and multiple assignments each periods require further innovations. For simplicity, consider the case with two suppliers at two orders to be assigned each period. Extending to more suppliers and more orders will be straightforward once the binary case is solved, as will become clear shortly.

Z's value function now can be written as:

$$V(w_1, w_2) = \max_{\substack{B \subseteq N, \\ w_1^f: \{0,1\}^2 \to \mathbb{R}, \\ w_2^f: \{0,1\}^2 \to \mathbb{R}}} \left\{ \sum_{\substack{(y_1, y_2) \in \{0,1\}^2 \\ w_2^f: \{0,1\}^2 \to \mathbb{R}}} P(y_1, y_2) \left[u(y_1, y_2) + \delta V(w_1^f(y_1, y_2), w_2^f(y_1, y_2)) \right] \right\}$$

subject to

$$P(y_1, y_2) = \prod_i \left[y_i(\theta_i + (1 - \theta_i)e_i) + (1 - y_i)(1 - \theta_i)(1 - e_i) \right]$$

$$e_i = \arg\max_e \left[b \mathbb{k}_{i \in B} + \delta \sum_{(y_1, y_2) \in \{0, 1\}^2} P(y_1, y_2) w_i^f(y_1, y_2) - c(e) |e_{-i} \right]$$

$$w_i = b \mathbb{k}_{i \in B} + \delta \sum_{(y_1, y_2) \in \{0, 1\}^2} P(y_1, y_2) w_i^f(y_1, y_2)$$

There are two inconveniences in this formulation. First, the optimization is over a 2^n dimensional set of continuous-valued continuation values. Second, conditional on any choice of these continuation values, there is an induced game between the suppliers on effort choice: each supplier's incentive is affected by the effort choice of the other. This can become extremely unwieldly when scaled to multiple suppliers - both theoretically and computationally.

But a reformulation of the problem can take care of these inconveniences. Consider the following reformulation:

$$V(w_1, w_2) = \max_{\Delta w_1, \Delta w_2} \left\{ \max_{(w_1^f, w_2^f) \in F(w_1, w_2, \Delta w_1, \Delta w_2)} \sum_{(y_1, y_2) \in \{0, 1\}^2} P(y_1, y_2) u(y_1, y_2) + \delta V(w_1^f(y_1, y_2), w_2^f(y_1, y_2)) \right\}$$

subject to

$$P(y_1, y_2) = \prod_i [y_i(\theta_i + (1 - \theta_i)e_i) + (1 - y_i)(1 - \theta_i)(1 - e_i)]$$
$$e_i = \arg\max_e [b + \delta e(1 - \theta_i)\Delta w_i - c(e)]$$

and $F(w_1, w_2, \Delta w_1, \Delta w_2)$ is the set of solutions to the following set of equations:

$$\begin{aligned} \Delta w_i = & [\theta_{-i} + (1 - \theta_{-i})e_{-i}][w_i^c(y_i = 1, y_{-i} = 1) - w_i^c(y_i = 0, y_{-i} = 1)] \\ & + [1 - (\theta_i + (1 - \theta_i)e_i)][w_i^c(y_i = 1, y_{-i} = 0) - w_i^c(y_i = 0, y_{-i} = 0)] \\ & w_i = & b + \delta \sum_{(y_1, y_2) \in \{0, 1\}^2} P(y_1, y_2) w_i^f(y_1, y_2) \end{aligned}$$

Fixing $\Delta w_1, \Delta w_2$, the effort choices of the suppliers are determined directly. Thus, the key here is that Z must promise them an expected variation in continuation values rather than the full array of continuation values, so that the supplier problems become decision theoretic rather than game theoretic. Given the solutions to these supplier problems, Z needs to pick the right array of continuation values that is consistent with those solutions.

Having multiple suppliers allows me to speak more directly to the underlying mechanisms that we will encounter in the empirical setting. In particular, we have the following results on how the buyer will select among the suppliers.

Proposition 3. Consider the case with two suppliers with known types $\theta_1 > \theta_2$. Fixing the higher type θ_1 , increasing θ_2

- 1. Increases the optimal value for the buyer
- 2. Increases the equilibrium effort of the higher type
- 3. Decreases the initial continuation value of the higher type

Moreover, when $\theta_1 >> \theta_2$, the higher type is always picked, and when $\theta_1 \approx \theta_2$, the buyer's optimal rule picks one supplier until they produce a high outcome and switches immediately to the other on a low outcome.

This proposition also speaks to the novel *competition effect* in this setting. Given an inside supplier pool, the buyer has an additional competition reason to explore from the outside pool, since having a higher second type makes the higher type exert more effort.

Finally, a note on tractability. Plugging this problem into a constrained optimizer can be computationally taxing for even a reasonably sized grid for value function iteration. However, computational burden can be eased using the following approach. The equations that determine $F(w_1, w_2, \Delta w_1, \Delta w_2)$ are linear in the variables, so it defines a convex polytope in 2^n dimensions. Computationally, it can be cheaper to calculate the extreme points of this polytope, which enables a reasonably efficient algorithm to do the inner optimization (fixing $\Delta w_1, \Delta w_2$) - just sample random weights for the extreme points and take their convex combination using those weights to generate points in the interior of the polytope.

