Underreaction in Publicly Traded Cross-Holdings *

Marius Savatier†

September 12, 2026

Abstract

To what extent do the pricing restrictions derived from no-arbitrage under perfect redundancy continue to hold when redundancy becomes partial? To answer this question, I examine publicly traded firms that hold a large stake in another publicly traded firm while also owning other assets that are not separately traded. In a frictionless market, where all assets are tradable, no-arbitrage implies that a one-dollar change in the value of the stake should translate into a one-dollar change in the parent's value. After isolating price movements specific to the subsidiary, I find that, on average, a one-dollar change in the value of the stake translates into only 38 cents in the parent's value. When the parent holds no other assets, however, the estimated transmission rate is close to one. The greater the uncertainty surrounding the parent's other assets, the stronger the underreaction. These findings indicate that no-arbitrage pricing restrictions weaken as redundancy becomes more partial, consistent with a limits-to-arbitrage mechanism in which non-fundamental demand shocks affecting the subsidiary price create a tendency toward underreaction, while risk associated with the parent's non-hedgeable other assets limits arbitrageurs' ability to trade against this force.

JEL Codes: G12, G14, G32

Keywords: Limits to Arbitrage, Law of One Price, Value Additivity

*I thank Fabrice Riva, Jérôme Dugast, Itay Goldstein, Juan Felipe Imbet, Florian Peters, Evgenia Passari, Serge Darolles, Ricardo Delao (discussant), Jean-Gabriel Cousin (discussant), Maximilian Schröder (discussant), Carole Gresse, Rémi Genet, Anouck Faverjon, and Alex Osberghaus for helpful comments, as well as participants at the 40th International AFFI Conference (2024), the 7th Dauphine Finance PhD Workshop (2024), the 18th International Behavioural Finance Conference (2025), the 9th HEC Paris Finance PhD Workshop in Memory of Denis Gromb (2025), and the 42nd International AFFI Conference (2026).

†PhD Candidate at University Paris Dauphine-PSL; email: marius.savatier@dauphine.psl.eu

"The usefulness of 'no arbitrage' hinges on additional assumptions such as the specific market structure, investors' beliefs and preferences. These auxiliary assumptions limit its empirical applicability and thus its testability."

— Brav, Heaton & Rosenberg (2004)

1. Introduction

Many of the achievements in financial economics come from extending no-arbitrage pricing restrictions beyond settings where assets are perfectly redundant.[1] No-arbitrage implies that the price of a redundant asset equals the sum of the prices of the assets in the portfolio that replicates it, leading to a common linear pricing rule across redundant assets. However, redundancy is rare in equity markets, as most firms have large idiosyncratic components (Roll, 1988; Shleifer and Vishny, 1997). Standard asset pricing extends the idea of a common linear pricing rule beyond perfect redundancy by relating prices to expected cash flows through a common stochastic discount factor that is compatible with the preferences of rational investors. This implies, for example, that the same exposure to a given risk factor should be priced the same across stocks, despite the presence of idiosyncratic components that make redundancy only partial.

The limits-to-arbitrage literature questions whether such common pricing restrictions continue to hold when redundancy is only partial. In particular, risk-averse arbitrageurs who cannot perfectly hedge and diversify their positions may be unable to fully absorb non-fundamental demand shocks, causing relative valuations to be affected by shocks unrelated to news about cash flows or discount rates and thus to depart from such a rational pricing rule (Shleifer and Vishny, 1997; Gromb and Vayanos, 2010). To what extent do the pricing restrictions derived from no-arbitrage under perfect redundancy continue to hold when redundancy becomes partial?

[1] As noted by Brav et al. (2004), many important results in financial economics can be viewed as incarnations of the no-arbitrage principle. At the same time, they argue that the concept, in its strictest form, "provides little to work with except in the most trivial of circumstances (the price of an identical loaf of bread in the same supermarket) or in the world's best developed and standardized derivatives markets."

A natural laboratory for studying this question is provided by publicly traded firms (parents) holding a stake in another publicly traded firm (subsidiary). In a frictionless benchmark where the parent's underlying assets are all tradable, no-arbitrage implies that the value of the parent equals the value of its stake plus the value of its other assets, net of liabilities. Real-world publicly traded parent–subsidiary pairs differ from this benchmark because the parent's other assets are not directly traded. Thus, when arbitrageurs believe that the parent's market value does not fully reflect the value of its stake, they can take a long position in the parent and a short position in the subsidiary, but cannot hedge their exposure to the parent's other assets. The limits-to-arbitrage literature predicts that this unhedgeable exposure limits arbitrageurs' ability to trade against price discrepancies, weakening the link between the parent's market value and the market value of its stake (Mitchell et al., 2002; Gromb and Vayanos, 2010).

A second consequence of the fact that the parent's other assets are not directly traded is that they do not have an observable market price. As a result, this no-arbitrage prediction cannot be tested directly in price levels.[2] The restriction nevertheless implies a testable prediction in returns: holding changes in the parent's other assets fixed, changes in the value of the stake should transmit one-for-one to the parent.

To test this prediction, I isolate movements in the subsidiary's stock price that are not explained by a large set of other traded-asset price movements and ask how much of these movements transmit to the parent. The identifying assumption is that observable price movements capture the common factors affecting both the subsidiary stake and the parent's other assets, such that the remaining subsidiary-specific variation is unrelated to contemporaneous changes in the parent's other assets. I limit the parent–subsidiary sample to cases where the subsidiary stake represents at least 10 percent of the parent's value and impose additional size and liquidity requirements. The resulting sample contains 187 U.S. parent–subsidiary pair-years between 1999 and 2024. The average daily transmission rate is 0.38, significantly below the frictionless benchmark of 1. Placebo tests replacing

[2] Brav et al. (2004) argue that extending no-arbitrage restrictions beyond settings of perfect redundancy requires auxiliary assumptions about investors' beliefs, preferences, and market structure. Yet these strong assumptions are often invoked precisely where the restriction becomes most difficult to test directly.

each subsidiary with the most correlated non-owned firm produce transmission rates close to zero, supporting the identification strategy.

I use the cross-section of transmission rates to test different mechanisms for the underreaction. The key cross-sectional determinant is the non-spanned variance: the parent's residual variance after removing the variance explained by the subsidiary stake. This measure captures the non-hedgeable risk associated with the parent's other assets. When this non-spanned variance is small because the parent holds few other assets and the stake therefore closely replicates the parent, the transmission rate is close to one. As this non-spanned variance increases, the transmission rate falls. This pattern is consistent with the limits-to-arbitrage prediction that imperfect hedgeability weakens the link between the parent's market value and the market value of its stake.

I investigate a range of alternative explanations for the measured underreaction. One possibility is that subsidiary shares held within the parent are worth less than otherwise identical shares held directly. Such a discount could arise from an additional layer of taxation within the parent or from poor management by the parent leading to inefficient use of the subsidiary's payoff stream. Studying seemingly mispriced internet carve-outs, Cornell and Liu (2001) argue that such mechanisms are less plausible when the parent controls the subsidiary. Control allows the parent to manage the subsidiary's payout policy to avoid double taxation, while poor parent management would plausibly affect all shares of the subsidiary, not only those embedded within the parent. There is some evidence that transmission rates are higher when the parent controls the subsidiary, although the result is not robust across specifications.

A second possibility is that subsidiary-specific price movements are associated with offsetting changes in the value of the parent's other assets. Given the estimated transmission rate of 0.38, a \$1 subsidiary-specific increase (decrease) would need to be associated with an approximately \$0.60 decrease (increase) in the value of the parent's other assets. Any omitted-variable explanation must therefore generate such opposite-signed comovement systematically within individual parent–subsidiary pairs, without being present in other observable price movements. A possibility is that such omitted variation operates through discount rates rather than cash-flow news. As redundancy is only partial, it is always possible to define a time-varying discount rate that matches the observed underreaction. However, doing so would require discount rates to vary systematically with subsidiary-specific price movements at the daily frequency.

Other observable characteristics, such as size and various liquidity measures, have little explanatory power for transmission rates. Taken together, a limits-to-arbitrage mechanism in which non-fundamental demand shocks to the subsidiary create pressure toward underreaction, while unhedgeable risk from the parent's other assets limits arbitrageurs' ability to correct it, provides the most parsimonious explanation for the evidence. I formalize this mechanism in a model in the spirit of Gromb and Vayanos (2010), where imperfect hedgeability limits arbitrageurs' ability to absorb non-fundamental demand shocks. The model predicts that the transmission rate depends on arbitrageurs' risk-bearing capacity and on the level of unhedgeable risk associated with the parent's other assets. When the parent has no other assets, the transmission rate equals one. As unhedgeable risk from the parent's other assets increases, arbitrage becomes more difficult and the transmission rate falls.

The combination of arbitrageurs' limited risk-bearing capacity and non-fundamental demand shocks can lead to long-horizon return predictability, a mechanism studied theoretically by Campbell and Kyle (1993) and documented empirically by Fama and French (1988); Poterba and Summers (1988); Campbell and Shiller (1989). However, parent–subsidiary relationships are often observed for only one or two years, making them less suitable for studying this type of long-horizon predictability.

At shorter horizons, I test whether subsidiary returns predict parent returns over the following days, as would be expected under delayed information diffusion. The evidence suggests only a small delayed reaction. Chang et al. (2022) document such predictability in a broad international sample of ownership-linked firms, without restricting the sample to cases where the subsidiary stake represents a large fraction of the parent's market value. They construct portfolio strategies that buy or sell the parent depending on the subsidiary's performance over the previous month.[3] They interpret this evidence as reflecting the slow incorporation of fundamental information related to the parent's internal capital market.

Publicly traded parent–subsidiary pairs with large stakes attracted particular attention during the internet bubble of the early 2000s, when some stakes exceeded the market value of their parent. Such situations are commonly referred to as negative stub values (Cornell and Liu, 2001; Mitchell et al., 2002; Lamont and Thaler, 2003). However, negative stub values do not in themselves demonstrate mispricing. They simply imply that the market assigns a negative value to the parent's other assets net of liabilities, which is not impossible. To address this issue, Mitchell et al. (2002) analyze trading strategies based on negative stub values and employ two different definitions of mispricing. One definition treats a negative stub value itself as evidence of mispricing. The other uses the book value of the parent's other net assets as a proxy for their value. By estimating a transmission rate between the stake and the parent's market value, this paper approaches the problem from another angle, allowing the analysis to extend beyond negative stub values to all parent–subsidiary pairs with a large stake.

In their analysis of trading strategies based on negative stub values, Mitchell et al. (2002) show that exposure to fundamental risk arising from the parent's other assets makes arbitrage risky and may limit arbitrageurs' positions. This echoes the argument in Shleifer and Vishny (1997) that arbitrageurs are specialized and hold concentrated portfolios, making idiosyncratic risk not only unhedgeable but also undiversified.[4] The evidence uncovered in this paper is consistent with this intuition.

In the sample, negative stub values do not behave differently from other parent–subsidiary pairs in general, except for the internet carve-outs of the early 2000s, where the variance of the subsidiary is extreme while transmission rates are very low.[5] Lamont and Thaler (2003) study a few such cases where the shares of the subsidiary are subsequently

[3]See also Li et al. (2016) for evidence among US parent–subsidiary firms.

[4]See Brav et al. (2010), who find weak support for idiosyncratic volatility as an explanation for positive-return anomalies such as value stocks, but stronger support for negative-return anomalies such as growth stocks.

[5]I do not include the famous 3Com–Palm case because I require at least one year of observation.

spun off to the parent's shareholders, increasing confidence that mispricing is present.[6] They explain the apparent mispricing through irrational frenzy combined with short-sale constraints arising from the limited amount of subsidiary shares available for trading. Cherkes et al. (2013) provide additional evidence consistent with the importance of these short-sale constraints. For transmission rate, however, the amount of subsidiary shares available for trading does not explain cross-sectional variation in transmission rates.[7]

This paper also relates to the twin-stock anomalies: firms entitled to identical cash flows but traded in different national markets exhibited large deviations from parity (Rosenthal and Young, 1990). Froot and Dabora (1999) attribute these deviations to international market segmentation. By contrast, in my sample, the parent and subsidiary in each pair trade in the same U.S. market and on the same exchange. Consistent with this difference, when the stake nearly perfectly replicates the parent, the transmission rate is close to one.[8]

More recent developments in the limits-to-arbitrage literature have emphasized additional frictions, including investor withdrawals (Hombert and Thesmar, 2014), margin constraints (Gromb and Vayanos, 2002), liquidity spirals (Brunnermeier and Pedersen, 2009), and dynamic convergence risk (Kondor, 2009). Gromb and Vayanos (2018) further show how arbitrageurs can propagate shocks across markets, while Davila et al. (2024) study the social value of arbitrage. This literature tends to focus on settings where assets are perfectly redundant. This paper instead uses a specific empirical setting to make visible a more widespread limit to arbitrage arising from limited risk-bearing capacity under partial redundancy (Gromb and Vayanos, 2010).

The rest of the paper is structured as follows. Section 2 introduces the transmission rate and derives its benchmark value under the frictionless benchmark. Section 3 describes the data. Section 4 presents the estimation framework. Section 5 reports the empirical

[6] Lamont and Thaler (2003) also note that Bergstresser and Karlan (2000) independently documented underreaction of parents to movements in their subsidiaries' values, although the manuscript does not appear to be publicly available.

[7] See also Cornell and Liu (2001), who study several such internet carve-outs and conclude that mechanisms commonly used to explain conglomerate or closed-end fund discounts do not apply well to publicly traded parent–subsidiary structures.

[8] See also de Jong et al. (2009), who show that arbitrage strategies involving twin stocks do not generate large Sharpe ratios.

results. Section 6 systematically investigates the potential sources of the underreaction. Section 7 develops a limits-to-arbitrage model that generates the observed underreaction. Section 8 concludes.

