Bargaining Power in the Market for Ideas*

Li Guo† Yang Yu‡

August 12, 2026

Abstract

We study the division of total surplus between buyers and sellers in patent transactions and its impact on the development of ideas. Using comprehensive data on U.S. patent transactions from 1975 to 2012, we estimate gains from trade for both parties based on how their stock prices react to patent transaction news. Additionally, we quantify bargaining shares as the proportion of their gains from trade within the total surplus. We find that firms' bargaining shares are heterogeneous and proportional to their size relative to their counterparties, starkly contrasting to the conventional assumption of homogeneous and fixed bargaining shares in the literature. We further show that the relationship between relative size and bargaining share is stronger among firms with weaker corporate governance. To study the implications of these findings, we develop a model with size-dependent bargaining power in the patent market. Our analysis shows that the substantial bargaining power of large firms diminishes the incentive for innovation among smaller firms.

Keywords: R&D, Patent trade, Bargaining share, Market power.

JEL Classification: O31, O34, L24, C78, G34.

*Accepted by *Management Science*. We are grateful to the Department Editor, Stylianos Kavadias, the associate editor, and three anonymous referees for constructive comments that substantially improved the paper. We also thank Linyi Cao, Jeremy Greenwood, Suting Hong, Dirk Krueger, and Yueyuan Ma, as well as conference and seminar participants, for helpful comments and suggestions. Yang Yu gratefully acknowledges financial support from the National Natural Science Foundation of China (Grant Nos. 72495153 and 72522013). All errors are our own.

†Li Guo: Department of Economics, University of Pennsylvania, US. gary16@sas.upenn.edu; ‡Yang Yu: Antai College of Economics and Management, Shanghai Jiao Tong University, China. yu.yang.econ@sjtu.edu.cn.

1 Introduction

New ideas are pivotal drivers of technological progress. However, those who generate ideas often differ from those who implement them. For instance, in April 2012, America Online (AOL) sold over 800 patents—covering areas such as advertising, search, content management, social networking, and multimedia—to Microsoft for $1.1 billion. This transaction accelerated Microsoft's technological development and enhanced its competitive edge in critical areas like search engines and ad targeting. Meanwhile, it allowed AOL to monetize intellectual properties underutilized within its core business. This transaction is not an anomaly. As Akcigit et al. (2016) documented, approximately 16% of U.S. patents registered between 1976 and 2006 were traded, a trend that continues to grow (Gans and Stern, 2000; Arora and Ceccagnoli, 2006; Galasso and Schankerman, 2010; De Marco et al., 2017; Ma, 2022). The key finding of the literature is that patent transactions create value by placing patents in the hands of firms that can better utilize them.

How large are the surpluses in the transaction of ideas, and what factors influence their magnitude? How do patent sellers and buyers share the surplus from mutually beneficial transactions? Understanding these questions has critical implications for fostering innovation and allocating ideas effectively. This paper aims to answer these questions and study their impact on technological growth.

By integrating United States Patent and Trademark Office (USPTO) data, news content extracted through Large Language Models, and the stock market data, we estimate buyer- and seller-side gains from trade using stock price reactions to patent transaction news for 26,509 patent trades between 1975 and 2012.[1] Notably, we observe a surge in trading volume around the patent transaction news dates, indicating that financial investors pay attention to and trade on this information. Hence, it is reasonable to posit that stock prices are informative about patent trades. To demonstrate the relationship between stock prices and gains from patent transactions, consider the case of AOL and Microsoft. On April 9th, 2012, they announced an agreement to trade a portfolio of patents. The market responded favorably, driving AOL's share prices up by 43% and Microsoft's by 5.2%. These price movements reflect both firms' gain

[1]We focus on patents traded as standalone assets and exclude patent transactions associated with mergers and acquisitions from our analysis.

from trade—precisely, each party's gain from trade, defined as the improvement in firm value from completing the transaction after accounting for the transaction price and transaction costs. Unlike traditional methods, our approach does not rely on knowing the exact value of the patent for either party or transaction prices; instead, under the efficient market hypothesis, we infer gains from trade from investor-perceived changes in firm value reflected in stock price reactions to patent transaction news. We compute the total surplus as the total gains from trade for both the seller and the buyer and measure their bargaining shares as the proportion of their gains within the total surplus.

With the constructed measures for total surplus and bargaining shares in patent trades, we uncover a set of novel findings. We document that bargaining shares exhibit significant heterogeneity across trades. Specifically, the distribution of bargaining shares is concentrated at both ends of the spectrum for both sellers and buyers, suggesting that the division of total surplus overwhelmingly favors one side of the trade rather than being evenly split. While the conventional literature assumes a fixed bargaining share in Nash bargaining between buyers and sellers (e.g., Akcigit et al., 2016), our empirical analysis reveals substantial dispersion in bargaining shares across patent transactions.

To understand this heterogeneity, we document that traders' bargaining shares are proportional to their relative size. Specifically, the output ratio between trading partners alone accounts for about eighty percent of the variation in bargaining shares across patent trades, implying that large traders disproportionately reap more surplus than their smaller counterparties.[2] Our findings suggest that managers of small firms should not assume that patents misaligned with their core business can be easily monetized through sales, as limited bargaining power may leave them able to capture only a small fraction of the surplus.

One institutional factor that may strengthen the size dependence of bargaining shares is incentive misalignment between shareholders and the managers responsible for patent transactions. Patent transactions often involve managerial discretion over negotiation, timing, and counterparty choice. When shareholder control is weaker, managers have greater scope to accept terms that do not maximize shareholder value. Smaller firms may be especially vulnerable

[2] A practical implication of this finding is that when direct estimation of bargaining shares is infeasible, calibrating them using relative sizes offers a reasonable alternative.

because they have fewer legal and negotiating resources to discipline managerial behavior and counter better-resourced counterparties. Weak governance may therefore make bargaining shares more sensitive to relative firm size. Using the Entrenchment Index (Bebchuk et al., 2009), we find that the positive relationship between relative size and bargaining share is stronger among firms with weaker shareholder control.

While the preceding results focus on how surplus is divided between trading partners, an innovating firm's payoff from developing a patent that may later be sold depends not only on its bargaining share, but also on the total surplus a transaction generates. We therefore examine how firm size relates to total surplus. We find that total surplus in patent transactions increases with the sizes of both buyers and sellers, with the effect being more pronounced for buyer size. The positive relationship between buyer size and total surplus aligns with a scaling effect: larger firms are able to leverage patents on a grander scale, thereby enhancing the overall gains. Meanwhile, the positive association between seller size and total surplus can be driven by lower transaction costs associated with large sellers due to their greater organizational efficiency.

Together, these findings create a tension for small firms. Trading with a larger counterparty can generate more surplus, but the larger firm also captures a greater share of that surplus. The empirical results alone do not reveal which force dominates a small firm's ex ante incentive to innovate. To study this trade-off, we develop a theoretical model and calibrate it with our empirical estimates. In the model, firms decide whether to develop a new blueprint based on their expectations about the technological propinquity between the blueprint and their core business, as well as the blueprint's market value in a secondary market. Our results indicate that the secondary market for ideas not only reallocates patents to their most efficient uses (Akcigit et al., 2016), but also enhances firms' incentives to engage in R&D by mitigating uncertainties inherent in the innovation process (Fleming, 2001). However, due to their limited bargaining power, smaller firms derive significantly fewer benefits from this market than larger firms. Compared to a counterfactual scenario with symmetric bargaining shares, size-dependent bargaining weakens the R&D incentives of smaller firms, potentially stifling creative destruction and impeding technological growth.

Our study is closely related to the literature on the secondary market for patents. Prior research has explored the conditions under which patents are more likely to be traded and the characteristics that influence these transactions (Serrano, 2010; Akcigit et al., 2016; Hegde and Luo, 2018; Figueroa and Serrano, 2019; Kwon et al., 2022; Arora et al., 2022; Ma and Yang, 2023).[3] Other work has examined the link between patent transactions and litigation (Galasso and Schankerman, 2010; Galasso et al., 2013), as well as the rising trend in patent exchanges (Gans and Stern, 2000; Arora and Ceccagnoli, 2006; De Marco et al., 2017; Hegde and Luo, 2018; Ma, 2022). Our paper is particularly related to Akcigit et al. (2016), who demonstrate that the patent market enhances social welfare by reallocating patents to their most efficient uses. Our contribution is to show that large firms have a disproportionately high bargaining share in patent trades, which discourages the development and implementation of ideas by small firms and may impede technological growth, especially if smaller firms exhibit higher R&D efficiency. Another closely related study is Serrano (2018), who estimates the total surplus from patent trades using a structural model of sales and renewals focused on patents' productive value. In contrast, we adopt a stock-market-based approach to measure gains from trade from stock market reactions to transaction news. This approach is designed to capture a broader measure of value, including strategic and reallocation benefits net of transaction costs, as assessed by investors. Importantly, our method allows us to decompose the total surplus and quantify its division between buyers and sellers, which is central to our findings on bargaining power.

Furthermore, our research intersects with studies on the varying motivations for innovation among firms of different sizes (Bena et al., 2016; Arora et al., 2018; Akcigit and Kerr, 2018; Hochberg et al., 2018; Farre-Mensa et al., 2020; Ma, 2022; Bustamante and Zucchi, 2023; Arora et al., 2023; Serrano and Ziedonis, 2024), as well as literature examining the role of patents in driving the innovation process (Fleming and Sorenson, 2004; Ziedonis and Galasso, 2019; Arora et al., 2023; Marx and Scharfmann, 2024). Our contribution is to show how size influences the incentives for innovation by affecting bargaining power within the patent market.

The remainder of the paper proceeds as follows. Section 2 describes the measurement of the bargaining share and the total surplus from the patent transaction. Section 3 presents the empirical facts. Section 4 develops a model. Section 5 discusses key assumptions underlying our

[3]Prior research suggests that patent transactions are more likely to occur when the parties share closer technological and geographic proximity, the patent holds high scientific and economic value, information is disclosed at an earlier stage, and the transaction involves smaller or science-based inventors. Conversely, transactions are less likely to occur when there is significant product-market overlap or when the original holder has strong technological capabilities.

empirical identification strategy and the limitations they imply. Finally, Section 6 concludes.

2 Measuring Gains from Patent Transactions and Bargaining Shares

In this section, we estimate two key components of patent transactions: the total surplus and the bargaining share. The total surplus represents the combined gains from trade received by both the buyer and seller, net of any trade costs. Specifically, the buyer's gain from trade is the difference between the private value of the patent to the buyer and the transaction price of the patent. Conversely, the seller's gain from trade is the difference between the transaction price of the patent and the private value of the patent to the seller. The bargaining share refers to the portion of the total surplus attributed either to the buyer's or the seller's gain from trade.

2.1 Data

We begin by describing the database we use and the process for matching these data.

USPTO Patent Assignment Dataset (UPAD). The core of our empirical analysis relies on UPAD, which contains detailed records of approximately 6 million patent-asset conveyances recorded at the USPTO from 1970 to 2014. This dataset provides information on the parties involved in each conveyance, the execution and recording dates, the relevant patent and application identification numbers, and the types of conveyances. Conveyance types indicate whether a transaction resulted from trade, merger, mortgage, license, joint research agreement, change of name, government interest assignment, employment assignment, among other reasons (Graham et al., 2018).

We restrict our sample to patents traded as standalone assets by retaining patent transactions categorized as *trade*, while discarding those associated with other types of conveyance.

To connect these patent transactions with firms' economic fundamentals and stock prices, we obtain GVKEYs (Global Company Key identifiers) for both buyers and sellers of patent trades between 1975 and 2012 from Arqué-Castells and Spulber (2022, 2023). Using these GVKEYs, we identify 361,507 patent transactions where both parties are publicly listed companies. To ensure that our sample does not involve any mergers and acquisitions, we integrate this sample with the SDC Platinum Mergers & Acquisitions Database by matching assignees and assignors in

UPAD with acquirers and targets in SDC, respectively. To account for cases where an acquirer retains the target's name following a merger, we also perform a reverse match, linking assignors to acquirers and assignees to targets. Transactions involving patent conveyance that occur within a two-year window around an acquisition are excluded.

Compustat North America Fundamentals. We rely on the Compustat database to access the fundamentals of both sellers and buyers, including information such as sales and industry affiliation.

Center for Research in Security Prices (CRSP) Database. We use the CRSP database to analyze market reactions to patent transaction news. The main variables we use include stock price, turnover rate, and market capitalization. To link patent transaction data with financial market information, we merge the CRSP database with a concordance table that maps PERMNOs (unique identifiers for securities) to GVKEYs provided by Wharton Research Data Services (WRDS).

Private Value of Patents. We obtain the initial value of patents (i.e., their value at the time of invention) from Kogan et al. (2017). Their dataset integrates stock market data with patent information for U.S. firms since 1926. Specifically, they analyze the stock market response within a three-day window following a firm's issuance of a new patent, while controlling for market returns during that period. Additionally, their dataset includes details on patent classes (CPC) and forward citations. To link this data, we utilize unique patent identifiers provided by the USPTO.

Corporate Governance Indices. To assess the degree of incentive misalignment between managers and investors, we employ a standard corporate governance index: the Entrenchment Index (E-index) from Bebchuk et al. (2009). It focuses on six provisions particularly linked to managerial entrenchment, such as poison pills and supermajority requirements, with higher values indicating weaker shareholder rights. The index is derived from governance provisions cataloged by the Investor Responsibility Research Center (IRRC).

Merged data. We merge the datasets described above, resulting in a final dataset of 26,509 standalone patent transactions between publicly listed companies that contain stock prices within the event window (as described in Section 2.3), turnover rates, sales, NAICS, and initial patent values. It serves as the primary sample for our empirical analysis, with 6,642 occurring within the same four-digit NAICS industry and the rest involving cross-industry transactions.

2.2 Empirical Strategy

To fix ideas, we first define the key variables and concepts. Consider a patent transaction denoted by j, involving a buyer firm denoted by b and a seller firm denoted by s. Each firm's valuation of the patent is denoted $v_{i,j}$ for $i \in \{b, s\}$, defined as the net change in firm value attributable to holding the patent, measured before accounting for transaction costs or the monetary transfer associated with the transaction. This valuation encompasses the direct profit from utilizing the patent in production, the strategic value it confers to the firm, and the option value of potential future resale.

Formally, let $d_{i,j} \in \{0, 1\}$ denote whether firm $i \in \{b, s\}$ holds patent j. We define each party's valuation as the marginal contribution of the patent to firm value:

$$v_{i,j} = V_i(d_{i,j} = 1) - V_i(d_{i,j} = 0), \qquad i \in \{b, s\}.$$

Here, $V_i(\cdot)$ represents the total value of firm i in a given state, which is forward-looking and incorporates strategic considerations and outside option values. For the buyer, $V_b(d_{b,j} = 0)$ reflects the value of not acquiring the patent (or acquiring an alternative), while $V_b(d_{b,j} = 1)$ includes the patent's productive value and the option value of future resale. For the seller, $V_s(d_{s,j} = 1)$ reflects the patent's productive value and the value of opportunities to sell to other potential buyers, whereas $V_s(d_{s,j} = 0)$ captures the firm's value after the patent is sold.

The transaction price of the patent, p_j, determines the gains from trade for the two firms. For the buyer, the gain is expressed as $\xi_{b,j} = v_{b,j} - p_j - tc_{b,j}$, and for the seller, it is given by $\xi_{s,j} = p_j - v_{s,j} - tc_{s,j}$. Here, $tc_{b,j}$ and $tc_{s,j}$ represent the transaction costs incurred by the buyer and the seller, respectively.

The total surplus of the patent transaction is the sum of these gains from trade: $TS_j = \xi_{b,j} + \xi_{s,j}$, which is equal to the difference between the two firms' valuations of the patent, net of the total transaction cost, $v_{b,j} - v_{s,j} - tc_j$, where $tc_j = tc_{b,j} + tc_{s,j}$. The bargaining shares of the two firms are the fraction of gains from trade in total surplus: $\tau_{b,j} = \xi_{b,j}/TS_j$ for the buyer and $\tau_{s,j} = \xi_{s,j}/TS_j$ for the seller. Our baseline framework treats observed patent transactions as trades that both parties expect to generate positive surplus, i.e., $\xi_{b,j} > 0$ and $\xi_{s,j} > 0$. This implies that $TS_j > 0$, or, $v_{b,j} - v_{s,j} > tc_j$, indicating that the buyer's valuation of the patent exceeds the seller's valuation by an amount greater than the transaction costs. This assumption reflects the idea that managers undertake transactions they expect to be value-increasing for the firm.

Note that neither $v_{b,j}$, $v_{s,j}$, $tc_{b,j}$, $tc_{s,j}$, nor p_j is directly observed in our data. The key to our strategy is to use both firms' stock price reactions to infer $\xi_{b,j}$ and $\xi_{s,j}$, and use them to form estimates of TS_j, $\tau_{b,j}$, and $\tau_{s,j}$. Our baseline specification adopts the standard approach in the literature by assuming the efficient markets hypothesis. Under this assumption, stock prices reflect investors' expectations of the transaction-related change in firm value. In Appendix E.1, we generalize the specification to allow for a transient divergence between managerial and market valuations, arising from investors' initial incomplete information. Our primary findings remain robust to this extension.

To isolate the surplus specifically attributable to the patent transaction from fluctuations in overall market capitalization, we implement a two-step process. First, we identify a narrow window following the date when the transaction is first recognized by the market, as outlined in Section 2.3. By narrowing the focus to this window, we reduce most factors unrelated to the patent transaction. Second, we estimate the portion of the firms' market capitalization change within this window that is attributable to the patent transaction. This methodology is detailed in Section 2.4.

2.3 Choosing The Event Window

In this section, our goal is to select an appropriately narrow window following the date when the transaction is first recognized by the market, referred to as the market awareness date.

Identifying Market Awareness Date. The USPTO Patent Assignment Dataset (UPAD) provides two key dates for each patent transaction: the *execution date*, when the conveyance was executed, and the *record date*, when it was recorded in the USPTO database. In the vast majority of cases (99.70% of our sample), record dates follow execution dates, with an average lag of 163 days and a median of 47 days. However, these dates do not necessarily correspond to the date when the market first became aware of the transaction, as the news could have been disclosed earlier via the company website or other relevant news outlets.

To extract the news disclosure dates, we utilize a large language model (LLM) trained on an extensive repository of publicly available historical news and information. Specifically, for each patent transaction in our sample, we query the GPT-4 model via its API, asking whether the trade occurred between the two companies during the execution year, which helps identify whether the patent transaction was reported in the news. The model is instructed to return either "1" (yes) or "0" (no). A response of "0" is interpreted as an indication that the transaction was likely not publicly disclosed or, at the very least, not documented in mainstream internet sources. Among the 26,509 patent transactions, 2,655 received "1". A possible explanation for the small fraction of "1" responses is that most patent transactions do not receive coverage in mainstream media.[4]

Next, we determine the earliest news dates for these identified transactions reported in the news. A common concern when using LLMs in academic research is their tendency to occasionally fabricate information. To address this issue, we manually searched Google for the 2,655 transactions labeled as "1" to locate the corresponding patent transactions. We then identified the date of the earliest publicly accessible online source reporting the transaction, referred to as the *news date*. In Appendix B.2, we conduct a textual analysis of the news coverage and find that the vast majority of articles portray the transactions as value-enhancing for both parties. In the remaining cases, the tone is negative for the seller, often reflecting distressed asset liquidations. Even under such circumstances, however, these transactions may still yield positive gains from trade, as the seller's counterfactual payoff in the absence of the deal would be even lower.