4 Data

I use a comprehensive dataset from a large buyer in the Indian manufacturing space. The buyer acts as an intermediary, contracting with mostly small and medium scale suppliers to deliver a range of products to its customers. The customers are often major construction MNCs, and the orders are for customer-ordered metal parts, such as steel beams and rods. The order records span from 2019 to mid-2022. These records detail the interactions between downstream customers, the buyer, and approximately 200 suppliers who fulfill various orders. Over this period, we observe repeated transactions, offering a rich temporal view of buyer-supplier dynamics.

Each order is composed of multiple parts, with a median of around 50 parts per order, amounting to a total of approximately 94,000 parts across 700 orders. In most cases, parts are individual manufactured items, such as a steel beam. Multiple parts are inspected and delivered in lots, which can be though of as one truckload worth of parts. Beginning end of 2019, the dataset captures detailed part-, lot-, and order-level information. This includes the supplier and customer numbers, contract numbers, and the managers associated with each transaction. We also observe the planned and actual dates for key stages in the production and delivery process, such as raw material acquisition, cutting, welding, polishing, painting, shipment from the supplier's factory, and final arrival at the destination.

The dataset also includes granular information about each part: its description, size, weight, and type, as well as its associated lot number. Quality check outcomes are recorded for each part, with binary indicators (0/1) specifying whether a part meets the technical specifications, along with the identity of the inspector conducting these quality checks. This binary check is what we use as a measure of the quality of that part. In addition to product specifics, the data tracks payments made to suppliers and the total order value, categorized into four broad buckets, and includes information on auxiliary services provided by the buyer, such as credit, cash upfront, or sourcing of primary inputs.

Suppliers in the dataset are geographically dispersed across many states (which we collapse into three self-contained regions) and exhibit varying degrees of specialization and certification. We also observe the contract terms between the buyer and suppliers, which are decided at the beginning of the relationship, without renegotiation ex post.

5 Preliminary empirical analysis

5.1 Summary stats

Since this is a paper about the relationship between the two parties, Figure 1 presents a histogram of the length of relationships, across all suppliers. Most relationships we see consist of a single order, which is consistent with the buyer learning about supplier types. Some relationships last quite long, indicating that the buyer settles on a set of standard suppliers.



Figure 1: A large percent of relationships end after the first order.

Finally, there is substantial heterogeneity in how suppliers fulfil orders, as Figure 2 shows. Supplier performance is averaged over all orders⁴.

⁴This is with a synchronization exercise using the common orders to make the two regimes consistent, but even restricted to just one regime, the histogram is essentially the same.



Figure 2: There is substantial heterogeneity in supplier performance.

5.2 Rewarding high performance

Next, we note that the buyer is responsive to performance. Tables Table 1 and Table 2 show that suppliers that have lifetime better performance get more orders, and higher average past performance leads to more orders today.

	Dependent variable:
	Total lifetime orders
Lifetime performance	3.0^{***}
	(1.0)
Constant	1.4**
	(0.7)
Observations	274
\mathbb{R}^2	0.03
Adjusted \mathbb{R}^2	0.03
Residual Std. Error	$4.0 \; (df = 272)$
F Statistic	9.2^{***} (df = 1; 272)
Note:	*p<0.1; **p<0.05; ***p<0.01

	Dependent variable:	
	Orders per month	Probability of getting an order
	OLS	Multinomial log-linear
	(1)	(2)
Past average performance	0.251***	2.845***
	(0.037)	(0.389)
Constant	-0.159^{***}	-6.925***
	(0.046)	(1.051)
Month FE	Yes	Yes
Observations	3,424	3,424

Table 2: Supplier monthly order statistics based on past average performance

Note:

*p<0.1; **p<0.05; ***p<0.01

PS: The observations are at the Supplier X Month level, recording numbers of order (and whether or not an order was received).

5.3 Evidence of moral hazard

Finding reduced form evidence of moral hazard is trickier, since it is an equilibrium choice object rather than a fundamental parameter (like type). But it is possible to present structure informed reduced form evidence of moral hazard here. For this, recall that the buyer should be more likely to drop a supplier after bad performance rather than good. As a result, suppliers that have a bad first order should strive to make up for it in their second order. This is exactly what Figure 3 demonstrates. The right hand panel in the figure shows that this result is not driven by a reversion to the mean phenomenon, since average performance in fact rises in the second order.

Another way to showcase this is to note that supplier performance is higher whenever the performance of existing competitors are higher, which is what Table 3 shows.

In fact, we can take this logic even further. Suppliers that do well initially are more likely to slack off later, since the optimal incentive scheme will involve some sort of "tenure" system, whereby the buyer incentivizes high effort in the beginning by promising a less future punishment for good performance today, which translates into a greater likelihood of getting orders in the future *even when* future performance is low. Figure 4 illustrates that



Figure 3: Suppliers who do badly on first orders try to make up for it on second orders, and vice versa.