2. Transmission Rate

In a frictionless market where all components of a portfolio are themselves tradable, no-arbitrage implies value additivity: the value of the portfolio equals the sum of the values of its constituent assets. For a parent holding a stake in a subsidiary, this implies

$$p_t^P = p_t^S + p_t^O, \tag{1}$$

where p_t^P denotes the parent's market value (number of shares outstanding times share price), p_t^S the value of its stake in the subsidiary (number of shares held times share price), and p_t^O the value of its other assets net of liabilities. Defining the simple return as $r_t \equiv \Delta p_t / p_{t-1}$, equation (1) implies that the parent's return equals the value-weighted return of its components:

$$r_t^P = r_t^{wS} + r_t^{wO}, \tag{2}$$

where $r_t^{wS} \equiv w_{t-1} r_t^S$ and $r_t^{wO} \equiv (1 - w_{t-1}) r_t^O$ denote the weighted returns of the subsidiary stake and the parent's other assets, respectively, and $w_{t-1} \equiv p_{t-1}^S / p_{t-1}^P$ is the beginning-of-period equity weight of the subsidiary stake.

Neither equation (1) nor equation (2) can be tested directly because the parent's other assets are not publicly traded as a stand-alone security, rendering p_t^O and r_t^O unobservable. Instead, I focus on one implication of equation (2): the sensitivity of the parent's return to the weighted return of its subsidiary stake, holding the return on the parent's other assets fixed:

$$\mathcal{T} \equiv \frac{\partial r_t^P}{\partial r_t^{wS}}. \tag{3}$$

Under the frictionless benchmark of equation (2), this sensitivity would equal one, so $\mathcal{T} = 1$. Intuitively, when the market value of the subsidiary stake increases by \$x, the market value of the parent should also increase by $x. Section 4 shows how to estimate the transmission rate.

For the publicly traded parent–subsidiary pairs studied here, the key departure from the frictionless benchmark is that the parent's other assets are not separately traded, making redundancy only partial. No-arbitrage alone therefore no longer requires a transmission rate of one. As emphasized by Brav et al. (2004), extending arbitrage-based pricing restrictions beyond perfect redundancy requires additional economic structure.

In the present setting, three forces could provide such structure and sustain a transmission rate of one. First, the parent itself can act: if its market value exceeds the combined value of its subsidiary stake and its other assets, it can issue equity and acquire additional shares in the subsidiary; if it falls short, it can sell shares in the subsidiary and repurchase its own stock. However, such actions require the parent to know the value of its other assets. Second, arbitrageurs can hedge systematic risk and diversify idiosyncratic risk, as in the Arbitrage Pricing Theory (APT) of Ross (1976). Third, if arbitrage is understood more broadly as trading by risk-bearing speculators who exploit advantageous opportunities, the conventional asset-pricing benchmark is one in which prices behave as if set by a rational representative agent who values assets as discounted conditional expectations of future cash flows. Because the expectation operator is linear, equations (1) and (2) hold, implying $\mathcal{T} = 1$ (see Appendix A). Intuitively, when new information changes the valuation of the stake, the representative agent adjusts valuations consistently both inside and outside the parent. Overall, a transmission rate of one remains a natural benchmark even when redundancy is only partial. Whether these forces are sufficient to enforce this benchmark in practice is ultimately an empirical question.

3. Data

The sample consists of 73 publicly traded parent–subsidiary pairs between 1999 and 2024, yielding 187 pair–year observations. All firms are headquartered in the United States and listed on either the NYSE or NASDAQ. Table 1 reports summary statistics. The mean subsidiary and parent market values are $12 billion and $17 billion, respectively, and the mean stake value is \$6.5 billion. The mean portfolio weight, defined as the stake value divided by the parent's market value, is 50%. The mean economic interest, measured as the parent's share of the subsidiary's cash flows, is 46%. In 56% of cases, the parent holds a majority of the subsidiary's voting rights. Appendix B reports the time-series evolution and sectoral composition of the sample. About one third of the pairs belong to the Energy sector (Fama–French 48: Oil and Gas, Utilities, and Transport), with the remainder distributed across 19 industries.

I construct the number of shares held by the parent in the subsidiary from Schedule 13D/G filings. Acquirers of more than 5% of a public company must normally file Schedule 13D, disclosing both their stake and intentions regarding control. Passive investors may instead use the shorter Schedule 13G, which some parent firms, despite holding large stakes, were permitted to file when planning to sell or spin off the subsidiary. A drawback of Schedule 13G is that, during most of the sample period, it was generally updated only annually, whereas Schedule 13D required prompt amendments following material changes.[9] From 2004 onward, I also use Form 3 and Form 4 filings to cross-check the ownership series. These forms require insiders, including beneficial owners holding more than 10%, to report ownership and subsequent changes promptly. Before 2004, I instead rely on Forms 10-K and 10-Q to cross-check ownership.

The number of shares outstanding for the parent is obtained from Compustat. Because this figure may be inaccurate when multiple classes of common shares are outstanding, I cross-check shares outstanding using 10-Q and 10-K filings. Stock prices, returns, bid and ask quotes, and trading volumes are obtained from CRSP. For control variables, I supplement these data with returns from Compustat not available in CRSP, primarily for exchange-traded funds (ETFs). The CPI is obtained from the U.S. Bureau of Labor Statistics, and the Fama–French factors are obtained from Kenneth French's website.

I partition each parent–subsidiary pair into consecutive, non-overlapping 252-trading-day periods beginning with the pair's first observation. I retain pair-years for which

[9]The SEC revised these reporting deadlines in 2023. Schedule 13D amendments are now due within two business days of a material change, while Schedule 13G amendment deadlines depend on the type of filer.

Table 1: Summary Statistics. The sample consists of 187 parent–subsidiary pair-year observations, each corresponding to a 252-trading-day period. For each variable, pair-year observations are constructed by averaging daily observations over the corresponding 252-trading-day period. Ownership is the proportion of economic interest held by the parent in the subsidiary; Weight is the ratio of Stake MV to Parent MV; #Trades is the number of trades (some values are missing in CRSP); Relative Spread is the bid–ask spread divided by the mid-quote, $(\text{Ask}_{i,t} - \text{Bid}_{i,t})/\left(\frac{1}{2}(\text{Ask}_{i,t} + \text{Bid}_{i,t})\right)$; Amihud is $|r_{i,t}|/\text{DollarVol}_{i,t}$; Dollar Volume is price times trading volume; Turnover is daily trading volume divided by shares outstanding.

Panel A. Subsidiary Characteristics

Variable	N	Mean	Median	SD	Min	Max
Market Value	187	11,604	5,493	18,470	1,001	122,669
# Trades	70	11,792	9,052	9,452	243	35,041
Relative Spread	187	0.002	0.001	0.003	0.000	0.023
Dollar Volume	187	81	32	127	5	745
Amihud	187	2.062	0.778	6.463	0.016	75.296
Turnover	187	0.009	0.006	0.010	0.001	0.077

Panel B. Parent Characteristics

Variable	N	Mean	Median	SD	Min	Max
Market Value	187	16,960	8,235	21,914	307	161,141
# Trades	61	18,311	10,825	30,503	339	216,902
Relative Spread	187	0.002	0.001	0.004	0.000	0.023
Dollar Volume	187	162	72	214	3	1,404
Amihud	187	0.942	0.245	2.928	0.011	27.271
Turnover	187	0.012	0.009	0.009	0.001	0.085

Panel C. Parent–Subsidiary Relationship

Variable	N	Mean	Median	SD	Min	Max
Stake MV	187	6,583	2,748	9,133	136	54,908
Ownership	187	0.467	0.496	0.229	0.050	0.935
Weight	187	0.504	0.364	0.421	0.102	1.955

Note: Dollar values are in millions; Amihud is reported in billions of USD^{-1}.

the mean subsidiary market value exceeds $1 billion, the mean subsidiary daily trading volume exceeds $5 million, the mean portfolio weight exceeds 10%, and both firms have a mean share price above $5. Applying these filters yields the sample described above, consisting of 73 parent–subsidiary pairs and 187 pair–year observations. Dollar thresholds are expressed in January 2000 dollars using the Consumer Price Index.

Table 1 shows that both parents and subsidiaries are actively traded. Mean daily dollar trading volumes are \$162 million and \$81 million, respectively, while relative bid–ask spreads average about 0.2% for both groups of firms. Mean turnover is 1.2% and 0.9% of shares outstanding per day, respectively, and the mean Amihud illiquidity measure is low. The mean number of daily trades is 18,311 for parents and 11,792 for subsidiaries, although this variable is available only for a subset of observations because it is missing for some firms in CRSP. Taken together, these statistics indicate that the sample comprises liquid securities, as intended by the sample selection criteria.

4. Estimation Framework

4.1. Specifications

Identifying the transmission rate requires separating subsidiary-specific movements from common factors that also affect the parent's other assets. To this end, I decompose the subsidiary's return as

$$r_t^S = F_t \gamma^S + e_t^S, \tag{4}$$

where F_t is a row vector of factors whose first element is 1, γ^S is the corresponding column vector of parameters, and e_t^S is the subsidiary-specific return, such that $E[F_t e_t^S \mid \mathcal{H}_{t-1}] = 0$, where $\mathcal{H}_{t-1}$ denotes the information set available at time $t-1$. For identification, the relevant factor structure is the one that also renders the subsidiary-specific residual orthogonal to the return on the parent's other assets: $E[e_t^S r_t^O \mid \mathcal{H}_{t-1}] = 0$.

If this factor structure were observable, the transmission rate could be estimated through the following regression:

$$r_t^P = \mathcal{T} r_t^{wS} + F_t^w \beta + u_t, \tag{5}$$

where $r_t^{wS} = w_{t-1} r_t^S$ denotes the weighted return of the stake, F_t^w is a row vector whose first element is 1 and whose remaining elements are given by $w_{t-1} F_t$, β is the corresponding column vector of regression coefficients, and u_t is a residual orthogonal to both r_t^{wS} and F_t^w. Under this specification, the transmission rate is identified from the weighted subsidiary-specific residual $w_{t-1}e_t^S$, which is orthogonal to the weighted return on the parent's other assets, $r_t^{wO} = (1 - w_{t-1})r_t^O$.[10] Therefore, if value additivity and return additivity, as defined in equations (1) and (2), respectively, hold, the transmission rate satisfies $\mathcal{T} = 1$.[11]

In practice, however, the exact common factor structure F_t is unobservable. The guiding principle of the empirical strategy is that these common factors can be approximated using the observable returns of other assets. Different empirical procedures can be used to organize these returns into a set of controls. For any such procedure, the transmission rate is identified from the component of the subsidiary's return that is not explained by the selected controls.

The key empirical question is whether a given set of controls removes enough common variation between the subsidiary and the parent's other assets. A placebo procedure, in which the subsidiary is replaced by its most correlated non-owned firm, serves as a falsification test of the procedure used to select these controls. If the procedure successfully approximates the relevant common factor structure, the same procedure should also be able to remove the common variation between the parent and the non-owned placebo firm. Because the parent does not hold the placebo firm, firm-specific movements in the placebo firm should not transmit to the parent, and the estimated transmission rate should therefore be zero. A nonzero placebo transmission rate indicates that the procedure has failed to remove common variation between the parent and the placebo firm, casting doubt on the procedure's ability to isolate firm-specific movements in general and, therefore, subsidiary-specific movements in the main specification.

A natural alternative would be to instrument subsidiary returns using shocks plausibly unrelated to the parent's other assets, such as earnings surprises or ownership-flow shocks. I do not follow this approach because parent–subsidiary structures are rare and observed over relatively short windows, making strong instruments difficult to construct across pairs. Moreover, instrumental-variable estimation would identify transmission from a

[10]Since $w_{t-1} \in \mathcal{H}_{t-1}$, $E\big[w_{t-1}e_t^S(1 - w_{t-1})r_t^O\big] = E\big[w_{t-1}(1 - w_{t-1})E\big[e_t^S r_t^O \mid \mathcal{H}_{t-1}\big]\big] = 0$.

[11]In Appendix A, I provide a more detailed derivation that also defines the factor structure of the parent's other assets, allowing the regression residual u_t to be defined explicitly.

narrower subset of subsidiary price movements, yielding a local estimate whose underlying mechanism may be specific to those movements. For example, earnings-announcement days are atypical trading days and attract unusually high investor attention, so the resulting transmission mechanism may differ from that operating on ordinary trading days. Finally, imperfect removal of common factors naturally biases the estimated transmission rate upward, as the placebo exercise shows: specifications that insufficiently remove this common variation produce positive estimated transmission rates even for firms not held by the parent. Because the main finding of the paper is underreaction rather than overreaction, this bias works against the main result.

Let X_t denote the set of controls constructed by a given empirical procedure. I estimate the following empirical specification at the daily frequency:

$$r_t^P = \mathcal{T} r_t^{wS} + X_t^w \beta + u_t, \tag{6}$$

where $X_t^w = w_{t-1} X_t$. The contents of X_t depend on the procedure, but always include $1/w_{t-1}$ and a constant. I consider three control procedures of increasing sophistication.

(1) No Controls. The parent's simple return is regressed on a constant, the previous period's portfolio weight, and the weighted simple return of the subsidiary. Formally,

$$X_t = [\,1/w_{t-1},\ 1\,].$$

(2) FF5 + Sector. This specification adds the five Fama–French factors and the sector portfolio(s) corresponding to the parent and the subsidiary, based on the Fama–French 48-industry classification. Formally,

$$X_t = \left[\,1/w_{t-1},\ 1,\ MKT_t,\ SMB_t,\ HML_t,\ RMW_t,\ CMA_t,\ SECTOR_t^P,\ SECTOR_t^S\,\right].$$

When the parent and the subsidiary belong to the same sector, the corresponding sector portfolio is included only once.

