Automotive Procurement Under Opaque Prices: Theory with Evidence from the BMW Supply Chain

Danko Turcic[1], Panos Markou[2], Panos Kouvelis[3], and Daniel Corsten[4]

[1]*A. Gary Anderson Graduate School of Management, University of California Riverside, Riverside, CA 92507, USA*

[2]*Darden School of Business, University of Virginia, Charlottesville, VA 22903, USA*

[3]*Olin Business School, Washington University in St. Louis, St. Louis, MO 63130, USA*

[4]*IE Business School, IE University, Madrid, 38006, Spain*

January 17, 2023

Abstract

Several features of automotive procurement distinguish it from the prototypical supply chain in the academic literature: pass-through pricing that reimburses suppliers for raw material costs, market frictions that prohibit cost transparency and imbue suppliers with pricing power, and contractual commitments that span multiple production periods. In this context, we formalize a procurement model by considering an automaker that buys components from an upstream supplier to assemble cars over several production periods in an environment where period demands and raw material costs are both stochastic. Our paper clarifies how information asymmetry, and market factors that amplify or weaken this asymmetry, affect the firms' procurement protocol preferences. Then, using proprietary contract and supplier data from BMW, we empirically validate this model and show that it reflects BMW's reality: the factors that should theoretically go into automotive procurement decisions do so. Our analysis also reveals that existing contracting protocols in this context are not optimal for procurement under asymmetric information, and so we propose an alternative contracting method. We calibrate our model and estimate an automaker's performance improvement from this optimal contract over the status quo.

1 Introduction

Several features of automotive procurement distinguish it from the prototypical procurement supply chains known from the academic literature (see, e.g., de Kok and Graves, 2003). The BMW supply chain is a case in point. First, procurement protocols are built around the fact that raw material costs are stochastic and represent a substantial portion of the total production costs of components and subassemblies purchased from auto part suppliers. For perspective, at the end of 2021, BMW reported a total raw material exposure of €6.9 billion (BMW Group, 2022). Under accepted industry norms, BMW's supply agreements pay Tier 1 suppliers' reported raw material costs, with an added transfer for R&D and tooling costs, as suppliers are expected to make significant recurrent investments into R&D and tooling. In return, automakers must make quantity commitments, requiring them to purchase some components over a specified period. Automakers such as BMW enter into supply agreements either via a fixed price (FP) or a floating price ("material-plus-surcharge"; MPS) protocol. The former (latter) pays the supplier a fixed (floating) unit cost for all raw materials used to produce a particular component or subassembly. Moreover, the automaker, not the supplier, chooses the procurement protocol.

Second, dozens of different types and grades of raw materials go into the production of an automobile, and so there are considerable concerns about information asymmetry regarding raw material costs. On the one hand, cars consist of commodities such as primary (pure) aluminum, copper, and lead. These commodities actively trade on international exchanges such as the London Metal Exchange (LME). Internationally accepted price benchmarks exist, and so automakers benefit from symmetric information by being able to verify a supplier's costs. Data from BMW suggest that about 20 percent of an average vehicle's weight consists of exchange-traded raw materials. (In this environment, Swinney and Netessine, 2009; Kouvelis et al., 2018 guide optimal procurement.) On the other hand, this also means that about 80 percent of an average vehicle's weight is from specialty materials (e.g., aluminum alloys, automotive-grade steels, and lithium) that trade in opaque, disaggregated markets. Suppliers often purchase such materials from Tier-2 suppliers for privately negotiated prices. The problem of asymmetric cost information arises because suppliers and supplier trade groups are the ultimate arbiters of this environment and exercise tight control over the flow of cost information. Suppliers' costs are, therefore, in such cases, unverifiable.

Savvy suppliers can (and do) exploit information asymmetries to extract rents. Figure 1 plots price movements from two aluminum indexes: primary aluminum traded on the LME and a specialized aluminum alloy used in automobile production with quoted prices from a trade group. In early 2008 and mid-2009, price increases in both primary aluminum and the alloy moved in lockstep: the cost of primary aluminum on the LME increased, and the alloy price quickly followed suit. However, between mid and late 2008, the LME price dropped, yet the alloy price was slow to react. Suppliers that produced components with this embedded alloy benefited by extracting information rents. Therefore, one of the main difficulties automakers such as BMW face is ex-ante raw material cost screening rather than matching supply with demand and supplier distress, which were the primary concerns in the earlier automotive procurement literature (see, e.g., Swinney and Netessine, 2009; Babich, 2010).

Against this backdrop, we answer the following questions. First, given the current automotive supply

Figure 1: LME and Trade Group Aluminum Prices

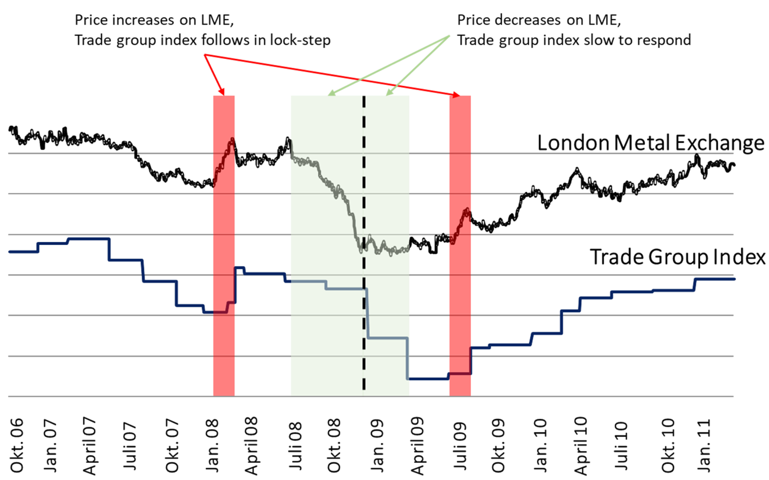

Source: BMW Internal Document

chain environment, which linear contract (MPS or FP) is better for purchasing components produced from raw materials characterized by symmetric vs. asymmetric price information? Second, how is this contract decision affected by various market factors? We answer these questions by following a top-down approach, developing a dynamic procurement model, deriving propositions and hypotheses from the model, and ultimately validating this model empirically using proprietary contract data from BMW. In developing our model, we collaborated closely with managers of the Raw Material Management and Economics Departments at BMW, who helped us capture the critical structural features and market factors in this context. Thus, although our model is stylized, it encompasses the economically relevant elements of BMW's (and other European automakers') procurement practices. In the course of our analysis, we also show that the linear FP and MPS contracts adopted by the industry are not the ones best suited to the screening task highlighted above. Therefore, we conclude our analysis by answering the question: what is the optimal contract for this environment?

Our results can be summarized as follows. When information about supplier costs is symmetric, automakers should offer the floating MPS contract. In such situations, the MPS contract maximizes the total trading surplus and allows the automaker to first-degree price-discriminate the supplier, something the FP contract cannot do. But, when raw materials are not exchange-traded, the automaker's preferred procurement protocol is contingent upon three additional market factors. These factors are the price elasticity of consumer demand, the supplier's R&D and tooling costs, and raw material cost variability. Our model generally reveals that when suppliers are better positioned to distort prices and extract rents, automakers should offer FP contracts.

Do these prescriptions reflect an automaker's reality? We empirically assess the validity of our theoretical model by leveraging a proprietary data set from BMW. The data set offers detailed contract data between the automaker and its hundreds of Tier 1 suppliers delivering parts and components to over 30 manufacturing and assembly plants across multiple continents and countries. We show that BMW tends to tailor its contract offerings depending on situations where suppliers are better positioned to distort prices and extract rents. In line with our model's findings, information asymmetry, price elasticity of demand, supplier R&D and tooling costs, and raw material cost variability are key determinants in BMW's contract selection. Overall, the data support the notion that our model is capturing an automaker's reality and that market factors which *theoretically* inform contract choices in fact *do so* in practice.

Up to this point, our analysis theorizes how certain factors influence the choice between MPS or FP contracts, and it empirically documents that these factors are taken into consideration by BMW. However, we also uncover that neither currently-implemented contractual arrangement can maximize the trading surplus under asymmetric information. This is because the two contracts cannot screen cost information. The question then becomes: what contract can overcome the shortcomings of the MPS and the FP in the case of asymmetric information? Unfortunately, even the existing second-best procurement contracts identified by the literature are unsuitable for the current context. They are static (one-shot) and primarily designed to screen demand information from the downstream buyer in the absence of pre-existing inventory commitments. By contrast, the optimal automotive screening contract must be dynamic, allow pre-existing inventory commitments, and screen cost information from an upstream supplier.

Therefore, we supplement our analysis of the current contracting environment by deriving a new, optimal contract. We show that the second-best contract is a modified MPS contract in both its ex-ante and ex-post terms. Ex-ante (before production begins), the automaker commits to purchasing some fixed quantity of components for a fixed price in each production period. Ex-post (at the beginning of each production/decision period), the automaker purchases additional quantity according to a pre-specified price-quantity schedule with an optimized upper bound on the reported raw material cost. The automaker only purchases the fixed quantity if the supplier reports a raw material cost that exceeds an upper bound. On the other hand, if the supplier reports a raw material cost less than the upper bound, the automaker purchases the ex-ante quantity for the ex-ante price, *plus* she buys additional quantity according to the pre-specified schedule.

Finally, we calibrate the model using our proprietary data and publicly available information. We estimate that the second-best contract could give the automaker a profit boost of 1%–6% over the existing industry contracts, depending on the situation. It is particularly effective when the automaker purchases components for price-elastic models (for example, entry-level models such as the BMW 3-series) produced from raw materials whose prices significantly fluctuate. In such cases, our numerical estimations suggest that an automaker's profit gain could exceed 6% compared to the existing contracts.

Our paper is organized as follows. We review related research in §2. In §3, we generate an abridged model of the current procurement process at BMW. Given the automotive industry norms, however, the model is general enough to apply to other (European) automakers. We solve the model and identify several propositions in §§4–6. §7 uses a proprietary data set from BMW to validate the model's insights. In §8, we deviate from current practice, develop an optimal contract, and estimate the profit boost the contract could deliver. §9 concludes.

2 Related Research

Our paper is related to three streams of literature. The first stream studies supply chain contracts; the second focuses on why and how firms manage risk; and the third analyzes contracts under asymmetric information.

From the contracting literature, the papers closest to ours are Bajari and Tadelis (2001) and Swinney and Netessine (2009). The setting in Bajari and Tadelis (2001) is the commercial construction industry, which heavily utilizes cost-plus and fixed-price contracts featured in this paper. The leading complication in that industry is contract re-negotiation due to randomly needed design modifications after the start of construction. The authors assume that construction costs without changes are verifiable, while the costs of modifications are not. The primary research question is which linear contract – fixed-price or cost plus – to use, which the authors answer by creating a multi-stage, static model. The setting in Swinney and Netessine (2009) is the automotive industry. The primary research goal is to find an optimal contract and consider the possibility that a supplier may be unable to fulfill their obligations if their production costs become too high. Multiple suppliers in their model have correlated stochastic production costs, and the automobile manufacturer can create contracts that last for one or multiple periods with one of the suppliers. It should be noted that there is no private information involved.

Relative to the above papers, our model is dynamic with multiple sources of uncertainty (prices and demand), inventory constraints, and contractually binding quantity commitments. We look for optimal linear and non-linear contracts in that setting. Finally, unlike the previous studies, we validate our model with data.

Changes in commodity prices or exchange rates frequently lead to fluctuations in profits. A large body of articles in operations and finance studies *how* firms can optimally lessen profit fluctuations by entering into financial transactions that complement their operational payoffs. For example, in Gaur and

Seshadri (2005), Caldentey and Haugh (2009), and Turcic et al. (2015), the financial transactions are trades in options and futures. Others study how to hedge operationally. In Yang et al. (2009) firms use dual sourcing. In Swinney and Netessine (2009); Kouvelis et al. (2018), firms use contracting strategies.

We contribute to this literature by highlighting that firms sourcing commodities regularly face two types of risks. The first is the commodity price risk that firms face in all markets; the second is a reporting risk, which exposes buyers to price distortions if they buy raw materials in opaque markets.

Finally, closely connected is the issue of optimal contracting in opaque markets. First-best outcomes are generally not possible; second-best is the best outcome firms can achieve. A comprehensive review of that literature presented in Shen et al. (2019) reveals that previous articles have derived second-best contracts while narrowly focusing on situations in which the buyer, not the supplier, is privy to private information about demand (e.g., Kostamis and Duenyas, 2011) or the inventory holding cost (e.g., Ha, 2001; Corbett et al., 2004; Kostamis and Duenyas, 2011). All models mentioned in the papers are static and do not allow for contractual quantity pre-commitments. Both are essential features of the automotive supply chains that are baked into the second-best contract we present in this paper.

3 An Automotive Component Procurement Model

In this section, we build an $\bar{n}$-period dynamic automotive procurement model.[1] Insights for building the model came from the Raw Material Management and Economics Department managers at BMW, who also provided us with data for our analysis. Prior academic research (Simchi-Levi et al., 2015; Markou et al., 2019) and industry publications (e.g., Zehender, 2015), however, suggest that many automakers utilize similar procurement protocols. Table 1 provides information on the vital notation and expressions developed over the subsequent sections.

In this model, a risk-neutral automaker (she) supplies $\mathbf{x} = (x_1, x_2, \ldots, x_{\bar{n}})$ cars to the market while facing a market potential, $\mathbf{V} = (V_1, V_2, \ldots, V_{\bar{n}})$, and an inverse demand curve, $g(\mathbf{x}, \mathbf{V})$. Each car that the automaker produces requires a sub-assembly (component) that the automaker purchases from a Tier-1 auto supplier (he) facing a unit raw material cost $\mathbf{C} = (C_1, C_2, \ldots, C_{\bar{n}})$.[2] We assume that $\{(V_n, C_n) : n = 1, 2, \ldots, \bar{n}\}$ follow a stochastic process defined on a filtered probability space $(\Omega, \mathcal{A}, \mathbb{F}, \mathbb{P})$. The process $\mathbf{V}$ has support $[\mathbf{v}^l, \mathbf{v}^h]$, so there is some minimum and maximum demand in each period. Unless stated otherwise, only the *supplier* is privy to the realization, c_n, of the random cost variable, C_n in each period n. However, following standard assumptions from the contracting literature (see, e.g., Bolton et al., 2005), all firms are privy to the probability measure, $\mathbb{P}$.

[1] The number $\bar{n}$ reflects the lifespan of a specific car model.

[2] Multiple suppliers/components add complexity with little contribution to the main idea.

Table 1: Key Notation and Expressions

Expression	Description			
v_n	Market potential in period $n = 1, 2, \ldots, \bar{n}$ †			
$g(x, v_n)$	Inverse demand in period n when the automaker outputs x units			
$r(x, v_n)$	Retail revenues in period n when the automaker outputs x units, i.e., $r(x, v_n) = x\, g(x, v_n)$			
c_n and $\widetilde{c}_n$	Actual raw material cost and supplier reported raw material costs in period n †			
$\mathcal{H}$	Unit holding cost			
$\mathcal{K}_i$	Supplier i's R&D and tooling cost, $i \in \{1, 2, \ldots, I\}$			
		MPS Contract with Symmetric Information	MPS Contract with Asymmetric Information	FP Contract
q	*Ex-ante* quantity commitment in period n	$q^*_{s,n}$	$q^*_{a,n}$	$q^*_{p,n}$
lq and hq	Minimum and maximum purchase quantities period n	$lq^*_{s,n}$ and $hq^*_{s,n}$	$lq^*_{a,n}$ and $hq^*_{a,n}$	$lq^*_{p,n}$ and $hq^*_{p,n}$
y	Supplemental *ex-post* purchase quantity in period n	$y^*_{s,n}$	$y^*_{a,n}$	$y^*_{p,n}$

† Upper case letters denote random variables; lower case letters denote realizations of random variables.

The automaker selects the supplier via an electronic second-price auction in which $I > 1$ suppliers participate. The automaker awards the contract to the lowest-bidding supplier, paying that supplier the second-lowest bidder's bid. Before the auction begins, the automaker announces (1) the contract duration, (2) the market potential for the car, [3] and (3) the supply contract type. The contract type remains unchanged during the contract's duration. Under the *current practice*, the automaker offers each supplier either a *material-plus-surcharge* (MPS) or a *fixed price* (FP) supply contract. Such contracts are prevalent in the European automotive industry and are the contracts behind our data.

Both contracts are two-part tariffs. The *fixed tariff* reimburses the supplier for his R&D and tooling cost in each period n, which is essential in preempting hold-up (Demski and Sappington, 1991). The fixed tariff size is determined in the auction. The *floating tariff* (i.e., unit cost) is a standard pass-through that reimburses the supplier for his raw material costs. If the contract type is MPS, then the period-n floating tariff is what the supplier reports to the automaker as his period-n raw material cost. If the contract type is FP, the period-n per-unit fee is fixed.

Internationally accepted prices of *exchange-traded* raw materials are published by the exchanges themselves, such as the London Metal Exchange and the London Platinum and Palladium Market. Benchmark prices for *non-exchange-traded* materials do not exist, so representative prices are typically published by industry trade groups, such as the Wirtschaftsvereinigung Stahl, or determined by the

[3]Before the auction, BMW, for example, shares a market potential forecast with its suppliers to inform them of the bid's capacity requirements. After production begins, the automaker shares demand forecast updates. Vehicle sales data, reported by vehicle type, and the handful of economic variables driving vehicle sales are available in annual reports and other industry/economic publications.

suppliers themselves (Markou et al., 2019). In practice, contracts also contain provisions restricting suppliers from reporting unrealistic prices. Whether explicitly negotiated or implicitly agreed upon, suppliers are effectively restricted to per-period raw material price updates within some range $[\mathbf{c}_n^l, \mathbf{c}_n^h]$. The range is typically based on some fixed multiple of a reference price.