Notes: This plot is for all suppliers that saw at least two orders in our full dataset.

Table 3: Suppliers perform better when competitor pool is better.

Dependent Variable:	Performance of supplier
Variables	
Performance of competitor pool	0.81***
	(0.08)
Fixed-effects	
Month	Yes
Supplier ID	Yes
Fit statistics	
Observations	682
\mathbb{R}^2	0.43
Within \mathbb{R}^2	0.06

Clustered (month, supplier ID) standard-errors in parentheses Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

this trend is observed in the data.



Figure 4: Suppliers who do well initially slack off later, and vice versa. **Notes:** This plot is for all suppliers that saw at least ten orders in the course of their relationship with the buyer. Suppliers in red have an average performance that is more than 0.1 less than their first order performance, and vice verse for those in green. The rest are grey.

6 Structural analysis

The buyer and suppliers interact in the equilibrium of a game, with a specific structure that dictates how the outcomes realize. As is clear from the preliminary empirical analysis, this equilibrium interaction makes it hard to separate the effect of types from the effect of incentivization, which necessitates a structural model that can take into account the game form.

6.1 Taking the theory to data

The theory cannot be taken directly to the data, however. The state space for the value function, which is the promised continuation value for each supplier and the type (or current belief) of each supplier, would be prohibitively large for a pool of suppliers that numbers in the hundreds. To overcome this, I will leverage a key fact and some important simplifications:

- 1. The buyer uses a tiered structure to classify suppliers. At any point in time, a supplier is classified as Gold, Silver, Bronze, or Blacklisted. These tiers are functions of history, but also connote a promise of future interactions. This naturally leads me to use these tiers to fix levels of CVs that the suppliers can be promised.
- 2. I will assume that the buyer tracks the type distribution at the tier level rather than the individual level, with a two step choice process for choosing the supplier for an order. In the first step, the buyer chooses a tier, given the distribution of types within each, and in the second, the buyer chooses from the available suppliers. This is mainly to make choice probabilities consistent across different supplier pool sizes.

With this tiered structure, fixing the promised CV for each tier implies that what is relevant at any point in time is the number of suppliers in that tier. Multiplying the number by the respective promise yields the total promise made to that tier, which can then be used to construct value functions. Thus, the number of suppliers in each tier and the average type of suppliers in each tier function as the state. The value from picking a supplier of type θ of a given tier T can then be written as:

$$V_{T,\theta}(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B}, \theta) = \max_{\Delta w} \max_{N_G^f, N_S^f, N_B^f} u(p(e(w_T, \Delta w), \theta)) + \delta V(N_G^f, N_S^f, N_B^f, H_{\theta_G}^f, H_{\theta_S}^f, H_{\theta_S}^f)$$

where

$$V(N_G, N_S, N_B, \theta_G, \theta_S, \theta_B) = \max_T \{ V_T(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B}) + \varepsilon_T \}$$
$$V_T(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B}) = \max_\theta \{ V_{T,\theta}(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B}, \theta) + \varepsilon_\theta \}$$

I discuss the ε shocks in the next section. These value functions can then be used to calculate the appropriate choice probabilities. This is implemented via a value function iteration algorithm, starting from some initial guess. For a fixed vector of promises per supplier in each tier, we get an upper bound on the number of possible suppliers in each tier since we know the maximum possible value each supplier can get, so N_G, N_S, N_B lie between 0 and this upper bound. For simplicity, I start with just two types in the support of F. Before the iteration, $(N_G^f, N_S^f, N_B^f, H_{\theta_G}^f, H_{\theta_S}^f, H_{\theta_B}^f)$ is obtained for the choice of each tier given $(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B})$ and Δw using an extreme points approach. I get the extreme points of the set and then during the VFI process, for each state I sample from the set to find the maximizer.

6.2 Choice probabilities

There are a lot of factors that go into selecting a supplier for an order, many of which I see, but there are still quite a few that I cannot. For example, at the time of a given order, a supplier might have other external orders pending which make her less likely to be available for the order. This motivates the addition of logit shocks to the current utility from picking a supplier, in order to account for unobserved reasons for picking suppliers.

Adding logit shocks to this model is not straightforward. To see this, note that the principal provides incentives to suppliers via variations in future continuation value subject to some promised value today. In general, starting with some ex ante promised value, it is possible that the principal could find it optimal to condition the ex post value (post realization of the preference shocks) on the preference shocks, in a way that in expectation the ex post values equal the ex ante value. This adds an additional layer of optimization into the principal's problem, which can make it intractable to calculate choice probabilities directly. As an alternative, we restrict the buyer's ex ante and ex post promised values to coincide, circumventing this issue⁵.

⁵Two additional source of randomness that might be important to incorporate (depending on fit of previous model) will be random shocks to availability of suppliers, and random relationship-level shocks. The former can be tempered with capacity information, which we have for a significant subset of the suppliers.