(3) Post-Double-Lasso. Following Belloni et al. (2014), this method uses lasso regression to select controls from a potentially high-dimensional set, whose size may exceed the number of observations. Applied to the estimation of the transmission rate, the procedure starts from a set of potential controls $\mathcal{P}$, assumed to contain a sparse subset sufficient to approximate the relevant common factor structure. The Post-Double-Lasso procedure then works as follows:

1. I use a lasso regression to select from $\mathcal{P}$ the controls that best predict the subsidiary's return. Let $\mathcal{I}_S$ denote the set of indices selected for the subsidiary.

2. I use a second lasso regression to select from $\mathcal{P}$ the controls that best predict the parent's return. Let $\mathcal{I}_P$ denote the set of indices selected for the parent.

3. I take the union of the two selected sets of controls and use it in the final OLS regression. The final set is $\mathcal{I} = \mathcal{I}_S \cup \mathcal{I}_P$, so that

$$X_t = \left[\, 1/w_{t-1},\ 1,\ \{x_{j,t}\}_{j\in\mathcal{I}} \,\right].$$

For both lasso regressions, I use the theoretically grounded penalty proposed by Belloni et al. (2014).

I construct the set of potential controls $\mathcal{P}$ from the combined universe of CRSP and Compustat firms in North America as follows:

1. I select the 100 assets most correlated with the subsidiary, subject to the requirement that no selected asset has a correlation above 0.70 with an asset already selected.

2. After removing these assets from the candidate universe, I select the 100 assets most correlated with the parent, subject to the same correlation restriction, yielding 200 distinct asset returns.

3. I augment these 200 returns with the first 100 principal components computed from them.

This procedure yields 300 potential controls for each parent–subsidiary pair. Reducing the dimensionality of the potential control set from the full CRSP–Compustat universe to 300 substantially reduces computation time. The correlation restriction further reduces redundancy among the candidate returns and helps ensure that the preselected assets span a broad range of common variation. Finally, augmenting the 200 asset returns with their principal components allows the potential control set to capture common factors that may not be well represented by any individual asset return.[12]

In sum, the No Controls specification provides a benchmark for the bias when not controlling, the FF5 + sector model reflects the standard approach in asset pricing, and the Post-Double-Lasso specification offers a data-driven method designed to approximate the factor structure with greater realism.

4.2. Testing the Transmission Benchmark

For each parent–subsidiary pair-year (hereafter, pair), I estimate one transmission rate. Under the return-additivity condition implied by value additivity, each estimate can be written as

$$\hat{\mathcal{T}}_i = 1 + \varepsilon_i, \tag{7}$$

where i indexes the parent–subsidiary pair and ε_i captures finite-sample noise. Letting n denote the number of parent–subsidiary pairs, the cross-sectional average is

$$\bar{\mathcal{T}} = \frac{1}{n}\sum_{i=1}^{n}\hat{\mathcal{T}}_i, \tag{8}$$

which estimates $E[\hat{\mathcal{T}}_i]$. Appendix A.5 provides conditions under which $E[\varepsilon_i \varepsilon_{i'}] = 0$ for $i \neq i'$. The asymptotic normality of $\bar{\mathcal{T}}$ then follows from the central limit theorem as $n \to \infty$. If return additivity holds, the following hypothesis should not be rejected:

$$H_0^S : E[\hat{\mathcal{T}}_i] = 1, \qquad H_1^S : E[\hat{\mathcal{T}}_i] \neq 1,$$

[12] Lasso regression often selects the first component but generally prefers to select firm returns directly.

where the superscript S refers to the parent–subsidiary sample.

4.3. Placebo Tests

I implement the placebo test described above by replacing each subsidiary with a publicly traded firm not held by the parent. To make the test demanding, I select the non-owned asset whose return is most highly correlated with the subsidiary's return.[13]

Formally, the placebo regression is

$$r_t^P = \mathcal{T} r_t^{wPL} + X_t^w \beta + u_t, \tag{9}$$

where $r_t^{wPL} = w_{t-1} r_t^{PL}$ and r_t^{PL} denotes the simple return of the placebo asset. For the Post-Double-Lasso specification, the full selection procedure is repeated on the parent–placebo pairs, so that the set of selected controls may differ from those obtained for the parent–subsidiary pairs. Under a correctly specified model, the estimated transmission rate in the placebo sample should satisfy

$$H_0^{PL} : E[\hat{\mathcal{T}}_i] = 0, \qquad H_1^{PL} : E[\hat{\mathcal{T}}_i] \neq 0,$$

where the superscript PL refers to the parent–placebo sample. A rejection of this hypothesis falsifies the corresponding control procedure.

4.4. Determinants of the Transmission Rate

Parent–subsidiary pairs differ in how closely the parent resembles a pure claim on the subsidiary. In some cases, the parent holds little besides the subsidiary stake, while in others it owns substantial additional assets. When a larger share of the parent's return variance is attributable to its other assets, arbitrageurs seeking to exploit underreaction in the parent face greater unhedgeable risk (Mitchell et al., 2002; Gromb and Vayanos, 2010). This weakens the corrective force of arbitrage, allowing underreaction to persist.

[13] One might worry that, in finite samples, replacing the subsidiary with its most correlated asset could overfit by chance. In large samples this should be negligible, and if anything, it makes the placebo test slightly too demanding.

Underreaction should therefore be stronger when unhedgeable risk is larger.

I proxy for unhedgeable risk by estimating what I refer to as the Unspanned Variance Ratio:

$$\widehat{\mathrm{UVR}}_i = \frac{Var(r_t^{P\perp S})_i}{Var(r_t^{wS})_i}, \tag{10}$$

where

$$Var(r_t^{P\perp S})_i = (1 - R^2_{\text{No Control},i})Var(r_t^P)_i.$$

The term $R^2_{\text{No Control},i}$ denotes the share of the parent's return variance explained by the weighted return on the subsidiary stake when no controls are included. Thus, $\widehat{\mathrm{UVR}}_i$ measures the amount of unspanned parent variance relative to the variance of the subsidiary stake. As shown in the model developed in Section 7, the transmission rate is determined by the amount of unhedgeable risk relative to hedgeable risk. $\widehat{\mathrm{UVR}}_i$ provides an empirical counterpart to this theoretical driver. The model further implies a highly nonlinear relation between the transmission rate and this ratio, which becomes substantially more linear when the ratio is expressed in logs. I therefore use $\log(\widehat{\mathrm{UVR}}_i)$ in the cross-sectional regressions.

To test if unhedgeable risk can explain the underreaction, I estimate cross-sectional regressions of the form

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\mathrm{UVR}}_i) + Z_i'\gamma + \varepsilon_i, \tag{11}$$

where $\hat{\mathcal{T}}_i$ denotes the estimated transmission rate for pair i, $\widehat{\mathrm{UVR}}_i$ is the estimated Unspanned Variance Ratio, and Z_i is a vector of control variables capturing alternative mechanisms affecting transmission rates across pairs.

The amount of unspanned variance could alternatively be measured by $(1-R^2_{\text{No Control},i})$. However, this measure scales unspanned variance by the parent's total return variance, which itself depends on the transmission of subsidiary returns to the parent and could therefore induce a mechanical relation with the transmission rate. For this reason, the baseline specification uses the Unspanned Variance Ratio defined above. As an additional robustness check, I construct an alternative UVR in which $R^2_{\text{No Control},i}$ is replaced by the

R^2 obtained after including controls. Both alternative measures yield similar results (see Appendix D.2).

Because both $\hat{\mathcal{T}}_i$ and $\widehat{\mathrm{UVR}}_i$ are estimated, estimation error could induce a mechanical relation between the two variables. To avoid this concern, I use cross-fitting. For each parent–subsidiary pair, I split the 252 trading days into two folds of 126 trading days each, with alternating trading days assigned to each fold. I then estimate one transmission rate and one Unspanned Variance Ratio in each fold, yielding two estimates of each measure. I cross-fit these estimates by matching the transmission rate estimated in each fold with the Unspanned Variance Ratio estimated in the opposite fold. Thus, each transmission rate is explained by an Unspanned Variance Ratio estimated from a different set of return observations. This procedure doubles the number of cross-sectional observations from 187 to 374, and I cluster standard errors at the parent–subsidiary pair level.

The vector Z_i contains several control variables capturing alternative explanations for the underreaction. The first such variable is $\textsc{Float}_i$, which measures the share of the subsidiary's stock available for trading after deducting the parent's holdings. Limited float can make it difficult for arbitrageurs to locate shares to short, as documented by Lamont and Thaler (2003) for negative stub values in Internet carve-outs with subsequent spin-offs.

I also include a dummy MC_i for majority control, equal to one when the parent holds more than 50% of the subsidiary's voting rights. If the parent wastes the subsidiary's cash flows, the shares embedded within the parent should trade at a discount. If the parent controls the subsidiary, this discount should be minimal because the mismanagement would affect all shares of the subsidiary (Cornell and Liu, 2001). Control also reduces potential double taxation of the subsidiary's cash flows. When the parent controls the subsidiary, it can influence payout policy, for instance by encouraging the subsidiary to repurchase its own shares rather than distribute dividends when these are tax-inefficient. Repurchases increase the parent's ownership stake and may eventually allow the parent to execute a tax-free spin-off under Section 355 of the IRS Code. Other mechanisms also exist to minimize taxes.[14]

Liquidity frictions may also influence transmission rates through trading costs or price-measurement errors. I therefore include two standard liquidity measures: Turnover_i, defined as average daily trading volume relative to shares outstanding, and Amihud_i, the price-impact measure proposed by Amihud (2002). Because parent and subsidiary liquidity are highly correlated in the sample, I average each measure across the two firms to obtain a single pair-level characteristic, thereby reducing multicollinearity while preserving the common liquidity component. I further include $\log(\text{Size}_i)$, the logarithm of the combined market capitalization of the parent and the subsidiary, as firm size can proxy for investor attention or the informativeness of stock prices (Dávila and Parlatore, 2025).

Finally, I include a dummy variable Post2002_i, equal to one for observations after the internet bubble. The pre-2002 sample is concentrated in the internet carve-outs that attracted considerable attention because of suspected mispricing (Lamont and Thaler, 2003; Cornell and Liu, 2001). In these cases, the subsidiary stake often exhibits very high return variance, sometimes exceeding that of the parent itself. Because the UVR_i measure uses the variance of the stake in the denominator, this mechanically affects its value. In Appendix C, I also examine whether transmission rates differ between energy-sector pairs, which account for roughly one-third of the sample, and the rest of the sample. I find no evidence that they do.

5. Empirical Results

5.1. The Average Transmission Rates

For each of the 187 parent–subsidiary pairs, and for the corresponding parent–placebo pairs, I estimate transmission rates under the three specifications introduced above, yielding six distributions in total.

[14]Under U.S. tax law, corporations receiving dividends from other corporations benefit from the dividends-received deduction, which depends on the ownership stake: 50% if ownership is below 20%, 65% if between 20% and 80%, and 100% if above 80% (prior to 2018: 70%, 80%, and 100%, but with a 35% corporate tax rate instead of 21%).

The first column of Table 2 reports the mean transmission rate when no controls are used. In Panel A, the mean for the parent–subsidiary pairs is 1.92 (s.e. 0.126). Although the significance levels reported in the table are relative to zero, this estimate is also significantly greater than the frictionless benchmark of one. Panel B reports the corresponding placebo mean transmission rate, which is similarly large at 1.83 (s.e. 0.139). Because the parents do not hold the placebo firms, the null hypothesis for the placebo sample is $H_0^{PL} : E[\hat{\mathcal{T}}_i] = 0$. This hypothesis is strongly rejected, falsifying the No Controls specification.

The positive placebo transmission rate shows that insufficient controls bias estimated transmission rates upward. Such a result is not surprising, as most firms' returns tend to be positively correlated, while parents and subsidiaries generally operate in the same sector, making their returns even more positively correlated. The large mean transmission rates in both samples are driven by pairs with small portfolio weights. When the portfolio weight is small, the subsidiary and placebo returns are scaled by a small w_{t-1} in r_t^{wS} and r_t^{wPL}. Any remaining common variation between these returns and the parent's other assets can therefore generate very large estimated transmission rates.

The second column of Table 2 reports the mean transmission rate when controlling for the Fama–French factors plus sector factors. In Panel A, the mean transmission rate is 0.848 (s.e. 0.069), significantly greater than zero but also significantly below the benchmark of one ($t = -2.20$). In Panel B, the mean transmission rate for the parent–placebo pairs is 0.607 (s.e. 0.079). The placebo null $H_0^{PL} : E[\hat{\mathcal{T}}_i] = 0$ is again strongly rejected, falsifying the FF5+Sector specification. Importantly, the positive placebo transmission rate shows that even after controlling for sector factors, the remaining bias is upward.

The third column of Table 2 reports the mean transmission rate when the Post-Double-Lasso procedure is used. In Panel A, the mean transmission rate for the parent–subsidiary pairs is 0.381 (s.e. 0.036). It is significantly greater than zero, consistent with the market recognizing the parent's ownership stake in the subsidiary. However, it is also significantly below the benchmark of one implied by return additivity. The subsidiary null $H_0^S : E[\hat{\mathcal{T}}_i] = 1$ is rejected by a large margin ($t = -17.1$). In Panel B, the mean placebo transmission rate is -0.022 (s.e. 0.044). The placebo null $H_0^{PL} : E[\hat{\mathcal{T}}_i] = 0$ cannot be rejected ($p = 0.611$). Thus, unlike the two preceding specifications, Post-Double-Lasso is not falsified by the placebo test. In Figure 1, I plot the distribution of transmission rates estimated using Post-Double-Lasso for parent–subsidiary pairs and parent–placebo pairs.