After the auction and before the start of production, the automaker *must commit* to procurement quantities $\mathbf{q} = (q_1, q_2, \ldots, q_{\bar{n}})$ with the proviso that her actual procurement quantities will be between $l\mathbf{q}$ and $h\mathbf{q}$, where $1 < h < \infty$ and $0 < l < 1$ are positive scalar multipliers. The automaker gets to decide the actual procurement quantity at the beginning of each production period n, after seeing the period n market potential and raw material cost.[4]

4 Procurement of Components Made From Exchange-Traded Raw Materials

In this section, the supplier produces a component manufactured from a raw material that trades on an exchange. So, the raw material price, C_n, is an index published by the exchange and made publicly available. We analyze the MPS protocol in this environment and refer to it as our *benchmark* protocol.

Before production begins (*ex-ante*), the suppliers engage in a second-price auction, which identifies the winning supplier $i \in \{1, 2, \ldots, I\}$ and establishes a vector of fixed transfers, $\boldsymbol{\tau}_s^i = (\tau_{s,1}^i, \tau_{s,2}^i, \ldots)$. (The first subscript, s, indicates that information is symmetric. The second subscript corresponds to the production period number.) Seeing the fixed transfer vector, $\boldsymbol{\tau}_s^i$, the automaker commits to procurement quantities, $\mathbf{q}_s = (q_{s,1},\ q_{s,2},\ \ldots)$, with the proviso that in period n she purchases at least $lq_{s,n}$ and no more than $hq_{s,n}$ units.

After production begins (*ex-post*), the automaker's optimization runs as follows. At the beginning of each production period, n, she sees the market potential and raw material cost realizations, v_n, and c_n. Based on this information, she chooses a level of output between $lq_{s,n}$ and $hq_{s,n}$ units. Then, period-n payoffs are realized.

We begin by deriving the automaker's optimization objective. Then, given the automaker's objective function, we examine her ex-ante and ex-post quantity decisions. Finally, we determine the suppliers' bids in the second-price auction and derive expected payoffs.

[4]In unreported results, our data reflects that automakers respond to market conditions when picking the final output in each production period.

4.1 Optimization Objective

At the beginning of each production period n, the automaker gets to observe the market realizations, v_n and c_n, and her cycle inventory, $z_{s,n}$, which accumulates at the beginning of period n whenever she had more input components on hand than needed in period $(n-1)$. With this information, she must decide how many components to buy, $y_{s,n} \in [0, (h-l)q_{s,n}]$, above the contractually required minimum of $lq_{s,n}$ units and how many cars to produce, $x_{s,n}$. Period n output cannot exceed the number of available components; unused components are held as cycle inventory at a unit holding cost of $\mathcal{H}$ per period.

Let $g : x_n \times v_n \to \mathbb{R}$ and $\nu : x_n \times v_n \to \mathbb{R}$ be the inverse demand curve and the corresponding price elasticity of demand, e.g., see Jehle and Reny (2001, Definition 1.6). (Consistent with Berry et al., 1995; Anderson et al., 1997, we assume $\nu(\cdot) < -1$, i.e., demand is relatively elastic.) The automaker's revenue and marginal revenue curves, $r(x_n, v_n)$ and $r'(x_n, v_n)$, are [5]

$$r(x_n, v_n) = x_n \, g(x_n, v_n), \tag{1a}$$

$$r'(x_n, v_n) = g(x_n, v_n) \left(1 + \frac{1}{\nu(x_n, v_n)}\right). \tag{1b}$$

Following Nasser and Turcic (2019), we impose a few assumptions on the marginal revenue curve. We do not, however, restrict attention to a particular revenue function.

Assumption 1. The marginal revenue curve (1b) is concave, decreasing in quantity, x_n, concave, increasing in market potential, v_n, and supermodular in x_n and v_n.

Assumption 1 implies unique optimal output, which increases with market potential. However, incremental increases in the market potential have a diminishing effect on marginal output. As examples, linear and constant elasticity demand curves with their corresponding revenue and marginal revenue curves satisfy the restrictions of Assumption 1.

In summary, the flow of components in cycle inventory is subject to a conservation constraint (2a) and two logical constraints (2b):

$$z_{s,n+1} = z_{s,n} + lq_{s,n} + y_{s,n} - x_{s,n}, \tag{2a}$$

$$0 \le y_{s,n} \le (h-l)q_{s,n}, \qquad \text{and} \qquad 0 \le x_{s,n} \le z_{s,n} + lq_{s,n} + y_{s,n}. \tag{2b}$$

The automaker's payoff in each period n is

$$\psi(\mathbf{a}_{s,n}, \mathbf{s}_{s,n}) = r(x_{s,n}, v_n) - c_n \, y_{s,n} - \mathcal{H}\left(z_{s,n} + lq_{s,n} + y_{s,n} - x_{s,n}\right)^+ - \tau^i_{s,n},$$

$$\mathbf{a}_{s,n} = (x_{s,n}, y_{s,n}), \ \mathbf{s}_{s,n} = (z_{s,n}, v_n, c_n), \tag{3}$$

[5] See Equation 4.5 in Jehle and Reny (2001).

where $\mathbf{a}_{s,n}$ and $\mathbf{s}_{s,n}$ are the decision and the state vectors. Equation (3) demonstrates the implementation of two-part tariff pricing: the raw material cost, c_n, is the floating tariff, and the R&D and tooling cost, $\tau^i_{s,n}$, is the fixed tariff. Also, notice that although the automaker is acquiring $(lq_{s,n} + y_{s,n})$ components in period n, in Equation (3), she only pays for $y_{s,n}$ units because the cost of the $lq_{s,n}$ units is sunk. (She must buy the $lq_{s,n}$ units no matter what.) The optimization of the procurement and output decisions corresponds to the following dynamic program:

$$\pi(\mathbf{s}_{s,n} \mid \mathbf{q}_s) = \max_{\mathbf{a}_{s,n}} \left[\psi(\mathbf{a}_{s,n}, \mathbf{s}_{s,n}) + \mathbb{E}_n \pi(\mathbf{S}_{s,n+1} \mid \mathbf{q}_s) : x_{s,n} \leq z_{s,n} + lq_{s,n} + y_{s,n},\ y_{s,n} \leq (h-l)\, q_{s,n}\right], \quad (4)$$

where the subscript on the expectation operator denotes filtration. The optimal commitment policy is

$$\pi(\mathbf{s}_{s,0}) = \max_{\mathbf{q}_s} \mathbb{E}_0 \left[\pi(\mathbf{S}_{s,1} \mid \mathbf{q}_s)\right] - \mathbb{E}_0 \left[l\mathbf{q}_s \cdot \mathbf{C}\right]. \quad (5)$$

4.2 JIT Production Policy

The decision rule many automakers implement in the dynamic procurement model we just proposed is just-in-time (JIT) production. If $z_{s,0} = 0$ (cycle inventory is zero when production starts), then JIT production implements the following quantity decisions:

Definition 1 (JIT Production Policy)**.**

$$x^*(q_{s,n}) = lq_{s,n} + y^*_{s,n}(q_{s,n}), \quad (6a)$$

$$y^*_{s,n}(q_{s,n}) = \arg\max_{y_{s,n}} \ \psi(y_{s,n}, q_{s,n}) \quad \text{s.t.} \quad 0 \leq y_{s,n} \leq (h-l)\, q_{s,n}, \quad (6b)$$

$$q^*_{s,n} = \arg\max_{q_{s,n}} \ \mathbb{E}_0 \left[\psi\big(y^*_{s,n}(q_{s,n}), q_{s,n}\big)\right], \quad \forall n = 1, 2, \ldots, \bar{n}. \quad (6c)$$

Veinott (1965) provides conditions under which a class of myopic policies similar to (6) is optimal in settings where costs and prices are constant, and the firm faces no contractual quantity commitments. Our environment, however, does not match that of Veinott (1965). So, we must derive a new set of conditions.

There are two primary reasons for deviating from the JIT policy in our environment. First, because the automaker is contractually required to procure $lq_{s,n}$ units in each period $n = 1, 2, \ldots, \bar{n}$, she may have to obtain components in larger quantities than immediately required in production, resulting in *cycle inventories*. Lemma A.2 asserts that cycle inventories will be zero in our model whenever the initial quantity commitment, $\mathbf{q}_s$, is sufficiently costly, and thereby low enough, relative to the market potential, $\mathbf{V}$.

Second, the automaker may wish to acquire *speculative inventory* in some period $k < n$ to hedge against demand and cost fluctuations in a later period n. The automaker is incentivized to make such a speculative purchase if the period-n's expected marginal profit, evaluated at the optimal feasible output, is positive *when viewed from the period k*. In Lemma A.2, we argue that the evolving sequence of best approximations to the period-n expected marginal payoff conditional on period k's filtration is a martingale (Lemma A.1). Lemma A.2 obtains a lower bound on the holding cost, $\mathcal{H}$, to stop the automaker from adding speculative inventory. The martingale property implies that the automaker is incentivized to commit to the JIT policy before production begins. Once the automaker implements JIT, she will have little access to storage, implying relatively high holding costs.

The conditions of Lemma A.2 are new to the literature and help us understand when using the JIT policy does not reduce profits.

Assumption 2. Henceforth, we will restrict attention to the class of JIT policies. (JIT production policy is also the policy behind our data.)

4.3 Supplier Selection via Auction

So far, we have taken it as given that there is a supplier $i \in \{1, 2, \ldots, I\}$ producing components in quantities that the automaker wants and for which she pays a unit price, C_n, *plus* a fixed transfer, $\tau^i_{s,n}$, in each period n. In this section, the automaker gets to pick the supplier i by letting $I > 1$ firms bid in a second-price auction. The auction outcome also decides the fixed transfer amount.

Each supplier i's costs associated with handling the supply contract are (1) supplier-specific, fixed R&D and tooling cost of $\mathcal{K}_i$, accrued over $\bar{n}$ production periods and (2) unit raw material cost of C_n accrued in each period n.

Lemma 1 (Jehle and Reny, 2001, Theorem 9.3)**.** *Let $\mathcal{K}_1 \leq \mathcal{K}_2 \leq \ldots \leq \mathcal{K}_I$. Then, each supplier $i = 1, 2, \ldots, I$ bids $\mathcal{K}_i/\bar{n}$. Supplier $i = 1$ wins the auction and the equilibrium fixed transfer is $t^*_{s,n} = \mathcal{K}_2/\bar{n}$.*

Using Lemma 1, the automaker's payoff can be written as:

$$\max_{\mathbf{q}} \; \mathbb{E}_0 \left[\sum_{n=1}^{\bar{n}} \max_{y} \; \Big(r(y + lq_n, V_n) - C_n \, (y + lq_n) \;\; \text{s.t. } \; y \leq (h - l) \, q_n \Big) \right] - \mathcal{K}_2. \tag{7}$$

The above program is an aggregate surplus maximization problem with a fixed transfer of $\mathcal{K}_2$, which ensures the winning supplier's participation. Thus, we report that the automaker's output level is the first-best. Also, since the automaker extracts all surplus above the supplier's R&D and tooling cost, the MPS contract allows her to engage in first-degree price discrimination.

5 Procurement of Components Made From Non-Exchange-Traded Raw Materials

Most supply contracts are written on materials that do not trade on an exchange. These materials include aluminum alloys, automotive-grade steels, lithium, and rare earth metals, used in batteries and hybrid drives. Tier 1 suppliers purchase these materials from various smelters, foundries, and other upstream suppliers for privately negotiated prices. The negotiated prices vary substantially across producers because of differences in refining processes, quality grades, regional economics, etc.

Frequently, because there is little other price data available, the prices of non-exchange-traded raw materials in automotive procurement contracts are pegged either to prices published via supplier price charts or to trade group indices (Markou et al., 2019). Individual suppliers publish the former. The latter is published by supplier trade groups and compiled from the supplier price charts. In Germany, one trade group that publishes reference prices is the *Wirtschaftsvereinigung Stahl*.

This system allows suppliers to report something different from their actual costs without negatively impacting their reputations. For automakers, the absence of widely accepted reference prices means that suppliers' raw material costs can neither be readily corroborated nor easily challenged. [6]

This section aims to offer a model of the current contracting practice at BMW in the environments we just described. The contracts we study in this section are not necessarily optimal or the best practice – they are the *current* practice. We address the optimality question later in §8, where we perform counterfactual experiments.

5.1 The Material-Plus-Surcharge Contract

In this section, we re-analyze the benchmark contract of §4 in an environment where the automaker cannot corroborate the supplier's actual raw material cost; the automaker only gets to see the price chart index, say $\widetilde{c}_n$ (e.g., the "Trade Group Index" line in Figure 1). We conjecture that the supplier treats the index as a decision variable, use Nash equilibrium as our solution concept, and re-derive the benchmark contract while assuming that the automaker follows a JIT policy.

In Lemma A.3, included in the Appendix, we argue that the automaker's propensity to hold cycle inventories is no greater than in the symmetric information case of Lemma A.2. However, the automaker may be more incentivized to speculate on future reported raw material costs with the inventory. Anand et al. (2008) were the first to observe that the supplier is incentivized to set lower wholesale prices when the buyer holds speculative inventory and production costs are deterministic. Lemma A.3 effectively

[6] Ramkumar (2018) also illustrate the difficulty with price discovery of non-exchange-traded metals with a story of the LME abandoning the launch of futures contracts for cobalt and lithium because the exchange struggled to secure a reliable source of price data.

extends their argument into our stochastic setting and identifies a holding cost sufficient to stop the speculation. Thus, if the automaker's holding cost exceeds the minimum levels identified in both Lemma A.2 and A.3, the JIT policy is optimal regardless of information symmetry, and Assumption 2 holds.

5.1.1 Supplier's Cost Reporting Decision

Using Equation (6), after seeing the index, $\widetilde{c}_n$, the automaker responds by choosing a top-up order quantity $y^*_{a,n}(\widetilde{c}_n) = \min\{(h-l)q_{a,n}, y\}$, where y is a number of units that equates the automaker's marginal revenue of Equation (1b) with marginal cost:

$$\widetilde{c}_n = g(y + lq_{a,n}, v_n, \gamma)\left(1 + \frac{1}{\nu(y + lq_{a,n}, v_n, \gamma)}\right). \tag{8}$$

Above, notice how we write that the marginal revenue depends on an additional parameter, γ, which affects the automaker's price elasticity of demand. (To illustrate the role of γ, suppose that the automaker faces linear demand. Then, γ would be one over the demand curve slope. Although the parameter γ already existed in the benchmark case of §4, for ease of presentation, we notationally suppressed it.) For further analysis, we assume that as γ increases, demand becomes more elastic, i.e., $\partial\nu(y + lq_{a,n}, v_n, \gamma)/\partial\gamma \leq 0$.

Seeing the raw material input cost, c_n, and correctly anticipating the automaker's quantity decision in (8), the supplier $i \in \{1, 2, \ldots, I\}$ reports an index value, $\widetilde{c}_n$, optimal in

$$\max_{c^l_n \leq \widetilde{c}_n \leq c^h_n} \left(y^*_{a,n}(\widetilde{c}_n) + lq_{a,n}\right)(\widetilde{c}_n - c_n) + \tau^i_{a,n}. \tag{9}$$

After substituting for $\widetilde{c}_n$ from (8) and dropping the term, $\tau^i_{a,n}$, which does not influence his choice in (9), the supplier's reporting problem reduces to that of choosing quantity, y:

$$\widetilde{\Pi}_a(\mathbf{s}_{a,n}) \equiv \left\{\max_{0 \leq y \leq (h-l)q_{a,n}} r(y + lq_{a,n}, v_n, \gamma)\left(1 + \frac{1}{\nu(y + lq_{a,n}, v_n, \gamma)}\right) - (y + lq_{a,n})\, c_n\right\}, \tag{10}$$

revealing how the supplier's reporting decision depends on the elasticity of market demand, $\nu(\cdot)$ and the supplier's actual raw material cost, c_n.

The supplier would then report index value, $\widetilde{c}_n$ equal to

$$\widetilde{c}^*_n(y) = \max\left\{\min\left\{g(y + lq_{a,n}, v_n, \gamma)\left(1 + \frac{1}{\nu(y + lq_{a,n}, v_n, \gamma)}\right), c^h_n\right\}, c^l_n\right\} \tag{11}$$

which he obtains by substituting the quantity decision, y, optimal in (10) into the right side of Equation (8), subject to the pre-determined minimum and maximum allowed index value.

Further exploration of the supplier's choices in (10) and (11) reveal that a sufficient condition for the supplier to increase (reduce) his reported index value when the automaker's demand becomes less (more) elastic is:

Assumption 3. The marginal revenue curve (1b) is increasing in γ and supermodular in Q and γ.

Taken together, (10) and (11) give a model of how the supplier, locked into a supply contract with the automaker, derives the price chart index value when the automaker cannot corroborate the supplier's actual raw material cost, c_n. However, the automaker can use this model to partially infer the value of c_n – up to the bounds on quantity and price in (10) and (11) – and use the inferred price history to estimate the probability measure $\mathbb{P}_c(\cdot)$ (see, e.g., Lai and Ying, 1991). We are familiar with this approach when estimating demand distribution from censored retail sales data in inventory management. In that case, censoring occurs because of out-of-stocks (see, e.g., Musalem et al., 2010).

Yet, because the standard MPS contract links the contract's floating tariff to the supplier's price chart index, the automaker's ability to partly guess the supplier's cost does not affect the pass-through payment. Instead, the defensive tactic that BMW often uses in these markets replaces the MPS contract with the FP contract, which links the floating tariff (unit cost) to the forward price of $\mathbb{E}_0[C_n]$. Because the price is fixed at time 0, it is impossible to manipulate it, as will be seen in §5.2.