Finally, we define the *market awareness date* as the earliest of three dates: the execution date, the record date, and the news date. In most instances, the news date precedes both the execution and record dates. As demonstrated in Appendix B, our key variable of interest, bargaining share, remains robust regardless of the choice among these dates.

Determining the Window Length. With the identified market awareness date, the next step is to determine the length for the event windows. The goal is to ensure that the stock price reactions predominantly reflect information related to the patent transaction. Specifically, we want the market capitalization change within this window to capture the majority of gains from trade

[4]Moreover, as the LLM itself notes, it was not trained on UPAD or companies' financial filings unless such information was made public on company websites or other internet sources included in its training data.

resulting from patent transactions while minimizing other sources of noise.

If the window is too narrow, the market may not have sufficient time to thoroughly incorporate the information regarding patent transactions. Consequently, market capitalization changes might inadequately reflect the gains from trade. On the other hand, if the window is too wide, market capitalization changes within the window may be influenced by a significant amount of noise, making it challenging to isolate the gains from trade.

We determine the length of the window of the patent transaction by analyzing the turnover pattern surrounding the market awareness date. The turnover ratio, which represents the daily trading volume relative to shares outstanding, serves as a measure of market attention on a given day. An increase in turnover indicates heightened market attention on the stock related to the patent transaction. Once this attention wanes, it suggests that the gains from trade have already been fully incorporated into the change in market capitalization.

To examine the dynamics of the turnover ratio surrounding the market awareness date of the patent transaction, we estimate the following regression

$$h_{i,d} = a_0 + \sum_l b_l I_{i,d-l} + cZ_{i,d} + \epsilon_{i,d}, l \in \{-1, 0, 1, 2, 3\}, \quad (1)$$

where the subscripts i and d denote the firm and the date of patent transaction, respectively. $h_{i,d}$ is the ratio of daily trading volume (CRSP: vol) to shares outstanding (CRSP: shrout) of firm i on day d. $I_{i,d-l}$ is an indicator function that equals one if day $d - l$ is the market awareness date of a patent transaction involving firm i. $Z_{i,d}$ is a vector that stacks firm-year and calendar-day fixed effects. We estimate regression (1) separately for patent buyers and sellers, with the results summarized in Fact 1 below.

Fact 1. The turnover ratio exhibits a one-day increase following the market awareness date of a patent transaction.

Figure 1 displays the point estimate for the coefficient b_l for $l \in \{-1, 0, 1, 2, 3\}$ with the 90% confidence intervals. Panels (a) and (b) illustrate changes in the turnover ratio following the market awareness date of a patent transaction for patent buyers and sellers, respectively. After patent transactions, both buyers and sellers experience an initial increase of approximately 0.02% in the turnover ratio within one day, followed by a subsequent decline starting from the second day. Notably, both sides of the transaction receive similar attention, suggesting that analyses

Figure 1: Turnover around Patent Transaction and Issuance

(a) Buyers' turnover ratio (b) Sellers' turnover ratio (c) Applicants' turnover ratio

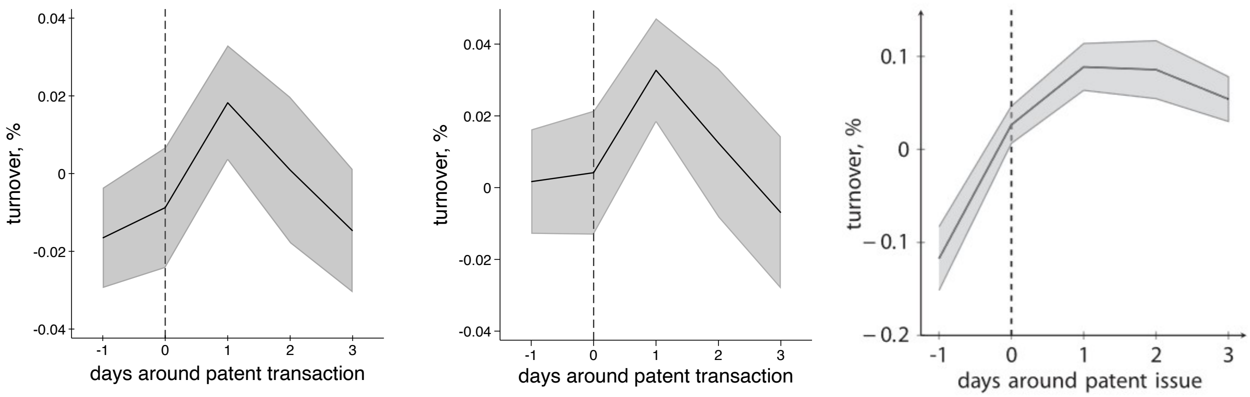

Note: Panels (a) and (b) depict the changes in the turnover ratio after the market awareness date of a patent transaction for buyers and sellers, respectively. Panel (c) (from Kogan et al., 2017) shows the changes in the turnover ratio after the announcement of patent issuance.

should not exclusively focus on the buyer's gains.

For comparison, Panel (c) (from Kogan et al., 2017) depicts turnover ratio changes after patent issuance. The rise in the turnover ratio after patent transactions is about 20% of the turnover ratio increase observed for patent issuances. This suggests that the market pays less attention to patent transactions than to patent issuances.

Another key difference in turnover ratio dynamics between patent transactions and issuances is that the turnover ratio increase following patent transactions (the first two panels) lasts only one day, whereas, for patent issuances (the third panel), it persists for three days. One plausible explanation is that publicly available information about traded patents allows for a quicker consensus on their market value.

In summary, trading activity in the stock of firms involved in patent transactions rises immediately after the transaction date. This pattern motivates a one-day benchmark window, which captures the period of concentrated market attention while limiting exposure to unrelated return shocks. In Appendix C, we extend the event window beyond the one-day benchmark using three-day, five-day, and seven-day windows around the market awareness date. The results demonstrate that estimates of gains from trade and bargaining share remain stable across all window lengths.

2.4 Estimating the Gains from Trade

The next step of our empirical strategy involves identifying the market capitalization change related to patent transactions. Specifically, we aim to isolate the portion of this change that arises from the transaction's surplus within a narrow window (one day, as determined in the previous subsection). During this window, a firm's stock price may also fluctuate for reasons unrelated to the patent transaction. To address this, we adjust for measurement error in stock returns following an approach similar to Kogan et al. (2017).

To filter out market-wide movements, we focus on the firm's idiosyncratic return, defined as the firm's return minus the return on the overall market portfolio. We decompose the idiosyncratic stock return $R_{i,j}$ for a given firm i around the date of patent transaction j as:

$$R_{i,j} = w_{i,j} + \epsilon_{i,j}, \tag{2}$$

where $w_{i,j}$ represents the value extracted by firm i from transaction j as a fraction of its market capitalization, and $\epsilon_{i,j}$ captures the residual return component unrelated to the patent transaction. Since each transaction j involves two distinct parties—a buyer b and a seller s—the subscript i can denote either the buyer or the seller, i.e., $i \in \{b, s\}$. In other words, the combination of i and j uniquely identifies a specific firm, either the buyer or the seller involved in the patent transaction j. For simplicity, with a slight abuse of notation, we will use i as the firm's index in the subsequent sections.

Firm i's gain from patent transaction j, denoted as $\xi_{i,j}$, is estimated as the change in market capitalization due to the surplus from the patent transaction:

$$\xi_{i,j} = E\left(w_{i,j} \mid R_{i,j}\right) M_{i,j} / N_{i,j}, \tag{3}$$

where $E\left(w_{i,j} \mid R_{i,j}\right)$ is the estimate of firm i's stock return due to the gain from the patent transaction j. $M_{i,j}$ is the market capitalization of firm i on the day preceding patent transaction j. Market capitalization is deflated on a monthly basis, with the base period in August 1982. $N_{i,j}$ represents the number of patent transactions related to firm i that occurred on the same day as patent transaction j.[5] By dividing the incremental market capitalization by the number of

[5]Figure 14 in Appendix H illustrates a negative relationship between $E\left(w_{i,j} \mid R_{i,j}\right)$ and $M_{i,j}$, indicating that gains from trade do not mechanically mirror market capitalization under identical distributions of market returns across firms.

patent transactions on the same day, we posit that these transactions contribute to the firm's stock return equally.

To implement equation (3), we need to make assumptions about the distributions of w (firm's stock return due to the gain from the patent transaction) and ϵ (stock returns unrelated to the patent transaction) in equation (2). Following the approach in Kogan et al. (2017), we allow both distributions to vary across firms and time periods t, which corresponds to the year of the patent transaction. To ensure that firms always extract positive surplus, we assume a normal distribution for $w_{i,j}$ truncated at 0: $w_{i,j} \sim \mathcal{N}^{+}(0, \sigma^2_{w,i,t})$. The residual return component $\epsilon_{i,j}$ is assumed to follow a normal distribution: $\epsilon_{i,j} \sim \mathcal{N}(0, \sigma^2_{\epsilon,i,t})$. As a result, the filtered value of $w_{i,j}$ as a function of the idiosyncratic stock return $R_{i,j}$ is determined by:

$$E\left(w_{i,j} \mid R_{i,j}\right) = \delta_{i,t} R_{i,j} + \sqrt{\delta_{i,t}}\sigma_{\varepsilon,i,t} \frac{\phi\left(-\sqrt{\delta_{i,t}}\frac{R_{i,j}}{\sigma_{\varepsilon,i,t}}\right)}{1 - \Phi\left(-\sqrt{\delta_{i,t}}\frac{R_{i,j}}{\sigma_{\varepsilon,i,t}}\right)} \tag{4}$$

where ϕ and Φ represent the standard normal probability density function (p.d.f.) and cumulative distribution function (c.d.f.), respectively, and $\delta_{i,t}$ is the signal-to-noise ratio represented by

$$\delta_{i,t} = \frac{\sigma^2_{w,i,t}}{\sigma^2_{w,i,t} + \sigma^2_{\varepsilon,i,t}}.$$

The detailed derivation of Equation (4) is provided in the Appendix A. The intuition behind equation (4) is that we adjust the observed changes in market capitalization based on the firm's stock price volatility prior to the market's first awareness of the patent transaction during period t. When a firm exhibits lower pre-event volatility (i.e., the term $\sigma_{\varepsilon,i,t}$ approaches zero), its stock market reaction during the event window is likely to contain less noise. In this case, $\delta_{i,t}$ approaches one, while $\phi\left(-\sqrt{\delta_{i,t}}R_{i,j}/\sigma_{\varepsilon,i,t}\right)$ approaches zero, so that $E\left(w_{i,j} \mid R_{i,j}\right)$ converges to $R_{i,j}$. This implies that the observed return $R_{i,j}$ is informative about the gain from the transaction. Conversely, as $\sigma_{\varepsilon,i,t}$ approaches infinity, $\delta_{i,t}$ approaches zero, and $E\left(w_{i,j} \mid R_{i,j}\right)$ converges to zero. In this scenario, the stock market reaction reveals little information about the gain from the patent transaction.

In Appendix E.2, we implement an alternative approach that bypasses the filtering procedure and instead relies directly on raw stock returns. This approach allows for negative gains from trade when announcement returns are negative. Despite this methodological difference, our main qualitative findings remain unchanged.

We estimate $\xi_{i,j}$ in three stages. In the first stage, we estimate $\delta_{i,t}$ by analyzing the changes in stock return volatility due to patent transactions. In the second stage, we estimate $\sigma_{\varepsilon,i,t}$ using a nonparametric method. Finally, we apply the estimated $\delta_{i,t}$ and $\sigma_{\varepsilon,i,t}$ to equation (4) to calculate $E\left(w_{i,j} \mid R_{i,j}\right)$, and then use equation (3) to obtain estimates of $\xi_{i,j}$.

Stage 1: Estimate $\delta_{i,t}$. To estimate the surplus resulting from the patent transaction using equation (4), we estimate the signal-to-noise ratios by analyzing the increase in the volatility of firms' stock returns around patent transaction days, and use them to recover the parameters $\sigma_{w,i,t}$ and $\sigma_{\epsilon,i,t}$. Following Kogan et al. (2017), we assume that the signal-to-noise ratio remains constant across transactions and years, while allowing it to differ between buyers and sellers. This restriction follows the empirical-Bayes logic of using pooled information to discipline the extraction of the transaction-related component from noisy event-window returns; we discuss its implications in Section 5. Specifically, we estimate the following panel regression separately for sellers and buyers:

$$ln\left(R_{i,d}^2\right) = \gamma I_{i,d} + cZ_{i,d} + u_{i,d},$$

where $R_{i,d}$ represents the idiosyncratic return of firm i of day d. $ln(R_{i,d}^2)$ is the log squared return that measures the stock return volatility. $I_{i,d}$ is an indicator function that equals one if firm i transacts a patent with another firm on day d. $Z_{i,d}$ is a vector of fixed effects that includes day-of-the-week and firm-by-year fixed effects to account for seasonality and time-variation in volatility fluctuations. The coefficient γ measures the percentage increase in firm's stock return volatility due to the patent transaction and reflects the relationship between $\sigma_{w,i,t}^2$ and $\sigma_{\varepsilon,i,t}^2$ given by

$$\gamma = ln\left[(\sigma_{w,i,t}^2 + \sigma_{\varepsilon,i,t}^2)/\sigma_{\varepsilon,i,t}^2\right], \tag{5}$$

where the right-hand side of this equation represents the log deviation of stock return volatility resulting from the increase in volatility caused by the patent transaction. Equation (5) further implies that $\delta = 1 - e^{-\gamma}$, which enables the calculation of the signal-to-noise ratio δ from γ. Specifically, for buyers, with an estimated $\widehat{\gamma}$ of 0.0229, we have $\widehat{\delta}_b = 0.0226$. For sellers, with an estimated $\widehat{\gamma}$ of 0.0297, it entails that $\widehat{\delta}_s = 0.0292$.

Stage 2: Estimate $\sigma_{\varepsilon,i,t}$. Next, we estimate the variance of the residual return component, $\sigma^2_{\varepsilon,i,t}$ using the nonparametric method described in Andersen and Teräsvirta (2009). This method estimates the conditional volatility of firm i in year t through the mean of squared realized idiosyncratic returns, $\sigma^2_{i,t}$. Since the estimation of $\sigma^2_{i,t}$ covers both the market awareness day and non-event days, it represents a combination of $\sigma^2_{\varepsilon,i,t}$ and $\sigma^2_{w,i,t}$ given by

$$\sigma^2_{i,t} = d_{i,t}\sigma^2_{w,i,t} + \sigma^2_{\varepsilon,i,t}, \tag{6}$$

where $d_{i,t}$ is the fraction of trading days that are market awareness days. Equation (6) shows that the total variance of the firm's stock return in a year with patent transactions is attributed to both the firm's intrinsic variance and the variance resulting from those transactions.

Note that the estimate of γ has revealed the ratio between $\sigma^2_{\varepsilon,i,t}$ and $\sigma^2_{w,i,t}$. We can derive the variance of the residual return component from equation (6) as

$$\widehat{\sigma}^2_{\varepsilon,i,t} = \widehat{\sigma}^2_{i,t}\left[1 + d_{i,t}\left(e^{\hat{\gamma}} - 1\right)\right]^{-1},$$

which shows that a higher fraction of market awareness days or higher relative stock return volatility associated with patent transactions (as captured by γ) contributes to a lower estimated volatility of the residual return component.

Stage 3: Estimate $\xi_{i,j}$. By applying the estimations of $\delta_{i,t}$ and $\sigma^2_{\varepsilon,i,t}$ to equation (4), we can gauge each firm's stock market return due to patent transactions. Finally, using equation (3) provides the estimates of firms' gain from patent transactions.

In estimating the gain from the patent transaction, $\xi_{i,j}$, a potential concern is ex-ante information leakage, i.e., the market may become aware of potential patent transactions before they are successfully executed. As a result, the observed market reaction may only capture a portion of the transaction's total impact on the firm's value, potentially reflecting the resolution of uncertainty post-announcement. When estimating the patents' market value at the time of issuance, Kogan et al. (2017) adjust for anticipated patent issuance by dividing the filtered changes in firms' market capitalization due to patent issuance by the average success probability of patent applications.[6] Due to data constraints, we cannot observe the probability of a successful

[6]The average probability of a successful patent application is about 56% according to Carley et al. (2015), hence Kogan et al. (2017) divide the filtered change in firms' market cap by 0.56.

patent transaction and adjust $\xi_{i,j}$ accordingly. Consequently, our estimated $\xi_{i,j}$ likely captures only a portion of the gains from patent transactions. However, it is important to note that our analyses focus on the ratio of gains from trade between counterparties and on log differences across transactions. If the probability of success is uniform across firms—as assumed in Kogan et al. (2017)—this constant cancels out in both specifications.[7] Therefore, our main results are not affected by the lack of adjustment for the average success probability.

Total surplus and bargaining shares. Our empirical method enables us to compute the gains from trade for both buyer ($\xi_{b,j}$) and seller ($\xi_{s,j}$) in each patent transaction. By summing up the two gains, we obtain the total surplus TS_j resulting from each patent transaction j as

$$TS_j = \xi_{s,j} + \xi_{b,j}.$$

The bargaining shares of the buyer ($\tau_{b,j}$) and the seller ($\tau_{s,j}$) in the patent transaction j are determined by dividing their respective gains from trade by the total surplus:

$$\tau_{b,j} = \frac{\xi_{b,j}}{TS_j}, \ \tau_{s,j} = \frac{\xi_{s,j}}{TS_j}.$$

By construction, we have $\tau_{b,j} + \tau_{s,j} = 1$. As $\tau_{b,j}$ approaches 1, the buyer receives nearly all the total surplus, while as $\tau_{b,j}$ approaches 0, the seller captures most of the total surplus. Table 1 summarizes the descriptive statistics for buyers' and sellers' gains from trade, their bargaining shares, and total surplus.

Table 1 reports the summary statistics for seller and buyer's gain from trade, buyer's bargaining share, and total surplus, respectively. The average buyer bargaining share is 0.42, suggesting a nearly even split of the total surplus. However, substantial dispersion is observed, with a standard deviation of 0.333. We further investigate the distribution of bargaining shares in Fact 2. The mean total surplus across transactions is \$11,103. For comparison, Serrano (2018) reports an average total surplus of \$15,240 in 2003 USD (\$7,993 in 1982 USD), about 28% lower than our estimate. This discrepancy likely reflects differences in estimation methodology, sample

[7] Specifically, the logarithm of adjusted gains equals the logarithm of unadjusted gains minus a constant (the log success probability). Hence, the constant does not affect the estimated coefficients in our regressions. See Appendix B for further discussion and empirical evidence.

coverage, and the broader scope of our gains-from-trade measure.[8]

Table 1: Descriptive Statistics

Note: This table presents descriptive statistics for buyers' gains from trade ($\xi_{b,j}$), sellers' gains from trade ($\xi_{s,j}$), buyers' bargaining share ($\tau_{b,j}$), and total surplus (TS_j). Monetary values are expressed in 1982 U.S. dollars.