Table 2: The Mean Transmission Rates. This table reports mean transmission rates across six distributions (three models × two samples: parent–subsidiary and parent–placebo). *No Controls* uses no controls. *FF5+Sector* uses the Fama–French five factors plus 48 sectors. *Post-Double–Lasso* selects controls via Lasso on both parent and subsidiary returns. Avg. R^2 is the mean R^2 across estimations; R^2 (sub) is the share of subsidiary return variance explained by the controls.

	No Controls	FF5+Sector	Post-Double-Lasso
Panel A: Parent–Subsidiary Pairs ($n = 187$)			
mean $\mathcal{T}$	1.922^{***}	0.848^{***}	0.381^{***}
SE of mean $\mathcal{T}$	(0.126)	(0.069)	(0.036)
t-stat	15.247	12.231	10.560
p-value	0.000	0.000	0.000
median $\mathcal{T}$	1.227	0.578	0.387
Avg R^2	0.385	0.581	0.684
Avg R^2 (sub)	–	0.336	0.459
Panel B: Parent–Placebo Pairs ($n = 187$)			
mean $\mathcal{T}$	1.826^{***}	0.607^{***}	-0.022
SE of mean $\mathcal{T}$	(0.139)	(0.079)	(0.044)
t-stat	13.114	7.638	-0.509
p-value	0.000	0.000	0.611
median $\mathcal{T}$	1.327	0.290	0.023
Avg R^2	0.293	0.491	0.628
Avg R^2 (sub)	–	0.501	0.589

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

As with an instrumental-variable approach, one concern is that the Post-Double-Lasso specification could identify the transmission rate from only a small fraction of the subsidiary's return variance. In particular, the controls could absorb most of the variation in subsidiary returns, leaving identification to rely on a narrow residual component. This is not the case: on average, the controls explain only 45.9% of the subsidiary's return variance, leaving more than half as subsidiary-specific variation. This is consistent with the well-known finding that a surprisingly large share of firm-level return variance is firm specific, first documented by Roll (1988).[15]

From this point on, I focus on transmission rates estimated using Post-Double-Lasso and investigate what drives the underreaction. Before turning to possible explanations, I examine whether this underreaction reverses in subsequent days.

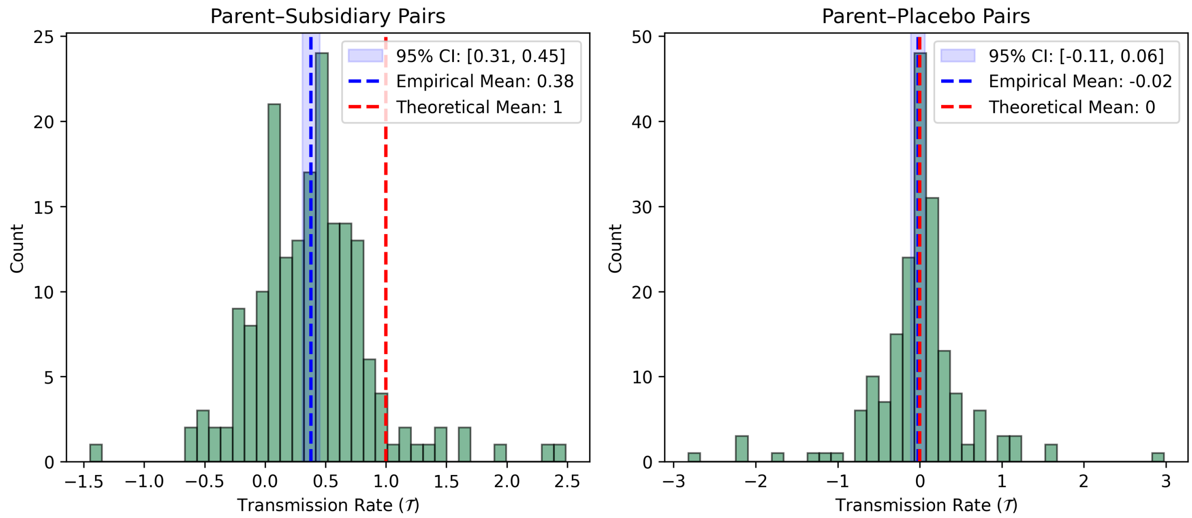

Figure 1: Distribution of estimated transmission rates using Post-Double-Lasso. On the left-hand side, the distribution for parent–subsidiary pairs should be centered at 1 (red vertical line), but is instead centered near 0.38 (blue vertical line). On the right-hand side, the transmission rates are estimated for parent–placebo pairs, where the subsidiary is replaced by the non-owned firm whose return is most highly correlated with the subsidiary's return. If the controls capture common factors, the distribution should be centered at 0, so the red and blue lines overlap, which is approximately the case.

5.2. Short-Term Reversal

The violation of value additivity suggested by the underreaction could be short-lived. If so, the delayed adjustment could potentially generate large profitable trading opportunities.

For each parent–subsidiary pair, I regress the parent's daily return on the contemporaneous weighted subsidiary return and its four daily lags, together with the controls selected by Double-Lasso and one lag of the parent's return. Panel A of Table 3 shows that

[15]For the same period, the 45.9% R^2 is slightly above the mean R^2 across all US stocks in Peters (2024), which is not surprising since R^2 tends to be higher for larger firms. The 33.6% for FF5+Sector can be compared to Roll (1988), who report about 20%. Aggregate R^2 declined until the early 2000s and has since started to increase, as documented by Campbell et al. (2001); Chiah et al. (2020); Campbell et al. (2023).

the contemporaneous transmission rate ($\bar{\mathcal{T}} = 0.386$) remains virtually unchanged, still significantly below one and above zero. The $t-1$ lag is 4.2% (s.e. 1.9%) and is significant at the 5% level, suggesting a slight delayed reaction of the parent to subsidiary news. The $t-2$ lag is 2.6% (s.e. 1.5%) and is not significant at the 10% level, while the two remaining lags have negative and insignificant coefficients. Summing all transmission rates yields a weekly cumulative transmission rate of 42.4% (s.e. 5%).[16] The cumulative transmission rate remains highly significantly below one. The subsequent adjustment is small and disappears after two days. This result does not rule out predictability in parent–subsidiary structures. It does, however, rule out the possibility that the large contemporaneous underreaction can be exploited through simple short-term trading strategies to generate correspondingly large profits. At monthly horizons, Chang et al. (2022); Li et al. (2016) obtain excess returns from strategies that buy parents whose subsidiaries had the highest returns over the previous months and sell parents whose subsidiaries had the lowest, consistent with slow information diffusion. On the other hand, if most of the underreaction is driven by noise, this noise might fade over the long term, leading to predictability in the subsidiary's share price over horizons of several years, consistent with the long-run return predictability documented in Fama and French (1988); Poterba and Summers (1988); Campbell and Shiller (1989).

Panel B of Table 3 reports the placebo estimates. Adding lags and the parent's lagged return barely changes the contemporaneous transmission rate, which remains very close to zero at -1.7% (s.e. 4.2%). Since the placebo's specific return should not predict the parent's return, the lag coefficients provide a useful benchmark. The first three lags are insignificant, and only the fourth lag is significant at the 5% level, consistent with a false positive arising from multiple tests.

[16]The standard error of the sum is computed as

$$\mathrm{SE}\left(\sum_{k=0}^{4} \hat{\mathcal{T}}_{i,t-k}\right) = \sqrt{\sum_{k=0}^{4} \mathrm{Var}(\hat{\mathcal{T}}_{i,t-k}) + 2\sum_{k<k'} \mathrm{Cov}(\hat{\mathcal{T}}_{i,t-k}, \hat{\mathcal{T}}_{i,t-k'})}.$$

If subsidiary returns are serially uncorrelated, the cumulative transmission rate corresponds to the transmission rate obtained from a regression at the lower frequency.

Table 3: Transmission Rates at Contemporaneous and Lagged Horizons. Within each pair, transmission rates are estimated using

$$r_t^P = \mathcal{T}_0 r_t^{wS} + \mathcal{T}_1 r_{t-1}^{wS} + \mathcal{T}_2 r_{t-2}^{wS} + \mathcal{T}_3 r_{t-3}^{wS} + \mathcal{T}_4 r_{t-4}^{wS} + \gamma r_{t-1}^P + X_t^w \beta + u_t,$$

where r_t^{wS} is the subsidiary's weighted return. The table reports the cross-sectional averages of $\hat{\mathcal{T}}_0, \hat{\mathcal{T}}_1, \hat{\mathcal{T}}_2, \hat{\mathcal{T}}_3, \hat{\mathcal{T}}_4$ for parent–subsidiary pairs and parent–placebo pairs.

	t	t-1	t-2	t-3	t-4
Panel A: Parent–Subsidiary Pairs ($n = 187$)					
mean $\mathcal{T}$	0.386^{***}	0.042^{**}	0.026^{*}	-0.020	-0.010
SE of mean $\mathcal{T}$	(0.038)	(0.019)	(0.015)	(0.017)	(0.014)
t-stat	10.248	2.234	1.798	-1.195	-0.687
p-value	0.000	0.026	0.072	0.232	0.492
median $\mathcal{T}$	0.359	0.045	0.014	-0.006	0.002
Panel B: Parent–Placebo Pairs ($n = 187$)					
mean $\mathcal{T}$	-0.017	0.006	0.000	-0.005	-0.033^{**}
SE of mean $\mathcal{T}$	(0.042)	(0.021)	(0.015)	(0.017)	(0.016)
t-stat	-0.394	0.261	-0.004	-0.293	-2.091
p-value	0.694	0.794	0.996	0.770	0.037
median $\mathcal{T}$	0.023	0.018	0.005	-0.003	-0.002

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. Standard errors are HC0.

5.3. Explaining the Cross-sectional Variation in Transmission Rates

The baseline WLS specification includes only $\log(\widehat{\mathrm{UVR}}_i)$ as a regressor. The coefficient is -0.110 with a standard error of 0.016, yielding a t-statistic of -6.875 and a p-value of 1.3×10^{-10}. Economically, this estimate implies that doubling the Unspanned Variance Ratio is associated with a decrease of about 7.5 percentage points in the transmission rate.[17]

The baseline OLS regression yields a smaller coefficient of -0.064 (s.e. 0.016, $t = -4$, $p = 6.3 \times 10^{-5}$), but the relation remains highly significant. Figure 2 plots $\log(\widehat{\mathrm{UVR}}_i)$ on the x-axis and $\hat{\mathcal{T}}_i$ on the y-axis. The relation becomes flatter as $\log(\widehat{\mathrm{UVR}}_i)$ increases, while the dispersion of the estimated transmission rates becomes larger. In the WLS regression, these noisier observations receive less weight, which helps explain why the estimated coefficient is further from zero.

Including the full set of controls leaves the coefficient on $\log(\widehat{\mathrm{UVR}}_i)$ largely unchanged

[17] A doubling of $\widehat{\mathrm{UVR}}_i$ implies $\Delta\hat{\mathcal{T}}_i = -0.11 \times \log(2) \approx -0.076$.

in both WLS and OLS specifications. In the WLS regression, the R^2 increases from 0.292 to 0.627, while in the OLS regression it rises from 0.058 to 0.105. Thus, $\log(\widehat{\mathrm{UVR}}_i)$ alone generates an R^2 nearly half as large as that of the full WLS specification, highlighting its importance in explaining variation in transmission rates. As a sanity check, Appendix D.1 reports the same regressions for placebo transmission rates, which show no meaningful relation. Appendix D.2 repeats the analysis using an idiosyncratic version of the UVR constructed from residual returns after partialling out the selected controls; the results are similar.

Among the control variables, $\log(\textsc{Size}_i)$ is significantly negatively related to transmission rates in the WLS specification, suggesting that larger parent–subsidiary pairs tend to exhibit lower transmission. However, this relation depends on the inclusion of $\textsc{Amihud}_i$ and the $\textsc{Post2002}_i$ dummy. When these variables are excluded, the size coefficient becomes statistically insignificant. The relation with $\textsc{Amihud}_i$ indicates that among firms of similar size, higher price impact is associated with lower transmission rates. The relation with $\textsc{Post2002}_i$ reflects the role of Internet carve-outs, for which larger subsidiaries are associated with lower transmission rates. Some of these subsidiaries reached extremely high valuations that did not transmit to the parent, generating negative stub values. In several cases, $\mathrm{Var}(r_t^{wS}) > \mathrm{Var}(r_t^{P})$, meaning that the variance of the subsidiary's portfolio-weighted return exceeds that of the parent. These observations appear in Figure 2 as outliers with low $\log(\widehat{\mathrm{UVR}}_i)$ but transmission rates close to zero. Their low $\widehat{\mathrm{UVR}}_i$ reflects the unusually large variance of the subsidiary stake in the denominator of the ratio. In contrast, $\textsc{Float}_i$ and $\textsc{Turnover}_i$ show little explanatory power in either specification.

The majority-control dummy MC_i is not significant in the WLS regression but becomes significant at the 5% level in the OLS specification. When the parent controls the subsidiary, the transmission rate increases by about 18 percentage points. This pattern is consistent with the possibility that agency frictions or tax considerations associated with not controlling the subsidiary create a discount on the subsidiary stake embedded in the parent, as suggested by Cornell and Liu (2001) and discussed in Section 4.4.

Overall, Float, Size, Turnover, and Amihud illiquidity provide relatively limited explanatory power, while the effect of majority control is unstable across specifications. The Unspanned Variance Ratio is the only strong explanatory variable that remains robust throughout.