Recall that the buyer makes tier-specific promises and tracks types at the tier level, so the buyer's choice is effectively at the tier level first i.e. which supplier tier to select at any given composition of each tiers, and then which supplier to select from within the tier. We add logit preference shocks to both of these stages. As a result, the choice probability of picking a supplier of tier T directly follows as a product of the probability of choosing tier T, and then choosing the supplier from tier T.

$$P(T) = \frac{V_T(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B})}{\sum_{T'} V_{T'}(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_S}, H_{\theta_B})} \frac{V_{s|T}(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B})}{\sum_{s'} V_{s'|T}(N_G, N_S, N_B, H_{\theta_G}, H_{\theta_S}, H_{\theta_B})}$$

Given the choice probabilities for every order, we can then construct a likelihood function:

$$\mathcal{L}(\theta|X) = \sum_{o \in X} \log(P_o(T_o))$$

where o denotes an order, T_o denotes the tier of the supplier that was chosen for that order, and P_o is the associated CCP of that tier at that time. Note that tiers of the suppliers update dynamically over time, as outcomes realize, so the choice probabilities keep track of whether a supplier's choice was justified given the history of observed outcomes.

6.3 Moment construction

I use the choice probabilities to construct a likelihood function from the observed actions. This process, however, is not straightforward and requires addressing several complexities inherent in the data and modeling assumptions.

At any given time, the states are not fully observed. While we may estimate the number of suppliers the buyer is considering, their tier assignments are unknown during each choice occasion. To address this, I track a probability distribution for each supplier over each tier, derived from the buyer's optimal choices. For example, if the buyer selects supplier A at a specific time, given a distribution over tiers for all suppliers, I update the suppliers' tiers for the next period using the buyer's optimal . This update operates at the tier level, so movements across tiers for suppliers not chosen are probabilistic. Ideally, both the value function and the movements would be at the supplier level rather than the tier level, but this approximation ensures computational feasibility.

Given this probabilistic tracking of states, the choice probability at any time is calculated as the expectation over choice probabilities across all possible states. This approach accounts for uncertainty in state assignments and integrates over the probabilistic structure to yield the overall likelihood of observed actions. To make this computation tractable, I assume a ceiling on the supplier pool size, which simplifies the construction of states from individual supplier tier probabilities and the associated moments.

Having the right supplier pool is crucial to accurately estimate the likelihood of a supplier being chosen. Many suppliers are observed only briefly before disappearing from the buyer's records, which I interpret as the supplier being dropped from active consideration. This interpretation aligns with suppliers' strong preference for regular loading and their aversion to idle capacity due to its shadow cost. When a supplier is observed as dropped, I incorporate the probability of this event as part of the choice probability, ensuring that these dynamics are reflected in the likelihood function.

6.4 Parameters

The estimation process will broadly recover two sides to the structural aspects of the buyersupplier relationship. First, the buyer's preferences: the share of order surplus β that flows to the supplier if the high outcome realizes, and the scale of the logit shocks σ that the buyer faces. Second, the supplier's preferences. The value to the supplier is treated as known, and the cost of effort c(e) required to achieve a high outcome will be estimated. These preferences influence the supplier's decisions and, ultimately, the outcomes observed.

In the process of estimating these parameters, other aspects of the buyer-supplier relationship will be estimated alongside them. The distribution of supplier types will be one. We know that suppliers are assigned to Gold, Silver, Bronze, or Blacklisted tiers based on average past performance. The parameters of the buyer's tier system will also be recovered in the process.

It is important at this stage to clarify that our interest is not in these parameters per se, but what their estimation will enable us to do. These parameters will be key to tease apart the effects of incentivization versus supplier types on order outcomes, and hence get a sense of how the buyer is adding value to this setting. It also sheds light on the optimal allocation rule for the buyer and how far the current system deviates from it. By quantifying the buyer's value to the system under assumptions about what would happen in the absence of relational constraints, we will also be able to speak to the quantitative impact of these constraints.

7 Identification

In standard dynamic models, identification is based on the fact that observing states and actions allows the construction of (CCPs). By inverting these probabilities, value functions can be determined. Assuming a distribution for the preference shocks, and because transition probabilities can be read off the data, fixing the discount factor leads to identification of utilities using the value functions.

In this model, there are some major departures that prevent directly applying identification results. First, there are two sets of agents that make choices in an interactive manner the buyer and the suppliers - instead of one entity making repeated choices. Furthermore, the choices of some of them - the effort that suppliers choose - is not directly observed. Second, the state w is not directly observed. The value w that is promised to a supplier is a function of both how many orders they will get in the future *and* how much effort they will be expected to put into each future order. Third, transition probabilities today are functions of future variation in the state, since they depend on effort, which depends on Δw . However, I show that identification can be approached in an appropriately modified manner.

The data essentially consists of the choices made by the buyer and the outcomes that realize as a result of these choices. Thus the choice probabilities $P(S|h_t)$ and the outcome probabilities $P(y|h_t, s)$ are the observables that will be used for formal identification. These probabilities may need to be constructed based over many independent settings of choices made by the Principal, but they can also be constructed within a single choice environment - if, for example, all suppliers are eventually replaced, the same history vector for active suppliers can repeat over time. The model parameters F_{θ} , p, β , c, u_Z will need to be identified off of the variation in these probabilities.