		(1)	(2)	(3)	(4)	(5)
Variable	Notation	#Observation	Mean	Std.Dev	Min	Max
Buyer's gain from trade	$\xi_{b,j}$	26,509	4,252	34,787	1.492	2,350,996
Seller's gain from trade	$\xi_{s,j}$	26,509	6,852	38,593	0.490	2,189,970
Total surplus	TS_j	26,509	11,103	54,819	9.266	2,354,355
Buyer's bargaining share	$\tau_{b,j}$	26,509	0.421	0.333	0.000350	1.000

3 Empirical Findings

In the previous section, we computed the total surplus and the bargaining share for each patent transaction. Next, we explore some important observations regarding the distributions of total surplus and bargaining shares, as well as their connection to firms' economic fundamentals. The first key observation regards the distribution of bargaining shares, as summarized by Fact 2.

Fact 2. The distribution of bargaining shares is concentrated at both ends of the spectrum.

Panels (a) and (b) in Figure 2 illustrate the distribution of buyers' and sellers' bargaining shares, respectively. Since the sum of both sides' bargaining shares for each specific patent transaction is one, the two panels are symmetric. The shape of the distributions resembles a U-curve, with concentration at both ends of the spectrum. This indicates that, in most cases, the total surplus of a transaction is significantly skewed toward one party rather than being evenly split. Interestingly, the density of bargaining shares around 0.5—the midpoint—is among the lowest. In other words, there are very few instances where the total surplus is equally divided. This finding is in stark contrast to the conventional assumption in the literature that bargaining shares are uniform across transactions and equal to 0.5.

A potential concern is that bargaining shares estimated to be close to zero may partly reflect negative stock returns during the event window. In Appendix D, Figure 9 shows the distribution

[8]Serrano (2018) employs a structural model with a sample of patents granted between 1988 and 1992 with transfer data through 2001, and focuses primarily on the productive value of patents.

of bargaining shares for the subsample of transactions in which both parties experienced positive stock returns. The distribution retains the same U-shaped pattern observed in the full sample. Furthermore, Appendix D confirms that all core findings remain robust when estimated on this more restricted sample.

Figure 2: Distribution of Bargaining Shares

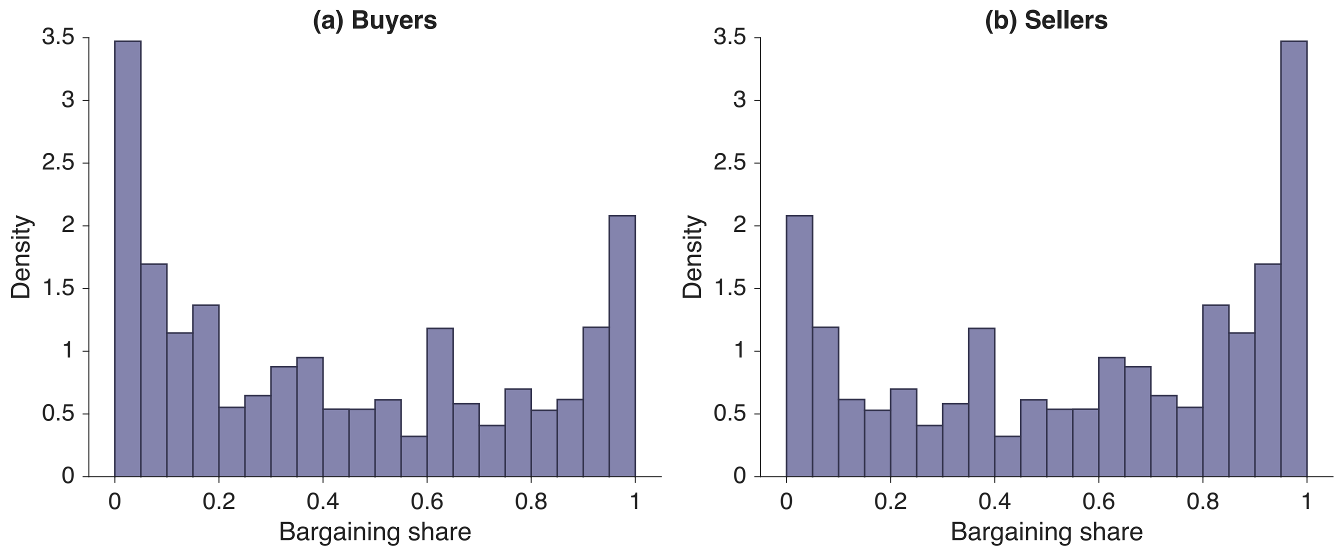

Note: This figure illustrates the distribution of buyers' and sellers' bargaining shares.

We next document that bargaining share is primarily determined by the relative size of traders, as summarized by Fact 3.

Fact 3. Traders' bargaining shares are proportional to their size, measured by output, relative to their counterparties.

To examine the role of firms' relative sizes in determining their bargaining shares, we estimate the following regression:

$$\tau_{b,j} = \beta \times RS_{b,j} + Z_{j,t} + \epsilon_j, \tag{7}$$

where $\tau_{b,j}$ represents the buyer's bargaining share in patent transaction j, and $RS_{b,j}$ denotes the buyer's relative size, measured by the ratio of its real sales to the combined real sales of both the buyer and seller in the year of the patent transaction.[9] The sellers' bargaining shares and relative sizes are symmetric to those of the buyers, with $\tau_{s,j} = 1 - \tau_{b,j}$ and $RS_{s,j} = 1 - RS_{b,j}$. Therefore,

[9]Specifically, $RS_{b,j} = \frac{Sale_{b,j}}{Sale_{b,j}+Sale_{s,j}}$, where $Sale_{b,j}$ and $Sale_{s,j}$ are the real sales of the buyer and seller, respectively, in the year of the patent transaction.

they are omitted in the regression. $Z_{j,t}$ stacks the transaction-year, industry-pair (four-digit NAICS), and patent-class fixed effects. Standard errors are clustered at the firm level.

The estimation results are shown in column (1) of Table 2. Relative size is positively correlated with bargaining share, indicating that firms with larger size disproportionately reap more surplus than their smaller counterparties.

To demonstrate the decisive role of relative size in determining the bargaining share, Column (2) of Table 2 presents the results of a univariate regression without fixed effects and intercept. Relative size alone accounts for approximately 78% of the variation in bargaining shares. In comparison, the inclusion of the fixed effects (as shown in Column (1)) raises the R^2 to 0.87. The limited improvement in R-squared despite adding those fixed effects suggests that relative size is the primary factor influencing bargaining shares. An important practical implication of this finding is that when estimating bargaining shares is challenging, calibrating them using relative sizes is a reasonable approach.

Appendix F analyzes the variation in bargaining shares across sector pairs. Manufacturing sectors tend to have higher bargaining shares, particularly in transactions with sectors such as transportation, information, and professional services.

Table 2: Bargaining Share is Positively Correlated with Relative Size

The dependent variable is the buyers' bargaining share. The independent variable is the buyers' relative size compared to the sellers. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)
Dependent Variable:	Bargaining Share	
Relative Size	0.763***	0.941***
	(0.033)	(0.024)
Year FE	Yes	–
Industry Pair and Patent Class FE	Yes	–
Observations	26,126	26,362
R-squared	0.869	0.782

Corporate governance and size-dependent bargaining share.

The preceding results show that relative firm size is strongly associated with bargaining share in patent transactions. A plausible institutional force behind the strength of this relationship is the incentive misalignment between shareholders and the managers who conduct these transactions. Patent transactions are complex and often involve managerial discretion over negotiation, timing, and counterparty choice. When shareholder control is weaker, smaller firms are particularly disadvantaged because they may lack the internal governance resources to discipline managerial opportunism (e.g., selling patents below fair value to preferred buyers) and the negotiating resources to counteract better-resourced and aggressive counterparties. Consequently, they are more likely to accept unfavorable terms and capture a smaller share of the surplus. Larger firms, by contrast, benefit from stronger outside options, greater transaction experience, and superior bargaining leverage. Weaker shareholder control can therefore steepen the size-bargaining relationship, making bargaining shares more sensitive to relative firm size.

To examine this possibility, we test whether corporate governance changes the strength of the relationship between firm size and bargaining share. To proxy for the quality of governance, we employ the Entrenchment Index (E-index) developed by Bebchuk et al. (2009). This index is constructed from six key provisions that restrict shareholder rights, with higher scores indicating weaker shareholder control and thus greater potential for managerial-shareholder misalignment.[10]

We examine this heterogeneity by estimating the following regression:

$$\tau_{i,j} = a + b \cdot IM_i \times RS_{i,j} + c \cdot RS_{i,j} + d \cdot IM_i + Z'_{i,j,t}\gamma + \epsilon_{i,j},$$

where $\tau_{i,j}$ is firm i's bargaining share in transaction j, $RS_{i,j}$ is its relative size, IM_i is the firm's governance-related incentive misalignment measured by the E-index, and $Z_{i,j,t}$ stacks the fixed effects and controls. The coefficient of interest is b on the interaction term $IM_i \times RS_{i,j}$. A positive coefficient means that the marginal relationship between relative size and bargaining share is stronger when shareholder control is weaker.

Estimation results are presented in Table 3. The result indicates that relative size is more strongly associated with bargaining share among firms with weaker shareholder control. The

[10] We match each firm-year observation in our sample to the corresponding E-index value using ticker identifiers. The underlying ISS/RiskMetrics data underwent a reporting methodology change, so this index is available directly only through 2006. For post-2006 observations, we carry forward each firm's last available score. This approach is supported by Li and Li (2018), who document the high persistence and slow evolution of corporate governance scores over time.

effects are economically meaningful. Moving from the minimum to the maximum E-index value, 0 to 6, raises the marginal effect of relative size by about 0.4 (0.066×6), equivalent to roughly 52% of the baseline coefficient reported in Table 2.

Table 3: Corporate Governance, Bargaining Share, and Relative Size

The dependent variable is the firm's bargaining share. The key independent variable is the interaction between the firm's relative size and governance-related incentive misalignment, proxied by the E-index. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, ** p<0.05, * p<0.1.

	(1)
Dependent Variable:	Firm's Bargaining Share
$IM_i \times RS_{i,j}$	0.066*
	(0.037)
$RS_{i,j}$	0.600***
	(0.105)
IM_i	-0.063**
	(0.026)
Year FE	Yes
Patent Class FE	Yes
Observations	24,249
R-squared	0.727

The preceding analysis examines how transaction surplus is divided between the two parties; however, the magnitude of surplus created is also important. A smaller firm may capture a lower share when transacting with a larger counterparty, but the larger counterparty may simultaneously generate greater gains from trade. We therefore turn to how total surplus varies with the sizes of both parties.

Fact 4. The total surplus of patent transactions increases with both the buyer's and the seller's firm sizes.

To examine the role of firm size in the determination of total surplus, we estimate the following regression:

$$ln(TS_{j,t}) = \beta_1 ln(Sale_{b,j,t}) + \beta_2 ln(Sale_{s,j,t}) + Z_{j,t} + \epsilon_{j,t}, \tag{8}$$

where $TS_{j,t}$ is the total surplus from patent transaction j occurring in year t. $Sale_{b,j,t}$ and $Sale_{s,j,t}$ are the real sales of the buyer and seller in the year of the patent transaction, respectively. $Z_{j,t}$ is a vector of control variables and fixed effects, such as the number of forward citations (using the HJT correction term), the patent's real market value at the time of issuance estimated by Kogan et al. (2017), transaction-year fixed effects, industry-pair (four-digit NAICS) fixed effects, and patent-class fixed effects. The inclusion of industry-pair and patent-class fixed effects helps to account for technology propinquity between the patent and the firms' core business lines (as revealed by the firms' industries), as patents closely aligned with a company's primary business may carry higher value (Akcigit et al., 2016).

Table 4: Total Surplus Increases with Firm Size

The dependent variable is the logarithm of total surplus. The independent variables are the sizes of the buyers and sellers, measured by the logarithm of their real sales. Standard errors, reported in parentheses, are clustered at the firm level. *** $p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

	(1)	(2)	(3)	(4)
Dependent Variable:	*Total Surplus*			
Buyer's Size	0.175***	0.171***	0.198***	0.247***
	(0.063)	(0.061)	(0.054)	(0.054)
Seller's Size	0.134*	0.125*	0.105	0.125**
	(0.073)	(0.070)	(0.066)	(0.053)
Controls				
Citation	-	Yes	Yes	Yes
InitialValue	-	Yes	Yes	Yes
Year FE	-	-	Yes	Yes
Industry Pair and Patent Class FE	-	-	-	Yes
Observations	26,275	26,275	26,273	26,038
R-squared	0.052	0.105	0.264	0.746

As shown in column (4) of Table 4, a 1% increase in buyer size is associated with a 0.25% increase in total surplus, while a 1% increase in seller size corresponds to a 0.13% increase. These effects are economically significant, highlighting the importance of firm sizes for overall surplus generation. The positive relationship for buyer size remains robust across specifications, though the effect for seller size is only marginally significant in columns (1) and (2).

The positive association between total surplus and buyer size is intuitive: the value of a patent often depends on the scale at which it can be utilized, making it more valuable to larger firms capable of broader deployment. Consider AOL's sale of its patent portfolio to Microsoft. AOL originally developed patents related to search and online advertising technologies, but its ability to commercialize them was limited by the scale of its own search platform. Upon acquisition, Microsoft integrated these patents into Bing, leveraging them across a substantially larger user base and a more established ecosystem.

The positive relationship between total surplus and seller size may stem from the greater organizational efficiency of larger firms, which enables them to navigate the market more effectively and identify suitable deals. A key component of this efficiency is the minimization of transaction costs, which represent a significant friction in patent markets. Using a structural model, Serrano (2018) estimates that the average transaction cost per traded patent is approximately \$6,680 in 2003 USD (about \$3,500 in 1982 USD). For context, the average total surplus in our sample is \$11,103 (also in 1982 USD), indicating that transaction costs absorb a substantial share of the potential gains from trade. Crucially, such costs tend to decline with firm size due to economies of scale in legal and administrative functions, reputation effects that lower due diligence costs for counterparties (Gans et al., 2008), and the ability to bundle assets and spread fixed costs. This scale-dependent efficiency, consistent with analogous findings in related settings such as university technology transfer (Siegel et al., 2003), enables larger sellers to generate greater surplus from patent transactions.

At the sector-pair level, Appendix F shows that total surplus is generally higher when patents transfer from less technology-dependent industries to more technology-dependent ones. This pattern aligns with the intuition that buyers in technology-intensive sectors are better positioned to commercialize and apply innovations, generating greater gains from trade.

Taking Stock. Fact 2 and Fact 3 indicate that bargaining shares exhibit substantial variation across patent transactions. Firm size plays a crucial role in determining these shares for both parties involved. Fact 4 shows that firm size also affects the amount of surplus created: larger firms can generate greater gains from trade through their scale and lower transaction costs. These two findings have opposing implications for smaller firms. A larger counterparty can create more value from a patent, but it also captures a greater share of that value. We next introduce a model that combines these two forces and studies their implications for firms' R&D decisions.

4 Model

To assess the net effect of size-dependent bargaining on innovation, we develop a parsimonious model that characterizes the key economic forces. We first derive the general analytical results of the model. We then calibrate the model using our empirical estimates to study the quantitative implications of the estimated bargaining rule. Finally, we conduct counterfactual analyses comparing the estimated size-dependent bargaining rule against two benchmarks: symmetric bargaining and the absence of a secondary market.

4.1 Theoretical Framework

4.1.1 Baseline Environment.

There are two firms, indexed by $i \in \{1, 2\}$, with different business types ϕ_i and sizes κ_i. Without loss of generality, we assume $\phi_1 = 0$, $\phi_2 = 1$, and $\kappa_i \in (0, 1]$. Firms need blueprints, corresponding to patents in the data, to produce output. Blueprints are heterogeneous in their technological types $\theta \in [0, 1]$.

The profit generated from a blueprint for firm i in production, denoted by $v_i(\theta)$, is negatively related to the technological distance between the blueprint's technological type and the firm's business type, $|\phi_i - \theta|$, and positively related to the firm's size, κ_i:

$$v_i(\theta) = f(\kappa_i)(1 - |\phi_i - \theta|), \tag{9}$$

where $f(\cdot)$ is an increasing function with $f(1) = 1$.

Figure 3: Technology Propinquity between the Blueprint and the Line of Business

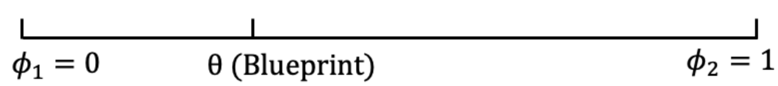

Note: This figure illustrates the technological types of blueprints and their technology propinquity with both firms.

Figure 3 illustrates the described setting. The line segment $[0, 1]$ represents the technological space of blueprints, with the two firms positioned at the endpoints. The distance of any given blueprint from a firm indicates its technological distance from that firm's business type. A blueprint located sufficiently to the left is closer to, and thus more valuable to, firm 1 than firm 2. Conversely, a blueprint located sufficiently to the right is more valuable to firm 2. Specifically, equation (9) implies that $v_1(\theta) = f(\kappa_1)(1-\theta)$ and $v_2(\theta) = f(\kappa_2)\theta$. We have $v_1(\theta) > v_2(\theta)$ if and only if $\theta < f(\kappa_1)/(f(\kappa_1) + f(\kappa_2))$.

A firm can produce output using multiple blueprints. For simplicity, we assume that there are no technological complementarities between blueprints, and the total production profit is the sum of the profits of each individual blueprint, $V_i(\Theta_i) = \sum_{\theta \in \Theta_i} v_i(\theta)$, where Θ_i is the set of blueprints possessed by firm i.

4.1.2 The Development and the Exchange of Blueprints.

Firms start with no blueprint. Each firm can choose to develop one blueprint, which incurs a fixed R&D cost of c. The technological type of the blueprint is uniformly drawn from $[0,1]$ and is realized and revealed after the R&D cost is incurred.[11] A firm i will engage in R&D if the R&D cost c is lower than the expected value of the blueprint, EV_i, derived below.

Firms can trade their developed blueprints in a secondary market. When the business type of firm i is sufficiently different from the realized technological type of the blueprint it developed, θ_i, where the subscript indicates the developer, there is a potential gain from trade because the blueprint is more valuable to the other firm, $-i$. Specifically, the blueprint's values to firms $-i$ and i are given by

$$v_{-i}(\theta_i) = f(\kappa_{-i})(1 - |\phi_{-i} - \theta_i|), \qquad v_i(\theta_i) = f(\kappa_i)(1 - |\phi_i - \theta_i|),$$

respectively.