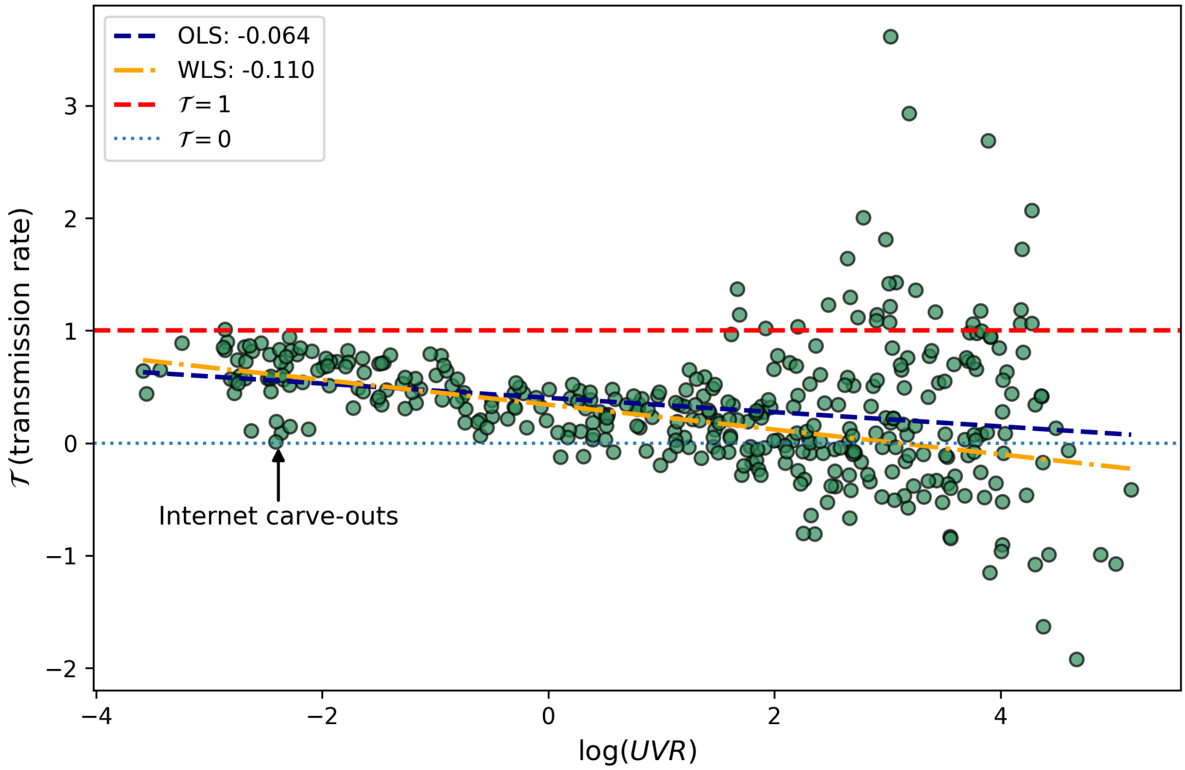

Figure 2: Scatter plots across parent–subsidiary pairs of $\mathcal{T}$ against $\log(UVR)$, where $UVR = \mathrm{Var}(r_t^{P\perp S})/\mathrm{Var}(r_t^{wS})$ denotes the Unspanned Variance Ratio. The number of observations is doubled because of the cross-fitting procedure. The blue dashed line reports the equal-weighted (OLS) regression, while the orange dashed line reports the weighted least squares (WLS) regression, weighted by the inverse squared standard errors of the estimated transmission rates. The red dotted line indicates the frictionless benchmark, $\mathcal{T} = 1$, and the light blue dotted line indicates the no-ownership benchmark, $\mathcal{T} = 0$.

Table 4: Explaining Transmission Rates Using Weighted Least Squares (WLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\mathrm{UVR}_i}) + \beta_2 \textsc{Float}_i + \beta_3 \mathrm{MC}_i + \beta_4 \textsc{Turnover}_i + \beta_5 \textsc{Amihud}_i + \beta_6 \log(\textsc{Size}_i) + \beta_7 \textsc{Post2002}_i + \varepsilon_i.$$

i denotes a parent–subsidiary pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. $\log(\widehat{\mathrm{UVR}_i})$ measures the ratio of the parent's unspanned variance to the subsidiary's variance, capturing uncertainty about the parent's other assets. $\textsc{Float}_i$ is the share of the subsidiary's stock available for trading after deducting the parent's holdings. MC_i is a dummy equal to one when the parent controls more than 50% of the voting rights. $\textsc{Turnover}_i$ is the average daily trading volume relative to shares outstanding, and $\textsc{Amihud}_i$ is the average price impact of trading, computed as the absolute return divided by the dollar volume. $\log(\textsc{Size}_i)$ is the logarithm of the combined market value of the parent and subsidiary, and $\textsc{Post2002}_i$ equals one for observations after the internet bubble. Cross-fitting ensures that $\log(\widehat{\mathrm{UVR}_i})$ and $\hat{\mathcal{T}}_i$ are estimated on separate folds. Standard errors are clustered by pair (HC3), and WLS weights are the inverse squared standard errors of $\hat{\mathcal{T}}_i$.

	Parent–Subsidiary (WLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–P02, –Am
Constant	0.342***	0.422***	0.474***	0.281**	0.244
	(0.023)	(0.137)	(0.158)	(0.134)	(0.202)
log(UVR)	−0.110***	−0.112***	−0.100***	−0.112***	−0.094***
	(0.016)	(0.017)	(0.026)	(0.018)	(0.031)
Float	–	0.077	0.254	0.027	0.228
	–	(0.223)	(0.379)	(0.229)	(0.431)
MC	–	−0.078	0.062	−0.052	0.157
	–	(0.121)	(0.144)	(0.125)	(0.142)
Turnover	–	−6.306*	−4.073	−5.493	−1.857
	–	(3.575)	(3.615)	(3.400)	(3.837)
Amihud	–	−0.179*	−0.279**	–	–
	–	(0.100)	(0.135)	–	–
log(Size)	–	−0.108***	−0.092***	−0.084***	−0.043
	–	(0.020)	(0.026)	(0.021)	(0.035)
Post-2002	–	0.352***	–	0.422***	–
	–	(0.134)	–	(0.131)	–
R^2	0.292	0.627	0.525	0.591	0.424
df	372	366	367	367	368

Notes: Weighted least squares (WLS) estimates with weights $1/\mathrm{SE}(\hat{\mathcal{T}}_i)^2$. Standard errors are clustered by pair (HC3). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 5: Explaining Transmission Rates Using Ordinary Least Squares (OLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\mathrm{UVR}_i}) + \beta_2 \textsc{Float}_i + \beta_3 \mathrm{MC}_i + \beta_4 \textsc{Turnover}_i + \beta_5 \textsc{Amihud}_i + \beta_6 \log(\textsc{Size}_i) + \beta_7 \textsc{Post2002}_i + \varepsilon_i.$$

i denotes a parent–subsidiary pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. $\log(\widehat{\mathrm{UVR}_i})$ measures the ratio of the parent's unspanned variance to the subsidiary's variance, capturing uncertainty about the parent's other assets. $\textsc{Float}_i$ is the share of the subsidiary's stock available for trading after deducting the parent's holdings. MC_i is a dummy equal to one when the parent controls more than 50% of the voting rights. $\textsc{Turnover}_i$ is the average daily trading volume relative to shares outstanding, and $\textsc{Amihud}_i$ is the average price impact of trading, computed as the absolute return divided by the dollar volume. $\log(\textsc{Size}_i)$ is the logarithm of the combined market value of the parent and subsidiary, and $\textsc{Post2002}_i$ equals one for observations after the internet bubble. Cross-fitting ensures that $\log(\widehat{\mathrm{UVR}_i})$ and $\hat{\mathcal{T}}_i$ are estimated on separate folds. Standard errors are clustered by pair (HC3).

	Parent–Subsidiary (OLS)				
Variable	Base	Full	No P02	No Ami.	–P02,–Am
Constant	0.402***	0.170	0.199	0.098	0.120
	(0.024)	(0.158)	(0.158)	(0.145)	(0.146)
log(UVR)	−0.064***	−0.057***	−0.055***	−0.056***	−0.053***
	(0.016)	(0.017)	(0.017)	(0.017)	(0.016)
Float	–	0.220	0.240	0.228	0.266
	–	(0.170)	(0.166)	(0.168)	(0.164)
MC	–	0.173**	0.182**	0.176**	0.194***
	–	(0.078)	(0.075)	(0.078)	(0.073)
Turnover	–	6.003	6.212	6.464	7.008
	–	(5.509)	(5.435)	(5.485)	(5.379)
Amihud	–	−0.078	−0.096	–	–
	–	(0.075)	(0.075)		
log(Size)	–	−0.045	−0.044	−0.035	−0.028
	–	(0.029)	(0.029)	(0.027)	(0.026)
Post-2002	–	0.059	–	0.090	–
	–	(0.071)		(0.067)	
R^2	0.058	0.105	0.104	0.103	0.100
df	372	366	367	367	368

Notes: Standard errors clustered by pair (HC3). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

6. Explaining the Underreaction

Let $\theta_{P,1}$, $\theta_{S,1}$, and $\theta_{O,1}$ denote the terminal-date payoffs at $t = 1$ of the parent, the stake, and the parent's other assets, respectively, where $t = 0$ denotes the current date, and where, for simplicity, we assume no intermediate payoffs. The ownership structure implies that

$$\theta_{P,1} = \theta_{S,1} + \theta_{O,1}.$$

Under standard asset pricing theory, the current market value of the parent should satisfy

$$p_{P,0} = E[m_0 \theta_{P,1} \mid \mathcal{I}_0] = E[m_0 \theta_{S,1} \mid \mathcal{I}_0] + E[m_0 \theta_{O,1} \mid \mathcal{I}_0],$$

where $\mathcal{I}_0$ denotes the information currently available and m_0 is the stochastic discount factor reflecting investors' preferences. Let u_0 denote the subsidiary-specific movement in the stake price $p_{S,0}$, after removing all other measurable stock-market movements. The transmission rate is then

$$\mathcal{T} = \frac{Cov(p_{P,0}, u_0)}{Var(u_0)} = \frac{Cov(E[m_0 \cdot \theta_{S,1} \mid \mathcal{I}_0], u_0)}{Var(u_0)} + \frac{Cov(E[m_0 \cdot \theta_{O,1} \mid \mathcal{I}_0], u_0)}{Var(u_0)},$$

where we should have

$$\frac{Cov(E[m_0 \cdot \theta_{S,1} \mid \mathcal{I}_0], u_0)}{Var(u_0)} = 1 \qquad \text{and} \qquad \frac{Cov(E[m_0 \cdot \theta_{O,1} \mid \mathcal{I}_0], u_0)}{Var(u_0)} = 0,$$

such that $\mathcal{T} = 1$. That is, u_0 is the component of $E[m_0 \theta_{S,1}]$ that cannot be explained by other measurable price movements. Under the value-additivity restriction, it should therefore transmit one for one to the parent. If the identification strategy is valid, it should also be orthogonal to $E[m_0 \theta_{O,1} \mid \mathcal{I}_0]$. Empirically, however, we obtain $\frac{Cov(p_{P,0}, u_0)}{Var(u_0)} \approx 0.38$. I discuss three mechanisms that can explain the underreaction.

First, the stake embedded within the parent may itself be discounted. This corresponds to the case where

$$\mathcal{T} = \frac{Cov(E[m_0\theta_{S,1} \mid \mathcal{I}_0], u_0)}{Var(u_0)} < 1,$$

so that $Cov(E[m_0\theta_{S,1} \mid \mathcal{I}_0], u_0) < Var(u_0)$, and the transmission rate reflects this discount. This interpretation may appear related to the conglomerate discount literature, where large conglomerates were found to trade at lower valuations than comparable standalone firms operating in the same industries (Lang and Stulz (1994); Berger and Ofek (1995); Servaes (1996)).[18] Agency problems and inefficient internal capital allocation arising from the management of multiple unrelated activities have often been proposed as a potential explanation for the conglomerate discount. However, in the setting considered here, the comparison is between shares embedded within another firm and the same shares traded directly on the market. Two frictions can still specifically affect the shares embedded within the parent: the parent may waste the subsidiary's cash flows, and an additional layer of taxation within the parent may warrant a discount. Both frictions are less likely to explain the underreaction when the parent controls the subsidiary, as it also controls the subsidiary's cash flows, and any wasteful behavior is more plausibly applied to all of the subsidiary's cash flows, not only to the shares held by the parent (Cornell and Liu (2001)). Similarly, taxes can be minimized when the parent controls the subsidiary's distribution policy (Cornell and Liu (2001)). The coefficient on the majority-control dummy is about 18% and statistically significant in the equal-weighted regression, consistent with the presence of such a discount. The WLS regression yields more mixed results.

Second, good (bad) subsidiary-specific news may be correlated with bad (good) news about the parent's other assets, such that removing all measurable price movements is not sufficient to identify the transmission rate. This corresponds to the case

$$\frac{Cov(E[m_0 \cdot \theta_{O,1} \mid \mathcal{I}_0], u_0)}{Var(u_0)} < 0, \tag{12}$$

[18]See also Lamont and Polk (2001), who show that conglomerates have higher expected returns, consistent with potential mispricing, and Villalonga (2004), who argues that measurement errors in segment data can account for the discount.

where, through the information set $\mathcal{I}_0$, a \$1 subsidiary-specific increase (decrease) would need to destroy (create) \$0.62 of value in the parent's other assets to account for the estimated transmission rate of 0.38 ($1 - 0.38 = 0.62$). This mechanism would require a firm-specific factor causing subsidiary-specific shocks to be systematically associated with opposite-signed shocks to the parent's other assets. Such a pattern appears implausible because parents and subsidiaries typically operate in related businesses and would therefore be expected to respond similarly to common firm- or sector-level information. Consistent with this, improving the controls in the empirical analysis decreases the estimated transmission rate rather than increasing it. Moreover, placebo firms selected to be the most correlated with the subsidiary have firm-specific returns that are uncorrelated with the parent, implying that such a factor would need to operate specifically at the parent–subsidiary level while remaining absent from other closely related firms. Thus, this alternative explanation itself implies a nontrivial empirical regularity requiring explanation.