The first step is to show that the supplier side can be identified under some restrictions. A key insight that makes this possible is the fact that the supplier problems can be separated out, given the observed choice and outcome probabilities. The suppliers value w can be rewritten as if she is chosen every period, precisely because the principal's promise does not condition on the realization of preference shocks. Even with this simplification, we need fairly strong restrictions to be able to say something useful about the supplier side. The reason behind these restrictions can be understood by examining the moments that will allow us to identify the supplier's parameters. Consider the effort FOC for the supplier $-p_2(\theta, e)\Delta w = c'(e)$ - and the outcome equation $P(y|h_t, s) = p(\theta, e)$. With e and θ both being unobserved, fully nonparametric estimation of p and c is a tall order. So I place some minimal restrictions on both, in a manner that is consistent with some ground truths. For

p, the type θ is supposed to represent some intrinsic capabilities of the supplier, with effort e enabling her to push beyond her capabilities. This motivates a functional form where $p = \theta$ when effort is 0, and then an additive effect of effort. For c, I restrict the effect of effort to a power form, capturing both increasingness and convexity, with both properties allowed to vary in sufficient generality to capture different ways that effort can affect costs.

Assumption 1. The outcome probability function $p(\theta, e)$ is strictly monotonic both arguments and $p(\theta, e) = \theta + \kappa q(\theta)e$ for some non-homogenous function q.

Assumption 2. The marginal cost function is log separable (e, X) i.e. $c'(e; X) = \gamma d(X)\tilde{c}(e)$ such that d is non-homogenous and bijective on $\mathbb{R}^+ \to \mathbb{R}^+$, and \tilde{c}^{-1} exists, is log separable, and is non-homogenous. Further, d satisfies an index restriction i.e. $\log(d(X)) = X_1 + \tilde{d}(X_{-1})$. The cost shocks X have full support on \mathbb{R}^+ .

Assumption 3. The type distribution F has full support on [0,1].

Proposition 4. Assume Assumption 1, Assumption 2, Assumption 3. Then κ , q, γ , d are identified, and hence p and c are identified. Moreover, $w(h_t)$ and F_{θ} are also identified.

Proof. Fix a t and T >> t. Select the set of histories H^0 such that $P_S(h_{t'}|h_t, 1) = P_S(h_{t'}|h_t, 0)$ for all t < t' < T i.e. the choice probability tree for a supplier is the same regardless of the outcome at t. For these histories, e = 0, so $\theta_S = P(y = 1|h_t, S)$. Denote these suppliers as $\Theta(S)$. Since F is full support, this is a non empty set, since there is a high enough type $\underline{\theta} < 1$ such that $(\underline{\theta}, 1) \subset \Theta(S)$.

Consider the set of histories H^1 where $P(s|h_t) > 0$ for some $s \in \Theta(S)$ i.e. there is a supplier in $\Theta(S)$ who could be selected. For these histories, take the supplier's FOC

$$p_2(\theta, e)\Delta w(h_t) = c'(e; X)$$

and substitute it into the outcome equation

$$P(y = 1) = \theta + \kappa q(\theta)e$$

$$\implies P(y = 1) - \theta = \kappa q(\theta)c'^{-1}(q(\theta)\Delta w(h_t); X)$$

$$\implies \log(P(y = 1) - \theta) = \log(\kappa q(\theta)) + \log(\tilde{c}'^{-1}(\frac{\kappa q(\theta)\Delta w(h_t)}{\gamma d(X)}))$$

Thus we get

$$\log(P(y=1) - \theta) = \log(\kappa) + \log(\tilde{c}'^{-1}(\kappa)) - \log(\tilde{c}'^{-1}(\gamma)) + \log(q(\theta)) + \log(\tilde{c}'^{-1}(q(\theta)))$$
(4)
$$- \log(\tilde{c}'^{-1}(X_1)) - \log(\tilde{c}'^{-1}(d(X_{-1}))) + \log(\tilde{c}'^{-1}(\Delta w(h_t))))$$

Fixing a supplier in $\Theta(S)$ and varying only X_1

$$\log(P(y=1|X_1') - \theta) - \log(P(y=1|X_1) - \theta) = -\log(\frac{\tilde{c}'^{-1}(X_1')}{\tilde{c}'^{-1}(X_1)})$$
(5)

Note that the LHS is known for the set of histories H^1 . This allows us to estimate \tilde{c} , since it is non-homogenous and invertible. Next, fixing Δw ,

$$\log(P(y=1|X'_{-i}) - \theta) - \log(P(y=1|X_{-i}) - \theta) = -\frac{1}{\rho}\log(\frac{\tilde{c}'^{-1}(d(X'_{-1}))}{\tilde{c}'^{-1}(d(X_{-1}))})$$

Since d is non-homogenous, the log ratio has sufficient variation, and hence, d is identified. Fixing the shock X and fixing Δw (which can be within order, or across orders where the choice probability tree is the same), for $\theta_1, \theta_2 \in \Theta(S)$

$$\log(P(y=1|\theta_1) - \theta_1) - \log(P(y=1|\theta_2) - \theta_2) = -(1+\frac{1}{\rho})\log(\frac{q(\theta_1)}{q(\theta_2)})$$

Since q is non-homogenous, the log ratio has sufficient variation, so at this step we can q within $\Theta(S)$.