If the blueprint is traded, the firms incur a transaction cost $tc(\kappa_i, \kappa_{-i})$ given by

$$tc(\kappa_i, \kappa_{-i}) = [1 - \eta(\kappa_i, \kappa_{-i})] \cdot [v_{-i}(\theta_i) - v_i(\theta_i)] \,. \tag{10}$$

In equation (10), $\eta(\kappa_i, \kappa_{-i}) \in [0,1]$ is an efficiency term that determines the fraction of value retained after accounting for transaction costs. The transaction cost captures, in reduced form, frictions in the trading process, such as search costs, informational barriers, negotiation and contracting expenses, and legal or operational complexities.

[11] This assumption is made for analytical convenience; Section 5 discusses how the model's implications extend under alternative distributional assumptions.

4.1.3 Value Functions.

The total surplus generated from a blueprint transaction is

$$\begin{aligned} TS(\theta_i) &= v_{-i}(\theta_i) - v_i(\theta_i) - tc(\kappa_i, \kappa_{-i}) \\ &= \eta(\kappa_i, \kappa_{-i}) \cdot [f(\kappa_{-i})(1 - |\phi_{-i} - \theta_i|) - f(\kappa_i)(1 - |\phi_i - \theta_i|)], \end{aligned} \tag{11}$$

which captures the net value created by reallocating a blueprint to a better-matched firm, accounting for transaction costs.

Equation (11) shows that total surplus is positive if and only if

$$|\phi_i - \theta_i| > \frac{f(\kappa_i)}{f(\kappa_i) + f(\kappa_{-i})}.$$

That is, trade occurs only when the blueprint's technological type is sufficiently far from the developer's business type. Otherwise, there is no gain from trade, and firm i keeps the blueprint it developed.

We assume that the total surplus is distributed between the firms by Nash bargaining. Firm i's bargaining share is denoted by $\tau_i(\kappa_i, \kappa_{-i})$, with $\tau_i(\kappa_i, \kappa_{-i}) + \tau_{-i}(\kappa_{-i}, \kappa_i) = 1$. The resulting gains from trade are $\tau_i(\kappa_i, \kappa_{-i})TS(\theta_i)$ for firm i and $\tau_{-i}(\kappa_{-i}, \kappa_i)TS(\theta_i)$ for firm $-i$. In the general model, we only require that firm i's bargaining share is weakly increasing in its own size and weakly decreasing in the counterparty's size. We introduce the specific empirical bargaining rule in the calibration subsection below.

The expected payoff from developing a blueprint for firm i is

$$\begin{aligned} EV_i = &\underbrace{\frac{f(\kappa_i)}{f(\kappa_i) + f(\kappa_{-i})} E\left[v_i(\theta_i) \mid |\phi_i - \theta_i| \leq \frac{f(\kappa_i)}{f(\kappa_i) + f(\kappa_{-i})}\right]}_{\text{keep}} \\ &+ \underbrace{\frac{f(\kappa_{-i})}{f(\kappa_i) + f(\kappa_{-i})} E\left[\tau_i(\kappa_i, \kappa_{-i})TS(\theta_i) + v_i(\theta_i) \mid |\phi_i - \theta_i| > \frac{f(\kappa_i)}{f(\kappa_i) + f(\kappa_{-i})}\right]}_{\text{sell}}. \end{aligned} \tag{12}$$

When the blueprint is sufficiently close to firm i's business type, firm i keeps it. When the blueprint is better matched to firm $-i$, firm i sells it in the secondary market.

Since θ_i is independently drawn from the uniform distribution on $[0, 1]$, we have

$$E\left[TS(\theta_i) \mid |\phi_i - \theta_i| > \frac{f(\kappa_i)}{f(\kappa_i) + f(\kappa_{-i})}\right] = \frac{\eta(\kappa_i, \kappa_{-i})f(\kappa_{-i})}{2}. \tag{13}$$

This expression shows that, conditional on trade, expected total surplus depends on the productive value of the buyer and on the transaction efficiency term.[12]

Substituting equation (13) into equation (12), we obtain

$$EV_i = \underbrace{\frac{f(\kappa_i)}{2}}_{\text{Own-use value}} + \underbrace{\tau_i(\kappa_i, \kappa_{-i})}_{\text{Bargaining share}} \cdot \underbrace{\frac{\eta(\kappa_i, \kappa_{-i}) f^2(\kappa_{-i})}{2\left[f(\kappa_i) + f(\kappa_{-i})\right]}}_{\text{Secondary-market option value}} . \tag{14}$$

Equation (14) separates the value of R&D into two components. The first term is the expected value of using the blueprint internally. The second term is the option value created by the secondary market. This option value depends on the total surplus generated by reallocating the blueprint and on the share of that surplus captured by the innovating firm.

Proposition 1. *Under the maintained assumptions that $f(\cdot)$ is increasing in firm size, $\eta(\kappa_i, \kappa_{-i})$ is weakly increasing in κ_i, and $\tau_i(\kappa_i, \kappa_{-i})$ is weakly increasing in κ_i, firm i's expected value of developing a blueprint increases with its own size.*

Proposition 1 captures how firm size raises the private value of an R&D outcome. A larger firm can generate greater value from using a blueprint, face weakly lower frictions when trading it, and capture a weakly larger share of transaction surplus. These forces increase the expected return from developing a blueprint.[13]

The effect of the potential buyer's size is more ambiguous. Holding firm i's size fixed, a larger potential buyer can create more surplus from using the blueprint and may reduce transaction costs. At the same time, if bargaining shares depend on relative size, a larger potential buyer reduces firm i's share of the surplus. The sign of this effect therefore depends on the relative strength of surplus creation and surplus capture. The calibrated model below quantifies this trade-off using the empirical estimates.

[12] The expected total surplus depends on κ_i solely through the transaction efficiency term. This is because we condition on trade occurring only when $|\phi_i - \theta_i| > f(\kappa_i)/(f(\kappa_i) + f(\kappa_{-i}))$. As the seller's size κ_i increases, the probability of trade decreases, selecting blueprints whose technological type is closer to the buyer's business type. These matched blueprints generate higher surplus ex post. The reduced likelihood of trade and improved matching offset each other, leaving expected total surplus independent of κ_i except through the transaction efficiency term.

[13] These channels reflect advantages in commercializing or trading technology once a blueprint exists, rather than greater productivity in generating the blueprint internally.

4.2 Calibration

We now calibrate the model to the empirical estimates from Section 3. As shown in Table 2, the relative-size term $\kappa_i/(\kappa_i + \kappa_{-i})$, referred to as RS in equation (7), closely approximates a firm's bargaining share, so we set

$$\tau_i = \frac{\kappa_i}{\kappa_i + \kappa_{-i}}. \tag{15}$$

We specify the productive value and transaction efficiency functions as

$$f(\kappa) = \kappa^{\alpha_0}, \qquad \eta(\kappa_i, \kappa_{-i}) = \kappa_i^{\alpha_1}\kappa_{-i}^{\alpha_2}.$$

The specification of $\eta(\cdot,\cdot)$ captures the idea that firm size may reduce transaction frictions by increasing visibility, reputation, and matching opportunities in patent trades (see Serrano 2010).

We draw on two empirical moments from column (4) in Table 4: the elasticity of total surplus with respect to buyer size is 0.25, and the elasticity of total surplus with respect to seller size is 0.13. In the model, these correspond to

$$\frac{\partial \ln\left[\eta(\kappa_i,\kappa_{-i})f(\kappa_{-i})\right]}{\partial \ln \kappa_{-i}} = 0.25, \qquad \frac{\partial \ln\left[\eta(\kappa_i,\kappa_{-i})f(\kappa_{-i})\right]}{\partial \ln \kappa_i} = 0.13, \tag{16}$$

where $\eta(\kappa_i,\kappa_{-i})f(\kappa_{-i})$ represents the expected total surplus conditional on trading, multiplied by two, as in equation (13).

Given the assumed functional forms, conditions (16) imply $\alpha_2 + \alpha_0 = 0.25$ and $\alpha_1 = 0.13$. By imposing symmetric contributions of buyer and seller size to the transaction efficiency term, we obtain $\alpha_0 = 0.12$ and $\alpha_1 = \alpha_2 = 0.13$.[14] Thus,

$$f(\kappa) = \kappa^{0.12}, \qquad \eta(\kappa_i,\kappa_{-i}) = \kappa_i^{0.13}\kappa_{-i}^{0.13}. \tag{17}$$

4.3 Numerical Analysis

In this subsection, we use the calibrated model to examine the implications of the empirical estimates for R&D incentives.

[14] We impose symmetry across buyers and sellers because we do not observe how trading partners meet and lack evidence of asymmetric size effects in the transaction efficiency term.

4.3.1 Counterparty Size.

We first examine how the size of a potential buyer affects firm i's R&D incentives in the calibrated model. Substituting equations (15) and (17) into equation (14) gives

$$EV_i = \frac{\kappa_i^{0.12}}{2} + \frac{\kappa^{0.37}}{(1+\kappa)(1+\kappa^{0.12})} \cdot \frac{\kappa_i^{0.38}}{2}, \tag{18}$$

where $\kappa \equiv \kappa_{-i}/\kappa_i$ is the size of firm $-i$ relative to firm i.

Equation (18) illustrates the trade-off described in the general model. A larger potential buyer increases the value created by reallocating a poorly matched blueprint and improves transaction efficiency. At the same time, it reduces firm i's bargaining share. As κ approaches zero, the potential buyer derives little value from the blueprint, causing the secondary-market option value to vanish. As κ tends to infinity, the potential buyer creates high gross value but also captures nearly all of the surplus, again driving firm i's option value toward zero.

Differentiating equation (18) with respect to κ gives the following calibrated implication.

Proposition 2. *Under the calibrated specification with $f(\kappa) = \kappa^{0.12}$ and $\eta(\kappa_i, \kappa_{-i}) = \kappa_i^{0.13}\kappa_{-i}^{0.13}$, firm i's expected value of developing a blueprint, EV_i, is non-monotone in the size of the potential buyer, κ_{-i}. Holding κ_i fixed, EV_i increases with κ_{-i} when $\kappa_{-i}/\kappa_i < \kappa^*$ and decreases with κ_{-i} when $\kappa_{-i}/\kappa_i > \kappa^*$, where κ^* is the positive solution to equation (19).*

The threshold κ^* solves

$$0.37\kappa^{-1.12} + 0.25\kappa^{-1} - 0.63\kappa^{-0.12} - 0.75 = 0, \tag{19}$$

where the positive root is approximately 0.46.[15]

Proposition 2 reflects the balance between two forces. A larger potential buyer creates more surplus from using the blueprint and improves transaction efficiency. At the same time, the empirically estimated size-dependent bargaining rule lowers the innovator's share of that surplus. Under the calibrated elasticities, the surplus-creation effect dominates when the potential buyer is relatively small, while the bargaining-share loss dominates when the potential buyer is sufficiently large.

Appendix G examines the ex post division of gains conditional on the seller having developed a blueprint and chosen to trade it. Conditional on trading, both buyers and sellers prefer

[15] Specifically, $\partial EV_i/\partial\kappa_{-i}$ has the same sign as $\partial EV_i/\partial\kappa$, because $\partial EV_i/\partial\kappa_{-i} = (\partial EV_i/\partial\kappa)/\kappa_i$. They are positive if $\kappa^{0.37}/[(1+\kappa)(1+\kappa^{0.12})]$ increases with κ.

counterparties significantly smaller than themselves, as the bargaining advantage from size asymmetry outweighs the efficiency losses from higher transaction costs. Specifically, a seller's expected gains are maximized when the buyer is one-third its size, while a buyer's expected gains peak when the seller is 13/87 of its size.

4.3.2 Comparison with No Secondary Market.

We next compare the expected value of developing a blueprint in the calibrated benchmark model to a counterfactual scenario in which a secondary market for blueprints does not exist. In the absence of a market, a firm must retain and use any blueprint it develops, regardless of its technological fit. Its expected value is therefore the average productive value of the blueprint:

$$E[v_i(\theta)] = f(\kappa_i) E\left[1 - |\phi_i - \theta|\right] = \frac{\kappa_i^{0.12}}{2}, \tag{20}$$

where $v_i(\theta)$ is the blueprint's productive value defined in equation (9).

A direct comparison of equation (20) with equation (18) shows that the secondary market raises the expected value of R&D by giving firms the option to sell poorly matched blueprints. The market therefore partially insures firms against the risk of low technological fit and increases the expected return to innovation.

4.3.3 Comparison with Symmetric Bargaining.

To isolate the role of size-dependent bargaining, we consider a counterfactual scenario in which the bargaining share is fixed at 0.5. This assumption leads to the following expected value of developing a blueprint:

$$\widetilde{EV}_i = \frac{f(\kappa_i)}{2} + \frac{1}{2} \cdot \frac{\eta(\kappa_i, \kappa_{-i}) f^2(\kappa_{-i})}{2[f(\kappa_i) + f(\kappa_{-i})]} = \frac{\kappa_i^{0.12}}{2} + \frac{\kappa^{0.37}}{2(1+\kappa^{0.12})} \cdot \frac{\kappa_i^{0.38}}{2}. \tag{21}$$

A direct comparison of equation (21) with equation (18) yields two implications. First, under the empirically estimated size-dependent bargaining rule, larger firms benefit disproportionately: $EV_i > \widetilde{EV}_i$ if and only if $\kappa_i > \kappa_{-i}$. That is, the estimated bargaining structure raises the innovation incentives of the larger party in a potential transaction while lowering those of the smaller one.

Second, under symmetric bargaining, $\widetilde{EV}_i$ is increasing in the size of the potential buyer κ_{-i}. Absent size-dependent bargaining, the expected value of innovation for small firms would increase with the size of the potential counterparty, because larger counterparties can create greater value from the blueprint and transact at lower cost. In contrast, size-dependent bargaining attenuates or can even reverse this effect for relatively small innovators.

4.3.4 Aggregate R&D Incentives.

To highlight the differences across the three scenarios, we simulate a continuum of potential innovators with heterogeneous sizes, assuming that the size of potential counterparties follows a uniform distribution over the unit interval. The left panel of Figure 4 displays the expected value of developing a blueprint as a function of a firm's own size, κ_i, for three cases: (1) no secondary market, corresponding to $E[v_i(\theta)]$ in equation (20); (2) a secondary market with symmetric Nash bargaining, given by $\int_0^1 \widetilde{EV}_i \, d\kappa_{-i}$, where $\widetilde{EV}_i$ is defined in equation (21); and (3) a secondary market with empirically estimated size-dependent bargaining, given by $\int_0^1 EV_i \, d\kappa_{-i}$, where EV_i is defined in equation (18). Note that a larger own size κ_i shifts the distribution of the counterparty's relative size $\kappa \equiv \kappa_{-i}/\kappa_i$ downward.

The right panel of Figure 4 incorporates heterogeneous R&D costs, where each firm's R&D efficiency, the inverse of cost, follows an i.i.d. Pareto distribution, consistent with evidence on the skewed distribution of research productivity (Kortum, 1997).[16] The curves plot the probability that a firm conducts R&D, defined as the fraction of cost draws for which the expected value exceeds the realized cost, as a function of firm size.

Across all three scenarios, both the expected value of developing a blueprint and the probability of innovation increase with firm size. Size-dependent bargaining yields the highest expected value and innovation probability for large firms with $\kappa_i > 0.5$. Symmetric bargaining yields the highest expected value and innovation probability for small firms with $\kappa_i \leq 0.5$, while the absence of a market for blueprints generates the lowest incentives.

Under the calibrated parameters, the aggregate share of firms developing blueprints is 20.2% without a secondary market, 29% under symmetric bargaining, and 28.8% under the empirically estimated size-dependent bargaining rule. These figures show that, in the calibrated model, a secondary market raises aggregate R&D relative to no market. They also show that the symmetric bargaining rule induces slightly higher aggregate innovation than the estimated size-dependent rule.[17]

[16] Scale and shape parameters are set to 1 and 2, respectively. Results are qualitatively similar with other parameter values.

[17] This result is based on a uniform firm-size distribution. The advantage of symmetric bargaining would be larger under a more realistic right-skewed distribution, where small firms are more prevalent.

Disaggregating by firm size illustrates the distributional effects of the bargaining rules. For small firms, $\kappa_i \in (0, 0.5]$, the R&D participation rates are 17.1% without a secondary market, 24.8% under symmetric bargaining, and 22.4% under size-dependent bargaining. For large firms, $\kappa_i \in (0.5, 1]$, the corresponding rates are 23.3%, 33.1%, and 35.2%. Two patterns emerge. First, large firms are more likely to innovate across all market structures, consistent with the upward-sloping relationships in Figure 4. Second, bargaining power has opposing effects by size: size-dependent bargaining provides the strongest incentive for large firms, while reducing R&D by small firms relative to the symmetric-bargaining counterfactual.

Figure 4: Innovation Incentives across Firm Size and Bargaining Rules

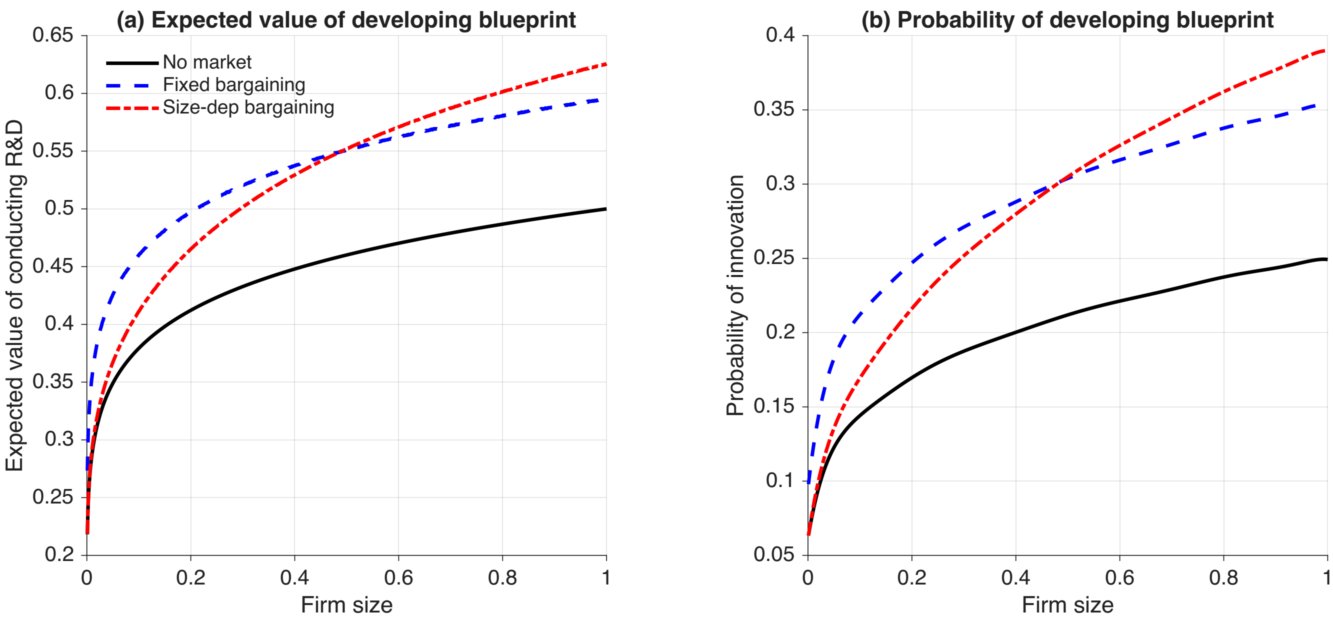

Note: The left panel shows the expected value of blueprint development for three market structures. The right panel displays the simulated probability of conducting R&D, conditional on R&D costs whose inverse follows a Pareto distribution. Both panels are plotted against firm size $\kappa_i \in (0, 1]$, with counterparty size drawn uniformly from $[0, 1]$.