Third, price changes can always be mechanically decomposed into payoff news and discount-rate news (Campbell and Shiller (1989)), such that one can always construct, ex post, a stochastic discount factor m_0 in Equation (12) capable of generating the underreaction. For example, one could assume investor preferences such that increases in the value of the stake make the parent's other assets appear more risky. However, rationalizing the underreaction in this way would require the pricing of idiosyncratic risk to vary systematically at the daily frequency, even though such risk is usually assumed to be diversifiable and therefore largely unrelated to discount rates in standard asset-pricing models.

Section 7 instead develops a model in the tradition of the limits-to-arbitrage literature, in which non-fundamental demand shocks affect prices because of arbitrageurs' limited risk-bearing capacity.

7. A Limits-to-Arbitrage Model

In this section, I develop a limits-to-arbitrage model in which non-fundamental demand shocks affecting the subsidiary generate underreaction when arbitrageurs have limited risk-bearing capacity and the parent holds additional non-traded assets with uncertain payoffs. The model explains the negative relation observed in the data between uncertainty surrounding the parent's other assets and the transmission rate.

This approach builds on the intuition of Gromb and Vayanos (2010), who study how arbitrageurs absorb non-fundamental demand shocks affecting one asset by trading a correlated hedging asset with an exogenous price. Unless the two assets are perfectly correlated, the arbitrage position remains risky and non-fundamental demand shocks are not fully absorbed. The exogeneity of the hedging asset's price is crucial: it allows arbitrageurs' trades in the hedging asset to absorb part of the shock without moving its price. When both prices are endogenous and a single group of investors trades both assets, arbitrage instead spreads the shock across the two prices as a function of the correlation between their payoffs.

To generate underreaction, I introduce market segmentation: some investors trade only in the parent, others trade only in the subsidiary, while arbitrageurs trade both assets. Without market segmentation, there is instead a single group of identical rational investors who trade both assets. These investors spread the shock across prices in a way that mirrors the overlap in payoffs, yielding a transmission rate of one. Similar segmentation assumptions, in which arbitrageurs with access to a broader investment universe trade against local investors, appear in Gromb and Vayanos (2002, 2018).

Before turning to the model, it is worth noting that a similar underreaction pattern could also arise if arbitrageurs face informational uncertainty about the parent's other assets. Because these assets do not have separately observed market prices, arbitrageurs, like the econometrician, cannot directly infer their value. Consequently, the uncertainty surrounding the parent's other assets may reflect not only non-hedgeable payoff risk but also imperfect information about their value. A model built around this informational friction generates analogous comparative statics: greater uncertainty leads to greater underreaction. Such a model is presented in Appendix E.

7.1. Environment

There are two periods, $t = 0, 1$. The market consists of one risk-free asset and two risky assets, P and S. Asset P denotes the parent company. I assume that the parent holds half of the subsidiary's outstanding shares. Asset S denotes the remaining floating subsidiary shares.

The gross risk-free rate is equal to one. The two risky assets, P and S, are traded at $t = 0$. The terminal payoff of S at $t = 1$ is

$$\theta_{S,1} \sim N(\mu_S, \tau_{\theta,S}^{-1}),$$

where μ_S denotes the expected payoff of S, and $\tau_{\theta,S}$ its precision, that is, the inverse of the payoff variance. Because the parent holds half of the subsidiary's shares, the payoff of the parent's stake is also $\theta_{S,1}$. The terminal payoff of the parent at $t = 1$ is

$$\theta_{P,1} = \theta_{S,1} + \theta_{O,1}, \qquad \theta_{O,1} \sim N(\mu_O, \tau_{\theta,O}^{-1}),$$

where $\theta_{O,1}$ denotes the payoff generated by the parent's other assets, μ_O its expected payoff, and $\tau_{\theta,O}$ its precision.

I assume that $\theta_{S,1}$ and $\theta_{O,1}$ are independent, reflecting the fact that the model is formulated at the firm-specific level.[19] Under this assumption,

$$\theta_{P,1} \sim N\left(\mu_S + \mu_O, \tau_{\theta,S}^{-1} + \tau_{\theta,O}^{-1}\right).$$

The payoffs $\theta_{P,1}$, $\theta_{S,1}$, and $\theta_{O,1}$ should be interpreted broadly as the terminal economic values relevant for investors and arbitrageurs. They may reflect future cash flows, liquidation values, or future market valuations. For example, $\theta_{S,1}$ and $\theta_{O,1}$ may capture the

[19]The results extend to the case where correlation is allowed.

future values of the stake and the parent's other assets following a spin-off of the stake.

7.2. Parent- and Subsidiary- investors

Markets are segmented, with two non-overlapping continua of unit mass investors: the P-investors only invest in the parent and the risk-free asset, while the S-investors only invest in the remaining floating subsidiary shares and the risk-free asset.

Each group of investors $k \in \{P, S\}$ has CARA preferences with risk-aversion parameter γ_k, and terminal wealth $W_{k,1}$. Because they are atomistic, they take prices as given. They solve

$$\max_{q^i_{k,0}} E\left[-e^{-\gamma_k W_{k,1}}\right]$$

subject to

$$W_{k,1} = W_{k,0} + q^i_{k,0}(\theta_{k,1} - p_{k,0}),$$

where $W_{k,0}$ is initial wealth, $q^i_{k,0}$ is the number of shares held, and $p_{k,0}$ is the price of asset k at $t = 0$. From CARA-normality, the demand of investors k is

$$q^i_{k,0} = \frac{\tau_{\theta,k}}{\gamma_k}(\mu_k - p_{k,0}),$$

7.3. Arbitrageurs

There is a unit mass continuum of arbitrageurs who can invest in both assets P and S. They have CARA preferences with risk-aversion parameter ξ, and terminal wealth $W_{\text{arb},1}$. Because they are atomistic, they take prices as given. They solve

$$\max_{q^i_{\text{arb},0}} E\left[-e^{-\xi W_{\text{arb},1}}\right]$$

subject to

$$W_{\text{arb},1} = W_{\text{arb},0} + (q^i_{\text{arb},0})'(\theta_1 - p_0),$$

where $W_{\text{arb},0}$ denotes initial wealth, and

$$q^i_{\text{arb},0} = \begin{pmatrix} q^i_{\text{arb},P,0} \\ q^i_{\text{arb},S,0} \end{pmatrix}, \qquad \theta_1 = \begin{pmatrix} \theta_{P,1} \\ \theta_{S,1} \end{pmatrix}, \qquad p_0 = \begin{pmatrix} p_{P,0} \\ p_{S,0} \end{pmatrix},$$

respectively denote the vector of arbitrage positions, terminal payoffs, and prices.

From CARA-normality, arbitrageur demand is

$$q^i_{\text{arb},0} = \frac{\Sigma_\theta^{-1}}{\xi}(\mu - p_0),$$

where

$$\mu = \begin{pmatrix} \mu_P \\ \mu_S \end{pmatrix}, \qquad \Sigma_\theta = \begin{pmatrix} \tau^{-1}_{\theta,S} + \tau^{-1}_{\theta,O} & \tau^{-1}_{\theta,S} \\ \tau^{-1}_{\theta,S} & \tau^{-1}_{\theta,S} \end{pmatrix}$$

is the variance–covariance matrix of the parent and subsidiary payoffs. This gives

$$q^i_{\text{arb},0} = \frac{1}{\xi} \begin{pmatrix} \tau_{\theta,O}\left(\mu_O - (p_{P,0} - p_{S,0})\right) \\ -\tau_{\theta,O}\left(\mu_O - (p_{P,0} - p_{S,0})\right) + \tau_{\theta,S}(\mu_S - p_{S,0}) \end{pmatrix}. \tag{13}$$

Arbitrageur demand in Equation (13) can be decomposed into two trades. The first allows arbitrageurs to access the payoff generated by the parent's other assets by taking a position in the parent, $\frac{\tau_{\theta,O}}{\xi}\left(\mu_O - (p_{P,0} - p_{S,0})\right)$, while hedging the payoff uncertainty associated with the subsidiary stake through the opposite position in the subsidiary, $-\frac{\tau_{\theta,O}}{\xi}\left(\mu_O - (p_{P,0} - p_{S,0})\right)$. The lower the uncertainty surrounding the payoff generated by the parent's other assets, that is, the greater $\tau_{\theta,O}$, the more aggressively arbitrageurs trade this long-parent, short-subsidiary portfolio. Independently, arbitrageurs also take a position in the subsidiary of $\frac{\tau_{\theta,S}}{\xi}(\mu_S - p_{S,0})$.

7.4. Equilibrium

I assume that both assets are in zero net supply, so that aggregate positions sum to zero in equilibrium. Assuming positive supply would only add a constant discount term to equilibrium prices to compensate investors for bearing aggregate risk, and therefore does not affect the transmission rate. The market-clearing condition for asset P is

$$\int_0^1 q^i_{P,0}\, di + \int_0^1 q^i_{\text{arb},P,0}\, di = 0$$
$$\frac{\tau_{\theta,P}}{\gamma_P}(\mu_P - p_{P,0}) + \frac{\tau_{\theta,O}}{\xi}\left(\mu_O - (p_{P,0} - p_{S,0})\right) = 0 \tag{14}$$

The market-clearing condition for asset S is

$$\int_0^1 q^i_{S,0}\, di + \int_0^1 q^i_{\text{arb},S,0}\, di + n_{0,S} = 0$$
$$\frac{\tau_{\theta,S}}{\gamma_S}(\mu_S - p_{S,0}) - \frac{\tau_{\theta,O}}{\xi}\left(\mu_O - (p_{P,0} - p_{S,0})\right) + \frac{\tau_{\theta,S}}{\xi}(\mu_S - p_{S,0}) + n_{0,S} = 0, \tag{15}$$

where $n_{0,S}$ denotes non-fundamental demand shocks affecting the subsidiary shares.

Solving the system formed by Equations (14) and (15), the equilibrium prices are

$$p_{P,0} = \mu_P + \frac{\tau_{\theta,O}\gamma_P\gamma_S\xi}{\tau_{\theta,P}\tau_{\theta,S}\gamma_S\xi + \tau_{\theta,P}\tau_{\theta,S}\xi^2 + \tau_{\theta,P}\tau_{\theta,O}\gamma_S\xi + \tau_{\theta,S}\tau_{\theta,O}\gamma_P\gamma_S + \tau_{\theta,S}\tau_{\theta,O}\gamma_P\xi} n_{0,S},$$
$$p_{S,0} = \mu_S + \frac{\gamma_S\xi\left(\tau_{\theta,P}\xi + \tau_{\theta,O}\gamma_P\right)}{\tau_{\theta,P}\tau_{\theta,S}\gamma_S\xi + \tau_{\theta,P}\tau_{\theta,S}\xi^2 + \tau_{\theta,P}\tau_{\theta,O}\gamma_S\xi + \tau_{\theta,S}\tau_{\theta,O}\gamma_P\gamma_S + \tau_{\theta,S}\tau_{\theta,O}\gamma_P\xi} n_{0,S}.$$

Before computing transmission rates, two limiting cases help clarify the mechanism of the model. First, when arbitrageurs are risk-neutral, $\xi \to 0$, or investors in the subsidiary are risk-neutral, $\gamma_S \to 0$, the non-fundamental demand shock is fully absorbed and prices coincide with fundamentals:

$$p_{P,0} = \mu_P, \qquad p_{S,0} = \mu_S. \tag{16}$$

Second, when investors in the parent become infinitely risk-averse, $\gamma_P \to \infty$, equilibrium prices satisfy

$$p_{P,0} = \mu_P + \frac{\gamma_S\xi}{\tau_{\theta,S}(\gamma_S + \xi)} n_{0,S},$$
$$p_{S,0} = \mu_S + \frac{\gamma_S\xi}{\tau_{\theta,S}(\gamma_S + \xi)} n_{0,S}.$$

Therefore, both prices move one-for-one in response to the demand shock, and the transmission rate is one. When $\gamma_P \to \infty$ and $\xi < \infty$, arbitrageurs freely set the relative price of the parent such that the transmission rate remains equal to one. Generating underreaction therefore requires $\gamma_P < \infty$, such that arbitrageurs trade against parent investors solving a different optimization problem. This underlies the market-segmentation assumption of the model, as without segmentation arbitrageurs and investors would collapse into a single group and the transmission rate would be equal to one.

7.5. Transmission Rate

Using the definition of the transmission rate introduced in Section 2, we obtain

$$\begin{aligned} \mathcal{T} &= \frac{\partial p_{P,0}}{\partial p_{S,0}} = \frac{\partial p_{P,0}/\partial n_{S,0}}{\partial p_{S,0}/\partial n_{S,0}} \\ &= \frac{\tau_{\theta,O}\gamma_P}{\tau_{\theta,P}\xi + \tau_{\theta,O}\gamma_P}, \end{aligned} \tag{17}$$

where $\tau_{\theta,P} = (\tau_{\theta,S}^{-1} + \tau_{\theta,O}^{-1})^{-1}$.

The transmission rate $\mathcal{T}$ is a function of the payoff precision of the parent's other assets, $\tau_{\theta,O}$, the payoff precision of the stake embedded in $\tau_{\theta,P}$, and the risk-aversion parameters of parent investors, γ_P, and arbitrageurs, ξ. The risk aversion of subsidiary investors, γ_S, does not determine the transmission rate directly; it only ensures that the demand shock is reflected in prices. As shown in Equation (16), if $\gamma_S \to 0$, subsidiary investors become risk-neutral and fully absorb the demand shock.