Finally, for the scale parameters κ and γ , going back to the FOC

$$\kappa q(\theta) \Delta w = \gamma d(X) e^{\rho} \tag{6}$$

$$\Longrightarrow \Delta w = \frac{d(X)e^{\rho}}{\frac{\kappa}{\gamma}q(\theta)\Delta w} \tag{7}$$

We can recover Δw as a function of $\frac{\kappa}{\gamma}$ using the following recursive method. First, fix any $h_t \in H^1$. Then, consider a $\theta \in \Theta(S)$. Let T >> t. For all t' > T, denote $\hat{w}_{t'}(h_{t'})$ as the expected value without accounting for the cost of effort i.e. defined recursively as

$$\hat{w}_{t'}(s, h_{t'}, o) = b_o \mathbf{1}_{\{P(s|h_{t'}, o) > 0\}} + \delta P(y = 1|s, h_{t'}, o) \hat{w}_{t'+1}(s, (h_{t'}, 1), o) + \delta(1 - P(y = 1|s, h_{t'}, o)) \hat{w}_{t'+1}(s, (h_{t'}, 0), o)$$

Thus, for any h_T , we can define $\Delta w_T(h_T) = \Delta \hat{w}_T(h_T) + \epsilon_{h_T}$ for some error ϵ_{h_T} . Approximating $\Delta w_T(h_T) = \hat{w}_T(h_T)$, we can derive $e_{t'-1}(\frac{\kappa}{\gamma})$ using 6. For t'' < T this allows us to

recursively get

$$\begin{split} \hat{w}_{t''}(s, h_{t''}, o) = &b_o \mathbf{1}_{\{P(s|h_{t''}, o) > 0\}} - c(e_{t''}(\frac{\kappa}{\gamma})) + \delta P(y = 1|s, h_{t''}, o) \hat{w}_{t''+1}(s, (h_{t''}, 1), o) \\ &+ \delta (1 - P(y = 1|s, h_{t''}, o)) \hat{w}_{t''+1}(s, (h_{t''}, 0), o) \end{split}$$

as a function of $\frac{\kappa}{\gamma}$, and hence get $e_{t''-1}$, with error that only depends on the no cost assumption made T onwards. Recursively applying this until t, we recover $\Delta w(\frac{\kappa}{\gamma})$ within the bounds of an error that vanishes as $T \to \infty$. Plugging this back into 4, we can identify κ and γ separately using the two resultant equations - utilizing the variation in the LHS as Δw varies, and using the constant in that equation.

To identify q outside of $\Theta(S)$, consider a supplier's FOC at the null history i.e. when they are first picked:

$$\kappa q(\theta) \Delta w_0 = \gamma d(X) c(e; X)$$

Note that at the null history, every supplier necessarily has the same Δw_0 . Thus, we know Δw_0 from the previous step, since we know it for suppliers in $\Theta(S)$. As a result, we can invert the equation to write

$$e_0^*(\theta) = \tilde{c}^{-1} \left[\frac{\kappa q(\theta) \Delta w_0}{\gamma d(X)} \right]$$

as the effort they will choose at the null history given their type. Plugging this back into the outcome equation, we get

$$P(y=1) = \theta + Aq(\theta)\tilde{c}^{-1}(q(\theta))$$

where A is a known number that depends on X. Fix a supplier. Since the X shocks are full support, there will be observations such that A is close to 0, which allows recovery of θ . Then, fixing any A > 0, we can recover $q(\theta)$. Doing this for different suppliers recovers q. The distribution $F(\theta)$ is also recovered in the process since θ is recovered for every supplier. Using the recovered q, we can now identify w at all histories for all types.

For the buyer's preferences, we will assume that the buyer also sees an estimate of $P(y|s, h_t)$ before t + 1, as happens in our data. Although this assumption is not necessary to get the net utility from choices, it enables us to separately identify the benefit from a high outcome and the cost from the low outcome, which is important for the counterfactuals that we want to conduct.

Proposition 5. Suppose p, c, w are identified. Then $u_s(w)$ is identified if δ is fixed.

Proof. Since w is identified, the transition probabilities F(w'|w) can be fixed by inspection. Choice probabilities P(w) are also known since w and $P(h_t)$ are known. Given logit shocks, we can use the standard CCP inversion method to obtain the relative conditional choice functions $V_s(w) - V_s(w)$ for any supplier s, relative to some choice S (which can be drawing from the outside pool). Fixing δ , this allows $u_s(w) - u_s(w)$ to be identified.