The above results suggest that bargaining shares skewed toward larger firms weaken the R&D incentives of smaller firms. This dynamic may stifle creative destruction and potentially impede technological growth. Note that our model is stylized and abstract from several factors that might magnify this concern. In particular, large and small firms may target different types of innovation (e.g., Bena et al., 2016; Arora et al., 2018; Akcigit and Kerr, 2018; Arora et al., 2023; Fernández-Villaverde et al., 2025). Large firms are more likely to focus on incremental and exploitative innovations, which refine existing technologies and protect them from potential competition. In contrast, small firms are more likely to pursue radical and exploratory innovations. This suggests that suppressing smaller firms' innovation efforts might lead to a more significant loss in technological growth.

5 Limitations and Future Research

Our analysis has a few limitations. First, our estimates are identified from stock price reactions to patent transactions. Under our baseline interpretation, event-window returns reflect investors' assessment of the transaction-related change in firm value, and completed transactions are treated as trades that both parties expect to generate positive gains. This interpretation may fail if investors misunderstand the transaction, if information leaks before the market awareness date, or if managers and investors systematically diverge in their valuation. We mitigate these concerns in Appendices B, C, D, and E.1 using alternative event dates, longer windows, positive-return subsamples, and a specification with transient mispricing.

Second, our filtering procedure imposes a constant signal-to-noise ratio across transactions, following Kogan et al. (2017). Because each patent transaction is observed only once, the transaction-specific signal variance cannot be separately estimated; pooling information across transactions is necessary to extract the transaction-related component from noisy event-window returns. However, if the informativeness of stock price reactions varies systematically with transaction characteristics, our estimates of gains from trade could be affected. Appendix E.2 mitigates this concern by replicating our core results using raw announcement returns, which bypass the filtering procedure entirely. Future research could refine these estimates by adopting a more flexible specification in which the signal-to-noise ratio varies with observable transaction-level features.

Third, our estimates are conditional on completed patent transactions. We do not observe failed negotiations, the number or identity of competing bidders, or the full information set available to each party. These unobserved factors may influence outside options and bargaining outcomes in addition to relative firm size. We therefore interpret our findings as describing the division of surplus among realized trades, rather than the allocation that would emerge across all potential trading opportunities. Future work could extend the analysis to unrealized trading opportunities by using data from settings where the negotiation process is observable, such as patent auctions or transactions with publicly disclosed bidding processes.

Fourth, the model assumes that blueprint types are uniformly distributed over the technological space. This assumption provides analytical tractability, but the shape of the distribution can affect the quantitative importance of the secondary market. When blueprint types are drawn from a distribution concentrated near a firm's core business, poorly matched blueprints arise less frequently, which reduces the option value of resale. Conversely, a more dispersed distribution generates greater scope for reallocation, making bargaining in the secondary market more consequential for R&D incentives. Future research could allow the distribution of blueprint types to vary across firms and examine how its shape interacts with firm size and bargaining power.

6 Conclusion

We introduce a method for estimating the total surplus generated by patent transactions between publicly listed companies and the bargaining shares received by both parties. Unlike previous approaches that focus primarily on the patent buyer, our approach uses stock-market reactions of both transacting firms to estimate the market-assessed gains on each side of the transaction.

Using this estimation approach, we analyze patent transactions among publicly traded U.S. companies from 1975 to 2012. Our empirical results reveal two key insights. First, bargaining shares vary significantly across transactions, with firm size playing a crucial role in determining each party's share, potentially influenced by the incentive alignment between shareholders and the managers overseeing these transactions. Second, larger firms are associated with greater total surplus from patent transactions, consistent with scale advantages in commercialization and lower transaction costs.

We then develop a model to organize these empirical findings and examine their implications for R&D incentives. The model shows that a secondary market for patents can raise aggregate R&D incentives by allowing firms to sell poorly matched inventions to better-matched users. However, when bargaining shares favor larger firms, the gains from this market are unevenly distributed. In the calibrated model, size-dependent bargaining shifts innovation incentives toward larger firms and weakens the R&D incentives of smaller firms relative to a symmetric-bargaining benchmark. Notably, the substantial bargaining power held by large firms diminishes the R&D incentives of smaller firms and may impede creative destruction.

These findings carry managerial implications for firms and entrepreneurs that use intellectual property as a potential source of monetization. Managers of smaller firms should not presume that patents misaligned with their core business can be readily monetized through sales, as limited bargaining power may confine them to capturing only a small fraction of the resulting surplus. This suggests that smaller firms may benefit from directing R&D toward technologies closely related to their core business, where the resulting patents can be deployed internally, rather than relying on the resale of less-related technologies to recover R&D investment. Although our empirical evidence draws on transactions between publicly listed firms, the same concern may extend to startups and private firms negotiating with substantially larger counterparties. When evaluating R&D projects that may ultimately be sold in the patent market, smaller firms should therefore consider not only the potential buyer's technological fit or commercialization capacity, but also how to strengthen their own bargaining position, for instance, by developing credible outside options and running structured, competitive bidding processes.

Our evidence on corporate governance also suggests that strengthening internal shareholder oversight can help smaller firms avoid managerial opportunism, such as selling to preferred counterparties without adequate competition.

For larger technology buyers, our results suggest that scale and commercialization advantages can make them valuable counterparties, but aggressive surplus capture may weaken smaller firms' incentives to generate future tradable inventions.

References

Akcigit, Ufuk and William R Kerr, "Growth through heterogeneous innovations," *Journal of Political Economy*, 2018, *126* (4), 1374–1443.

_, Murat Alp Celik, and Jeremy Greenwood, "Buy, keep, or sell: Economic growth and the market for ideas," *Econometrica*, 2016, *84* (3), 943–984.

Andersen, Torben G and Timo Teräsvirta, "Realized volatility," in "Handbook of financial time series," Springer, 2009, pp. 555–575.

Arora, Ashish and Marco Ceccagnoli, "Patent protection, complementary assets, and firms' incentives for technology licensing," *Management Science*, 2006, *52* (2), 293–308.

_, Sharon Belenzon, and Andrea Patacconi, "The decline of science in corporate R&D," *Strategic Management Journal*, 2018, *39* (1), 3–32.

—, —, and Jungkyu Suh, "Science and the market for technology," *Management Science*, 2022, *68* (10), 7176–7201.

—, —, Matt Marx, and Dror Shvadron, "When does patent protection spur cumulative research within firms?," *Journal of Law, Economics, and Organization*, 2023, p. ewad006.

Arqué-Castells, Pere and Daniel F Spulber, "Measuring the private and social returns to R&D: Unintended spillovers versus technology markets," *Journal of Political Economy*, 2022, *130* (7), 1860–1918.

— and —, "Firm Matching in the Market for Technology: Business Stealing and Business Creation," *Journal of Industrial Economics*, 2023, *71* (4), 961–1003.

Bebchuk, Lucian, Alma Cohen, and Allen Ferrell, "What matters in corporate governance?," *Review of Financial Studies*, 2009, *22* (2), 783–827.

Bena, Jan, Lorenzo Garlappi, and Patrick Grüning, "Heterogeneous innovation, firm creation and destruction, and asset prices," *Review of Asset Pricing Studies*, 2016, *6* (1), 46–87.

Bondt, Werner FM De and Richard Thaler, "Does the stock market overreact?," *Journal of Finance*, 1985, *40* (3), 793–805.

Bustamante, Maria Cecilia and Francesca Zucchi, "Innovation, industry equilibrium, and discount rates," *ECB Working paper*, 2023.

Carley, Michael, Deepak Hedge, and Alan Marco, "What is the probability of receiving a US patent," *Yale JL & Tech.*, 2015, *17*, 203.

Cohen, Randolph B, Paul A Gompers, and Tuomo Vuolteenaho, "Who underreacts to cash-flow news? Evidence from trading between individuals and institutions," *Journal of Financial Economics*, 2002, *66* (2-3), 409–462.

Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist, "What is a patent worth? Evidence from the US patent "lottery"," *Journal of Finance*, 2020, *75* (2), 639–682.

Fernández-Villaverde, Jesús, Yang Yu, and Francesco Zanetti, "Defensive Hiring and Creative Destruction," Technical Report, National Bureau of Economic Research 2025.

Figueroa, Nicolás and Carlos J Serrano, "Patent trading flows of small and large firms," *Research Policy*, 2019, *48* (7), 1601–1616.

Fleming, Lee, "Recombinant uncertainty in technological search," *Management science*, 2001, *47* (1), 117–132.

— and Olav Sorenson, "Science as a map in technological search," *Strategic Management Journal*, 2004, *25* (8-9), 909–928.

Galasso, Alberto and Mark Schankerman, "Patent thickets, courts, and the market for innovation," *The RAND Journal of Economics*, 2010, *41* (3), 472–503.

—, —, and Carlos J Serrano, "Trading and enforcing patent rights," *The RAND Journal of Economics*, 2013, *44* (2), 275–312.

Gans, Joshua S and Scott Stern, "Incumbency and R&D incentives: Licensing the gale of creative destruction," *Journal of Economics & Management Strategy*, 2000, *9* (4), 485–511.

—, David H Hsu, and Scott Stern, "The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent litigation," *Management Science*, 2008, *54* (5), 982–997.

Graham, Stuart JH, Alan C Marco, and Amanda F Myers, "Patent transactions in the marketplace: Lessons from the USPTO patent assignment dataset," *Journal of Economics & Management Strategy*, 2018, *27* (3), 343–371.

Hegde, Deepak and Hong Luo, "Patent publication and the market for ideas," *Management Science*, 2018, *64* (2), 652–672.

Hochberg, Yael V, Carlos J Serrano, and Rosemarie H Ziedonis, "Patent collateral, investor commitment, and the market for venture lending," *Journal of Financial Economics*, 2018, *130* (1), 74–94.

Hong, Harrison and Jeremy C Stein, "A unified theory of underreaction, momentum trading, and overreaction in asset markets," *Journal of Finance*, 1999, *54* (6), 2143–2184.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological innovation, resource allocation, and growth," *Quarterly Journal of Economics*, 2017, *132* (2), 665–712.

Kortum, Samuel S, "Research, patenting, and technological change," *Econometrica: Journal of the Econometric Society*, 1997, pp. 1389–1419.

Kwon, Jung H, Haemin Dennis Park, and Shu Deng, "When do firms trade patents?," *Organization Science*, 2022, *33* (3), 1212–1231.

Li, Di and Erica XN Li, "Corporate governance and costs of equity: Theory and evidence," *Management Science*, 2018, *64* (1), 83–101.

Ma, Yueyuan, "Specialization in a knowledge economy," *Available at SSRN 4052990*, 2022.

— and Shaoshuang Yang, "Technology Driven Market Concentration through Idea Allocation," Technical Report, Working Paper 2023.

Marco, Antonio De, Giuseppe Scellato, Elisa Ughetto, and Federico Caviggioli, "Global markets for technology: Evidence from patent transactions," *Research Policy*, 2017, *46* (9), 1644–1654.

Marx, Matt and Emma Scharfmann, "Does Patenting Promote the Progress of Science?," 2024.

Ottaviani, Marco and Peter Norman Sørensen, "Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal," *American Economic Review*, 2015, *105* (1), 01–34.

Serrano, Carlos J, "The dynamics of the transfer and renewal of patents," *The RAND Journal of Economics*, 2010, *41* (4), 686–708.

— , "Estimating the gains from trade in the market for patent rights," *International Economic Review*, 2018, *59* (4), 1877–1904.

— **and Rosemarie H Ziedonis**, "Do patent assets have a second life when startups fail? An analysis of the redeployment likelihood and mode of transfer," *Strategic Management Journal*, 2024.

Siegel, Donald S, David Waldman, and Albert Link, "Assessing the impact of organizational practices on the relative productivity of university technology transfer offices: an exploratory study," *Research Policy*, 2003, *32* (1), 27–48.

Stagni, Raffaele Morandi and Juan Santalo, "Investors' attention and the paradox of technologically related diversification: Evidence of stock market mispricing," *Strategic Management Journal*, 2025, *46* (10), 2432–2466.

Ziedonis, Rosemarie H and Alberto Galasso, "Patent rights and innovation: evidence from the semiconductor industry," in "Research Handbook on the Economics of Intellectual Property Law," Edward Elgar Publishing, 2019, pp. 423–444.

Online Appendix: Material not Considered for Publication

A Proof of Equation (4)

Proof. To derive $E[w_{i,j} \mid R_{i,j}]$, we first determine the joint distribution of $w_{i,j}$ and $R_{i,j}$. Since $w_{i,j}$ and $\epsilon_{i,j}$ are independent, their joint density function is:

$$f_{w_{i,j},R_{i,j}}(w,r) = f_{w_{i,j}}(w) f_{R_{i,j}|w_{i,j}}(r \mid w),$$

where $f_{w_{i,j}}(w)$ is the density function of the truncated normal distribution $\mathcal{N}^{+}(0, \sigma^2_{w,i,t})$, and $f_{R_{i,j}|w_{i,j}}(r \mid w)$ is the density function of $R_{i,j}$ given $w_{i,j}$, which follows $\mathcal{N}(w, \sigma^2_{\epsilon,i,t})$.

The density function of $w_{i,j}$ is:

$$f_{w_{i,j}}(w) = 2\frac{1}{\sqrt{2\pi\sigma^2_{w,i,t}}} \exp\left(-\frac{w^2}{2\sigma^2_{w,i,t}}\right) \text{ for } w \geq 0.$$

where ϕ denotes the standard normal density function and Φ denotes the standard normal cumulative distribution function (CDF).

While the density function of $R_{i,j}$ given $w_{i,j}$ is:

$$f_{R_{i,j}|w_{i,j}}(r \mid w) = \frac{1}{\sqrt{2\pi\sigma^2_{\epsilon,i,t}}} \exp\left(-\frac{(r-w)^2}{2\sigma^2_{\epsilon,i,t}}\right).$$

Given $R_{i,j}$, the distribution of $w_{i,j}$ follows:[18]

$$w_{i,j} \mid R_{i,j} \sim \mathcal{N}(\mu, \sigma^2),$$

[18]We use the result of the conditional normal distribution to derive the distribution of $w_{i,j} \mid R_{i,j}$ If X and Y follow a bivariate normal distribution:

$$\begin{pmatrix} X \\ Y \end{pmatrix} \sim \mathcal{N}\left(\begin{pmatrix} \mu_X \\ \mu_Y \end{pmatrix}, \begin{pmatrix} \sigma^2_X & \sigma_{XY} \\ \sigma_{XY} & \sigma^2_Y \end{pmatrix}\right),$$

then the conditional distribution $X \mid Y$ is:

$$X \mid Y = y \sim \mathcal{N}\left(\mu_X + \frac{\sigma_{XY}}{\sigma^2_Y}(y-\mu_Y), \sigma^2_X - \frac{\sigma^2_{XY}}{\sigma^2_Y}\right).$$

where:

$$\mu = \frac{\sigma^2_{w,i,t}}{\sigma^2_{w,i,t} + \sigma^2_{\epsilon,i,t}} R_{i,j},$$

$$\sigma^2 = \frac{\sigma^2_{w,i,t}\sigma^2_{\epsilon,i,t}}{\sigma^2_{w,i,t} + \sigma^2_{\epsilon,i,t}}.$$

Since $w_{i,j}$ is truncated at zero, the conditional expectation $E[w_{i,j} \mid R_{i,j}]$ for a normally distributed random variable $X \sim \mathcal{N}(\mu, \sigma^2)$ truncated at zero is:

$$E[X \mid X \geq 0] = \mu + \sigma \frac{\phi\left(-\frac{\mu}{\sigma}\right)}{1 - \Phi\left(-\frac{\mu}{\sigma}\right)}.$$

Substituting $\delta_{ft} = \frac{\sigma^2_{w,i,t}}{\sigma^2_{w,i,t}+\sigma^2_{\epsilon,i,t}}$ into these formulas, we obtain:

$$E[w_{i,j} \mid R_{i,j}] = \delta_{ft} R_{i,j} + \sqrt{\delta_{ft}} \sigma_{\epsilon,i,t} \frac{\phi\left(-\sqrt{\delta_{ft}} \frac{R_{i,j}}{\sigma_{\epsilon,i,t}}\right)}{1 - \Phi\left(-\sqrt{\delta_{ft}} \frac{R_{i,j}}{\sigma_{\epsilon,i,t}}\right)}.$$

□

B The Choice of Market Awareness Date

In this appendix, we discuss the implication of our choice of the market awareness date to our empirical findings.

B.1 Robustness to Alternative Market Awareness Dates

Our benchmark market awareness date is defined as the earliest among three dates: the execution date and record date from the USPTO Patent Assignment Dataset, and the news date identified through a two-step process involving a large language model and manual verification. While we consider this a reasonable proxy for the actual market awareness date, certain measurement errors may arise due to (1) incomplete news coverage and (2) the slow diffusion of information.

In the first case, investors may be aware of a potential transaction even if it is not confirmed in mainstream news. This leads to a decline in measured total surplus. Formally, the measured gain from trade following the transaction confirmation date (our market awareness date) captures only a $1-\pi_j$ fraction of the actual gain from trade $\xi_{i,j}$, where π_j denotes the probability investors assign to a successful transaction based on information not captured by our analysis.[19] However, this issue only affects the estimates of total surplus, while do not affect the estimates of bargaining shares, as the same perceived probability of a successful transaction π_j applies to the two sides of the transaction. As the result, the estimated bargaining share

$$\widehat{\tau}_{s,j} = \frac{\xi_{s,j} \cdot (1-\pi_j)}{\xi_{s,j} \cdot (1-\pi_j) + \xi_{b,j} \cdot (1-\pi_j)}$$

is equal to the actual bargaining share

$$\tau_{s,j} = \frac{\xi_{s,j}}{\xi_{s,j} + \xi_{b,j}}$$

In the second case, investors may overlook the patent transaction even after its announcement, leading to an incomplete incorporation of trade gains into stock prices following the market awareness date. As we demonstrate in the next section, this issue is unlikely to be significant, as we do not observe significant momentum in stock prices post-market awareness date, an outcome that would be expected if underpricing were substantial. Moreover, as in the previous case, slow information diffusion impacts only the estimated gains from trade and total surplus, while having minimal effect on the estimated bargaining share. Since lack of attention to a transaction affects both parties equally, its influence on bargaining dynamics remains limited.