5.1.2 Auction

As in the benchmark case of §4, the equilibrium transfer, $\tau^*_{a,n}$, is determined through a second-price auction described in §4.3. For previously explained reasons, supplier $i = 1$ wins, and the rules of the second-price auction assign him a transfer of $\tau^*_{a,n}$ that is given in the following lemma.

Lemma 2. *Assume $\mathcal{K}_1 \leq \mathcal{K}_2 \leq \ldots \leq \mathcal{K}_I$ and let $\widetilde{\Pi}^*_a(\mathbf{S}_{a,n})$ is given by (10). Then, the winning supplier's equilibrium transfer is* $\tau^*_{a,n} = \frac{1}{\bar{n}}\left(\mathcal{K}_2 - \sum_{n=1}^{\bar{n}} \mathbb{E}_0\left[\widetilde{\Pi}^*_a(\mathbf{S}_{a,n})\right]\right)^+$.

Let us compare the transfer of Lemma 2 to that of Lemma 1. The part $\mathcal{K}_2/\bar{n}$ and the reason for its presence are the same: it compensates the supplier for his R&D and tooling costs. The expression $\mathbb{E}_0\left[\widetilde{\Pi}^*_a(\mathbf{S}_{a,n})\right]$ is information rent, which the supplier earns by inflating his raw material cost as per Equation (10).

Seeing that the supplier does not get to keep the information rent, why does he inflate his raw material cost? It is an issue of time consistency (Kydland and Prescott, 1977): without an incentive to accurately report raw material cost after production begins, the MPS contract used in the automotive industry today makes it difficult for the supplier to commit to truthful reporting at the time of the second-price auction. The "$(\ldots)^+$" operator reflects that paying the automaker to enter into a supply agreement with him is not individually rational for the supplier because the automaker could, if she wanted to, extract the supplier's surplus by strategically under-ordering. Using Lemma 2, the automaker's expected payoff is:

$$\max_{\mathbf{q}} \mathbb{E}_0 \left[\sum_{n=1}^{\bar{n}} \max_{y} \Big(r(y + lq_n, V_n) - \widetilde{c}_n^*(C_n)\,(y + lq_n) \;\; \text{s.t.} \;\; y \leq (h-l)\, q_n \Big) \right] - \left(\mathcal{K}_2 - \sum_{n=1}^{\bar{n}} \mathbb{E}_0 \left[\widetilde{\Pi}_a^*(\mathbf{S}_{a,n}) \right] \right)^+, \quad (12)$$

where $\widetilde{c}_n^*(C_n)$ is obtained by substituting the value of y optimal in (10) into the constraint of (8).

Because $\widetilde{c}_n(C_n) \geq C_n$, the automaker's optimal output is distorted downward compared to the baseline case of §4 and MPS contract no longer maximizes the total surplus.

5.2 The Fixed Price Contract

The FP contract is an MPS contract with unit price pegged to the forward price of $\mathbb{E}_0[C_n]$ and an equilibrium transfer of $\tau_{p,n}^* = \mathcal{K}_2/\bar{n}$. (Since the FP contract is a special case of the MPS contract, Lemma A.2 also applies to the FP contract. Assuming that the conditions of Lemma A.2 are satisfied, we persist in our focus on the JIT policy.) The automaker's expected payoff with the FP contract is

$$\max_{\mathbf{q}} \mathbb{E}_0 \left[\sum_{n=1}^{\bar{n}} \max_{y} \Big(r(y + lq_{p,n}, V_n) - \mathbb{E}_0[C_n]\,(y + lq_{p,n}) \;\; \text{s.t.} \;\; y \leq (h-l)\, q_{p,n} \Big) \right] - \mathcal{K}_2. \quad (13)$$

By comparing (7) and (13), we see that the automaker can match the FP contract's expected payoff by using the MPS contract and constraining output to the FP contract's quantities. Because adding such a constraint reduces the automaker's payoff, we report that the FP contract fails to maximize the total surplus. However, the FP contract's redeeming quality is that the supplier cannot distort its fixed unit price, whether or not the raw material is exchange-traded.

6 Procurement Contract Choice Predictions

This section formally represents the automaker's preferences by the binary relation $\succsim$. If $MPS \succsim FP$, we say that "the MPS contract is at least as good for the automaker as the FP contract."

Proposition 1. *Let the raw material used by the supplier be exchange-traded. Then, $MPS \succsim FP$.*

When the raw material is *exchange-traded*, prices are public information, in which case the MPS contract dominates the FP contract because the former allows the automaker to first-degree price discriminate the supplier – something the latter does not allow. So, provided that the raw material is exchange-traded, a change in market conditions does not flip her procurement protocol preference.

In contrast, when the raw material is *not exchange-traded*, prices are private, and the automaker's procurement protocol preference depends on additional market factors. We formally describe these market factors by considering markets 1 and 2. In each market $j = 1, 2$, the elasticity parameter introduced in §5.1.1 is γ_j. The R&D and tooling cost of the second-lowest-cost supplier is $\mathcal{K}_{2,j}$. The reported raw material cost, $\widetilde{C}_{n,j}$, has support $[c^l_{n,j}, c^h_{n,j}]$ and variance $Var(\widetilde{C}_{n,j})$.

Proposition 2 (Price Elasticity of Demand)**.** *Let the raw material used by the supplier be non-exchange-traded, and Assumption 3 hold. There exists a threshold $\hat{\gamma}$ such that if $\gamma_1 \leq \hat{\gamma} \leq \gamma_2$, then $FP_1 \gtrsim MPS_1$ and $MPS_2 \gtrsim FP_2$.*

Proposition 2 orders the prospects in terms of price elasticity of demand. When demand is relatively elastic (i.e., when γ is sufficiently high – see §5.1.1), the percentage change in quantity demanded is more significant than in price. Hence, when the reported raw material cost is raised, the total revenue falls, and vice versa. So, a supplier selling to an automaker facing elastic demand has the incentive to report raw material costs, $\widetilde{c}_n$, close enough to the actual costs, c_n, and the payoffs of the MPS contracts of §§4 and 5.1 converge. In such circumstances, the automaker prefers the MPS contract §5.1 to the FP contract for reasons that we already described in our discussion of Proposition 1. (The reverse is true when demand is sufficiently inelastic.)

The car model effectively dictates the price elasticity of demand and can change over time for competitive reasons. Berry et al. (1995), for example, report elasticity values that vary widely, depending on the make and model. Economy cars tend to have more elastic demand than their luxury counterparts. So, we predict that a higher proportion of non-exchange-traded raw materials used in the production of the lower-priced (higher-priced) models, e.g., the 3-series (7-series), are procured via MPS (FP) contracts.

Proposition 3 (Supplier R&D and Tooling Cost)**.** *Let the raw material used by the supplier be non-exchange-traded. If $MPS_1 \gtrsim FP_1$ and $\mathcal{K}_{2,1} \leq \mathcal{K}_{2,2}$, then $MPS_2 \gtrsim FP_2$.*

Proposition 3 reflects the fact that the transfer payment to the supplier, $\tau^*_{a,n}$, in Lemma 2 is non-linear in $\mathcal{K}_2$. In particular, as $\mathcal{K}_2$ increases, the supplier returns a larger portion of his information rent to the automaker, thereby increasing the automaker's payoff from the MPS contract.

Proposition 4 (Reported Raw Material Price Variability)**.** *Let the raw material used by the supplier be non-exchange-traded, $\left[c^l_{n,1}, c^h_{n,1}\right] \subseteq \left[c^l_{n,2}, c^h_{n,2}\right]$, and $\mathbb{E}_0[\widetilde{C}_{n,1}] = \mathbb{E}_0[\widetilde{C}_{n,2}]$. (1) If $MPS_1 \gtrsim FP_1$ and the payoff is convex in $\widetilde{C}_{n,j}$, $j = 1, 2$, then $MPS_2 \gtrsim FP_2$ and $Var(\widetilde{C}_{n,1}) \leq Var(\widetilde{C}_{n,2})$. (2) If $FP_1 \gtrsim MPS_1$ and the payoff is concave in $\widetilde{C}_{n,j}$, $j = 1, 2$, then $FP_2 \gtrsim MPS_2$ and $Var(\widetilde{C}_{n,1}) \leq Var(\widetilde{C}_{n,2})$.*

Recall from §3 that suppliers are restricted to raw material price updates, $\widetilde{c}_n$, within some range $[\mathbf{c}^l_n, \mathbf{c}^h_n]$. Proposition 4, links the contract type preference to the variability of those updates. This reported cost variability is non-decreasing in the variability of the actual raw material cost, c_n, provided that an increase in the actual raw material cost's variability is not accompanied by tightening of the range $[\mathbf{c}^l_n, \mathbf{c}^h_n]$.

An increase in the reported cost's variability skews the automaker's preference towards the MPS (FP) protocol whenever her expected payoff is concave (convex) in the reported raw material cost. (Whether the expected payoff function is convex or concave depends on the revenue function of Equation 1.)

The standard approach from decision theory is to order prospects in terms of stochastic dominance and decide which one is preferable. In Part (1)[(2)], the automaker's expected payoff is convex [concave] in the reported raw material cost. This property translates into second-order stochastic dominance and a higher [lower] expected payoff under the MPS agreement in market 2 than in market 1. If MPS [FP] is the preferred agreement in market 1, it is also the preferred agreement in market 2. Finally, notice how the contract's preference coincides with the higher raw material cost variability in market 2 than in market 1. The standard result is that second-order stochastic ordering implies variance order, not necessarily vice versa (Shaked and Shanthikumar, 2007, Result 3.A.2).

7 Empirical Analysis of the Current Contracting Practice

Having developed a procurement model in the automotive industry, we now evaluate our model's predictions using the contract, supplier, and raw material data from BMW. In summary, Propositions 1 through 4 identify information asymmetry, price elasticity of demand, supplier R&D and tooling costs, and raw material price support, respectively, as key factors that determine contract choice in this context. We use these propositions to derive hypotheses, which we test in this section.

Hypotheses. *(H.1: Asymmetric Information.) All else equal, the proportion of FP contracts is greater when information is asymmetric than when information is symmetric.*

(H.2: Elasticity of Demand.) All else equal, the proportion of MPS contracts is greater when the price elasticity of demand is high than when elasticity is low.

(H.3: R&D and Tooling Cost.) All else equal, the proportion of MPS contracts is greater when supplier R&D and tooling cost is high than when R&D and tooling cost is low.

(H.4a: Price Support.) All else equal, the proportion of MPS contracts is greater when the raw material price support is greater than when it is smaller.

(H.4b: Price Support.) All else equal, the proportion of FP contracts is greater when the raw material price support is greater than when it is smaller.

Note that although our analytical model identifies a directional relationship between contract type and information asymmetry, price elasticity of demand, and supplier R&D and tooling costs, the empirical relationship between contract type and price support (Proposition 4) is *ex-ante* ambiguous. It will depend on whether the automaker's expected payoff is convex or concave in the raw material cost.

Therefore, although a wider range of permissible prices leads to a greater preference for MPS or FP contracts, *which* of the two protocols is preferred will depend on the convexity of the payoff for reasons already explained in our discussion of Proposition 4.

7.1 Our Data Sets

We test our hypotheses on data from three distinct data sets concerning BMW's Contracts, Suppliers, and Raw Material Prices. We describe the three data sets below. We also refer the reader to Appendix B, where we provide a data dictionary and additional details about the information and data fields contained in the data sets.

7.1.1 Contracts Data Set

BMW's supply chain involves hundreds of Tier 1 suppliers delivering parts and components to over 30 manufacturing and assembly plants across multiple continents and countries. The procurement strategy of the entire BMW Group is set by the Raw Material Management Department (RMD), which interfaces with various procurement, production, and finance functions. The RMD keeps a central database of all raw material transactions with its suppliers. We were provided with a subset of this database containing contract information between BMW and its Tier 1 suppliers from 2013–2015.

Automotive suppliers are required to disclose the complete raw material makeup (in weight) of the components they supply. Each contract between BMW and a supplier is at the raw material level: contracts are written on a raw material within a component. Each observation in the Contracts data set is a quarterly transaction for a particular raw material m within component a delivered to the manufacturing plant p by supplier i. We aggregate these quarterly transactions and collapse them to a single point because although order quantities change over time, the contract itself does not (see also Section 3). In other words, once contract terms are set, they remain the same until the contract's end (typically after six years, coinciding with the vehicle model life cycle), and there is no variation in contract type over time.

7.1.2 Suppliers Data Set

We were also provided with information on each Tier 1 supplier: identifying information consisting of the supplier's name and headquarters location and BMW's annual aggregate purchase volume from each supplier from 2013–2015. The purchase volume data cover the total payment made by BMW to each supplier per year. We refer to this as the "total transfer" from BMW to a supplier. This total transfer is composed of both raw material costs and other transfers.

7.1.3 Raw Material Prices Data Set

Finally, we were provided monthly raw material prices from Jan. 2010 to Dec. 2015. These are the same raw material index prices on which the contracts at BMW were written, and they come directly from organized exchanges and other supplier price charts and trade group publications.

7.2 Variables

7.2.1 Contract Type

BMW offers either one of two types of contracts: an FP contract which reimburses raw material costs at a fixed rate, and an MPS contract which reimburses raw material costs at a floating rate. We observe the contract type from the Contracts data set, and so we define our dependent variable as $MPS = 1$ if a contract is MPS, and $MPS = 0$ if the contract is FP.

7.2.2 Independent Variables

Information Asymmetry. We use two measures to capture information asymmetries in the various raw material markets. For the first, BMW provided a proprietary measure of its market opacity assessment in the raw materials it purchases. The variable $Opaque_m$ ranges from 1 to 3 and captures increasing market opacity (i.e., higher information asymmetry).

However, because BMW constructs this measure, it is likely endogenous to contract choice. We, therefore, also include an exogenous (to BMW) measure of information asymmetry based on the classification from Rauch (1999). This organizes materials according to whether they are sold on an organized exchange, are referenced in trade publications, or are neither exchange-traded nor referenced. Exchange-traded goods are generally homogeneous products with quoted and widely available reference prices. Conversely, goods without reference prices are highly differentiated and heterogeneous in quality and functional specifications. We set the binary variable $Heterogeneous_m = 0$ if the raw material in the supply contract is exchange-traded (homogeneous), and $Heterogeneous_m = 1$ if the raw material is not exchange-traded (heterogeneous). Goods in the heterogeneous category are indicative of high informational asymmetries.

In our data set, 53.9% of contracts are written on materials that are not exchange-traded and, therefore, do not have widely accepted reference prices. As Figure 2 also shows, there is considerable variation and significant overlap in contract value for exchange and non-exchange-traded raw materials.

Price Elasticity of Demand. Automakers manufacture multiple car models and cater to several consumer types. For example, BMW's 1- through 5-series vehicles are high-production automobiles meant to cater to a significant portion of the (upper-) middle-class population. These accounted for

Figure 2: Contracts for Exchange and Non-Exchange-Traded Materials

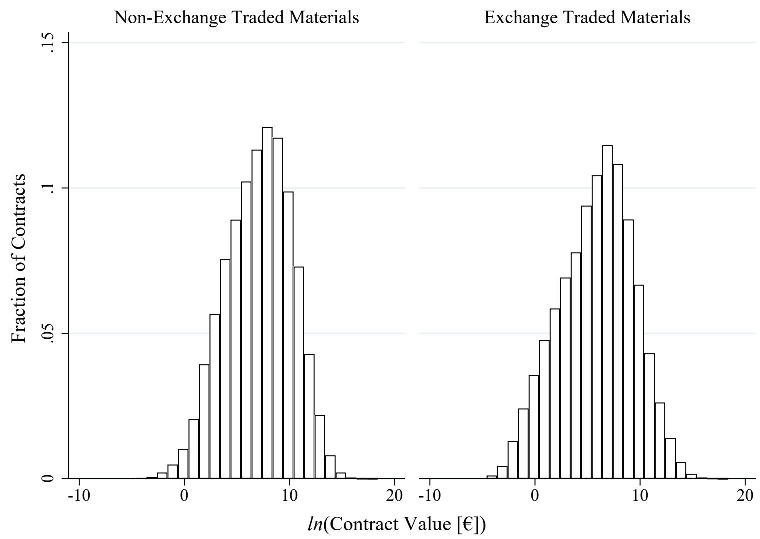

around 67% of BMW's unit sales volume in 2015, with the 3-series seeing more than 400,000 units delivered (BMW Group, 2016). An entry-level 5-series had an MSRP of €39,900. On the other hand, the 6- and 7-series models are more typical of "luxury goods" and cater to a smaller market of higher-net-worth individuals. Both models accounted for around 3% of BMW's 2015 unit sales volume. The 6- and 7-series models had MSRPs of more than €70,000 and achieved combined sales of 57,326 units in 2015 (BMW Group, 2016). Table 2 tabulates unit sales and MSRP's of several models by quarter, highlighting production and price differences between the 1- through 5-series and 6-/7-series vehicles. All models contain raw materials and commodities with symmetric and asymmetric price information, and suppliers often provide parts and components for multiple models.

Because these groupings capture two different consumer types, we can test whether the end model's price elasticity of demand affects the contract BMW offers its suppliers. We expect that consumers of the 1- through 5-series vehicles are more sensitive to price changes than consumers of the 6- and 7-series cars; the 1- through 5-series models are more price elastic than the 6- and 7-series models. In line with this intuition, Berry et al. (1995, see Table V, pg. 879) find that smaller, less-expensive cars display greater price elasticity than larger, more expensive vehicles. The BMW 735i exhibits one of the lowest price elasticities (i.e., it is an inelastic good).

We dichotomize parts in our data set according to whether they are installed in 1- through 5-series vehicles or in 6- and 7-series automobiles by performing a textual search on the part descriptions in the Contract data set and look for car model development codes. Development codes are used to disambiguate models and are also publicly known.[7] For example, in our data, we know that parts with

[7] For a list, see https://www.bimmerworld.com/About-Us/BMW-Chassis-Engine-Codes/ (accessed July 1, 2020).