Next, taking any two suppliers s_1, s_2 in $\Theta(S)$ such that $\theta_{s_1} \neq \theta_{s_2}$, we know $u_{s_1}(w) - u_S(w)$ and $u_{s_2}(w) - u_S(w)$. We also know $P(y = 1|s_1, w)$ and $P(y = 1|s_2, w)$. Using the fact that

$$u_s(w) = P(y = 1|s, w)b_z + (1 - P(y = 1|s, w))c_z$$

we can write

$$u_{s_1}(w) - u_{s_2}(w) = [P(y=1|s_1, w) - P(y=1|s_2, w)](b_z - c_z)$$

Thus $(b_z - c_z)$ is identified.

Finally, note that we observe variation in order value to the supplier. Taking two order values b_1 and b_2 , we know that $\frac{\beta}{1-\beta}b_1 + \phi$ and $\frac{\beta}{1-\beta}b_2 + \phi$ are identified. Thus, β and ϕ are identified.

Turning to the specific data that this paper works with, I delve deeper into the sources of identification. First, the buyer's cost of a low outcome is identified by the extent to which the buyer punishes a higher-type supplier for poor performance. To illustrate, if the buyer faces a high cost from a low outcome and their best supplier performs poorly on an order, the buyer may hesitate to assign the next order to a lower-type supplier, even though incentivization requires it. Next, the split of surplus between the buyer and suppliers is identified by the buyer's preference for assigning higher-value orders to higher-type suppliers. For example, if the buyer's share of the surplus is large, the relative value of higher-value orders increases more significantly, doubling the marginal benefit of giving high-value orders to higher-type suppliers. The cost of effort is identified by how order performance responds to shocks in the cost of effort, such as changes in steel or labor prices. If larger shocks result in a sharper decline in performance, we infer that the function governing effort cost is larger. The type distribution is identified by the outcome distribution under the buyer's incentivization scheme. Since the probability of success is a function of both type and effort, and effort is known under the scheme, the type distribution can be inferred. Finally, the cutoffs are identified by how the buyer selects suppliers based on observed performance.

8 Estimation

Given that the requirement that choice probabilities and outcome probabilities be known for every possible history is onerous, I work with some assumptions to make the estimation tenable. As we saw in the previous section, the buyer chooses a tiered rule, which allows me to restrict the continuation value w that suppliers can be assigned to. I also fix $\kappa = 1$, $q(\theta) = 1$, $c(e) = e^{\rho}$, and $d(X) = X_1 = \frac{1}{p_{\text{steel}}}$ for the first pass, which reduces the estimation burden.

With these restrictions, I implement the estimation in the following manner:

1. Cost parameters initial choice: (Imperfect version of Step 3) I set $\rho = 1$ and run a regression of the following equation at the lot level

$$P(y=1) = \alpha_p \left[\frac{E[\Delta w]}{p_{\text{steel}}^{-1}}\right]^{0.5}$$

where I use an empirical estimate of the average Δw in the sample. This recovers $\gamma = \frac{2}{\alpha_p}$.

- 2. Given the cost parameters, I solve for the value function and then run GMM to select the bargaining parameter β and logit shock scale σ that rationalizes choice probabilities at each history.
- 3. Cost parameters update: Given the estimated θ for each supplier and Δw at each history from the previous step, I then run a lot-level regression derived from Equation 5

$$\log(P(y=1) - \theta) = \alpha_{\rho} \log(p_{\text{steel}}^{-1}) + \alpha_{o}$$

which recovers $\rho = \frac{1}{\alpha_{\rho}}$. This regression incorporates order level fixed effects to ensure that the residuals are orthogonal to the cost shocks, since Δw , which is the only possible confounder as can be seen in Equation 4, is at the order level. Then I estimate γ by matching the FOCs at the order level:

$$P(y=1) - \theta = \alpha_{\gamma} \left[\frac{\Delta w}{p_{\text{steel}}^{-1}} \right]^{\frac{1}{\rho}}$$

which recovers $\gamma = \frac{1}{\alpha_{\rho}^{\rho}}$. I perform this regression with supplier level fixed effects in order to capture unobserved relationship specific variations in the FOCs.

Steps 2 and 3 can be re-run until a desired level of convergence. The cost parameter regressions are displayed in Table 4 and Table 5.

Dependent Variable: Model:	$\log(\hat{P}(Y) - \hat{\theta}) $ (1)
$\begin{array}{c} Variables\\ log(p_{\rm steel}^{-1}) \end{array}$	0.42^{*} (0.23)
<i>Fixed-effects</i> Order number	Yes
$\begin{array}{c} Fit \ statistics\\ Observations\\ R^2\\ Within \ R^2 \end{array}$	2,680 0.382 0.002

Table 4: Regression for cost exponent

Clustered (at the order level) standard-errors in parentheses $\hat{P}(Y)$ calculated at the lot level.