To demonstrate that the choice of market awareness date has minimal impact on the estimated bargaining shares, for the subset of transactions with an available news date, we re-estimate the bargaining share twice-once using the execution date and once using the news date as the market awareness date. Figure 5 demonstrates that the estimated bargaining shares remain highly consistent across these two specifications, with a correlation coefficient of 0.992. This confirms the theoretical prediction that the measure of bargaining share is robust to variations in the timing of market awareness.

To summarize, Facts 2 and 3 remain robust to the choice of market awareness date that they mainly discuss the bargaining shares. In comparison, the robustness of Fact 4, related to total

[19]Patent transactions often involve significant uncertainty and may not always be finalized due to factors such financing challenges, regulatory hurdles, or strategic shifts.

Figure 5: Robustness to the Choice of Market Awareness Date

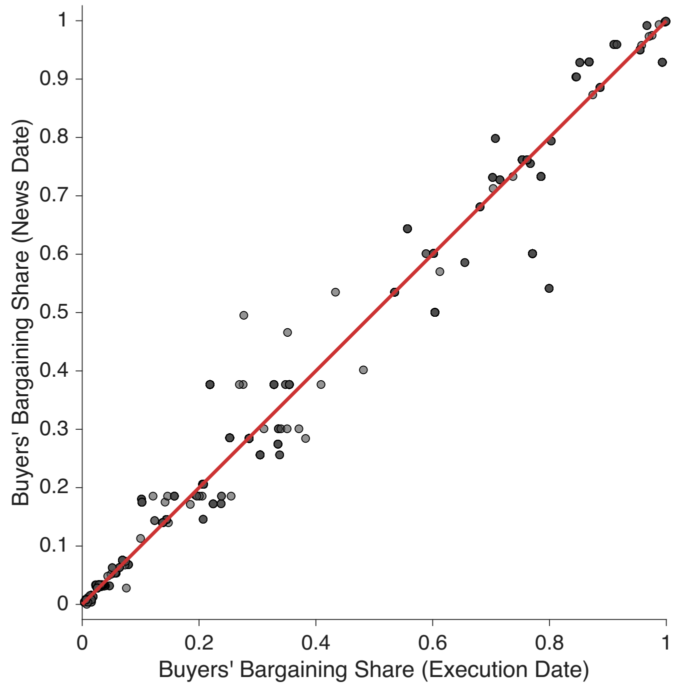

Notes This figure plots the buyers' estimated bargaining share using the execution date (x-axis) against the same share estimated using the news date (y-axis) for patent transactions where a news date is available. The 45-degree line represents perfect agreement between the two estimates. The points lie closely along the line, indicating that the estimated bargaining shares are highly consistent regardless of the choice of market awareness date.

surplus, depends on the variations in perceived success probabilities and speed of information diffusion across transactions. If these variations are minimal, as assumed to be constant in Kogan et al. (2017), Fact 4 remains largely unaffected: a constant $1 - \pi_j$ does not impact our point estimates once we take the logarithm of $\xi_{i,j} \cdot (1 - \pi_j)$.

B.2 Transactions with Explicit News Coverage

In this section, we estimate the core empirical results corresponding to Facts 2 and Fact 3 using this restricted sample covered by news. As shown in Figure 6 and Table B.1, the qualitative patterns documented in the main text remain robust.

Following recent literature on automated content analysis, we utilized a Large Language Model (Gemini) to categorize the strategic tone and rationale of these deals. For each transaction, we provided the model with the specific news URL and a standardized prompt designed to extract the strategic rationale and the directional impact of the deal. Specifically, the model was instructed to: "Analyze the following news article regarding a patent-related transaction.

Figure 6: Distribution of Bargaining Shares, News Covered Sample

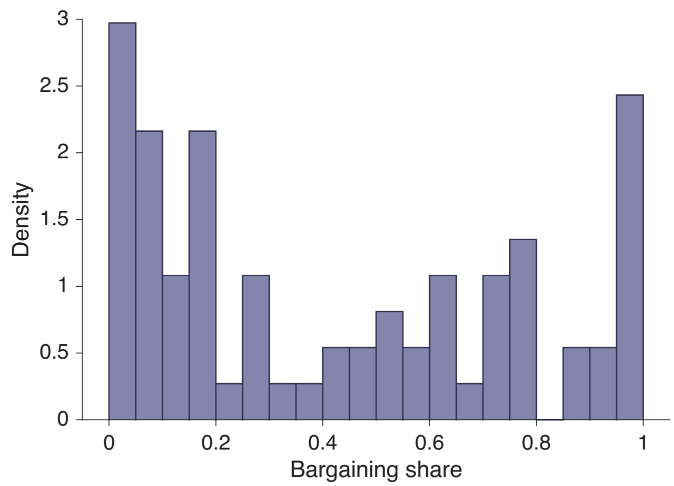

Note: This figure illustrates the distribution of buyers' bargaining shares estimated from patent transactions with explicit news coverage. Because some transactions are reported multiple times on the same day, we collapse duplicate buyer–seller observations so that each transaction contributes only once to the density. Without this adjustment, repeated mentions would mechanically create spikes at specific values. Applying the same de-duplication procedure to the full sample yields robust results.

Table B.1: Bargaining Share is Positively Correlated with Relative Size, News Covered Subsample

The dependent variable is the buyers' bargaining share, estimated from patent transactions that are explicitly covered by news reports. The independent variable is the buyers' relative size compared to the sellers. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)
Dependent Variable:	*Bargaining Share*	
Relative Size	0.662***	0.944***
	(0.229)	(0.120)
Year FE	Yes	-
Industry Pair and Patent Class FE	Yes	-
Observations	4,601	4,658
R-squared	0.994	0.760

Identify the primary buyer and seller, determine whether the tone of the report suggests a positive or negative strategic outcome for each party." This standardized prompting protocol allowed us to transform qualitative journalistic descriptions into a structured categorical dataset, providing a consistent and scalable way to validate the underlying economic assumptions of our bargaining model.

The analysis indicates that 88.2% of the articles describe the transactions as value-enhancing for both parties, which provides direct qualitative evidence for our fundamental assumption that participants engage in transactions that generate positive gains from trade. These deals typically fall into two broad strategic categories: strategic reallocation of patent portfolios—either through the integration of complementary technologies or the divestiture of non-core patents—and litigation settlements, where patent transfers resolve legal disputes and enhance business certainty.

The remaining 11.8% of articles report a negative outlook for the seller side, typically citing distressed asset liquidations or forced divestitures due to industry downturns. Crucially, even in these fire-sale scenarios, the gain from trade is likely to remain positive. In these instances, the seller's payoff in the counterfactual scenario—holding onto the patent during bankruptcy or a market collapse—would be even lower than the realized payoff.

Table B.2 provides a representative overview of the transactions identified through our systematic textual analysis, categorized by their strategic rationale and economic impact. The first column identifies the parties involved in the transaction, while the subsequent two columns indicate the direction of the strategic impact on the buyer and seller, respectively. The final two columns provide the economic reasoning for the transfer and its broader strategic classification. The first four entries illustrate the most frequent outcome in our sample: a "positive-positive" impact where the trade facilitates growth or operational independence for both parties. In contrast, the remaining two entries represent cases involving negative impacts, typically occurring when a seller is forced to liquidate assets due to broader financial distress or when a project fails to reach commercial viability.

C Different Window Length

A natural concern in our empirical strategy is the reliance on short-term stock price movements to estimate gains from trade and, subsequently, the bargaining share. Financial markets may not always perfectly price the economic impact of news upon arrival. In particular, prior literature documents that markets may overreact to firm-specific events, leading to subsequent price reversals, or underreact, resulting in delayed price adjustments over time (De Bondt and Thaler, 1985; Hong and Stein, 1999; Cohen et al., 2002; Ottaviani and Sørensen, 2015). If this occurs, our benchmark one-day window may not fully capture the total valuation effects of the patent transaction.

Table B.2: Representative Patent Transactions and Strategic Impact Analysis

l c c X l

Transaction	**Buyer**	**Seller**	**Reasoning**	**Category**
A.O. Smith / MagneTek	Pos	Pos	[l]Acquisition of motor technology to expand core industrial product line.	Core Strategic Growth
Microsoft / Connectix	Pos	Pos	[l]Integration of virtualization tech into Windows ecosystem to compete with VMware.	Core Strategic Growth
Qualcomm / Nokia	Pos	Pos	[l]Long-term settlement and cross-license ending multi-year global patent litigation.	Litigation Settlement
RIM / Motorola	Pos	Pos	[l]Strategic cross-license of wireless 2G/3G/4G IP to ensure freedom to operate.	Litigation Settlement
Alcatel / Avanex	Pos	Neg	[l]Forced divestiture of optical assets at low valuation due to telecom market crash.	Distressed Liquidation
SGI / Graphics IP	Pos	Neg	[l]Liquidation of high-end 3D graphics patents to raise cash during financial decline.	Distressed Liquidation

A related concern is whether public markets can fully price the gains from trade when private information–such as the value of a firm's outside options if the patent deal does not occur–is not observable. Although we cannot directly test for such information asymmetry between insiders and investors, its empirical relevance can be indirectly assessed. If public investors lack sufficient information to fully price the transaction upon initial disclosure, prices may subsequently adjust as more information becomes available or is better understood. In this case, we would expect to observe systematic price drifts or reversals over longer horizons.

To address these concerns, we conduct a robustness check by extending the length of the event window beyond the one-day benchmark used in the main analysis. Specifically, we re-estimate all key results using three-day, five-day, and seven-day windows around the market awareness date. This approach allows us to capture slower market adjustments and potential corrections after the initial reaction.

Figure 7 compares total surplus estimates using 3-, 5-, and 7-day windows against the benchmark result obtained with a 1-day window. Our estimates of gains from trade and bargaining share remain highly stable across different window lengths: The estimates across different window lengths align almost perfectly, suggesting minimal short-term over- or under-pricing.

Figure 7: Total surplus estimated using different window length

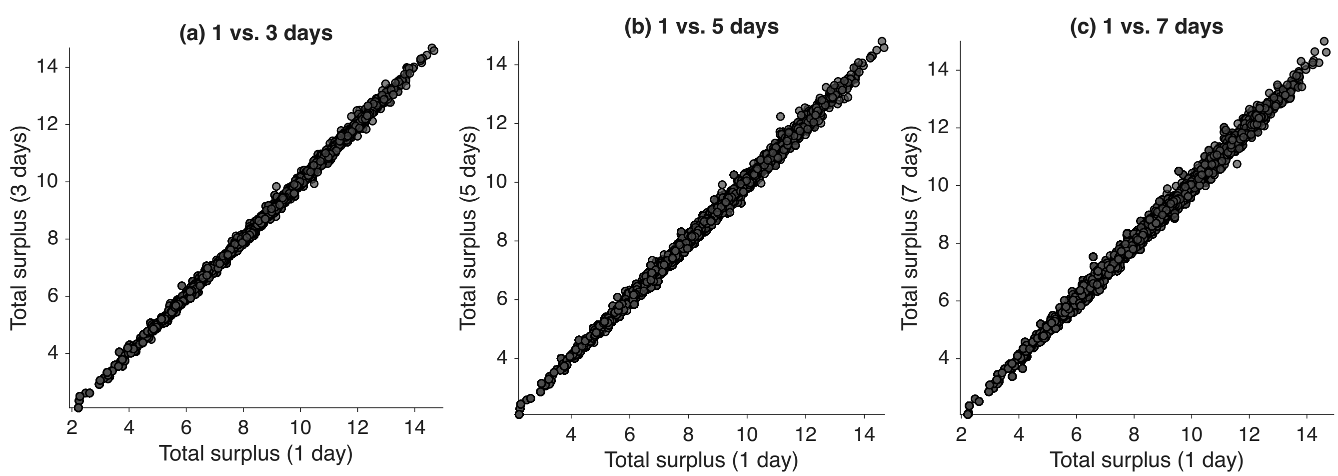

Notes This figure compares total surplus estimates using window lengths of 3, 5, and 7 days against the benchmark result obtained with a 1-day window.

Figure 8 shows that Fact 2 remains robust to different window length by comparing bargaining share estimates using 3-, 5-, and 7-day windows against the benchmark result obtained with a 1-day window (identical to the left panel of Figure 2). The distributions of bargaining shares remain nearly unchanged across different window lengths, indicating that the U-shaped pattern in bargaining share distribution is robust to variations in window length.

Figure 8: Bargaining shares estimated using different window length

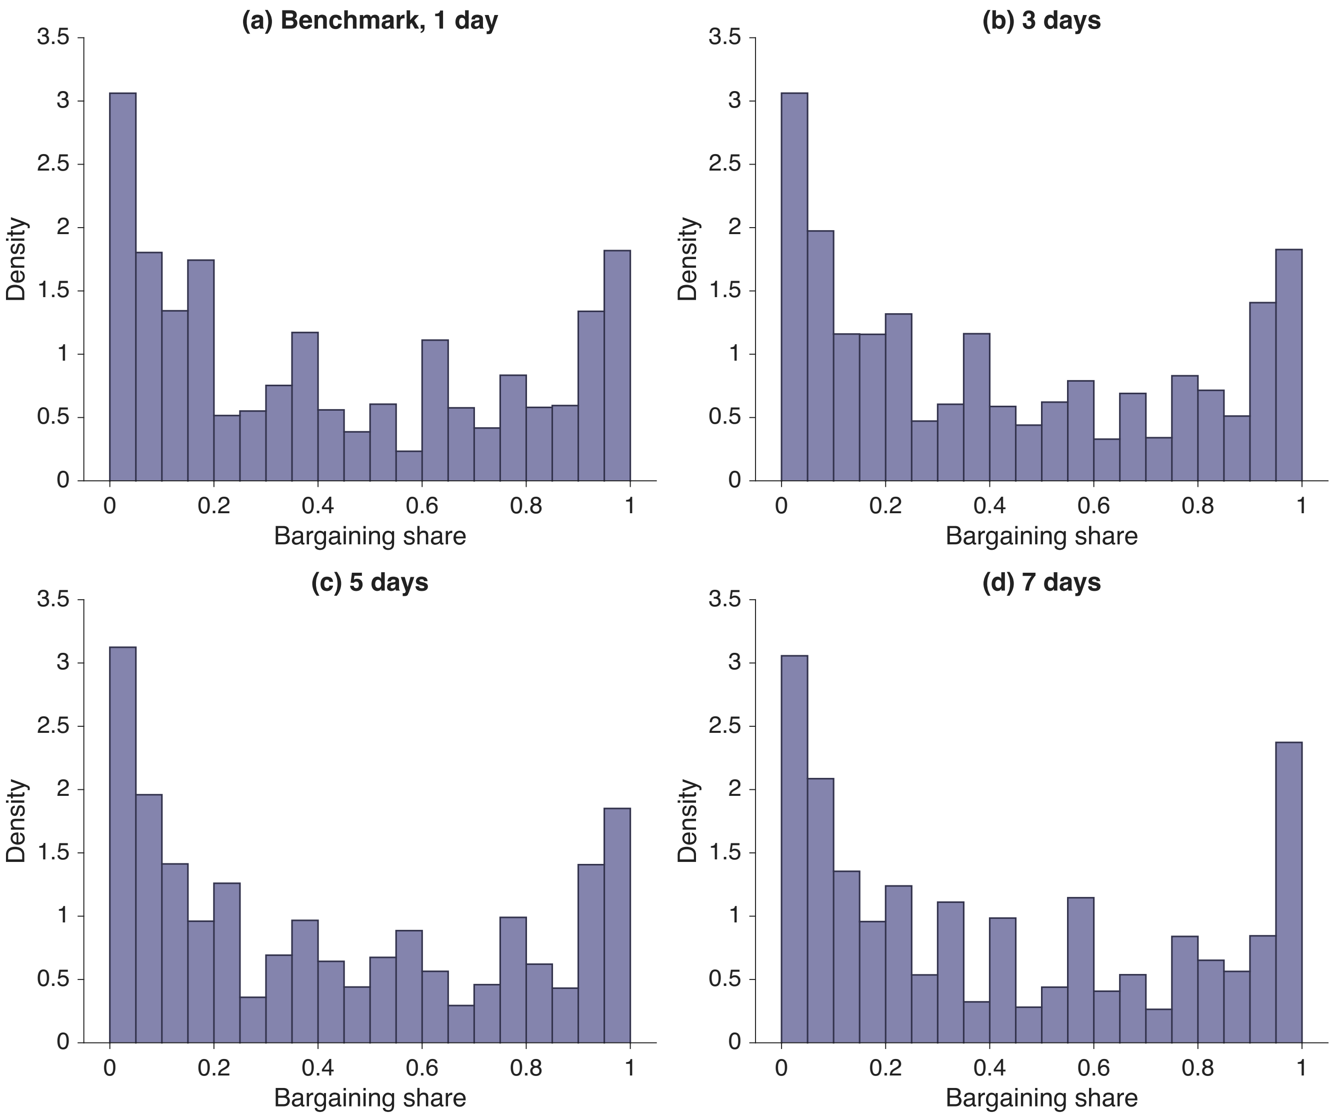

Notes This figure compares bargaining share estimates using window lengths of 3, 5, and 7 days against the benchmark result obtained with a 1-day window.

Table C.1 and Table C.2 show that Facts 3 and 4 remain robust to different window length. These results suggest that our conclusions are not driven by short-term mispricing or delayed market reactions, and that the stock price-based estimation strategy remains reliable even in the presence of potential information frictions.

Table C.1: Bargaining Share is Positively Correlated with Relative Size, Different Window Length

The dependent variable is the bargaining share of the buyers, estimated using different length of windows. The independent variable is the buyers' relative size compared to the sellers. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable:	*3 days*		*5 days*		*7 days*	
Relative Size	0.760***	0.942***	0.759***	0.938***	0.759***	0.934***
	(0.032)	(0.024)	(0.034)	(0.025)	(0.034)	(0.025)
Year FE	Yes	-	Yes	-	Yes	-
Industry Pair and Patent Class FE	Yes	-	Yes	-	Yes	-
Observations	26,122	26,358	26,122	26,358	26,122	26,358
R-squared	0.870	0.784	0.867	0.780	0.866	0.779

Table C.2: Total Surplus Increases with Firm Size, Different Window Length

The dependent variable is the total surplus, estimated using different length of windows. The independent variables are the sizes of the buyers and sellers, measured by the logarithm of their real sales. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)	(3)
Dependent Variable:	*3 days*	*5 days*	*7 days*
Buyer's Size	0.248***	0.245***	0.247***
	(0.055)	(0.054)	(0.055)
Seller's Size	0.121**	0.123**	0.127**
	(0.053)	(0.053)	(0.054)
Controls			
$Citation$	Yes	Yes	Yes
$InitialValue$	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Industry Pair and Patent Class FE	Yes	Yes	Yes
Observations	26,034	26,034	26,034
R-squared	0.745	0.743	0.743

D Estimation based on transactions with positive returns

The idiosyncratic stock return $R_{i,j}$ defined in Equation (4) can be either positive or negative. A negative idiosyncratic return suggests that the firm's gain from trade through the patent transaction is outweighed by unrelated negative factors affecting its stock price. In such cases, our empirical implementation yields a small but positive gain from trade. Approximately 25% of transactions result in positive returns for both the buyer and the seller. The remaining 75% involves at least one party exhibits a non-positive return.[20] Table D.1 reports the idiosyncratic stock return of buyers $R_{b,j}$ over the one-day window following the transaction; the filtered return component for buyers $E[w_{b,j} \mid R_{b,j}]$ attributed to gains from trade (based on Equation (4)); and the buyer's gain from trade $\xi_{b,j}$ as defined in Equation (3), deflated to 1982 dollars using the Consumer Price Index (CPI). The table also presents the corresponding three variables for sellers: $R_{s,j}$, $E[w_{s,j} \mid R_{s,j}]$, and $\xi_{s,j}$. Our baseline estimation involve all transactions, including the ones with negative returns.