Table 2: Quarterly Model Sales and MSRP.

Quarter	1-/2-Series	3-/4-Series	5-Series	6-Series	7-Series
2013 Q1	53,906	109,309	85,731	6,174	12,390
2013 Q2	55,797	128,391	94,102	7,838	14,710
2013 Q3	53,383	128,935	91,069	6,348	15,345
2013 Q4	50,525	148,460	96,090	7,327	13,556
2014 Q1	52,786	134,380	91,600	8,223	12,670
2014 Q2	56,083	148,940	101,960	5,511	13,708
2014 Q3	55,207	151,634	84,919	4,499	9,900
2014 Q4	66,995	164,840	94,574	5,755	12,241
2015 Q1	69,471	143,828	88,621	6,977	8,803
2015 Q2	80,843	154,892	85,607	4,416	10,521
2015 Q3	91,707	147,087	84,614	3,808	5,863
2015 Q4	97,281	150,921	88,254	5,761	11,177
MSRP (entry-level)	€20,200	€29,650	€39,900	€75,300	€73,500

Source: Production numbers are from BMW's quarterly reports and MSRP numbers are for base-level models from AutoBild.de.

the model code "F30" go into a 3-series vehicle; "F01" model codes belong to a 7-series. Yet, not all parts specify which vehicle series they belong to, which restricts our sample (i.e., no model code is specified). Because of component sharing across car models, we cannot know whether such parts were intended for a price-elastic or price-inelastic model. However, our approach ensures that the parts are not shared across multiple components and models; parts are model-specific. [8] We set $Elastic = 1$ if the contracted part is for the 1- through 5-series models and $Elastic = 0$ if the part is for the 6- or 7-series models.

R&D and Tooling Cost. We proxy R&D and tooling costs by taking the total transfer from BMW to a supplier (from the Suppliers data set) minus the raw material portion of the transfer (from the Contracts data set). Specifically, we first take the raw material cost (in €) of each contract and sum these up to the level of the supplier ($RawMatTransfer_i = \sum_{a=1}^{A} \sum_{m=1}^{M} RawMatCost_{m,a,i}$, where m denotes a specific raw material within component a). This sum is the *raw material transfer* BMW paid supplier i in a given year. We subtract this from the *total transfer* BMW paid the supplier in a year (i.e., $TotalTransfer_i$). This difference is directly comparable to $\mathcal{K}$ in our model due to the auction format that BMW uses to award contracts (see §4.3.) As we are only provided with a range of the total transfer in our Suppliers data set, we calculate the variable $Tooling$ as the ln-transformed difference between the ceiling of the provided total transfer minus raw material costs. [9] In order words, $Tooling_i = ln(TotalTransfer_i - RawMatTransfer_i)$. Figure 3 shows a large variance in supplier R&D and tooling cost. More importantly, the figure highlights how suppliers of varying R&D and tooling cost

[8] For example, in our data set, one part is named "MD Belt force limiter US F33." The "F33" notes that this part goes into a 4-series vehicle and *only* into a 4-series vehicle. However, a different part in our data set is named "Right roof trim, spacegray metallic." No identifying information here would allow us to link it to a specific car model, so we cannot know whether this part was intended for a price-elastic or price-inelastic car model.

[9] The results are qualitatively similar if we use the floor of the provided total transfer instead.

Figure 3: Supplier R&D and Tooling Cost

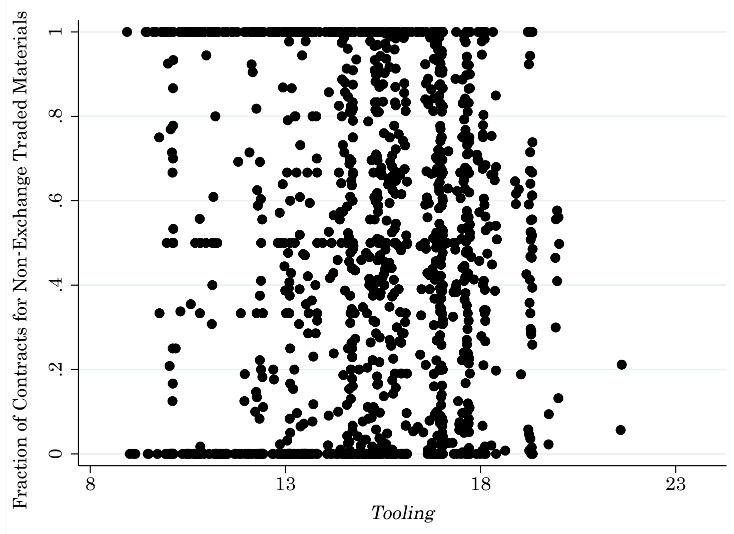

supply materials with symmetric and asymmetric price information. [10]

Raw Material Price Support. Our model predicts that contract type also depends on a raw material's price support. Due to data limitations, we cannot directly observe $\mathbf{c}^l$ and $\mathbf{c}^h$, so we proxy for this price support with price variability. Using the Raw Material Prices data set, we first calculate the arithmetic range of a raw material's reported prices (i.e., the maximum reported price minus the minimum reported price) between January 2010 and the contract start. Then, because prices of raw materials are reported in different currencies and units, we normalize the range by the average price over the same period. Our variable $PriceSupport$ captures the range of a raw material's reported cost variability over the time period and is related to the range $[\mathbf{c}^l, \mathbf{c}^h]$. Figure 4 plots the price trajectories of three raw materials in our data set: aluminum and lead (exchange-traded on the LME), and hot-dipped galvanized steel (which is not exchange-traded). The three materials all exhibit relatively high levels of variability.

7.2.3 Controls

We control for additional factors that may affect BMW's contract type to the extent possible. First, we control for the amount of raw material contracted by including the variable $Tonnage$, which is the ln-transformed weight of the material delivered under the contract.

Second, we look to control for supplier bargaining power, which may affect contract negotiations and decisions. $NumSuppliers$ is the natural logarithm of the number of suppliers supplying the same component to BMW in a given year. Sourcing a component from a single supplier is likely due to a lack

[10] Econometrically, because $Tooling$ is only a proxy, it will necessarily imprecisely measure the supplier's *true*, unobservable R&D and tooling cost. This leads to endogeneity by creating the well-known errors-in-variables problem and biasing the coefficient on $Tooling$ toward zero (attenuation bias). This makes it less likely for us to detect a statistically significant effect *a priori*. We address this concern using instrumental variables in Appendix C.

Figure 4: Raw Material Prices

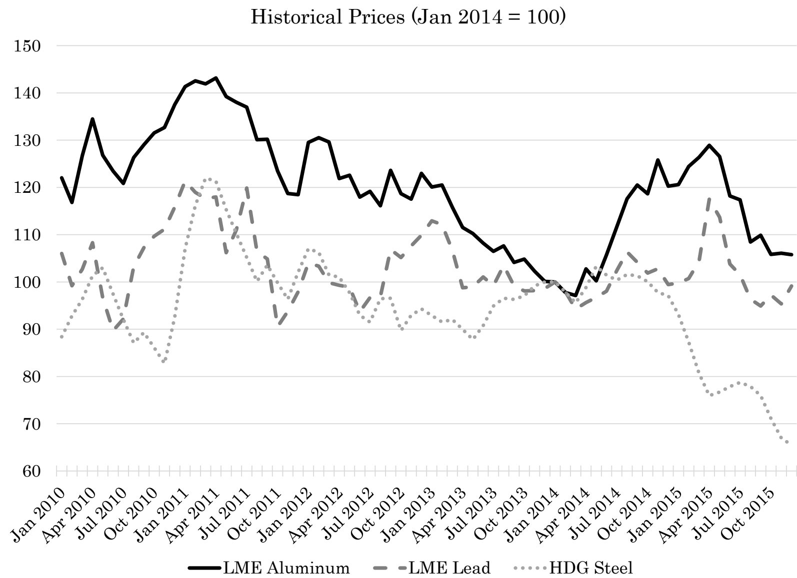

of a competitive supplier base. We would expect greater bargaining power for the supplier. Yet, when BMW procures a component from multiple suppliers, this should be indicative of greater competition in its supplier base for that component and BMW's ability to procure a component from several different suppliers. In such cases, we would expect greater bargaining power for BMW and lesser bargaining power for any one supplier of that component. [11]

Third, we include a set of indicator variables corresponding to the quarter in which a contract was entered into. This controls for any contemporaneous time effects or other "cohort" effects.

Fourth, we include a set of indicator variables corresponding to the manufacturing plant to which a component was delivered. Manufacturing plants assemble various car models and may require different suppliers, components, and raw materials.

Finally, we include component fixed effects to control for heterogeneity in component-specific factors. Certain components (e.g., an electrical wiring harness) may be more complex and require greater know-how and expertise than others (e.g., a door trim panel). They may also require different materials: a wiring harness must be made of copper and cannot be substituted for steel. Such factors will likely be correlated with both contract type and our independent variables. For example, components with different materials will have different levels of information asymmetry. Complex components may also go in more high-end vehicle models and/or require greater supplier investments in R&D. Kim and Netessine

[11] In our data, we see that the component with the greatest number of suppliers is listed as "Stamped part – Steel." This refers to a generic, low-tech part stamped out of steel that many companies can most likely manufacture. On the contrary, most of the components supplied to BMW by a single supplier are for electronics parts, especially for vehicles' electrical wiring harnesses. These are much more technical and complex to manufacture, typically requiring specialized expertise and capabilities.

Table 3: Descriptive Statistics.

	Obs.	mean	sd	Correlation Matrix (1)	(2)	(3)	(4)	(5)	(6)	(7)
(1) *MPS*	325,109	0.363	0.481							
(2) *Opacity*	307,124	1.592	0.546	-0.805						
(3) *Heterogeneous*	259,668	0.539	0.498	-0.795	0.869					
(4) *Elastic*	45,021	0.774	0.418	0.037	-0.060	-0.015				
(5) *Tooling*	324,862	17.908	1.721	0.313	-0.244	-0.293	0.038			
(6) *PriceSupport*	268,479	0.595	0.467	-0.369	0.532	0.457	-0.017	-0.040		
(7) *Tonnage*	308,139	5.854	3.770	-0.246	0.195	0.302	0.051	-0.162	-0.079	
(8) *NumSuppliers*	325,109	1.238	1.127	-0.286	0.209	0.265	-0.388	0.122	0.019	0.335

Notes: All correlations are significant at the $\alpha = 0.01$ level.

(2013) further note that there is also more cost uncertainty around complex parts. Component fixed effects account for important characteristics such as component complexity and raw material makeup.

Descriptive statistics for our variables are provided in Table 3 and Appendix Tables B2 and B3. Our sample consists of 325,109 contracts, but because we cannot match some raw materials and suppliers due to BMW's coding and classification scheme, observation numbers across models vary.

7.3 Econometric Specification

We test our hypotheses by estimating a model of the following form:

$$MPS_{maip} = \beta_0 + \beta_1 InfoAsymmetry_m + \beta_2 Elastic_a + \beta_3 Tooling_i + \beta_4 PriceSupport_m + \zeta \mathbf{X} + P_p + Q_q + A_a + \epsilon_{mcip}. \tag{14}$$

Each contract is for a raw material m in component a, which is delivered by supplier i to manufacturing/assembly plant p. In the equation above, $InfoAsymmetry$ denotes either one of *Opaque* or *Heterogeneous*, and $\mathbf{X}$ is the set of control variables mentioned in the previous section. Plants, quarters, and components are controlled for through a set of indicators denoted by P, Q, and A, respectively. All models include heteroskedasticity-robust standard errors clustered at the level of the component.

We estimate a Linear Probability Model (LPM) rather than a binary choice model for two reasons. First, the size of the data set and the necessity for component fixed effects make maximum likelihood computation cumbersome, and conditional logit models with component-level fixed effects are not estimable. Second, a substantial number of components have contract types that are perfectly predicted – that is, components are for one raw material or part, and contract type does not vary – leading a large number of observations to drop out of logit and probit models when we include component fixed effects. A well-specified LPM produces valid estimates of the direction of the effect (Wooldridge, 2015); however, the magnitudes of the effects are generally not interpretable. Thus, we comment only on the directions and significance of β_{1-4} in the analysis. (In Appendix D, we compare several probit models to LPMs.)

Finally, since the model makes clear directional predictions for β_1, β_2, and β_3, we test these coefficients under one-tailed tests. The relationship between reported price variability and contract type (H.4a/b) is an empirical question, so β_4 is tested under a two-tailed test.

7.4 Empirical Results of Contract Choice

Our objective with this empirical analysis is to assess the validity of our model from Sections 3–5. Although our model is somewhat stylized, the data from BMW allow us to empirically test whether it captures reality. In this section, we, therefore, seek to validate whether the market factors identified by the model as influencing contract choice do, in practice, matter.

The estimates from Equation 14 are displayed in Table 4. We exclude *Elastic* in Models I and II, whereas Models III and IV include this variable. This is due to the number of observations available for *Elastic*. (See Section 7.2.1.) Models I and II allow us to estimate the effects of the other variables on a larger sample of contracts but at the risk of model misspecification from an omitted variable. On the other hand, Models III and IV eliminate the misspecification concern but lead to a much more limited sample from which to draw conclusions. For transparency and completeness, we, therefore, report estimates from both kinds of models.

We observe a contracting pattern that is in line with our theoretical predictions. All coefficient signs in Models I–IV are directionally consistent with the model predictions and hypotheses, and our variables of interest are generally statistically significant. As hypothesized, information asymmetry, price elasticity of demand, supplier R&D and tooling costs, and raw material price support (as proxied by reported price variability) seem to be important factors for contract selection.

Moreover, as price variability increases, BMW tends to offer a greater proportion of FP contracts. This result corroborates Hypothesis 4b. Accordingly, the empirical results suggest that BMW's expected payoff is concave in the reported raw material cost (as per Proposition 4).

Overall, our results suggest that BMW tends to tailor its contract offerings depending on situations where suppliers are better positioned to distort prices and extract rents. Moreover, the data support the idea that our model is capturing reality and that market factors which *theoretically* inform contract choices in fact *do* so in practice. The results, therefore, lend confidence that our model is valid, allowing us to identify a better contract for this context.

Table 4: Baseline Analysis of Contract Choice.

	(I)	(II)	(III)	(IV)
Opaque	-0.296***		-0.084***	
	(0.043)		(0.032)	
Heterogeneous		-0.503***		-0.189***
		(0.036)		(0.060)
Elastic			0.007**	0.008**
			(0.003)	(0.004)
Tooling	0.006***	0.004**	0.004*	0.004*
	(0.002)	(0.002)	(0.003)	-0.003
PriceSupport	-0.057*	-0.111***	-0.154***	-0.191***
	(0.032)	(0.021)	(0.049)	(0.051)
Tonnage	0.001*	0.001***	-0.000	-0.000
	(0.000)	(0.001)	(0.000)	(0.000)
NumSuppliers	-0.028**	-0.021	0.011	0.012
	(0.014)	(0.014)	(0.014)	(0.012)
Constant	0.840***	0.706***	0.925***	0.865***
	(0.065)	(0.045)	(0.060)	(0.070)
Quarter FE	Yes	Yes	Yes	Yes
Plant FE	Yes	Yes	Yes	Yes
Component FE	Yes	Yes	Yes	Yes
# Contracts	247,167	237,195	39,646	39,098
# Components	3,741	3,656	479	453
Within-R^2	0.280	0.456	0.167	0.294

Notes: The dependent variable is MPS. Significance on *Opaque*, *Heterogeneous*, *Tooling*, *Elastic* is tested with one-sided hypothesis tests; *PriceSupport* with two-sided tests. Robust standard errors are in parentheses and are clustered at the component. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$

8 An Optimal Contract and Numerical Estimation

This section proposes a second-best procurement protocol à la Myerson (1979), which the automaker can unilaterally offer to any supplier in the environment of §5.1. This is notable because suppliers are organized in trade groups that collectively safeguard price information (§5). As such, it may be tricky for an individual supplier to agree to voluntarily post prices on a blockchain, which would require implementing when using smart contracts surveyed in Zheng et al. (2020). Although several studies designed second-best procurement contracts before us (e.g., Corbett et al., 2004; Özer and Wei, 2006; Kostamis and Duenyas, 2011), the existing protocols are static (one-shot) and primarily designed to screen demand information from the downstream buyer in the absence of pre-existing inventory commitments. (Here, we need to screen upstream cost information with commitments.)

The proposed procurement protocol is non-linear with two sets of terms: *ex-ante* and *ex-post*. Ex-ante, the automaker commits to purchasing $lq_{r,n}$ units for a unit cost of $w_{r,n}$ in each period n. This ex-ante commitment is analogous to that of §§4 and 5.

Ex-post, the automaker purchases additional quantity according to the following pre-specified price-quantity menu with a fixed transfer (tooling cost), $\tau_{r,n}$: $\{(\widehat{c}_n, y_{r,n}) \mid \widehat{c}_n \leq \widehat{c}_n^{max}\}$. Following a procedure well documented in the literature (see, e.g., Jehle and Reny, 2001, §8), the menu is designed to maximize the automaker's expected profit subject to the supplier's individual rationality (IR) and incentive compatibility (IC) constraints.

Seeing the menu in a particular period n and knowing the period n's raw material cost, c_n, the better-informed supplier picks the most favorable price-quantity pair, $(\widehat{c}_n, y_{r,n})$. Then, he reports raw material cost $\widehat{c}_n$ to which the automaker responds by augmenting her initial quantity commitment of $lq_{r,n}$ units with $y_{r,n}$ units. Note, however, that $y_{r,n} > 0$ only if the supplier's price is low enough.