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

m 1 '	1	-	D	•	C		1
liah		h.	- 8	orroggion	tor	COST	CCO D
LaD.	IC.	υ.	т (Cgr Coolon	IOI	COSU	scare
				()			

Dependent Variable: Model:	$\hat{P}(Y)$ - $\hat{ heta}$ (1)
$\begin{array}{c} Variables\\ (\hat{\Delta w} \ / p_{\text{steel}}^{-1})^{(1/\hat{\rho})} \end{array}$	$8.52e-5^{*}$ (4.2e-5)
$\begin{array}{c} Fit \ statistics\\ Observations\\ R^2\\ Within \ R^2 \end{array}$	$336 \\ 0.016 \\ 0.004$

Clustered (at the supplier level) standard-errors in parentheses. $\hat{P}(Y)$ calculated at the order level.

Signif. Codes: ***: 0.01, **: 0.05, *: 0.1

The value function iteration in Step 2 is run as detailed in the previous section. This gives me a vector of conditional choice probabilities $p_{i,\theta}(n_G, n_S, \theta_G, \theta_S)$ in two stages - first, the probability of choosing different type θ , fixing the group *i* that is chosen, and then with the implied group level value function, the probability of choosing the groups $i \in \{G, S, O\}$ where O is the outside option of drawing from the external pool. This gives me $3 \times |\Theta| - 1$ moments that I can potentially construct, based on which group and which type was observed selected at each observation. There is, however, an unusual complication. Since there is a distribution over possible states today, depending on the probability of high outcomes yesterday, each supplier has a distribution over possible groups she could be in (or having been dropped completely). Thus, the moment conditions must be averaged over this distribution of states for each observation. Moreover, I also add a moment condition for the probability with which a supplier that is selected today has already been dropped. This probability should be zero for the correct parameter.

I implement a two step GMM with these moments. Since the routine expects an unbounded parameters space, I use the transformation:

$$\beta_{\text{norm}} = G(\beta_{\text{unbounded}}) = \frac{\beta_{\text{unbounded}}^2}{1 + \beta_{\text{unbounded}}^2}$$

Then, to calculate standard errors, I run a Bayesian bootstrap procedure.

The full set of results are in Table 6.

Category	Parameter	Method	Estimate
Buyon side	Bargaining weight β	GMM	0.77
Duyer side			(0.4)
	Shock scale σ	GMM	33.78
			(25.4)
Sollon aida	Cost function exponent ρ	OLS	2.38
Seller side			(1.3)
	Cost function scale γ (1000 INR)	OLS	2.6
			(1.23)
Common	Discount Rate δ	Exogenous	0.95

Table 6: Parameter Estima

9 Counterfactuals

A major motivation for performing the estimation is to enable counterfactual analyses that address key dimensions of this setting. One crucial aspect is the value of intermediation. The buyer, by fostering competition among suppliers, improves outcomes by increasing equilibrium effort. Bilateral relational contracts, in contrast, would feature weaker incentives and consequently lower effort. Estimation results will quantify the value that the buyer adds through this mechanism.

I present some results in Figure 5 and Figure 6 that speak to this aspect. This is incomplete however, since a full counterfactual simulation remains to be run.



Figure 5: The buyer selects a high performing long term pool.

Another important question concerns the sufficiency of simple rules. With estimates of the underlying parameters, we can evaluate how much the buyer could improve outcomes by transitioning from their current simple rule to an optimal rule. This comparison will highlight the trade-offs between simplicity and optimality in rule design.

Finally, the analysis will address the surplus loss arising from contracting frictions. In this setting, incentives depend on surplus-destroying punishments along the equilibrium path due to these frictions. By running the optimal mechanism without such frictions, we can measure the extent of surplus loss and better understand the cost of these constraints.

References

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



Figure 6: Incentivization improves performance, with a more marked improvement for lower types.

- Clive Bull. Optimal enforcement of performance in the absence of binding contracts. *Econometrica: Journal of the Econometric Society*, pages 1443–1469, 1987.
- Drew Fudenberg, David Levine, and Eric Maskin. The folk theorem in repeated games with imperfect public information. *Econometrica: Journal of the Econometric Society*, pages 997–1039, 1994.
- V Joseph Hotz and Robert A Miller. Conditional choice probabilities and the estimation of dynamic models. *The Review of Economic Studies*, 60(3):497–529, 1993.
- Jonathan Levin. Relational incentive contracts. American Economic Review, 93(3):835–857, 2003.
- Rocco Macchiavello and Ameet Morjaria. The value of relationships: evidence from a supply shock to kenyan rose exports. *American Economic Review*, 105(9):2911–45, 2015.
- Rocco Macchiavello and Ameet Morjaria. Competition and relational contracts in the rwanda coffee chain. *The Quarterly Journal of Economics*, 136(2):1089–1143, 2021.
- W Bentley MacLeod and James M Malcomson. Implicit contracts, incentive compatibility,

and involuntary unemployment. *Econometrica: Journal of the Econometric Society*, pages 447–480, 1989.

- Martin Pesendorfer and Philipp Schmidt-Dengler. Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies*, 75(3):901–928, 2008.
- John Rust. Optimal replacement of gmc bus engines: An empirical model of harold zurcher. Econometrica: Journal of the Econometric Society, pages 999–1033, 1987.
- Stephen E Spear and Sanjay Srivastava. On repeated moral hazard with discounting. *The Review of Economic Studies*, 54(4):599–617, 1987.