Next, we show that our results are robust to focusing on the transactions with positive returns for both the buyer and seller.

Fact 2. Figure 9 illustrates the distribution of bargaining shares across transactions in the restricted sample, revealing a pattern similar to that observed in the full sample (Figure 2). The distributions exhibit a U-shaped curve, with concentrations at both extremes, suggesting that in most cases, the total surplus from a transaction is heavily skewed toward one party rather than being evenly divided.

Fact 3. Table D.2 replicates the estimation results from Fact 3 using transactions in the restricted sample. The results are similar to the full-sample case: traders' bargaining shares are proportional to their size relative to their counterparties.

Fact 4. Finally, Table D.3 replicates the estimation results from Fact 4 using transactions in the restricted sample. The results are similar to the full-sample case: The total surplus of patent transactions increases with both the buyer's and the seller's firm sizes.

[20] Similar pattern is found in Kogan et al. (2017) in analyzing the patent values.

Table D.1: Detailed Descriptive Statistics

Note: The table reports the distribution of the following variables across the patents in our sample: the idiosyncratic stock return of buyers $R_{b,j}$ on the one-day window following the patent transaction date; the filtered component of returns of buyers $E[w_{b,j}|R_{b,j}]$ related to the gain from patent transactions using Equation (4); the buyers' gains from trade ($\xi_{b,j}$) using Equation (3) deflated to 1982 dollars using the CPI; the idiosyncratic stock return of sellers $R_{s,j}$ on the one-day window following the patent transaction date; the filtered component of returns of sellers $E[w_{s,j}|R_{s,j}]$ related to the gain from patent transactions using Equation (4); the sellers' gains from trade ($\xi_{s,j}$) using Equation (3) deflated to 1982 dollars using the CPI.

Moment	$R_{b,j}$ (%)	$E[w_{b,j}\|R_{b,j}]$ (%)	$\xi_{b,j}$	$R_{s,j}$ (%)	$E[w_{s,j}\|R_{s,j}]$ (%)	$\xi_{s,j}$
Mean	2.23	0.36	4,251.51	0.23	0.35	6,851.86
Std. dev.	16.07	0.38	34,787.04	2.96	0.20	38,592.60
Percentiles						
p1	-7.18	0.09	3.04	-6.49	0.11	1.17
p5	-3.44	0.12	7.87	-3.75	0.14	8.50
p10	-2.21	0.15	11.89	-2.58	0.17	18.88
p25	-0.91	0.19	44.62	-1.38	0.21	64.50
p50	-0.04	0.27	250.43	0.03	0.29	386.11
p75	1.30	0.39	1,085.72	1.53	0.45	1,644.06
p90	3.44	0.61	4,115.72	3.30	0.57	7,969.98
p95	10.01	0.91	11,461.76	3.90	0.71	28,094.12
p99	146.52	3.39	73,517.88	6.70	0.98	145,589.30

Figure 9: Distribution of Bargaining Shares: Positive-Return Subsample

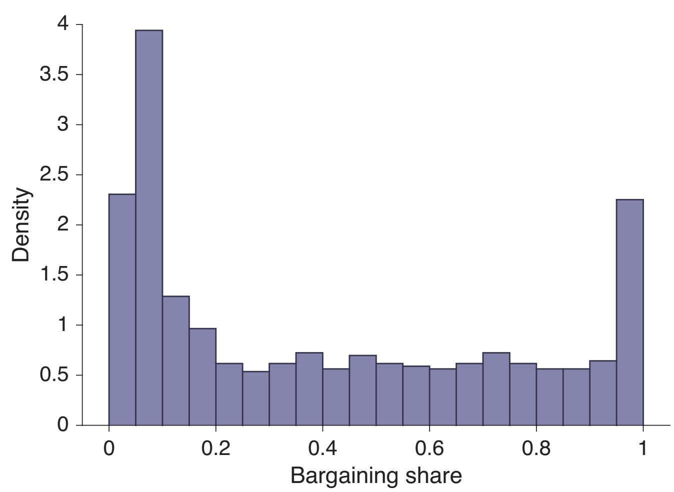

Notes This figure illustrates the distribution of buyers' bargaining shares, using only transactions where both parties have positive idiosyncratic returns.

E Alternative Econometric Approaches

This appendix demonstrates the robustness of our results to alternative methods for estimating gains from trade using stock price data. We consider two key departures from our baseline

Table D.2: Bargaining Share is Positively Correlated with Relative Size, Positive Return Subsample

The dependent variable is the bargaining share of the buyers, using only transactions where both parties have positive idiosyncratic returns. The independent variable is the buyers' relative size compared to the sellers. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)
Dependent Variable:	*Bargaining Share*	
Relative Size	0.781***	0.964***
	(0.069)	(0.035)
Year FE	Yes	-
Industry Pair and Patent Class FE	Yes	-
Observations	6,929	7,069
R-squared	0.944	0.817

Table D.3: Total Surplus Increases with Firm Size, Positive Return Subsample

The dependent variable is the total surplus, using only transactions where both parties have positive idiosyncratic returns. The independent variables are the sizes of the buyers and sellers, measured by the logarithm of their real sales. Standard errors are in parentheses, *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)	(3)	(4)
Dependent Variable:	*Total Surplus*			
Buyer's Size	0.226**	0.228**	0.185**	0.333***
	(0.092)	(0.088)	(0.088)	(0.076)
Seller's Size	0.191*	0.169	0.227**	0.136*
	(0.107)	(0.107)	(0.089)	(0.072)
Controls				
Citation	-	Yes	Yes	Yes
InitialValue	-	Yes	Yes	Yes
Year FE	-	-	Yes	Yes
Industry Pair and Patent Class FE	-	-	-	Yes
Observations	7,068	7,068	7,067	6,928
R-squared	0.107	0.150	0.281	0.865

framework. First, we extend the model to allow for a transient wedge between managerial and market valuations, capturing potential mispricing due to investors' initial incomplete information. Second, we investigate a more parsimonious approach that bypasses the filtering procedure and relies exclusively on raw announcement returns. In both cases, our main qualitative findings remain unchanged.

E.1 Transient Divergence Between Managerial and Market Valuations

In this subsection, we relax the assumption of efficient market by introducing a transient divergence between managerial and market valuations stemming from investors' incomplete information at the time of a transaction announcement. As in the benchmark analysis, managers correctly assess the gain from trade and only engage in transactions with a positive expected value. However, investors may initially misprice these transactions, as mentioned in Stagni and Santalo (2025). We assume that such mispricing gradually decays as investors learn over time, causing stock prices converge toward managerial valuations.

Econometric framework. Formally, let t denote calendar days, and let t_0 denote the trading day immediately *before* the market becomes aware of transaction j (the market awareness date). Thus $t_0 + 1$ is the first trading day on which the news is incorporated into stock prices.

Let $R_{i,j}(t)$ denote the cumulative idiosyncratic return of firm i over the event window around transaction j, measured from t_0 through t (assume $t \geq t_0 + 1$). We model it as

$$R_{i,j}(t) = w_{i,j} + \theta_{i,j}\, e^{-\eta\,(t-t_0-1)} + \sum_{t'=t_0+1}^{t} \epsilon_{i,j,t'}, \tag{E.1}$$

where $w_{i,j}$ is the manager-perceived value impact of the transaction, $\theta_{i,j}$ captures the initial wedge between investor and managerial valuations, and $\eta > 0$ governs the rate at which this wedge decays over time. The term $\epsilon_{i,j,t'}$ represents idiosyncratic return innovations unrelated to the transaction on calendar day t'.

If $\theta_{i,j}$ is sufficiently negative, indicating that investors are initially over-pessimistic about the trade, the announcement-day stock price reaction may be negative; this initial divergence is subsequently corrected at a rate determined by η.

Define daily idiosyncratic returns as first differences of cumulative returns,

$$r_{i,j,t} \equiv R_{i,j}(t) - R_{i,j}(t-1), \quad \text{with } R_{i,j}(t_0) \equiv 0, \quad t \geq t_0 + 1.$$

Using equation (E.1), daily returns satisfy

$$r_{i,j,t} = \begin{cases} w_{i,j} + \theta_{i,j} + \epsilon_{i,j,t}, & t = t_0 + 1, \\ \theta_{i,j}\left[e^{-\eta(t-t_0-1)} - e^{-\eta(t-t_0-2)}\right] + \epsilon_{i,j,t}, & t \geq t_0 + 2, \end{cases} \tag{E.2}$$

which shows that the manager-perceived component $w_{i,j}$ loads only on the announcement day $t_0 + 1$, while the mispricing component $\theta_{i,j}$ generates a deterministic decaying pattern in subsequent returns.

Our specification restricts the shape of the mispricing correction to follow a common exponential profile governed by η, while allowing the initial mispricing amplitude $\theta_{i,j}$ to vary freely across events. Observing returns for each transaction over multiple post-announcement days therefore provides the variation needed to separately identify the common decay rate η from event-specific mispricing magnitudes.several days, provides the necessary variation to separately identify the common decay rate η from the event-specific mispricing magnitudes.

Estimating η. For a given firm-transaction pair (i, j), stacking post-announcement returns yields the matrix form representation:

$$\mathbf{r}_{i,j} = \theta_{i,j}\,\mathbf{g}(\eta) + \boldsymbol{\varepsilon}_{i,j},$$

the vector $\mathbf{g}(\eta) \equiv \left(e^{-\eta(t-t_0-1)} - e^{-\eta(t-t_0-2)}\right)_{t \geq t_0+2}$ captures the deterministic exponential decay profile and $\boldsymbol{\varepsilon}_{i,j}$ collects the idiosyncratic noise terms $(\epsilon_{i,j,t})_{t \geq t_0+2}$. For a given decay rate η, the least-squares estimator of the mispricing value $\theta_{i,j}$ is

$$\hat{\theta}_{i,j}(\eta) = \frac{\mathbf{g}(\eta)^\top \mathbf{r}_{i,j}}{\mathbf{g}(\eta)^\top \mathbf{g}(\eta)}.$$

Substituting $\hat{\theta}_{i,j}(\eta)$ into the sum of squared residuals gives a *profile objective function* that depends only on η. This function is obtained by profiling out the event-specific amplitudes $\theta_{i,j}$. For a single pair (i, j), the sum of squared residuals conditional on η is

$$\mathrm{SSR}_{i,j}(\eta) = \left\|\mathbf{r}_{i,j} - \hat{\theta}_{i,j}(\eta)\,\mathbf{g}(\eta)\right\|^2 = \mathbf{r}_{i,j}^\top \mathbf{r}_{i,j} - \frac{\left(\mathbf{g}(\eta)^\top \mathbf{r}_{i,j}\right)^2}{\mathbf{g}(\eta)^\top \mathbf{g}(\eta)}.$$

Aggregating across all transactions yields the profile objective

$$Q(\eta) = \sum_{i,j} \mathrm{SSR}_{i,j}(\eta) = \sum_{i,j} \left[\mathbf{r}_{i,j}^{\top} \mathbf{r}_{i,j} - \frac{\left(\mathbf{g}(\eta)^{\top} \mathbf{r}_{i,j}\right)^2}{\mathbf{g}(\eta)^{\top} \mathbf{g}(\eta)} \right].$$

The aggregate decay parameter η is then estimated by minimizing the total residual variance across all transactions:

$$\hat{\eta} = \arg\min_{\eta > 0} \; Q(\eta) = \arg\min_{\eta > 0} \; \sum_{i,j} \left\| \mathbf{r}_{i,j} - \hat{\theta}_{i,j}(\eta)\, \mathbf{g}(\eta) \right\|^2.$$

The estimation employs daily returns from $t = t_0 + 2$ to $t = t_0 + 30$ after the transaction announcement, excluding the announcement-day return ($t = t_0 + 1$) because it contains the unknown term $w_{i,j}$. To control for market-wide shocks affecting all firms on a given calendar day, we also include date fixed effects.

Applying this estimation procedure yields decay parameters of $\hat{\eta} \approx 5.28$ for buyer firms and $\hat{\eta} \approx 3.70$ for seller firms. These estimates imply a strikingly rapid dissipation of initial mispricings. Under an exponential decay process, the transitory mispricing component falls to approximately $e^{-\hat{\eta}} \approx$ 0.5–3% of its initial magnitude after a single trading day, consistent with a highly efficient stock market. This rapid adjustment aligns closely with the abnormal trading volume dynamics found in this paper: volume spikes sharply at the announcement and largely returns to baseline by the second trading day. Together, the parallel decay in price divergence and normalization of trading activity indicate that any gap between market prices and managerial valuations is almost entirely transitory.

Interestingly, the estimated decay rates are notably similar for buyers and sellers, indicating that the market's learning dynamics are broadly symmetric on both sides of the transaction. If anything, mispricing decays marginally more slowly for buyers, which is consistent with greater investor uncertainty about how an acquired patent will integrate into the buyer's existing technology portfolio, compared to the relatively clearer signal of a divestment decision by a seller.

Estimating $\theta_{i,j}$. Using the estimated decay parameter $\hat{\eta}$, we recover the event-specific amplitude of transient mispricing, $\hat{\theta}_{i,j}$, by projecting post-announcement daily idiosyncratic returns on the estimated exponential decay profile:

$$\hat{\theta}_{i,j} = \frac{\mathbf{g}(\hat{\eta})^\top \mathbf{r}_{i,j}}{\mathbf{g}(\hat{\eta})^\top \mathbf{g}(\hat{\eta})}.$$

This projection is performed separately for each firm–transaction pair (i, j), yielding an estimate $\hat{\theta}_{i,j}$ and its associated standard error.

Estimating gains from trade $\omega_{i,j}$ and $\xi_{i,j}$. We then implement a transaction-level correction based on statistical significance. For each event (i, j), if we cannot reject the null hypothesis $H_0 : \theta_{i,j} = 0$ at the 10% significance level, we set $\hat{\theta}_{i,j} = 0$. Otherwise, we interpret $\hat{\theta}_{i,j}$ as evidence of statistically significant transient mispricing and subtract it from the observed idiosyncratic return. This yields an adjusted idiosyncratic return:

$$\widetilde{R}_{i,j} = R_{i,j} - \hat{\theta}_{i,j},$$

where $R_{i,j}$ is the original idiosyncratic return from the baseline analysis.

We then re-estimate the entire surplus and bargaining-share framework exactly as in the main text, substituting the original return $R_{i,j}$ with its adjusted counterpart $\widetilde{R}_{i,j}$. This generates corrected estimates of firm-level gains from trade, total surplus, and bargaining shares, while maintaining all other assumptions and procedures from the baseline analysis.

Empirical results. We replicate the empirical exercises corresponding to Facts 2 and Fact 3 using these adjusted estimates. As shown by Figure 10 and Table E.1, the results remain quantitatively similar to those reported in the main text.

E.2 Using Raw Stock Price Reactions Without Filtering

In this section, we investigate a more parsimonious approach that bypasses the baseline filtering procedure and relies exclusively on raw announcement returns.

Our baseline specification, which follows Kogan et al. (2017), decomposes a firm's idiosyncratic return around a patent transaction into a transaction-related component and residual noise. Under parametric assumptions, we recover the expected transaction-related return, $E[w_{i,j} \mid R_{i,j}]$,

Figure 10: Distribution of Buyer Bargaining Shares, Adjusted for Potential Mispricing

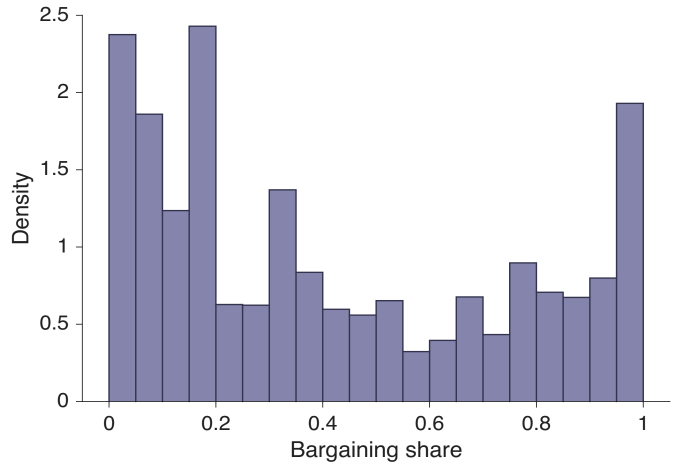

Notes: This figure plots the distribution of buyers' bargaining shares.

Table E.1: Bargaining Share is Positively Correlated with Relative Size

The dependent variable is the buyers' bargaining share. The independent variable is the buyers' relative size compared to the sellers. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)	(2)
Dependent Variable:	*Bargaining Share*	
Relative Size	0.759***	0.940***
	(0.033)	(0.024)
Year FE	Yes	-
Industry Pair and Patent Class FE	Yes	-
Observations	26,126	26,362
R-squared	0.869	0.783

and construct firm i's gain from trade as

$$\xi_{i,j} = E[w_{i,j} \mid R_{i,j}] \cdot \frac{M_{i,j}}{N_{i,j}}, \tag{E.3}$$

where $M_{i,j}$ denotes market capitalization and $N_{i,j}$ the number of patent transactions occurring on the same day. This approach requires a filtering procedure that relies on the underlying distribution of key variables.

As an alternative, we now construct gains from trade by directly substituting the observed idiosyncratic return for the filtered component:

$$\zeta_{i,j}^{\text{raw}} = R_{i,j} \cdot \frac{M_{i,j}}{N_{i,j}}.$$

This reduced-form specification makes no attempt to separate transaction-related gains from residual noise and therefore permits negative stock price reactions at the firm level. Consequently, a firm's estimated gain from trade is negative when its announcement return is negative.

Under this construction, the bargaining share is defined analogously to the baseline case:

$$\tau_{b,j}^{\text{raw}} = \frac{\zeta_{b,j}^{\text{raw}}}{\zeta_{b,j}^{\text{raw}} + \zeta_{s,j}^{\text{raw}}}.$$

Although we allow individual gains to be negative, we require the total surplus $\zeta_{b,j}^{\text{raw}} + \zeta_{s,j}^{\text{raw}}$ to be positive for the bargaining share to be interpretable. We therefore retain only transactions with positive total surplus and discard those where $\zeta_{b,j}^{\text{raw}} + \zeta_{s,j}^{\text{raw}} \leq 0$.