The analysis proceeds by backward induction while assuming that the automaker follows the JIT production policy. (Lemma A.4 summarizes the conditions under which the JIT policy can be implemented without loss of optimality. They are analogous to those of Lemma A.3 with the provision that the lower bound on the holding cost sufficient to halt speculative inventory purchases is either lower or higher than that of Lemma A.3, depending on the model parameters.)

We begin by characterizing the optimal ex-post menu, taking the automaker's ex-ante commitment and the supplier's R&D and tooling cost as given (§8.1). Then, in §8.2, we characterize the ex-ante commitment and the equilibrium R&D and tooling cost. Proposition 5 and Corollary 1 characterize the contract's ex-post and ex-ante terms. Lemma 3 describes the equilibrium transfer, compensating the supplier for R&D and tooling cost.

The analysis concludes with a numerical experiment that leverages our data, estimating the profit boost that an automaker might expect from this section's contract.

8.1 Ex-Post Price-Quantity Schedule Design

The problem of designing the optimal menu is a calculus of variation problem that seeks to find functions $\widehat{c}_n(c_n)$, $y_{r,n}(c_n)$ subject to $c_n \leq c_n^{max}$. The optimal menu is a function of the supplier's actual raw material cost, c_n. The max price that the supplier is allowed to report in period n is $\widehat{c}_n(c_n^{max})$. Mathematically:

$$\max_{\widehat{c}_n(c_n), y_{r,n}(c_n), c_n^{max}} \quad \mathbb{E}_0\left[r(y_{r,n}(C_n) + lq_{r,n}, v_n) - \widehat{c}_n(C_n)\, y_{r,n}(C_n) \mid C_n \leq c_n^{max}\right] \mathbb{P}\{C_n \leq c_n^{max}\}$$

$$+\mathbb{E}_0\left[r(lq_{r,n}, v_n) \mid C_n > c_n^{max}\right] \mathbb{P}\{C_n > c_n^{max}\} - \tau_{r,n} \qquad (15a)$$

$$\text{s.t.} \qquad (h - l)\, q_{r,n} - y_{r,n}(c_n) \geq 0 \qquad (15b)$$

$$(y_{r,n}(c_n) + lq_{r,n})\, (\widehat{c}_n(c_n) - c_n) \geq 0 \qquad (15c)$$

$$(y_{r,n}(c_n) + lq_{r,n})\, (\widehat{c}_n(c_n) - c_n) - \left(y_{r,n}(c_n') + lq_{r,n}\right) \left(\widehat{c}_n(c_n') - c_n\right) \geq 0, \qquad (15d)$$

for all $c_n \in \mathbb{R}^+$ and $c_n' \in \mathbb{R}^+$. Although the automaker is acquiring $(lq_{r,n} + y_{r,n})$ components in period n, in Equation (15), she only pays for $y_{r,n}$ units because the cost of the $lq_{r,n}$ units is sunk.

The constraint (15b) restrains the automaker's total purchase quantity to $hq_{r,n}$ units. The inequality (15c) is an IR constraint. It reflects that the supplier facing cost c_n can choose to supply nothing but prefers to supply $(y_{r,n}(c_n)+lq_{r,n})$ units. Inequality (15d) is an IC constraint. It accounts for the fact that the supplier can report some cost $\widehat{c}_n(c'_n)$, that is different from the intended cost, $\widehat{c}_n(c_n)$, but prefers to report $\widehat{c}_n(c_n)$. As such, the IC constraint must hold for all possible pairs of costs, c_n and c'_n. This feature makes the IC constraint complicated to evaluate. Without the IR and IC constraints, the problem of designing the optimal contract reduces to the first-best problem of §4.2.

Proposition 5. *Let $F_n(\cdot)$ denote the distribution function of C_n and $\lambda_n(\cdot) \equiv dF_n(\cdot)/F_n(\cdot)$. The price-quantity menu optimal in (15) is*

$$(\widehat{c}^*_n(c_n), y^*_{r,n}(c_n)) = \left\{ \left(c_n + \frac{1}{y^*_{r,n}(c_n)} \int_{c_n}^{c_n^{max\,*}} y^*_{r,n}(x)dx,\ y^*_{r,n}(c_n) \right) \;\middle|\; c_n \leq c_n^{max\,*} \right\}, \ \textit{where} \tag{16a}$$

$$\begin{aligned}\left(y^*_{r,n}(c_n), c_n^{max\,*}\right) = \arg\max_{y,c}\ & \mathbb{E}_0\left[r(y + lq_{r,n}, v_n) - y\left(C_n + \frac{1}{\lambda_n(C_n)}\right) \mid C_n \leq c \right] \mathbb{P}\{C_n \leq c\} \\ & + \mathbb{E}_0\left[r(lq_{r,n}, v_n) \mid C_n > c \right] \mathbb{P}\{C_n > c\} \ \ s.t.\ \ y \leq (h-l)\, q_{r,n}. \end{aligned} \tag{16b}$$

Equations (16) reveal how the automaker designs a contract that does not distort the total output whenever the supplier faces low raw material cost and leaves no surplus for the high-cost supplier. The quantity distortion is seen in Equation (16b), where $(1/\lambda_n(c_n)) \to 0$ as $c_n \to 0$. The price distortion is seen in Equation (16a), where $\frac{1}{y^*_{r,n}(c_n)} \int_{c_n}^{c_n^{max\,*}} y^*_{r,n}(x)dx \to 0$ as $c_n \to c_n^{max\,*}$. (As $c_n \to c_n^{max\,*}$, the unit price converges to the first-best unit price.)

8.2 Ex-Ante Commitment and Auction

Corollary 1 reviews the automaker's initial quantity commitment. The intuition behind Corollary 1 is analogous to that of §§4 and 5.1: The automaker commits to a quantity that maximizes her expected payoff for the most negligible unit price that satisfies the supplier's individual rationality constraint.

Corollary 1. *Let $F_n(\cdot)$ denote the distribution function of C_n and $\lambda_n(\cdot) \equiv dF_n(\cdot)/F_n(\cdot)$. The automaker's equilibrium commitment price and quantity, $(w^*_{r,n}, lq^*_{r,n}) = \left(\mathbb{E}_0[C_n],\ lq^*_{r,n}\right)$, where:*

$$\begin{aligned} q^*_{r,n} = \arg\max_{q}\ & \mathbb{E}_0\left[r(y^*_{r,n} + lq, v_n) - y^*_{r,n}\left(C_n + \frac{1}{\lambda_n(C_n)}\right) \mid C_n \leq c_n^{max\,*} \right] \mathbb{P}\{C_n \leq c_n^{max\,*}\} \\ & + \mathbb{E}_0\left[r(lq, v_n) \mid C_n > c_n^{max\,*} \right] \mathbb{P}\{C_n > c_n^{max\,*}\} - lq\mathbb{E}_0[C_n]. \end{aligned} \tag{17}$$

Now that we derived the automaker's ex-ante commitment, the last moving part to be determined is the supplier's fixed transfer, $\tau_{r,n}$. As in §§4 and 5, the equilibrium transfer, $\tau^*_{r,n}$, is determined through second-price auction that we have already described in §4.3.

The expression for the transfer is equivalent to that of Lemma 2 with the proviso that the expression for information rent is $\mathbb{E}_0\left[\frac{y^*_{r,n}(C_n)}{\lambda_n(C_n)} \mid C_n \leq c_n^{max\,*}\right]\mathbb{P}\{C_n \leq c_n^{max\,*}\}$. (For details, see our discussion of Lemma 2 in §5.1.2.)

Lemma 3. *Let $F_n(\cdot)$ denote the distribution function of C_n and $\lambda_n(\cdot) \equiv dF_n(\cdot)/F_n(\cdot)$. The winning supplier's equilibrium transfer is:*

$$\tau^*_{r,n} = \frac{1}{\bar{n}}\left(\mathcal{K}_2 - \mathbb{E}_0\left[\frac{y^*_{r,n}(C_n)}{\lambda_n(C_n)} \mid C_n \leq c_n^{max\,*}\right]\mathbb{P}\{C_n \leq c_n^{max\,*}\}\right)^+ .$$

In summary, we can see the automaker's expected payoff under the contract that we derived in this section:

$$\sum_{n=1}^{\bar{n}} \mathbb{E}_0\left[r(y^*_{r,n} + lq^*_{r,n}, V_n) - C_n\left(y^*_{r,n} + lq^*_{r,n}\right) \mid C_n \leq c_n^{max\,*}\right]\mathbb{P}\left\{C_n \leq c_n^{max\,*}\right\}$$
$$+\,\mathbb{E}_0\left[r(lq^*_{r,n}, V_n) - lq^*_{r,n}\, C_n \mid C_n > c_n^{max\,*}\right]\mathbb{P}\left\{C_n > c_n^{max\,*}\right\} - \mathcal{K}_2. \quad (18)$$

The automaker's effective marginal cost is C_n (i.e., she buys the components at cost – exactly what she paid in the benchmark case of §4). In contrast to the benchmark case, however, she purchases additional components in the ex-post stage only if the realized raw material cost is below the cutoff cost of $c_n^{max\,*}$.

8.3 Numerical Estimation

We use our data and publicly available information to derive parameter values for our model. We use these to estimate the performance of an automaker's current contracting practice against the optimal contract presented above while making approximations where parameters are proprietary and difficult to obtain. Our analysis would estimate the potential performance improvement if an automaker changed her procurement protocol from the best of FP and MPS, to the menu contract of §8.

We leverage our Raw Materials Prices data set, the Contract Info data set, academic publications, and other publicly available sources, to calibrate our estimation. These inputs are described in Table 5. We base our parameter values on the BMW 3-series because the consumer demand for this model is more elastic than the demand for the higher-priced, luxury models (Berry et al., 1995; Anderson et al., 1997). Elasticity is a powerful driver of the protocol choice, and focusing on a vehicle with such characteristics means that our estimates are more generalizable to similar high-volume automobiles. We approximate production by pretending that each car is produced from two aggregate components and the first (second) component is produced from (non-)exchange-trade raw materials.

For the components made from exchange-traded raw materials, we assume that the automaker pays the raw material cost $C_n^{\text{symmetric}}$ plus the tooling cost. This is equivalent to the automaker using the

Table 5: Estimation Parameter Inputs.

Aggregate Revenue Information		
Parameter	**Value**	**Source**
MSRP	€29,650 per vehicle	Table 2; AutoBild.de
Unit Sales	600,000 per year	BMW Quarterly Reports
Gross Margin (per vehicle)	22%	BMW Group (2016)
Price Elasticity Range	Between -2.6 and -1.2	Berry et al. (1995); Anderson et al. (1997)

Aggregate Cost Information (per vehicle)		
Parameter	**Value**	**Source**
Total Production Cost (%-age of total vehicle cost)	60%	Fries et al. (2017); Kallstrom (2020)
Aggregate Tooling Cost (%-age of total vehicle cost)	16.6%	Contracts data set
Aggregate %-age weight of exchange-traded raw material inputs	21%	Contracts data set
Baseline distribution of aggregate exchange-traded raw material costs	$LN(7.526, 2.467 \times 10^{-3})$	Raw Material Prices data set
Aggregate %-age weight of non-exchange-traded raw material inputs	79%	Contracts data set
Baseline distribution of aggregate non-exchange-traded raw material costs	$LN(7.135, 5.913 \times 10^{-2})$	Raw Material Prices data set

Notes: $LN(\mu, \sigma)$ denotes log-normal distribution with parameters μ and σ. It means that log of the raw material cost is normally distributed with a mean of μ and an annual standard deviation of σ.

MPS contract of §4.

For the components made from non-exchange-traded raw materials, we assume that the supplier faces the raw material cost of $C_n^{\text{asymmetric}}$. The automaker pays a unit cost of G_n and the tooling cost. If the contract is MPS, then $G_n = \widetilde{c}_n(C_n)$. The supplier computes $\widetilde{c}_n(C_n)$ in the manner that we describe in §5.1. If the contract is FP, then $G_n = \mathbb{E}_0[C_n^{\text{asymmetric}}]$. With the optimal contract of §8, $G_n = C_n^{\text{symmetric}} + \frac{1}{\lambda_n(C_n^{\text{symmetric}})}$. For simplicity, we compare the automaker's expected payoffs for a single production period.

We assume that demand for the vehicle model is isoelastic

$$\text{MSRP} = k \, \text{total output}^{\frac{1}{\nu}}, \tag{19}$$

where k is a scalar (baseline demand level) parameter and ν is the price elasticity of demand. We calculate the value of k from the fact that sales are 600,000 units when buyers face price elasticity of ν and a selling price (MSRP) of €29,650 per vehicle. Estimates for automobile price elasticities differ across sources (see, e.g., Berry et al., 1995; Anderson et al., 1997), ranging between (-1.6) and (-6), depending on the make and the model of the car. We take $\nu = -2$ as our baseline level and let it vary between (-1.2) and (-2.6). Elasticities less than (-2.6) are irrelevant because BMW is a premium product;

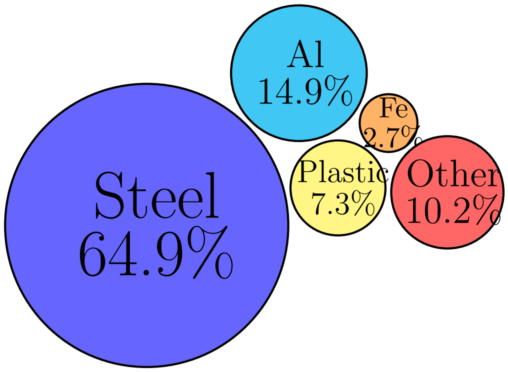

(a) Estimated Raw Material Composition of BMW 3-Series Model

(b) Changes in Efficiency Gain vs. Changes the Demand Elasticity

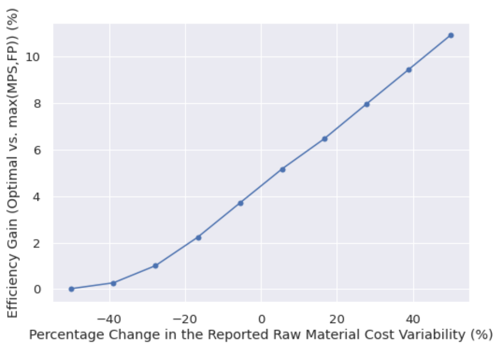

(c) Changes in Efficiency Gain vs. Changes in the Reported Raw Material Cost Variability †

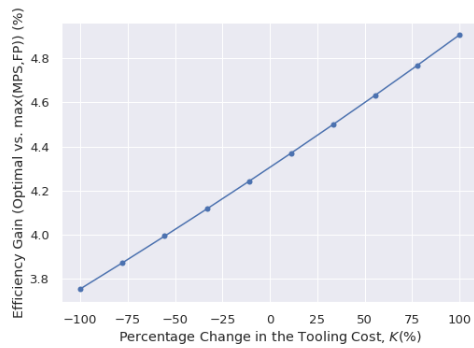

(d) Changes in Efficiency Gain vs. Changes in the Tooling Cost

Notes: † Changes in the reported raw material cost variability are assessed by varying the baseline variance of the reported raw material cost distribution included in Table 5.

Figure 5: Efficiency Gain from Using the Contract of §8 to Procure Components Manufactured From Non-Exchanged-Traded Raw Materials vs. Using the FP Contract of §5.2

likewise, elasticity values greater than (-1.2) are meaningless to study. No studies on the price elasticity of cars report values greater than (-1.2).

With these assumptions, the automaker's expected payoff from the single production period is:

$$\mathbb{E}_0\left[\text{total output}\left(\text{MSRP} - \left(C_n^{\text{symmetric}} + G_n + \text{overhead cost}\right)\right) - \text{tooling cost}\right], \tag{20}$$

where MSRP is given by Equation (19).

To compare the optimal contract's performance, we compare it against the max of FP and MPS contract payoffs. In line with our analytical predictions, the expected gain from choosing the price-quantity menu of §8 to procure the components produced from non-exchange-traded raw materials increases as price elasticity of demand decreases and variability of reported raw material and tooling costs increase. For a representative 3-series model, we see an increase between 1% and 10%. We separate the roles of price elasticity, reported raw material cost variability, and tooling cost by performing the same exercise while varying one parameter at a time. Figures 5b, 5c, and 5d present the results of these experiments. For the 3-series, the price-quantity menu appears to deliver the most benefit boost when price elasticity or reported raw material cost variability increases. A change in the tooling costs is relatively less impactful than price elasticity or reported raw material cost variability changes.

9 Conclusion

The institutional details of the procurement process crucially influence contract choice. Focusing on the procurement norms prevalent in the (European) automotive industry, this paper analyzes how cost information asymmetries and various market factors drive an automaker's decision between two commonplace contracts. This setting is dynamic, with demands and costs that fluctuate over time. The automaker chooses the supplier through an auction, and the automaker – not the supplier – picks the procurement contract type. However, information asymmetry in raw material costs means that the supplier controls the contract pricing.

Furthermore, each procurement period's inventory decisions happen in two stages: before production begins (ex-ante), the automaker must commit to some procurement quantity; then, at the beginning of each production stage (ex-post), the automaker can adjust her quantity commitment after she gets to see costs and demands. To the best of our knowledge, no procurement model from the literature fits the description above. Although this model is new to the literature, we assess the validity of our model using a unique data set of contractual arrangements provided to us by BMW. We empirically show that the factors that *theoretically* drive contract choice, in fact, *do so in practice*.

We conclude by identifying an optimal contract in this context and estimating the performance gain that could be had by switching to this contract. The optimal contract has a two-part tariff, MPS-like structure, but with refinements on how it executes its ex-ante and ex-post terms. Ex-ante, the automaker commits to a minimum purchase quantity but at an agreed price (expected cost) rather than the supplier reported price. Ex-post, the contract specifies a smart upper bound on the component cost, and reported supplier costs above that threshold trigger the minimum purchase quantity at the pre-specified price. The automaker adds to the committed minimum purchase quantity according to a given price (reported cost) and quantity menu for all other instances. Our numerical estimations suggest that adopting the optimal contract could boost an automaker's profit by an estimated 6%, depending on the situation.