The transmission rate converges to one in four corner cases. As shown in Section 7.4, when parent investors become infinitely risk-averse, $\gamma_P \to \infty$, arbitrageurs effectively determine relative prices, and the transmission rate converges to one. When $\xi \to 0$, arbitrageurs become risk-neutral and absorb the demand shock entirely, while simultaneously enforcing consistency between the parent and the stake, so the transmission rate converges to one. When all uncertainty associated with the parent's other assets vanishes, $\tau_{\theta,O} \to \infty$, hedging becomes perfect, the Law of One Price applies, and the transmission rate converges to one. Lastly, when the uncertainty associated with the stake becomes arbitrarily large, $\tau_{\theta,S} \to 0$, the uncertainty associated with the parent's other assets becomes negligible in relative terms, and the transmission rate converges to one.

The partial derivative of $\mathcal{T}$ with respect to $\tau_{\theta,O}^{-1}$ is

$$\frac{\partial \mathcal{T}}{\partial \tau_{\theta,O}^{-1}} = -\frac{\gamma_P \xi \tau_{\theta,S} \tau_{\theta,O}^2}{\left(\gamma_P(\tau_{\theta,O} + \tau_{\theta,S}) + \tau_{\theta,S}\xi\right)^2} \leq 0. \tag{18}$$

Holding all other parameters fixed, an increase in the uncertainty associated with the parent's other assets decreases the transmission rate. The UVR introduced in Section 4.4, by proxying the ratio $\tau_{\theta,S}/\tau_{\theta,O}$, allows us to test how the transmission rate evolves when uncertainty about the parent's other assets increases while uncertainty about the stake is held fixed. Holding the uncertainty of the stake fixed is important because $\mathcal{T}$ also increases with $\tau_{\theta,S}^{-1}$, such that it is really the ratio $\tau_{\theta,S}/\tau_{\theta,O}$ that yields an unambiguous prediction across parent–subsidiary pairs.

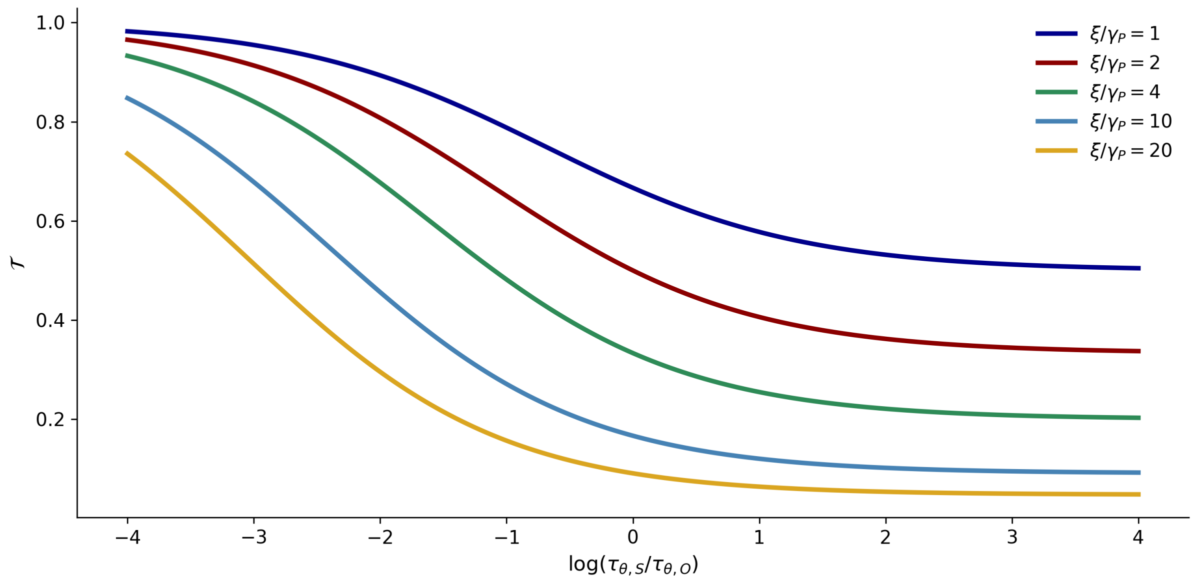

Figure 3: Transmission rate as a function of $\log(\tau_S/\tau_O)$.

Figure 3 is the theoretical counterpart of Figure 2. As $\log(\frac{\tau_{\theta,S}}{\tau_{\theta,O}})$ increases, the transmission rate decreases. In Figure 2 the relation between the transmission rate and the $\log(\widehat{\text{UVR}}_i)$ becomes flatter for high values of the UVR. Similarly, as $\log\left(\frac{\tau_{\theta,S}}{\tau_{\theta,O}}\right)$ increases, the relationship between the transmission rate and $\log\left(\frac{\tau_{\theta,S}}{\tau_{\theta,O}}\right)$ also becomes flatter. In both cases, the inflection point is around zero, that is, when the uncertainty associated with the parent's other assets becomes as large as the uncertainty associated with the stake. In level form rather than log form, both imply a decrease of about 5–10% in the transmission rate when the uncertainty associated with the parent's other assets relative to the stake doubles. In sum, when the parent has no other assets, the transmission rate is equal to one. As uncertainty associated with the parent's other assets progressively increases, the transmission rate initially falls sharply, but the marginal effect becomes smaller as uncertainty accumulates. Analytically, we can see from Equation (18) that the partial derivative of the transmission rate with respect to the payoff uncertainty associated with the parent's other assets decreases in absolute value as uncertainty increases.

The ratio of risk-aversion parameters $\frac{\xi}{\gamma_P}$ that best matches the empirical relationship in Figure 2 is approximately 10, implying that arbitrageurs behave as if they were substantially more risk-averse than investors in the parent. The fact that the empirical relationship is better matched by such a parametrization may reflect the presence of additional limits to arbitrage beyond limited risk-bearing capacity alone. In particular, while the model allows arbitrageurs to short both assets freely, short selling in practice requires capital to be posted as margin and involves additional stock-borrowing fees.

8. Conclusion

This paper studies publicly traded parent firms holding large stakes in publicly traded subsidiaries. In a frictionless market, changes in the value of the subsidiary stake should transmit one-for-one to the parent's market value. Empirically, however, after isolating subsidiary-specific price movements, I estimate an average transmission rate of only 38 cents on the dollar across U.S. parent–subsidiary pairs between 1999 and 2024.

The key cross-sectional determinant of transmission rates is the parent's variance not spanned by the subsidiary stake, which proxies uncertainty surrounding the parent's other assets. When the parent holds few or no other assets, the transmission rate is close to one. As uncertainty surrounding the parent's other assets increases, the transmission rate falls. This pattern is difficult to reconcile with explanations based solely on limited subsidiary float, liquidity frictions, or broad discount-rate variation. Instead, it is consistent with a limits-to-arbitrage mechanism in which non-fundamental demand shocks affecting the subsidiary are only partially transmitted to the parent because imperfect hedgeability limits arbitrageurs' ability to trade against them.

To formalize this intuition, the paper develops a limits-to-arbitrage model in the spirit of Gromb and Vayanos (2010) in which parent investors and subsidiary investors trade separately, while arbitrageurs can trade both assets. In the model, non-fundamental demand shocks affecting the subsidiary price are only partially transmitted to the parent when arbitrageurs have limited risk-bearing capacity and the parent holds additional non-hedgeable assets. The model also matches the negative relation observed in the data between transmission rates and uncertainty surrounding the parent's other assets.

Ross (1987) argues that what distinguishes finance from economics is that supply and demand should be largely irrelevant for asset prices, which should instead be determined by information and investor preferences. More recent developments in the literature have emphasized that equity prices can be substantially affected by non-fundamental demand shocks (Lou, 2012; Gabaix and Koijen, 2021; Haddad et al., 2025).

The findings of this paper suggest that non-fundamental demand shocks also matter for relative prices. Even when two securities are linked by a direct ownership relation, shocks to one security are only partially transmitted to the other. This result is consistent with limits-to-arbitrage theories in which non-hedgeable risk prevents arbitrageurs from fully correcting relative mispricing. A deeper question, however, is how much of this limited transmission reflects an informational friction. When the parent holds other assets without separately observed market prices, arbitrageurs may themselves be uncertain about fundamental values and therefore unable to observe the potential mispricing in the first place (Brav and Heaton, 2002, 2003; Brav et al., 2004).

References

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.

Angrist, J. D. and Pischke, J.-S. (2009). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *Review of Economic Studies*, 81(2):608–650.

Berger, P. G. and Ofek, E. (1995). Diversification's effect on firm value. *Journal of Financial Economics*, 37(1):39–65.

Bergstresser, D. and Karlan, D. (2000). The market valuation of cross-corporate equity holdings. Manuscript, Massachusetts Institute of Technology, Cambridge.

Brav, A. and Heaton, J. B. (2002). Competing theories of financial anomalies. *Review of Financial Studies*, 15(2):575–606.

Brav, A. and Heaton, J. B. (2003). Market indeterminacy. *Journal of Corporation Law*, 28(4):517–536.

Brav, A., Heaton, J. B., and Li, S. (2010). The limits of the limits of arbitrage. *Review of Finance*, 14(1):157–187.

Brav, A., Heaton, J. B., and Rosenberg, A. (2004). The rational-behavioral debate in financial economics. *Journal of Economic Methodology*, 11(4):393–409.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238.

Campbell, J. Y. and Kyle, A. S. (1993). Smart money, noise trading and stock price behaviour. *Review of Economic Studies*, 60(1):1–34.

Campbell, J. Y., Lettau, M., Malkiel, B. G., and Xu, Y. (2001). Have individual stocks become more volatile? an empirical exploration of idiosyncratic risk. *Journal of Finance*, 56(1):1–43.

Campbell, J. Y., Lettau, M., Malkiel, B. G., and Xu, Y. (2023). Idiosyncratic equity risk two decades later. *Critical Finance Review*, 12(1–4):203–223.

Campbell, J. Y. and Shiller, R. J. (1989). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, 1(3):195–228.

Chang, R., Gonzalez, A., Sarkissian, S., and Tu, J. (2022). Internal capital markets and predictability in complex ownership firms. *Journal of Corporate Finance*, 74:102219.

Cherkes, M., Jones, C. M., and Spatt, C. S. (2013). A solution to the palm–3com spin–off puzzles. Technical Report 12/52, Columbia Business School Research Paper Series. Working paper.

Chiah, M., Gharghori, P., and Zhong, A. (2020). Has idiosyncratic volatility increased? not in recent times. *Critical Finance Review*, 9(2):207–248.

Cochrane, J. H. (2005). *Asset Pricing*. Princeton University Press.

Cornell, B. and Liu, Q. (2001). The parent company puzzle: When is the whole worth less than one of the parts? *Journal of Corporate Finance*, 7(4):341–366.

Davila, E., Graves, D. D., and Parlatore, C. (2024). The value of arbitrage. *Journal of Political Economy*, 132(6):1947–1993.

Dávila, E. and Parlatore, C. (2025). Identifying price informativeness. *Review of Financial Studies*, 38(6):1947–1993.

de Jong, A., Rosenthal, L., and van Dijk, M. A. (2009). The risk and return of arbitrage in dual-listed companies. *Review of Finance*, 13(3):495–520.

Fama, E. F. and French, K. R. (1988). Permanent and temporary components of stock prices. *Journal of Political Economy*, 96(2):246–273.

Froot, K. A. and Dabora, E. M. (1999). How are stock prices affected by the location of trade? *Journal of Financial Economics*, 53(2):189–216.

Gabaix, X. and Koijen, R. S. J. (2021). In search of the origins of financial fluctuations: The inelastic markets hypothesis. NBER Working Paper 28967, National Bureau of Economic Research.

Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2-3):361–407.

Gromb, D. and Vayanos, D. (2010). Limits of arbitrage: the state of the theory. *Annual Review of Financial Economics*, 2(1):251–275.

Gromb, D. and Vayanos, D. (2018). The dynamics of financially constrained arbitrage. *Journal of Finance*, 73(4):1713–1750.

Haddad, V., Huebner, P., and Loualiche, E. (2025). How competitive is the stock market? theory, evidence from portfolios, and implications for the rise of passive investing. *American Economic Review*, 115(3):975–1018.

Hombert, J. and Thesmar, D. (2014). Overcoming limits of arbitrage: Theory and evidence. *Journal of Financial Economics*, 111(1):26–44.

Kondor, P. (2009). Risk in dynamic arbitrage. *American Economic Review*, 99(3):631–655.

Lamont, O. A. and Polk, C. (2001). The diversification discount: Cash flows versus returns. *Journal of Finance*, 56(5):1693–1721.

Lamont, O. A. and Thaler, R. H. (2003). Can the market add and subtract? mispricing in tech stock carve-outs. *Journal of Political Economy*, 111(2):227–268.

Lang, L. H. P. and Stulz, R. M. (1994). Tobin's q, corporate diversification, and firm performance. *Journal of Political Economy*, 102(6):1248–1280.

Li, J., Tang, Y., and Yan, A. (2016). Corporate equity ownership and expected stock returns. Working paper.

Lou, D. (2012). A flow-based explanation for return predictability. *Review of Financial Studies*, 25(12):3457–3489.

Mitchell, M., Pulvino, T., and Stafford, E. (2002). Limited arbitrage in equity markets. *Journal of Finance*, 57(2):551–584.

Peters, F. S. (2024). Optimal peers. Working paper.

Poterba, J. M. and Summers, L. H. (1988). Mean reversion in stock prices: Evidence and implications. *Journal of Financial Economics*, 22(1):27–59.

Roll, R. (1988). R2. *Journal of Finance*, 43(3):541–566.

Rosenthal, L. and Young, C. (1990). The seemingly anomalous price behavior of royal dutch/shell and unilever nv/plc. *Journal of Financial Economics*, 26(1):123–141.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory*, 13(3):341–360.

Ross, S. A. (1987). The interrelations of finance and economics: Theoretical perspectives. *American Economic Review*, 77(2):29–34.

Servaes, H. (1996). The value of diversification during the conglomerate merger wave. *Journal of Finance*, 51(4):1201–1225.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.

Villalonga, B. (2004). Diversification discount or premium? new evidence from the business information tracking series. *Journal of Finance*, 59(2):479–506.