An important feature of this approach is that bargaining shares are no longer restricted to the unit interval. A negative $\tau_{b,j}^{\text{raw}}$ indicates that investors perceive the buyer as overpaying, while a value greater than one implies that the seller is perceived to have accepted a price below its valuation.

Empirical results. Using this alternative measure, we replicate the main empirical patterns concerning bargaining shares documented in Facts 2 and 3 and show that our core empirical findings do not depend on the specific filtering procedure or on mechanically excluding negative firm-level stock price reactions.

Figure 11 shows that when bargaining shares are constructed directly from raw idiosyncratic returns, the majority of observations still lie between zero and one. This indicates that in most cases, investors perceive patent transactions as value-enhancing for both parties. A small fraction of observations fall outside this range, reflecting instances in which market participants view the transaction as unfavorable for one side—precisely the type of behavior that the baseline filtering procedure abstracts from.[21]

[21] For expositional clarity, Figure 11 presents a kernel density plot of bargaining shares winsorized at the 5th and 95th percentiles. This winsorization is applied solely for visualization purposes; the regression results underlying Fact 3 are virtually unchanged when estimated on the full, unwinsorized sample.

Figure 11: Distribution of Buyer Bargaining Shares, Unfiltered

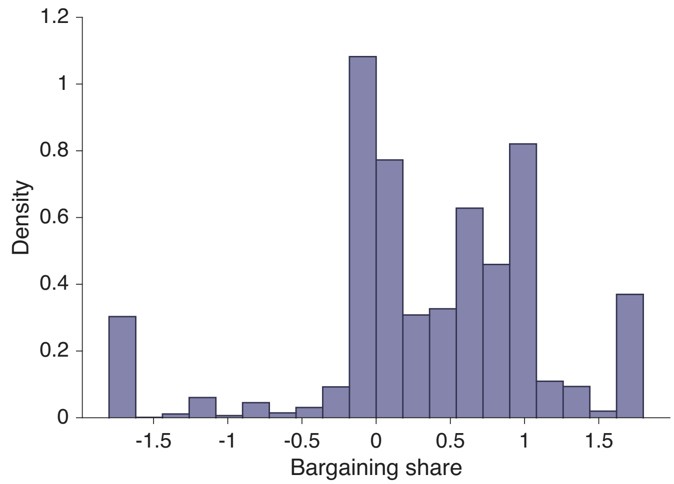

Notes: This figure plots the distribution of buyers' bargaining shares constructed directly from raw idiosyncratic stock returns, without applying the filtering procedure described in the main text. For visualization purposes, bargaining shares are winsorized at the 5th and 95th percentiles.

Table E.2: Raw Bargaining Share is Positively Correlated with Relative Size

The dependent variable is the buyers' raw bargaining share, winsorized at the 5th and 95th percentiles. The independent variable is the buyers' relative size compared to the sellers. Standard errors, reported in parentheses, are clustered at the firm level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

	(1)	(2)
Dependent Variable	*Raw Bargaining Share*	
Relative Size	0.653^{***}	0.814^{***}
	(0.124)	(0.128)
Year FE	Yes	-
Industry Pair and Patent Class FE	Yes	-
Observations	13,461	13,662
R-squared	0.837	0.212

Importantly, within the zero–one interval, the distribution of bargaining shares remains concentrated near the extremes. In particular, substantial mass continues to accumulate near zero and near one, mirroring our baseline finding that surplus division strongly favors one side of the transaction rather than being evenly split.

We further replicate the estimation of Fact 3 using these raw bargaining shares. The estimated coefficients are strikingly close to our benchmark results: 0.65 with fixed effects and 0.81 without fixed effects, compared to 0.76 and 0.94 in the filtered benchmark specification.

The explanatory power, as measured by R-squared, remains similarly high in the specification with fixed effects (0.83 versus 0.87 in the benchmark). In contrast, the R-squared falls substantially when fixed effects are omitted (0.21 versus 0.78 in the benchmark). This decline likely reflects substantial industry- and firm-level shocks unrelated to patent transactions, which are embedded in raw stock returns and dilute the measured relationship between bargaining shares and relative firm size. Our baseline filtering approach structurally and explicitly addresses this issue by separating transaction-related returns from residual noise under specific parametric assumptions, which helps to isolate the signal of interest.

F Sectoral differences in bargaining shares and total surplus

In this section, we examine sector-pair-level heterogeneity in total surplus and bargaining shares. Figures 12 and 13 present heatmaps depicting the average buyer's bargaining share and log total surplus across sector pairs, with the x-axis representing the buyer's industry and the y-axis representing the seller's industry. Shades of blue indicate higher values, while shades of yellow represent lower values.

As shown in Figures 12, the manufacturing sectors (NAICS codes 31–33) tend to have higher bargaining shares, particularly when engaging with counterparties from sectors such as transportation, information, and professional services (NAICS codes 48–62). However, this pattern may simply reflect the fact that manufacturing firms are, on average, larger, as firm size has been shown to strongly correlate with bargaining power.

As shown in Figures 13, total surplus exhibits significant variation across sector pairs. A natural hypothesis is that when a patent transfers from a seller in a less technology-dependent industry to a buyer in a more technology-dependent one, the buyer is better positioned to commercialize and apply the innovation effectively. This reallocation of intellectual property to a more productive use can generate greater gains from trade, resulting in a larger total surplus. Accordingly, we expect total surplus to be positively associated with the buyer's technological dependence and negatively associated with the seller's, controlling for other transaction- and firm-level factors.

Following a standard approach in the literature, we use per capita income within each industry as a proxy for technological dependence, given that technology-dependent industries

Figure 12: Heterogeneity in bargaining shares across sector-pairs

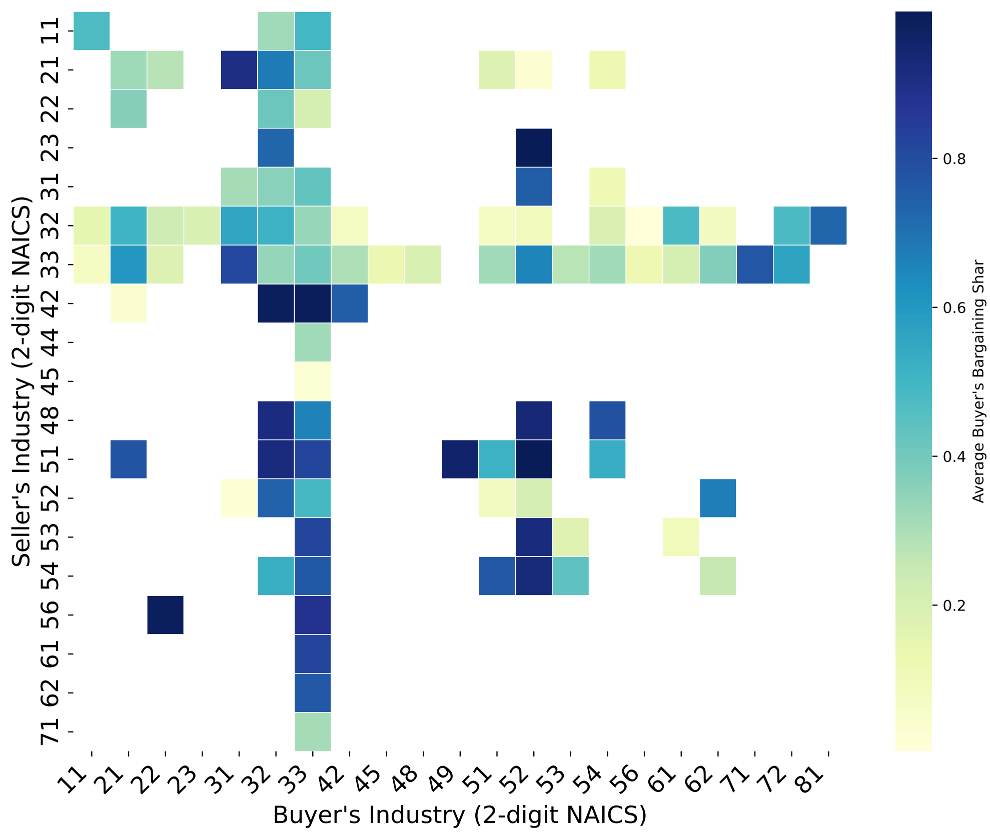

Notes This figure displays the average buyers' bargaining shares across industry pairs. Industries are categorized using two-digit NAICS codes, with the buyer's industry on the x-axis and the seller's industry on the y-axis.

tend to exhibit higher labor productivity and, consequently, higher average wage. Specifically, we estimate the following regression:

$$TS_{j,t} = \beta_1 Techdep_{b,j} + \beta_2 Techdep_{s,j} + Z_{j,t} + \epsilon_{j,t}$$

where $Techdep_{b,j}$ ($Techdep_{s,j}$) represents the average wage in 2016 for the industry in which the buyer (seller) operates. $Z_{j,t}$ denotes a vector of control variables and fixed effects, including the number of forward citations (adjusted using the HJT correction term), the patent's real market value at issuance estimated by Kogan et al. (2017), and transaction-year fixed effects.

The result is shown in Table F.1. The total surplus from patent transactions increases as the seller's technological dependence decreases, while it also rises as the buyer's technological dependence increases. This finding suggests that patents are more valuable when they transition

Figure 13: Heterogeneity in total surplus across sector-pairs

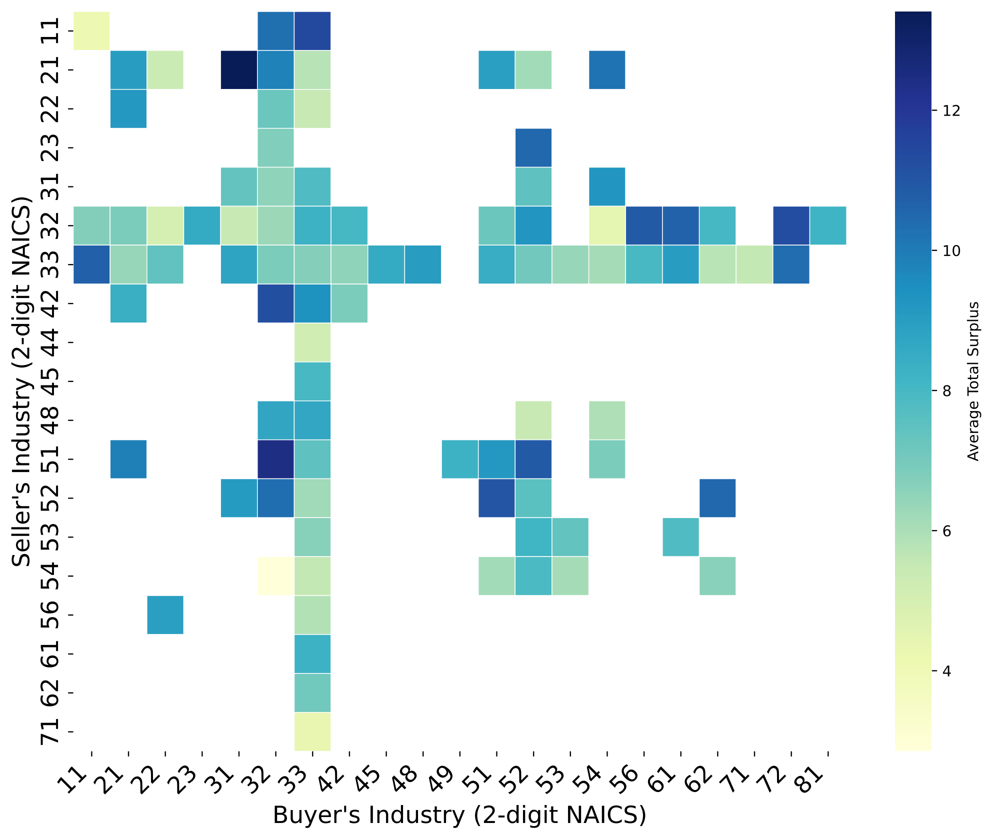

Notes This figure displays the average log total surplus across industry pairs. Industries are categorized using two-digit NAICS codes, with the buyer's industry on the x-axis and the seller's industry on the y-axis.

to sectors with greater reliance on technology.

G The effect of firm size on the expected gains from trade

Having established how counterparty size affects the ex-ante incentive to conduct R&D in the main text, in this appendix, we now examine its role in the ex-post division of gains, conditional on the seller having developed a blueprint and chosen to trade it.

The effect of buyer size on the seller's expected gains from trade. We begin with the seller's perspective. From equation (13) and the calibrated functional forms, the expected gain of firm i from selling a blueprint to firm $-i$, denoted by $\xi_{s,i}$, is given by

$$\xi_{s,i} = E\left[\tau_i TS\left(\theta_i\right) \mid |\phi_i - \theta_i| > \frac{f(\kappa_i)}{f(\kappa_i) + f(\kappa_{-i})}\right] = \frac{\kappa^{0.25}}{2(1+\kappa)}\kappa_i^{0.38}, \tag{G.1}$$

Differentiating equation (G.1) with respect to κ, we find that firm i's expected gains from selling

Table F.1: Total Surplus And Technology-dependence

The dependent variable is the logarithm of total surplus. The independent variables are the technology-dependence of the buyer and the seller, measured by the per capita income in 2016 in the industry that the buyer (seller) belongs to. Standard errors are reported in parentheses. *** $p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

	(1)
Dependent Variable:	*Total Surplus*
Tech-dependence of buyer	0.000***
	(0.000)
Tech-dependence of seller	-0.000***
	(0.000)
Controls	
Citation	Yes
InitialValue	Yes
Year FE	Yes
Observations	23,354
R-squared	0.261

a blueprint is maximized when

$$\left[\kappa^{0.25}/(1+\kappa)\right]' = 0, \tag{G.2}$$

which holds when $\kappa \equiv \kappa_{-i}/\kappa_i = 1/3$.

The result is summarized in the following proposition.

Proposition 3. *Given its own size, firm i's expected gain from selling a blueprint is maximized when the relative size of the potential buyer satisfies $\kappa_{-i}/\kappa_i = 1/3$.*

Proposition 3 suggests that firms prefer smaller buyers when selling blueprints. Although smaller buyers often incur higher transaction costs and may struggle to scale the use of acquired technologies–both of which can diminish total surplus–these disadvantages are offset by the seller's ability to claim a larger share of the surplus. In other words, the bargaining advantage gained from engaging with a smaller counterparty outweighs the efficiency losses associated with lower total surplus.

Notably, the optimal buyer size here is lower than that in Proposition 2, which characterizes the optimal buyer size for incentivizing seller-side innovation *prior* to the decision to sell the blueprint. Once the seller has decided to sell, their preference for smaller buyers becomes much stronger.

The effect of seller size on the buyer's expected gains from trade. We now shift to the buyer's perspective, exploring which seller size is most likely to maximize buyer's gains from a patent transaction.

From equation (13) and the calibrated functional forms, the expected gain of firm i from purchasing a blueprint from firm $-i$, denoted by $\xi_{b,i}$, is given by

$$\xi_{b,i} = E\left[\tau_i TS(\theta_{-i}) \mid |\phi_{-i} - \theta_{-i}| > \frac{f(\kappa_{-i})}{f(\kappa_i) + f(\kappa_{-i})}\right] = \frac{\kappa^{0.13}}{2(1+\kappa)}\kappa_i^{0.38}, \tag{G.3}$$

which depends on the firm's own size, κ_i, and the relative size of the seller, κ.[22] Notably, equation (G.3) characterizes the expected gains from trade, $\tau_i TS(\theta_{-i})$, conditional on the presence of trading opportunities., which occurs when the seller's business operates at a sufficient distance from the patent it develops, determined by $|\phi_{-i} - \theta_{-i}| > \frac{f(\kappa_{-i})}{f(\kappa_i)+f(\kappa_{-i})}$.

Differentiating equation (G.3) with respect to κ, we find that firm i's expected gains from purchasing a blueprint is maximized when

$$\left[\kappa^{0.13}/(1+\kappa)\right]' = 0,$$

which holds when $\kappa \equiv \kappa_{-i}/\kappa_i = 13/87$.

The result is summarized in the following proposition.

Proposition 4. *Given its own size, firm i's expected gain from purchasing a blueprint is maximized when the relative size of the potential seller satisfies* $\kappa_{-i}/\kappa_i = 13/87$.

Proposition 4 suggests that firms favor smaller sellers. Although small sellers entail higher transaction costs, partially borne by the buyer due to Nash bargaining, the advantage of securing a larger bargaining share outweighs these costs from the buyer's perspective.

Propositions 3 and 4 indicate that, conditional on seeking a trade, firms—whether as buyers or sellers—tend to prefer counterparties significantly smaller than themselves.

[22] Specifically, the LHS equation (G.3) is equal to $\frac{\kappa_i}{\kappa_i+\kappa_{-i}} \frac{\eta(\kappa_i,\kappa_{-i})f(\kappa_i)}{2}$, which can be derived as the RHS of equation (G.3).

H Additional Results

Figure 14 presents a scatter plot of log market values versus filtered stock returns within the event windows, illustrating that larger firms generally exhibit lower stock returns.

Figure 14: Market cap and stock returns

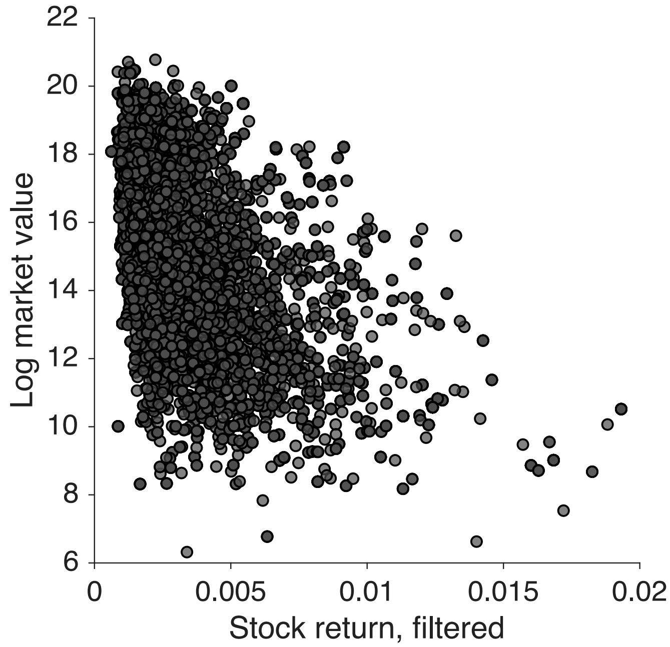

Table H.1: Bargaining Share is Positively Correlated with Relative Size: Small-to-Big Transactions

The dependent variable is the buyers' bargaining share. The independent variable is the buyers' relative size compared to the sellers. The sample is restricted to transactions in which the buyer's sales are at least three times those of the seller. Standard errors, reported in parentheses, are clustered at the firm level. *** p<0.01, **p<0.05, *p<0.1.

	(1)
Dependent Variable:	*Bargaining Share*
Relative Size	0.901***
	(0.017)
Observations	5,547
R-squared	0.951