The automotive procurement protocol involves many players, moving parts, and decisions. Our model captures the most salient features of this context faced by an automaker, one being information asymmetry in raw material costs. In practice, however, there is information asymmetry across many dimensions and in favor of either party. We are hopeful that future research will explore these avenues.

References

Anand, K., Anupindi, R., and Bassok, Y. (2008). Strategic inventories in vertical contracts. *Management Sci.*, 54(10):1792–1804.

Anderson, P. L., McLellan, R. D., Overton, J. P., and Wolfram, G. L. (1997). Price elasticity of demand. https://www.mackinac.org/1247. Accessed Nov 11, 2020.

Babich, V. (2010). Independence of capacity ordering and financial subsidies to risky suppliers. *Manufacturing Service Oper. Management*, 12(4):583–607.

Bajari, P. and Tadelis, S. (2001). Incentives versus transaction costs: A theory of procurement contracts. *Rand journal of Economics*, pages 387–407.

Beil, D. R. (2010). Supplier selection. *Wiley Encyclopedia of Operations Research and Management Science*.

Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.

BMW Group (2016). BMW annual report 2015. https://www.bmwgroup.com/content/dam/grpw/websites/bmwgroup_com/ir/finanzberichte/pdf/en/12784_GB_2015_en_Finanzbericht.pdf, Accessed July 31, 2018.

BMW Group (2022). BMW group report 2021. file:///C:/Users/markoup/Downloads/BMW-Group-Report-2021-en.pdf.

Bolton, P., Dewatripont, M., et al. (2005). *Contract theory*. MIT press.

Caldentey, R. and Haugh, M. B. (2009). Supply contracts with financial hedging. *Oper. Res.*, 57(1):47–65.

Corbett, C. J., Zhou, D., and Tang, C. S. (2004). Designing supply contracts: Contract type and information asymmetry. *Management Sci.*, 50(4):550–559.

de Kok, A. d. and Graves, S. C. (2003). *Supply chain management: Design, coordination and operation*. Elsevier.

Demski, J. S. and Sappington, D. E. M. (1991). Resolving double moral hazard problems with buyout agreements. *The RAND Journal of Economics*, 22(2):232–240.

Fries, M., Kerlera, M., Rohra, S., Schickrama, S., Sinninga, M., and Lienkamp, M. (2017). An overview of costs for vehicle components, fuels, greenhouse gas emissions and total cost of ownership update 2017. Institute of Transportation Studies, University of California, Davis, Research Report UCD-ITS-RR-17-45.

Gaur, V. and Seshadri, S. (2005). Hedging inventory risk through market instruments. *Manufacturing Service Oper. Management*, 7(2):103–120.

Ha, A. Y. (2001). Supplier-buyer contracting: Asymmetric cost information and cutoff level policy for buyer participation. *Naval Research Logistics (NRL)*, 48(1):41–64.

Jehle, G. A. and Reny, P. J. (2001). *Advanced Microeconometric Theory*. Addison-Wesley, Boston, MA, 02116, 2 edition.

Kallstrom, H. (Accessed: 07/13/2020). Raw materials – the biggest cost driver in the auto industry. https://marketrealist.com/2015/02/raw-materials-biggest-cost-driver-auto-industry/.

Kim, S.-H. and Netessine, S. (2013). Collaborative cost reduction and component procurement under information asymmetry. *Management Sci.*, 59(1):189–206.

Kostamis, D. and Duenyas, I. (2011). Purchasing under asymmetric demand and cost information: When is more private information better? *Oper. Res.*, 59(4):914–928.

Kouvelis, P., Turcic, D., and Zhao, W. (2018). Supply chain contracting in environments with volatile input prices and frictions. *Manufacturing Service Oper. Management*, 20(1):130–146.

Kydland, F. E. and Prescott, E. C. (1977). Rules rather than discretion: The inconsistency of optimal plans. *Journal of political economy*, 85(3):473–491.

Lai, T. L. and Ying, Z. (1991). Estimating a distribution function with truncated and censored data. *The Annals of Statistics*, 19(1):417–442.

Markou, P., Corsten, D., Carduck, C., and Koblbauer, M. (2019). Commodity risk management at BMW: Price indices and contracts. In Tate, W. L., Bals, L., and Ellram, L., editors, *Supply Chain Finance: Risk Management, Resilience and Supplier Management*, chapter 9, pages 177–194. Kogan Page, London.

Maskin, E. and Riley, J. (1984). Monopoly with incomplete information. *The RAND Journal of Economics*, 15(2):171–196.

Musalem, A., Olivares, M., Bradlow, E. T., Terwiesch, C., and Corsten, D. (2010). Structural estimation of the effect of out-of-stocks. *Management Sci.*, 56(7):1180–1197.

Myerson, R. B. (1979). Incentive compatibility and the bargaining problem. *Econometrica: journal of the Econometric Society*, pages 61–73.

Nasser, S. and Turcic, D. (2019). Temporary contract adjustment to a retailer with a private demand forecast. *Management Sci.*, 65(1):209–229.

Özer, O. and Wei, W. (2006). Strategic commitments for an optimal capacity decision under asymmetric forecast information. *Management Sci.*, 52(8):1238–1257.

Ramkumar, A. (2018). Want to invest in lithium and cobalt? good prices are hard to find. *The Wall Street Journal – Jun 19. WSJ.*

Rauch, J. E. (1999). Networks versus markets in international trade. *Journal of International Economics*, 48(1):7–35.

Ross, S. M. and Peköz, E. A. (2007). *A second course in probability.* www. ProbabilityBookstore. com.

Shaked, M. and Shanthikumar, J. G. (2007). *Stochastic Orders.* Springer, New York, NY 10013, 1 edition.

Shen, B., Choi, T.-M., and Minner, S. (2019). A review on supply chain contracting with information considerations: information updating and information asymmetry. *International Journal of Production Research*, 57(15-16):4898–4936.

Simchi-Levi, D., Schmidt, W., Wei, Y., Zhang, P. Y., Combs, K., Ge, Y., Gusikhin, O., Sanders, M., and Zhang, D. (2015). Identifying risks and mitigating disruptions in the automotive supply chain. *Interfaces*, 45(5):375–390.

Staiger, D. and Stock, J. H. (1997). Instrumental variables regression with weak instruments. *Econometrica*, 65(3):557.

Stock, J. H. and Yogo, M. (2005). Testing for weak instruments in linear IV regression. In Andrews, D. W. K. and Stock, J. H., editors, *Identification and Inference for Econometric Models: Essays in Honor of Thomas Rothenberg*, pages 80–108. Cambridge University Press.

Swinney, R. and Netessine, S. (2009). Long-term contracts under the threat of supplier default. *Manufacturing Service Oper. Management*, 11(1):109–127.

Turcic, D., Kouvelis, P., and Bolandifar, E. (2015). Hedging commodity procurement in a bilateral supply chain. *Manufacturing Service Oper. Management*, 17(2):221–235.

Veinott, A. F. J. (1965). Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Sci.*, 12(3):206–222.

Wooldridge, J. M. (2015). *Introductory Econometrics: A Modern Approach*. Cengage Learning, 6 edition.

Yang, Z., Aydin, G., Babich, V., and Beil, D. (2009). Supply disruptions, asymmetric information, and a backup production option. *Management Sci.*, 55(2):192–209.

Zehender, K. (2015). Mercedes-Benz cars procurement & supplier quality strategy. https://www.daimler.com/dokumente/investoren/kapitalmarkttage/daimler-ir-mercedes-benzcarscapitalmarketdaydrklauszehender-20150611.pdf, Accessed July 31, 2018.

Zheng, Z., Xie, S., Dai, H.-N., Chen, W., Chen, X., Weng, J., and Imran, M. (2020). An overview on smart contracts: Challenges, advances and platforms. *Future Generation Computer Systems*, 105:475–491.

A Technical Appendix

Lemma A.1 (Ross and Peköz, 2007, Result 3.11, p.86). *Let Y be an arbitrary random variable, $\mathcal{F}_n$ be a filtration, and define: $Z_n = \mathbb{E}_n[Y]$. Z_n is a (Doob) martingale with respect to the filtration $\mathcal{F}_n$.*

Lemma A.2. *Suppose that information is symmetric. If $\mathbb{E}_0[V_n - C_n] \leq v_n^l$, then given $\varepsilon > 0$, there exists a holding cost $\underline{\mathcal{H}}$ such that if $\mathcal{H} \geq \underline{\mathcal{H}}$, then the probability of deviating from the JIT policy of Definition 1 is less than ε.*

Proof of Lemma A.2. We use the lemma's conditions to show that the automaker will carry either (a) cycle inventories or (b) speculative inventories with the probability of ε or less. (For definitions of the cycle and speculative inventories, see §4.2.)

For Part (a), we begin by assuming that the JIT policy picks and chooses to purchase $lq_{s,n}$ components – the contractually required minimum – in each period n. Gradient inequality and Assumption 1 yield

$$\frac{\partial r(lq_{s,n}, v_n)}{\partial q_{s,n}} \leq \frac{\partial r(lq_{s,n}, v_n^l)}{\partial q_{s,n}} + \underbrace{\frac{\partial^2 r(lq_{s,n}, v_n)}{\partial q_{s,n} \partial v_n}}_{\equiv M(q_{s,n}, v_n)} \Bigg|_{v_n = v_n^l} \left(v_n - v_n^l\right). \tag{A.1}$$

Above, $M(\cdot) > 0$ because of Assumption 1. Also, $M(\cdot) = \left(1 + \frac{1}{\nu(x_n, v_n)}\right) \frac{\partial}{\partial v_n} g(x_n, v_n) - \frac{g(x_n, v_n) \frac{\partial}{\partial v_n} \nu(x_n, v_n)}{\nu^2(x_n, v_n)} < \left(1 + \frac{1}{\nu(x_n, v_n)}\right) \frac{\partial}{\partial v_n} g(x_n, v_n) \leq \frac{\partial}{\partial v_n} g\left(x_n, v_n\right) \leq 1$.

Next, subtract lc_n from both sides of (A.1) to obtain

$$\frac{\partial r(lq_{s,n}, v_n)}{\partial q_{s,n}} - lc_n \leq \frac{\partial r(lq_{s,n}, v_n^l)}{\partial q_{s,n}} + l \underbrace{\frac{\partial^2 r(q_{s,n}, v_n)}{\partial q_{s,n} \partial v_n}}_{\equiv M(q_{s,n}, v_n)} \Bigg|_{v_n = v_n^l} \left(v_n - v_n^l\right) - lc_n,$$

Finally, by taking expectation over V_n and C_n, we get

$$0 \leq \mathbb{E}_0 \left[\frac{\partial r(lq_{s,n}, V_n)}{\partial q_{s,n}} - lC_n\right] \leq \frac{\partial r(lq_{s,n}, v_n^l)}{\partial q_{s,n}} + l\mathbb{E}_0 \left[V_n - C_n\right] - lv_n^l. \tag{A.2}$$

The left side of (A.2) is a first-order condition with respect to $q_{s,n}$ when $y_{s,n} = 0$. The left side must be non-negative for all values of $q_{s,n}$ that the automaker can possibly pick in equilibrium. So, the right side of (A.2) must be non-negative also. Given that $\mathbb{E}_0[V_n - C_n] \leq v_n^l$, it must be the case that $\left(\partial r(lq_{s,n}, v_n^l)/\partial q_{s,n}\right) \geq 0$. Positive marginal revenue evaluated at $V_n = v_n^l$ implies that the automaker has an incentive to output at least $lq_{s,n}$ units in each period n, even when the market potential is at its lowest possible level, i.e., she accumulates no cycle inventories.

For Part (b), let I denote speculative inventory that the automaker might add in some period $k < n$

to use it in period n. Adding inventory pays if

$$\frac{\partial}{\partial I}\left(\mathbb{E}_k\left[\psi(\mathbf{a}_{s,n},\mathbf{S}_{s,n})\right]\Big|_{\mathbf{a}_{s,n}=(y_{s,n},lq_{s,n}+I)}\right)>0.$$

The above inequality can be written as follows:

$$-(n-k)\,\mathcal{H}+Z_k>0,$$

where $(n-k)\,\mathcal{H}$ is the marginal holding cost and Z_k is the remainder that does not contain the marginal holding cost, $(n-k)\,\mathcal{H}$. Using the expressions above, adding speculative inventory, I, in period $k<n$ is profitable whenever

$$Z_k\geq(n-k)\,\mathcal{H},\quad\forall k=1,2,\ldots,(n-1),\ n=1,2,\ldots,\bar{n},$$

where Z_k, $k=1,2,\ldots,(n-1)$ is a (Doob) martingale as per Lemma A.1. Azuma's martingale inequality (Ross and Peköz, 2007, Theorem 3.24) implies

$$\mathbb{P}\{Z_k\geq(n-k)\mathcal{H}\}\leq\exp\left\{-\frac{2\,(n-k)^2\,\mathcal{H}^2}{k\delta_n^2}\right\}\leq\exp\left\{-\frac{2\mathcal{H}^2}{(n-1)\delta_n^2}\right\},\tag{A.3}$$

where $\delta_n\equiv L(\mathbf{a}_{s,n},\mathbf{S}_{s,n}^{+})-L(\mathbf{a}_{s,n},\mathbf{S}_{s,n}^{-})$, $\mathbf{S}_{s,n}^{+}\equiv(V_n=v_n^l,C_n=c_n^h)$, and $\mathbf{S}_{s,n}^{-}\equiv(V_n=v_n^l,C_n=c_n^h)$. Seeing that the right side of (A.3) strictly decreases in $\mathcal{H}$, we define $\underline{\mathcal{H}}$ implicitly as follows:

$$\varepsilon\equiv\exp\left\{-\frac{2\underline{\mathcal{H}}^2}{(n-1)\delta_n}\right\},\ \forall n=1,2,\ldots,\bar{n}\quad\Rightarrow\quad\mathbb{P}\Big\{Z_k\geq(n-k)\mathcal{H}\Big\}\leq\varepsilon,\quad\forall\mathcal{H}\geq\underline{\mathcal{H}}.\tag{A.4}$$

□

Proof of Lemma 1. See Theorem 9.3 in Jehle and Reny (2001, §9.2.3). □

Lemma A.3. *Suppose the contract is MPS and the raw material is not exchange-traded. If the conditions of Lemma A.2 hold, then given $\varepsilon>0$, there exists a holding cost $\widetilde{\underline{\mathcal{H}}}$ such that if $\mathcal{H}\geq\widetilde{\underline{\mathcal{H}}}$, then the probability of deviating from the JIT policy is less than ε.*

Proof of Lemma A.3. With asymmetric information, the condition that prevents cycle inventories is $\mathbb{E}_0[V_n-\widetilde{C}_n]\leq v_n^l$. The condition holds whenever the condition of Lemma A.2 is satisfied because $\mathbb{E}_0[\widetilde{C}_n]\geq\mathbb{E}_0[C_n]$. To rule out speculative inventories, replace c_k and $\widetilde{c}_n(c_k)$. The expression for $\widetilde{c}_n(c_k)$ is given by Equation (11). Then, in Equation (A.4), replace $\underline{\mathcal{H}}$ with $\widetilde{\underline{\mathcal{H}}}$. □

Proof of Lemma 2. Theorem 9.3 in Jehle and Reny (2001, §9.2.3) implies that the supplier should earn

$\mathcal{K}_2$ on expectation. In the environment of §5, the supplier earns

$$\mathbb{E}_0\left[\sum_{n=1}^{\bar{n}}\left(y_{a,n}^* + lQ_{a,n}^*\right)\left(\widetilde{c}_n(C_n) - C_n + \frac{\tau_{a,n}}{y_{a,n}^* + lQ_{a,n}^*}\right)\right] = \sum_{n=1}^{\bar{n}}\left[\mathbb{E}_0\left[\widetilde{\Pi}_a^*(\mathbf{S}_{a,n})\right] + \tau_{a,n}^*\right].$$

The solution $\tau_{a,n}^*$ is the smallest positive constant for which the expression sums to $\mathcal{K}_2$, or more. □

Proof of Proposition 1. Omitted. □

Proof of Proposition 2. Let MPS_s and MPS_a denote the expected payoffs under the MPS contract when the raw material is and is not exchange-traded. When we write $MPS_a \to MPS_s$, we mean that the expected payoff of the non-exchange-traded contract converges to the expected payoff of the exchange-traded one. Clearly, if $MPS_s \approx MPS_a$, then $MPS_a \succsim FP$ because $MPSs \succsim FP$ as per Proposition 1.

In (8), $y_{a,n}^*$ increases if and only if $\widetilde{c}_n$ decreases. Moreover, $\left(\left(y_{a,n}^* \to y_{s,n}^*\right) \Leftrightarrow (\widetilde{c}_n \to c_n)\right) \Rightarrow (MPS_a \to MPS_s)$. In Part (i), we show that if γ increases, then $\left(y_{a,n}^* \to y_{s,n}^*\right)$; $\hat{\gamma}$ the level of γ for which $(MPS_a \approx FP)$.

Let $r'(y + lq, v_n, \gamma)$ be given by Equation (1b). The first-best top-up quantity, $y_{s,n}^*$ is given by: $y_{s,n}^*(q) = \max\{\min\{(h-l)q, y\}, 0\}$, where y solves

$$c_n = r'(y + lq, v_n, \gamma).$$

The supplier's most preferred quantity in (10) is given by: $y_{a,n}^*(q) = \max\{\min\{(h-l)q, y\}, 0\}$, where

$$c_n = r'(y + lq, v_n, \gamma) + (y + lq)\, r''(y + lq, v_n, \gamma).$$

Above, $r''(y + lq, v_n, \gamma)$ is the second derivative of the revenue function, $r'(\cdot)$, with respect to quantity. By concavity, $r''(y + lq, v_n, \gamma) \leq 0$. Supermodularity condition of Assumption 3 asserts $\partial r''(y + lq_{a,n}, v_n, \gamma)/\partial\gamma \geq 0$. So, as γ increases, $r''(y + lq_{a,n}, v_n, \gamma) \to 0$ and $y_{a,n}^* \to y_{s,n}^*$, completing the proof of Part (i). Part (ii) follows from the definition of elasticity. □

Proof of Proposition 3. The proof is included in §6 (main body of the paper). □

Proof of Proposition 4. We use an argument from Shaked and Shanthikumar (2007, §3). For Part (i), take fixed a and define functions $\phi_a^c \equiv (\widetilde{C}_{n,i} - a)^+$, $\phi_a^v \equiv (V_n - a)^+$, and $\phi_a \equiv (V_n + \widetilde{C}_{n,i} - a)^+$, $i = 1, 2$. All three functions are convex. The automaker's payoff function – can be written as a linear combination of the functions ϕ_a^c, ϕ_a^v, and ϕ_a for various choices of a's. Now, $[c_{n,1}^l, c_{n,1}^h] \subseteq [c_{n,2}^l, c_{n,2}^h]$ implies $\mathbb{E}[(\widetilde{C}_{n,1} - a)^+] \leq \mathbb{E}[(\widetilde{C}_{n,2} - a)^+]$ and $\mathbb{E}[(V_n + \widetilde{C}_{n,1} - a)^+] \leq \mathbb{E}[(V_n + \widetilde{C}_{n,2} - a)^+]$. Together, with the fact that $\mathbb{E}[\widetilde{C}_{n,1}] = \mathbb{E}[\widetilde{C}_{n,2}]$, the last two inequalities imply that the automaker's expected MPS contract payoff in market 2 is larger in the convex order than the expected MPS contract payoff in market 1.

So, per (Shaked and Shanthikumar, 2007, Theorem 3.A.1), $MPS_2 \succsim MPS_1$. Since the FP contract is the same in markets 1 and 2, transitivity implies $MPS_2 \succsim FP_2$, since $MPS_1 \succsim FP_1$. The variance comparison is a result (3.A.4) in Shaked and Shanthikumar (2007). The proof of Part (ii) is analogous to that of Part (i). □

Proof of Proposition 5. We derive the result for continuous costs by introducing a partition on $[0, c_n^{max}]$ such that $c_i = \Delta i$, where $\Delta = \frac{c_n^{max}}{m}$, $i = 0, 1, \ldots, m$ and then letting $m \to \infty$. Under this partition, we have a finite number of cost values $0 < c_1 < c_2 < \ldots < c_n^{max}$, which must satisfy the following (IR) and (IC) constraints:

$$(\widehat{c}_n(c_i) - c_i)y_{r,n}(c_i) \geq 0 \qquad \text{and} \qquad (\widehat{c}_n(c_i) - c_i)y_{r,n}(c_i) \geq (\widehat{c}_n(c_j) - c_i)y_{r,n}(c_j),\ \forall i \neq j.$$

This is a complicated problem, especially as m grows large: There are a total of m (IR) constraints and another $m(m-1)$ (IC) constraints. However, using a result from Maskin and Riley (1984), the problem reduces as follows:

$$(\widehat{c}_n(c_n^{max}) - c_n^{max})y_{r,n}(c_n^{max}) \geq 0, \quad (\widehat{c}_n(c_i) - c_i)y_{r,n}(c_i) \geq (\widehat{c}_n(c_{i+1}) - c_i)y_{r,n}(c_{i+1}), \quad \text{and} \quad y_{r,n}(c_i) \geq y_{r,n}(c_{i-1}),$$

for all $i = 1, 2, \ldots$. On top of that, Maskin and Riley (1984) show that all the above constraints are binding in equilibrium. Maskin and Riley (1984) result leads to the following expression on unit price:

$$\widehat{c}_n(c_i) = c_i + \frac{1}{y_{r,n}(c_i)} \sum_{k=i}^{m-1} y_{r,n}(c_{k+1})\Delta.$$

By letting $m \to \infty$, the above expression becomes

$$\widehat{c}_n(c_n) = c_n + \frac{1}{y_{r,n}(c_n)} \int_{c_n}^{c_n^{max}} y_{r,n}(x)dx,$$

which is (16a).

Next, we compute the automaker's cost of buying $y_{r,n}$ units if the supplier's cost is c_n and the unit price is $\widehat{c}_n$ and take expectation

$$\int_0^{c_n^{max}} z y_{r,n}(z)f(z) + \int_0^{c_n^{max}} \left(y_{r,n}(z) \frac{1}{y_{r,n}(z)} \int_z^{c_n^{max}} y_{r,n}(x)dx \right) f(z)dz.$$

By changing the order of integration, we obtain

$$\int_0^{c_n^{max}} z y_{r,n}(z)f(z)dz + \int_0^{c_n^{max}} \left(\int_z^{c_n^{max}} y_{r,n}(x)dx \right) f(z)dz = \int_0^{c_n^{max}} z y_{r,n}(z)f(z)dz + \int_0^{c_n^{max}} y_{r,n}(z) \left(\int_0^z f(x)dx \right) dz.$$

The above expression simplifies to

$$\int_0^{c_n^{max}} \left(z + \frac{1}{\lambda(z)} \right) y_{r,n}(z)f(z)dz,$$

where $\lambda(z) = f(z)/F(z)$. The last expression is the expected cost in (17). □

Proof of Corollary 1. The chosen terms maximize the automaker's expected payoff in the commitment stage while ensuring the supplier's participation. □

Proof of Lemma 3. Use the same argument as in the proof of Lemma 2. □

Lemma A.4. *Suppose the contract is the price-quantity menu of §8. If the conditions of Lemma A.2 hold, then given $\varepsilon > 0$, there exists a holding cost $\underline{\widetilde{\mathcal{H}}}$ such that if $\mathcal{H} \geq \underline{\widetilde{\mathcal{H}}}$, then the probability of deviating from the JIT policy is less than ε.*

Proof of Lemma A.4. With the second-best contract information, the condition that prevents cycle inventories is $\mathbb{E}_0[V_n - \widehat{C}_n] \leq v_n^l$. The condition holds whenever the condition of Lemma A.2 is satisfied because $\mathbb{E}_0[\widehat{C}_n] \geq \mathbb{E}_0[C_n]$. To rule out speculative inventories, replace c_k and $\widehat{c}_n(c_k)$. The expression for $\widehat{c}_n(c_k)$ is given by Equation (16a). Then, in Equation (A.4), replace $\underline{\mathcal{H}}$ with $\underline{\widehat{\mathcal{H}}}$. □

B Data Dictionary and Descriptives

Our final data set is the agglomeration of three distinct data sets of BMW's Contracts, Suppliers, and Raw Material Prices (see Section 7.1 in the paper). We describe the contents of these data sets in Table B1, along with descriptions, definitions, examples, and the sources of the variables used in Section 7. Tables B2 and B3 provide additional descriptive statistics.

Table B1: Data Dictionary.

Variable Name	Description	Example(s)	Source
File 1 - Contract Info			
Contract_id	Unique identifier per contract	C000001, C000002	BMW
Contract_type	MPS or FP	MPS, FP	BMW
MPS	1 if Contract_type = MPS; 0 if Contract_type = FP	1, 0	Authors
Commodity_id†	Commodity codename	STEEL01, STEEL02, ALU01	BMW
Opaque	Info. asymmetry [1-3]	3	BMW
Heterogeneous	Info. asymmetry based on Rauch (1999) [0/1]	1	Authors
YearQuarter	Year-Quarter contract was entered into	2013Q1	BMW
Component_id	Component name	MD Belt force limiter US F33	BMW
Series_Model	Automobile series into which component goes	Series 1/2, Series 3/4, Series 7	Authors
Elastic	Contracted part in a price-elastic model [0/1]	1	Authors
Plant	Manufacturing plant identifier	A1	BMW
Supplier_name†	Supplier name	ABC Inc.	BMW
RawMatCost	Total commodity cost to BMW in year [€ th.]	12,345.67	BMW
RawMatTransfer	$\sum_{c=1}^{C}\sum_{m=1}^{M} RawMatCost_{m,c,i}$	12,345.67	Authors
Tons	Total tonnage of material under contract [kg]	12,345.67	BMW
Tonnage	$ln(Tons)$	1.234	Authors
Tooling	$ln(TotalTransfer - RawMatTransfer)$	37,654.33	Authors
NumSuppliers	ln(# Suppliers per Component)	1.238	Authors
File 2 - Supplier Info			
Supplier_name†	Supplier name	ABC Inc.	BMW
Supplier_city	Supplier HQ city	Berlin	BMW
Supplier_country	Supplier HQ country	Germany	BMW
Supplier_TA	Total assets in year of contract start [€ mil.]	800.00	BvD Amadeus
lnTA	$ln(Supplier\_TA)$	6.68	Authors
Supplier_YearInc	Year of incorporation	1976	BvD Amadeus
Supplier_Age	Year of contract start - Supplier_YearInc	44	Authors
lnAge	$ln(Supplier\_Age)$	3.78	Authors
TotalTransfer	BMW's total transfer to supplier in year [€]	50,000	BMW
File 3 - Raw Material Prices			
Commodity_id†	Commodity codename	STEEL01, STEEL02, ALU01	BMW
Date	Month and year	Jan 2010	BMW; Index providers
Price	Price of commodity [€]	12,345.67	BMW; Index providers
PriceSupport	$\frac{MaxMonthlyPrice - MinMonthlyPrice}{AveragePrice}$	0.1625	Authors

Notes: Variables in *italics* are included in our econometric models. † denotes how files are linked.

Table B2: Further Descriptive Statistics.

Full Sample	Obs.	1%	5%	25%	50%	75%	95%	99%	Skewness	Kurtosis
(1) *MPS*	325,109	0.000	0.000	0.000	0.000	1.000	1.000	1.000	0.569	1.324
(2) *Opaque*	307,124	1.000	1.000	1.000	2.000	2.000	2.000	3.000	0.154	2.024
(3) *Heterogeneous*	259,668	0.000	0.000	0.000	1.000	1.000	1.000	1.000	-0.157	1.025
(4) *Elastic*	45,021	0.000	0.000	1.000	1.000	1.000	1.000	1.000	-1.311	2.720
(5) *Tooling*	324,862	12.345	14.710	16.947	18.085	19.248	19.836	20.688	-1.112	4.805
(6) *PriceSupport*	268,479	0.211	0.301	0.375	0.416	0.525	1.420	2.714	3.084	14.684
(7) *Tonnage*	308,139	-3.101	-0.693	3.303	6.003	8.551	11.815	13.805	-0.181	2.720
(8) *NumSuppliers*	325,109	0.000	0.000	0.000	1.099	1.946	3.434	4.533	0.771	3.030

$MPS = 0$	Obs.	1%	5%	25%	50%	75%	95%	99%	Skewness	Kurtosis
(2) *Opaque*	189,496	1.000	1.000	2.000	2.000	2.000	2.000	3.000	-0.582	6.347
(3) *Heterogeneous*	141,828	0.000	0.000	1.000	1.000	1.000	1.000	1.000	-2.675	8.154
(4) *Elastic*	7,166	0.000	0.000	0.000	1.000	1.000	1.000	1.000	-1.088	2.183
(5) *Tooling*	206,833	12.093	14.590	16.737	17.619	19.218	19.940	20.688	-0.833	4.373
(6) *PriceSupport*	155,316	0.211	0.304	0.415	0.431	0.963	1.820	3.137	2.291	9.238
(7) *Tonnage*	195,226	-1.661	0.759	4.165	6.730	9.041	11.942	13.770	-0.198	2.746
(8) *NumSuppliers*	207,058	0.000	0.000	0.693	1.386	2.197	3.638	4.533	0.597	2.805

$MPS = 1$	Obs.	1%	5%	25%	50%	75%	95%	99%	Skewness	Kurtosis
(2) *Opaque*	117,628	1.000	1.000	1.000	1.000	1.000	1.000	2.000	5.323	30.737
(3) *Heterogeneous*	117,840	0.000	0.000	0.000	0.000	0.000	1.000	1.000	2.591	7.715
(4) *Elastic*	37,855	0.000	0.000	1.000	1.000	1.000	1.000	1.000	-1.358	2.844
(5) *Tooling*	118,041	13.501	15.793	18.076	19.239	19.269	19.836	20.015	-2.007	9.280
(6) *PriceSupport*	113,163	0.211	0.215	0.366	0.375	0.412	0.466	1.074	4.494	31.337
(7) *Tonnage*	112,913	-3.863	-2.048	1.849	4.730	7.250	11.441	13.871	0.058	2.691
(8) *NumSuppliers*	118,051	0.000	0.000	0.000	0.693	1.386	2.708	3.296	1.026	3.417

Notes: Descriptive statistics for the variables used in our estimation models. The top chart contains descriptives for the full sample, the middle chart for the subsample where $MPS = 0$, and the bottom chart for the subsample where $MPS = 1$.

Table B3: Descriptives for Indicator Controls.

Plant	# FP	# MPS	Quarter	# FP	# MPS
Plant 1	156,396	88,868	2013Q1	62,666	28,416
Plant 2	74	83	2013Q2	7,025	5,219
Plant 3	3,134	3,045	2013Q3	9,640	5,790
Plant 4	2,789	1,852	2013Q4	6,151	5,772
Plant 5	3,505	1,126	2014Q1	58,215	30,726
Plant 6	12	6	2014Q2	10,780	6,293
Plant 7	21,190	13,025	2014Q3	11,534	5,616
Plant 8	72	0	2014Q4	9,202	6,881
Plant 9	12,037	6,798	2015Q1	12,065	7,528
Plant 10	731	309	2015Q2	5,770	4,440
Plant 11	3,903	1,261	2015Q3	7,980	4,926
Plant 12	3,215	1,678	2015Q4	6,030	6,444

Notes: The left chart contains MPS and FP contract counts by manufacturing plant, and the right chart contains counts for the quarter in which the contract was written.

C Robustness with 2SLS Analysis

We assess the robustness of our main results by implementing instruments in a two-stage least squares (2SLS) framework in this section and also in the next Appendix Section C.1. Although *Tooling* is in the hypothesized direction in each model in Table 4 in the paper, it only obtains strong significance in Models (I)–(II). As explained in §7, we attribute this to measurement error in *Tooling*. Because we are unable to capture each supplier's *true* R&D and tooling cost $\mathcal{K}$, this measurement error causes an attenuation bias: β_2 is biased towards zero and underestimates the effect of *Tooling* on contract choice (Wooldridge, 2015). After correcting for this measurement error, we expect the coefficient on *Tooling* to increase, suggesting that the "true" relationship between R&D and tooling costs is likely *stronger* than that estimated in our baseline models.

Before we proceed, we note that we do not use instrumental variables to imply a causal channel. Rather, this analysis is only meant to provide further confidence in our baseline results: absent this analysis, we nevertheless find evidence in line with our theory. Moreover, although our instruments statistically satisfy the instrument relevance and exclusion restriction conditions, and they are the best instruments available to us, they come with limitations.

We turn to the Bureau van Dijk Amadeus database to supplement our data set with additional supplier-level financial data that could function as instruments for *Tooling*. We match suppliers in our data set to firms in Amadeus based on the supplier's name and the city and country in which it is based. We can match approximately 82% of the suppliers in our data set to the Amadeus universe. Yet, because Amadeus does not contain complete financial information for all firms, this effectively reduces our sample size to 40% of the original.

We use suppliers' size (natural logarithm of total assets) and age (natural logarithm of the number of years since incorporation) as instruments for their R&D and tooling costs. To the extent that R&D and tooling costs differ across (1) larger/smaller firms and (2) older/younger firms, our instruments will be correlated with *Tooling* (instrument relevance). We also expect these variables to be uncorrelated with the error in Equation (14) (instrument exogeneity) for two reasons. First, contracts are *given* by BMW to the supplier, and suppliers do not receive different contracts just because they are larger and/or older. This was validated in discussions with the purchasing managers at BMW. Second, contract types are decided by BMW *prior* to suppliers being selected: conditional on having been invited to bid, suppliers are not offered different contract types based on their size, and age.[12] Nevertheless, because we have

[12] It could be argued that the size or age of a supplier may have some information on the stability of the company, which would affect the choice of contract. Yet, before the auction process even begins, automakers perform a supplier qualification screening to ensure that all potential suppliers are financially sound and able to meet technical specifications and quality standards (Beil, 2010). This is no different at BMW, where suppliers are screened before being invited to participate in the auction. If a company was financially unstable or incapable of engineering the necessary components, then it would not

Table C1: 2SLS Analysis of Contract Choice.

	(C.I)	(C.II)	(C.III)	(C.IV)	(C.V)	(C.VI)
Opaque	-0.374***		-0.149**		-0.149**	
	(0.044)		(0.065)		(0.065)	
Heterogeneous		-0.584***		-0.273***		-0.273***
		(0.041)		(0.100)		(0.100)
Elastic			0.003	0.017*	0.003	0.018*
			(0.004)	(0.012)	(0.005)	(0.013)
Tooling	0.017***	0.020***	0.011*	0.029*	0.014*	0.038**
	(0.006)	(0.007)	(0.009)	(0.023)	(0.010)	(0.020)
PriceSupport	-0.047	-0.112***	-0.082	-0.122**	-0.082	-0.122**
	(0.039)	(0.027)	(0.052)	(0.050)	(0.052)	(0.050)
Tonnage	0.001	0.002***	-0.002*	-0.003**	-0.002*	-0.003**
	(0.001)	(0.001)	(0.001)	(0.002)	(0.001)	(0.002)
NumSuppliers	0.011	0.017	0.055	0.037	0.056	0.038
	(0.036)	(0.030)	(0.045)	(0.055)	(0.045)	(0.056)
Quarter Indicators	Yes	Yes	Yes	Yes	Yes	Yes
Plant Indicators	Yes	Yes	Yes	Yes	Yes	Yes
Component Indicators	Yes	Yes	Yes	Yes	Yes	Yes
# Contracts	95,756	90,176	5,360	4,770	5,364	4,774
# Components	2,278	2,246	235	222	236	223
Within-R^2	0.365	0.551	0.206	0.318	0.205	0.314
First Stage F-statistic	70.81***	47.36***	18.07***	13.35***	21.78***	27.66***
First Stage Coefficients						
lnTA	0.435***	0.457***	0.409***	0.701***	0.377***	0.725***
lnAge	-0.274***	-0.203**	-0.087	0.431*		
Hansen J-statistic	0.135	2.876	2.740	0.653	-	-
Hansen J $\chi^2(1)p$-value	0.713	0.090	0.098	0.419	-	-

Notes: The dependent variable is *MPS*. *Tooling* is instrumented with *Size* and/or *Age*. Significance on *Opaque*, *Heterogeneous*, *Tooling*, *Elastic* is tested with one-sided hypothesis tests; *PriceSupport* with two-sided tests. Robust standard errors are in parentheses and are clustered at the component. $***p < 0.01$; $**p < 0.05$; $*p < 0.10$

more instruments (two) than endogenous regressors (one), we can statistically test the exogeneity of our variables and their validity as instruments.

Table C1 displays the results of our 2SLS models. Although we rely on a smaller subset of our initial data set, the results are still useful to assess whether we obtain evidence in line with our model after dealing with measurement error issues. Before analyzing the results, we remark on the validity of our instruments. The bottom of Table C1 shows that the first stages of Models V–X all have F statistics that are larger than the rule-of-thumb value of 10 (Staiger and Stock, 1997) and are generally greater than the critical values for estimation with two instruments (Stock and Yogo, 2005): our instruments are relevant and strong. Hansen J statistics are non-significant, further suggesting that our instruments are uncorrelated with the error and are exogenous. Therefore, we are confident in our instruments' validity and will proceed with the analysis.

be invited to bid in the auction in the first place. Thus, we are capturing that *conditional on having been invited to bid*, suppliers of different sizes or ages are not offered different contracts.

Models C.I–C.IV qualitatively mirror those from the baseline Models I–IV in Section 7. In Models C.V and C.VI, we re-estimate Models C.III and C.IV without *Age* as an instrument. All variables of interest are in the hypothesized direction, and they generally obtain statistical significance. Moreover, the extent to which the coefficients on *Tooling* in Models I–IV suffer from attenuation bias is readily visible. In Models C.I–C.VI, these coefficients are several times larger in magnitude than those from the baseline results in Table 4, providing evidence that the imprecise measurement of *Tooling* drove our coefficient estimates toward zero. After correcting for this measurement error, we obtain similar evidence as above that the supplier's R&D and tooling cost is a driver of contract choice. Similar to the models in Table 4, the inclusion of *Elastic* restricts our data set, so we lose some significance on *Tooling*. We thus attribute the instability of statistical significance in *Elastic* and *Tooling* to the nature of our data set and not the lack of empirical reality.

C.1 R&D Spending as an Instrument

The ideal instrument for our *Tooling* variable would be the supplier's investment in his capability to produce an individual component. However, such detailed information is difficult (if not impossible) to obtain. Barring that, another plausible measure would be related to R&D, such as R&D expenditure. A supplier's R&D spending is likely correlated with its tooling cost. However, the contract that BMW offers the supplier probably does not depend on how much the supplier spends on R&D, except for through R&D spending's direct effect on the supplier's tooling cost; R&D spending should be unrelated to the contract the supplier is offered. Therefore, as a robustness check for our main results, we use R&D expenditures as instruments for R&D and tooling costs by merging in data from the Amadeus database.

Before we proceed, we note one disadvantage of this approach. Unfortunately, even though Amadeus contains information on firm R&D, such data are only available for public companies: private firms are not required to disclose R&D expenditure information. In our data set, less than 4% of the suppliers are public, and so R&D expenditure information is relatively rare.

In Table C2, we re-estimate the models from above, this time using the natural logarithm of R&D expenditures (*lnRD*) as an instrument for *Tooling*. Reassuringly, our variable *Tooling* is in the hypothesized direction in all models and obtains significance in most, once again providing support for H.3. Yet, we note that Models C.IX and C.X are estimated on less than 0.2% of our baseline sample.

As a complementary approach and further robustness check, we also use *predicted* R&D expenditures to boost our sample size. Specifically, we first estimate a model where we regress *lnRD* on the natural logarithm of total assets, operating revenue scaled by assets, and a set of year- and firm-fixed effects (results available upon request). We then use the coefficients from this model to predict out-of-sample

Table C2: R&D Expenditure as an Instrument for R&D and Tooling Cost.

	(C.VII)	(C.VIII)	(C.IX)	(C.X)	(C.XI)	(C.XII)	(C.XIII)	(C.XIV)
Opaque	-0.284***		-0.265***		-0.369***		-0.148**	
	(0.049)		(0.090)		(0.044)		(0.065)	
Heterogeneous		-0.274***		-0.175*		-0.586***		-0.271***
		(0.066)		(0.113)		(0.042)		(0.100)
Elastic			-0.005	0.009			0.001	0.015
			(0.012)	(0.009)			(0.005)	(0.014)
Tooling	0.292**	0.453**	0.062	0.081**	0.011***	0.010**	0.013*	0.028**
	(0.132)	(0.261)	(0.247)	(0.042)	(0.004)	(0.005)	(0.009)	(0.017)
PriceSupport	-0.104***	-0.231***	-0.132***	-0.220***	-0.049	-0.112***	-0.082	-0.120**
	(0.022)	(0.012)	(0.030)	(0.035)	(0.039)	(0.027)	(0.051)	(0.050)
Tonnage	0.001	0.001	-0.005	0.004	0.001	0.002***	-0.002*	-0.003*
	(0.002)	(0.002)	(0.004)	(0.006)	(0.001)	(0.001)	(0.001)	(0.002)
NumSuppliers	0.743*	0.823			-0.001	0.003	0.057	0.034
	(0.447)	(0.748)			(0.039)	(0.029)	(0.046)	(0.055)
Quarter Indicators	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Plant Indicators	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Component Indicators	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
# Contracts	6,648	6,068	453	369	92,557	87,929	5,188	4,623
# Components	231	215	30	21	2,215	2,184	219	208
Within-R^2	0.455	0.419	0.550	0.480	0.371	0.563	0.206	0.317
First Stage F-statistic	3.82*	2.84*	11.31***	19.70***	114.40***	97.38***	19.58***	34.73***
First Stage Coefficients								
lnRD	0.174*	0.117*	-0.504***	-0.641***				
lnRD_predicted					0.548***	0.595***	0.502***	1.075***

Notes: The dependent variable is *MPS*. *Tooling* is instrumented with R&D expenditures in Models (C.VII)–(C.X), with predicted R&D expenditures in Models (C.X)–(C.XIV). Significance on *Opaque*, *Heterogeneous*, *Tooling*, *Elastic* is tested with one-sided hypothesis tests; *PriceSupport* with two-sided tests. Robust standard errors are in parentheses and are clustered at the component. ***$p < 0.01$; **$p < 0.05$; *$p < 0.10$

R&D expenditures (*lnRD_predicted*). Finally, we use *lnRD_predicted* as an instrument for *Tooling* in Models C.XI–C.XIV. The drawback of this approach is that now more than 90% of our estimation sample is predicted, and so we urge caution in drawing conclusions from these results. Nevertheless, the results provide evidence in favor of H.3.

Altogether, when we use R&D expenditures as an instrument for *Tooling*, we again find evidence that supplier R&D and tooling costs are associated with a greater proportion of MPS contracts.

D Comparing Probit and LPM Models

In Section 7, we estimate linear probability models (LPM). This is driven by the model's ability to control for a large number of component fixed effects. Since component fixed effects perfectly predict contract type for a large number of contracts, a probit (or logit) model would exclude those component groups which do not receive both contract types, leading to estimation on a severely restricted sample (approximately 23% of our original sample).

In this section, we validate the results of our linear model by comparing its performance to analogous probit models. Specifically, we estimate several probits and LPMs without and with component effects. These auxiliary estimations are reported in Tables D1 and D2.

Table D1: Probit and LPM with Component Indicators/FE.

	(D.I) Probit	(D.II) Probit	(D.III) LPM	(D.IV) LPM
Opaque	-0.341***		-0.296***	
	(0.020)		(0.043)	
Heterogeneous		-0.419***		-0.503***
		(0.047)		(0.036)
Tooling	0.008**	0.007**	0.006***	0.004**
	(0.004)	(0.004)	(0.002)	(0.002)
PriceSupport	-0.226***	-0.441***	-0.057*	-0.111***
	(0.076)	(0.072)	(0.032)	(0.021)
Tonnage	0.005***	0.006***	0.0009*	0.0014***
	(0.001)	(0.002)	(0.0005)	(0.0005)
NumSuppliers	-0.130**	-0.156**	-0.028**	-0.021
	(0.052)	(0.066)	(0.014)	(0.014)
Constant	3.947***	0.395	0.840***	0.706***
	(1.037)	(0.983)	(0.065)	(0.045)
Quarter Indicators	Yes	Yes	Yes	Yes
Plant Indicators	Yes	Yes	Yes	Yes
Component Indicators/FE	Yes	Yes	Yes	Yes
# Contracts	58,181	54,564	247,167	237,195
# Components	612	687	3,741	3,656
Pseudo-R^2	0.764	0.739		
log-likelihood	-9,342	-9,810		
Within-Component R^2			0.280	0.465
% Correctly Classified	92.7	92.0	93.1	91.2
True Positive	23,427	23,036	105,880	102,141
True Negative	30,530	27,163	124,170	114,255
False Positive	3,170	2,906	13,650	13,281
False Negative	1,054	1,459	3,467	7,518

Notes: The dependent variable is *MPS*. Significance on *Opaque*, *Heterogeneous*, *Tooling*, *Elastic* is tested with one-sided hypothesis tests; *PriceSupport* with two-sided tests. Estimates for probit Models D.I and D.II are average marginal effects. Robust standard errors are in parentheses and are clustered at the component. Classification statistics are against a cut-off greater than or equal to the proportion of MPS contracts in the sample. ***$p < 0.01$; **$p < 0.05$; *$p < 0.10$

In Table D1, we estimate probit models with component indicators (Models D.I–II) and compare these against LPMs with component fixed effects (Models D.III–IV). Across the different models, we find that the probit and LPM models are congruous: coefficients on the variables testing H.1 (*Opaque* and *Heterogeneous*), H.3 (*Tooling*), and H.4 (*PriceSupport*) are in the same, hypothesized direction and obtain similar levels of statistical significance. This provides evidence that estimates from the LPMs are directionally and statistically valid. However, as noted above, there are two disadvantages to the probit models. First, more than 75% of our sample drops out of these models due to the problem of perfect prediction within components. Second, and related, there are too few observations to estimate models, including the variable *Elastic*, and so we cannot test H.2.

Therefore, as a second check, we also compare in Table D2 probits without component indicators (Models D.V–VI), LPMs without component fixed effects (Models D.VII–VIII), and LPMs with component fixed effects (Models D.IX–X; also reported in our baseline Table 4 in the paper). Across Models D.V–VIII, coefficient estimates are at similar levels of significance and in the same direction (except

Table D2: Probit and LPM Comparison.

	(D.V) Probit	(D.VI) Probit	(D.VII) LPM	(D.VIII) LPM	(D.IX) LPM	(D.X) LPM
Opaque	-0.118***		-0.456***		-0.084***	
	(0.021)		(0.083)		(0.032)	
Heterogeneous		-0.130***		-0.737***		-0.189***
		(0.023)		(0.062)		(0.060)
Elastic	0.003	0.004	0.023*	0.024*	0.007**	0.008**
	(0.007)	(0.006)	(0.017)	(0.015)	(0.003)	(0.004)
Tooling	0.009***	0.006**	0.065***	0.011	0.004*	0.004*
	(0.003)	(0.003)	(0.013)	(0.010)	(0.003)	(0.003)
PriceSupport	-0.396***	-0.330***	-0.338**	-0.249***	-0.154***	-0.191***
	(0.075)	(0.061)	(0.132)	(0.082)	(0.049)	(0.051)
Tonnage	-0.006***	-0.005***	-0.013***	-0.006***	-0.0003	-0.0003
	(0.002)	(0.002)	(0.003)	(0.002)	(0.0002)	(0.0002)
NumSuppliers	-0.012**	-0.014***	(-0.098***)	(-0.098***)	(0.011)	(0.012)
	(0.005)	(0.005)	(0.034)	(0.036)	(0.014)	(0.012)
Constant			0.360	0.884***	0.945***	0.865***
			(0.243)	(0.197)	(0.060)	(0.070)
Quarter Indicators	Yes	Yes	Yes	Yes	Yes	Yes
Plant Indicators	Yes	Yes	Yes	Yes	Yes	Yes
Component Indicators/FE	No	No	No	No	Yes	Yes
# Contracts	39,392	38,845	39,646	39,098	39,646	39,098
# Components	477	451	479	453	479	453
Pseudo-R^2	0.823	0.814				
log-likelihood	-2,787	-2720				
Within-Component R^2			0.188	0.165	0.167	0.294
% Correctly Classified	98.3	98.4	91.5	92.5	97.9	98.5
True Positive	33,645	33,691	30,862	31,343	33,741	33,899
True Negative	5,092	4,533	5,415	4,820	5,079	4,607
False Positive	297	309	228	275	564	488
False Negative	358	312	3,141	2,660	262	104

Notes: The dependent variable is *MPS*. Significance on *Opaque*, *Heterogeneous*, *Tooling*, *Elastic* is tested with one-sided hypothesis tests; *PriceSupport* with two-sided tests. Estimates for Models D.V and D.VI are average marginal effects. Robust standard errors are in parentheses and are clustered at the component. Classification statistics are against a cut-off greater than or equal to the proportion of MPS contracts in the sample. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$

for *Elastic*, which is anyway not statistically different from 0). Once again, estimates from analogous probits and LPMs provide evidence that the LPM estimator is valid.

Models D.IX and D.X build on the previous LPM models by including component fixed effects. We first note that the coefficients on all variables between models D.VII–VIII and D.IX–X change substantially, evidence that components provide important variation that should be controlled for and that estimates absent component controls may be biased. We also find that the coefficient on *Elastic* has increased and is statistically significant. Across component groups, price elasticity is not associated with different contract types. However, after controlling for component-specific heterogeneity, price-elasticity (*Elastic*) is associated with a higher incidence of MPS contracts.

Altogether, the results provide evidence that LPM estimators are (1) necessary to control for component effects and remove bias from these and (2) valid.

E Including Supplier Fixed Effects

Suppliers may supply multiple components to BMW. By augmenting our models with supplier fixed effects, we can analyze the *within*-supplier variance in contract type. In our context, this means two things.

On the one hand, this allows us to test whether – for the *same* supplier – a supplier receives different contracts depending on the components or raw materials they are supplying. If (again, for the same supplier) information asymmetries in raw materials, price elasticity, and raw material price variability are significantly associated with contract type, this would provide further evidence that these factors drive contracting choices.

On the other hand, however, because R&D and tooling costs are also at the level of the supplier, the effect of tooling costs on contract choice will partially be absorbed by the suppliers' fixed effects. Because the supplier fixed effects would capture similar information as our tooling variable, we should not expect significance on $Tooling$.

We estimate a similar model to Equation 14 from the paper, this time including supplier fixed effects, and Table E1 displays the results. We find that information asymmetry (H.1) and raw material price variability (H.4) are negatively and significantly associated with MPS contracts. Also, although price elasticity (H.2) is significant in only one of the models, it is in the hypothesized direction in both. The results of these models are, therefore, qualitatively similar to those from Tables 4 and C1 in the paper. We also note that, as expected, the effects of tooling costs on contract type are absorbed by the supplier fixed effects. The results overall suggest that *even for the same supplier*, information asymmetry, price elasticity, and raw material price variability are strong predictors of BMW's contract choice.

Table E1: Models with Supplier Fixed Effects.

	(E.I)	(E.II)	(E.III)	(E.IV)
Opaque	-0.287***		-0.081**	
	(0.049)		(0.044)	
Heterogeneous		-0.504***		-0.182**
		(0.049)		(0.091)
Elastic			0.005*	0.003
			(0.003)	(0.003)
Tooling	0.003	0.003	0.004	0.008
	(0.006)	(0.006)	(0.004)	(0.007)
PriceSupport	-0.061	-0.114***	-0.151*	-0.189**
	(0.038)	(0.032)	(0.080)	(0.075)
Tonnage	0.0005	0.0009**	-0.0003	-0.0003
	(0.0004)	(0.0004)	(0.0003)	(0.0003)
NumSuppliers	-0.022*	(-0.009)	(-0.002)	(-0.0002)
	(0.012)	(0.012)	(0.002)	(0.003)
Constant	0.876***	0.715***	0.928***	0.797***
	(0.117)	(0.112)	(0.105)	(0.136)
Quarter Indicators	Yes	Yes	Yes	Yes
Plant Indicators	Yes	Yes	Yes	Yes
Component Indicators	Yes	Yes	Yes	Yes
Supplier FE	Yes	Yes	Yes	Yes
# Contracts	246,787	236,840	39,558	39,020
# Components	1,259	1,234	169	154
Within-R^2	0.291	0.480	0.163	0.287

Notes: The dependent variable is MPS. Significance on *Opaque*, *Heterogeneous*, *Tooling*, *Elastic* is tested with one-sided hypothesis tests; *PriceSupport* with two-sided tests. Robust standard errors are in parentheses and are clustered at the component. $^{***}p < 0.01$; $^{**}p < 0.05$; $^{*}p < 0.10$