A. Transmission Rate: Proofs

In this Appendix, I formally show that the transmission rate $\mathcal{T}$ is equal to one under value additivity and that controlling for the common factors affecting both the subsidiary and the parent's other assets permits unbiased estimation of the transmission rate.

A.1. Deriving Value Additivity

Let an asset with payoff x_{t+1}, consisting of its next-period cash flow and ex-dividend price, be priced according to the standard pricing equation (Cochrane, 2005)

$$p_t = \mathbb{E}(m_{t+1}x_{t+1} \mid \mathcal{F}_t)\,, \tag{19}$$

where m_{t+1} is the stochastic discount factor and $\mathcal{F}_t$ the information set at time t. Iterating (19) forward yields the infinite-horizon valuation formula

$$p_t = \mathbb{E}\left(\sum_{k=1}^{\infty} M_{t,t+k}\, d_{t+k} \;\middle|\; \mathcal{F}_t \right), \tag{20}$$

where d_{t+k} is the cash flow at $t+k$ and $M_{t,t+k} = m_{t+1}m_{t+2}\cdots m_{t+k}$ is the cumulative discount factor.

For a publicly traded parent firm whose total payoff stream is the sum of the payoff from its subsidiary stake and the payoff from its other net assets,

$$d^{p}_{t+k} = d^{s}_{t+k} + d^{o}_{t+k}. \tag{21}$$

Substituting (21) into (20) gives

$$p^{P}_t = \mathbb{E}\left(\sum_{k=1}^{\infty} M_{t,t+k}(d^{S}_{t+k} + d^{O}_{t+k}) \;\middle|\; \mathcal{F}_t \right). \tag{22}$$

By linearity of expectation and conditioning on the same information set,

$$\begin{aligned} p^{P}_t = {} & \mathbb{E}\left(\sum_{k=1}^{\infty} M_{t,t+k} d^{S}_{t+k} \;\middle|\; \mathcal{F}_t \right) \\ & + \mathbb{E}\left(\sum_{k=1}^{\infty} M_{t,t+k} d^{O}_{t+k} \;\middle|\; \mathcal{F}_t \right). \end{aligned} \tag{23}$$

Recognizing each term as the market value of the corresponding payoff stream gives

$$p^{P}_t = p^{S}_t + p^{O}_t. \tag{24}$$

Thus, under the standard asset pricing framework, asset prices equal the discounted expectation of future cash flows under a common stochastic discount factor, implying a common linear pricing rule under which the market value of the parent equals the sum of the market values of its subsidiary stake and its other net assets.

A.2. Deriving Return Additivity

Start with value additivity in levels,

$$p_t^P = p_t^S + p_t^O, \tag{25}$$

$$\Delta p_t^P = \Delta p_t^S + \Delta p_t^O, \tag{26}$$

$$\frac{\Delta p_t^P}{p_{t-1}^P} = \frac{p_{t-1}^S}{p_{t-1}^P}\frac{\Delta p_t^S}{p_{t-1}^S} + \frac{p_{t-1}^O}{p_{t-1}^P}\frac{\Delta p_t^O}{p_{t-1}^O}, \tag{27}$$

$$r_t^P = w_{t-1}r_t^S + (1-w_{t-1})r_t^O, \tag{28}$$

where

$$w_{t-1} = \frac{p_{t-1}^S}{p_{t-1}^P}.$$

Using the notation $r_t^{wS} = w_{t-1}r_t^S$ and $r_t^{wO} = (1-w_{t-1})r_t^O$, we obtain

$$r_t^P = r_t^{wS} + r_t^{wO}. \tag{29}$$

Combining this return-additivity relation with the definition of the transmission rate in equation (3), and holding the return on the parent's other assets fixed, gives

$$\mathcal{T} \equiv \left.\frac{\partial r_t^P}{\partial r_t^{wS}}\right|_{r_t^{wO}} = \left.\frac{\partial\left(r_t^{wS} + r_t^{wO}\right)}{\partial r_t^{wS}}\right|_{r_t^{wO}} = 1. \tag{30}$$

Thus, value additivity implies return additivity and, in turn, a transmission rate of one. The reverse implication does not generally hold.

A.3. Regression with Common Factors

Let us assume that each simple return can be decomposed into two parts: (i) a component driven by common factors, consisting of a constant and factors shared across assets, and (ii) an idiosyncratic term. Formally,

$$r_t^j = F_t\beta^j + e_t^j, \qquad j \in J, \quad \{s, o\} \subseteq J, \tag{31}$$

where F_t is the row vector of common factors (including a constant), β^j is the column vector of associated factor loadings, and e_t^j is the idiosyncratic component of the return. The index set J denotes the universe of assets under consideration and includes in particular the subsidiary and the parent's other net assets.

The idiosyncratic component is defined to satisfy mean-independence conditions:

$$E[\, e_t^j \mid \mathcal{H}_t \,] = 0, \qquad \text{with } \mathcal{H}_t = \{\Omega_{t-1},\ F_t,\ e_t^{j'}\ (j' \neq j)\}. \tag{32}$$

That is, each idiosyncratic return e_t^j has zero conditional mean given the current common factors and the contemporaneous idiosyncratic returns of other assets. The information set Ω_{t-1} includes at least past prices and therefore contains w_{t-1}.

A.4. From Return Additivity to a Regression Specification

Starting from return additivity and plugging in the factor structure,

$$r_t^P = r_t^{wS} + r_t^{wO}, \tag{33}$$

$$r_t^P = w_{t-1} r_t^S + (1 - w_{t-1}) r_t^O, \tag{34}$$

$$r_t^P = w_{t-1}(F_t \gamma + e_t^S) + (1 - w_{t-1})(F_t \beta^O + e_t^O), \tag{35}$$

$$e_t^P = e_t^{wS} + e_t^{wO}, \tag{36}$$

where $e_t^{wS} = w_{t-1} e_t^S$ and $e_t^{wO} = (1 - w_{t-1}) e_t^O$. By (32),

$$E[e_t^{wS} e_t^{wO}] = E\big[E[e_t^{wS} e_t^{wO} \mid \mathcal{H}_t]\big] = E\big[w_{t-1}(1 - w_{t-1}) E[e_t^S e_t^O \mid \mathcal{H}_t]\big] = 0.$$

The equation $e_t^P = e_t^{wS} + e_t^{wO}$ is what results from partialling out r_t^p and r_t^s with respect to F_t. By the Frisch–Waugh–Lovell theorem, estimating $\mathcal{T}$ with

$$e_t^P = \mathcal{T} e_t^{wS} + e_t^{wO} \tag{37}$$

is equivalent to estimating τ with

$$r_t^P = \mathcal{T} r_t^{wS} + F_t^w \beta + e_t^{wO}. \tag{38}$$

The issue with this specification is that $F_t^w = w_{t-1} F_t$ changes only slowly. However, it is sufficient to estimate

$$r_t^P = \mathcal{T} e_t^{wS} + v_t, \tag{39}$$

where $v_t = F_t(\beta^O - w_{t-1}\beta) + e_t^{wO}$ and, by (32),

$$E[e_t^{wS} v_t] = E\big[E[e_t^{wS} v_t \mid \mathcal{H}_t]\big] = 0.$$

This shows that it is enough to partial out the regressor of interest, a result known as the *regression anatomy theorem* Angrist and Pischke (2009). Equivalently,

$$r_t^P = \mathcal{T} r_t^{wS} + F_t^w \beta + u_t, \tag{40}$$

where $u_t = F_t\beta^O - F_t^w\beta + e_t^{wO}$ and β is chosen such that $E[F_t^w u_t] = 0$, which in turn implies that u_t is orthogonal to r_t^{wS}, since $\mathcal{T}$ is identified only from the orthogonal component e_t^{wS}.

A.5. Convergence to Normality of $\bar{\mathcal{T}}$

For each pair i (one-year window),

$$\hat{\mathcal{T}}_i = \mathcal{T} + \varepsilon_i, \qquad \varepsilon_i = (\tilde{F}_i' \tilde{F}_i)^{-1} \tilde{F}_i' u_i, \qquad u_{i,t} = F_t \beta_i^O - \tilde{F}_t \beta_i + \tilde{e}_{i,t}^O, \tag{41}$$

with $\tilde{F}_t = w_{t-1} F_t$ and $\tilde{F}_i$ stacking $\tilde{F}_t$ for $t \in T_i$. Under H_0^s, $\mathcal{T} = 1$, so $\hat{\tau}_i = 1 + \varepsilon_i$. Define

$$\bar{\mathcal{T}} = \frac{1}{n} \sum_{i=1}^{n} \hat{\mathcal{T}}_i. \tag{42}$$

A central limit theorem for $\bar{\mathcal{T}}$ follows once $E[\varepsilon_i \varepsilon_{i'}] = 0$ for $i \neq i'$ and second moments are bounded.

Non-overlapping windows. Assume $T_i \cap T_{i'} = \varnothing$. Let $\mathcal{H}_t$ denote the filtration up to time t, including past factors and weights. Since

$$\varepsilon_i = (\tilde{F}_i' \tilde{F}_i)^{-1} \tilde{F}_i' u_i = \sum_{t \in T_i} a_{i,t} u_{i,t}, \tag{43}$$

where $a_{i,t}$ is the t-th element of $(\tilde{F}_i' \tilde{F}_i)^{-1} \tilde{F}_i'$, I have from (32) that

$$E[a_{i,t} u_{i,t} \mid \mathcal{H}_t] = 0 \quad \text{for all } i, t. \tag{44}$$

For any $t \in T_i$ and $t' \in T_{i'}$, conditioning on the later sigma-field $\mathcal{H}_{\max(t,t')}$ gives

$$E[a_{i,t} u_{i,t} \, a_{i',t'} u_{i',t'}] = E\big[E[a_{i,t} u_{i,t} \, a_{i',t'} u_{i',t'} \mid \mathcal{H}_{\max(t,t')}]\big] = 0, \tag{45}$$

and therefore

$$E[\varepsilon_i \varepsilon_{i'}] = 0. \tag{46}$$

Overlapping windows. When $T_i \cap T_{i'} \neq \varnothing$, I fix the current common information $\mathcal{H}_t$. It is sufficient that the factor-adjusted components be conditionally orthogonal across pairs on overlapping dates:

$$E\Big[(F_t \beta_i^O - \tilde{F}_t \beta_i)\,(F_t \beta_{i'}^O - \tilde{F}_t \beta_{i'}) \mid \mathcal{H}_t\Big] = 0 \quad \text{for } t \in T_i \cap T_{i'}. \tag{47}$$

Because $\tilde{F}_t = w_{t-1} F_t$ and w_{t-1} varies slowly within a year,

$$F_t \beta_i^O - \tilde{F}_t \beta_i = F_t(\beta_i^O - w_{t-1} \beta_i) \tag{48}$$

is typically small. I assume that any remaining deviations are orthogonal across pairs as in (47). Under this condition and bounded moments, the same linear-combination argument yields

$$E[\varepsilon_i \varepsilon_{i'}] = 0. \tag{49}$$

With $E[\varepsilon_i \varepsilon_{i'}] = 0$ for $i \neq i'$ and finite second moments, a central limit theorem gives

$$\sqrt{n}\,(\bar{\mathcal{T}} - E[\hat{\mathcal{T}}_i]) \xrightarrow{d} \mathcal{N}(0, \sigma^2), \qquad \sigma^2 = \mathrm{Var}(\hat{\mathcal{T}}_i). \tag{50}$$

B. Supplementary Information on the parent–subsidiary pairs data

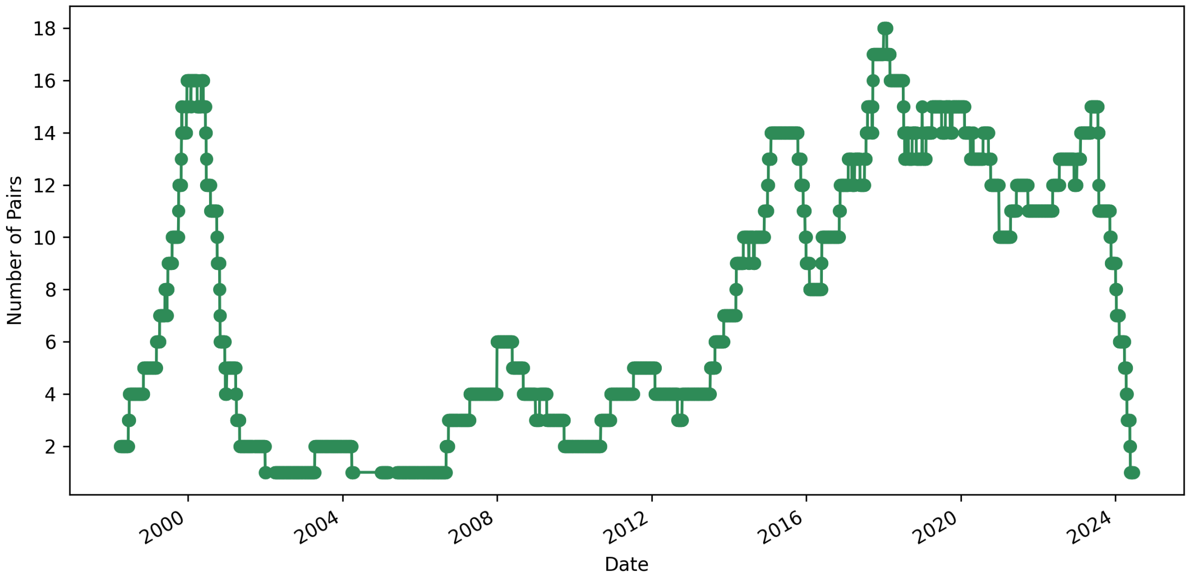

Figure 4: Time-series of the number of parent–subsidiary pairs across time.

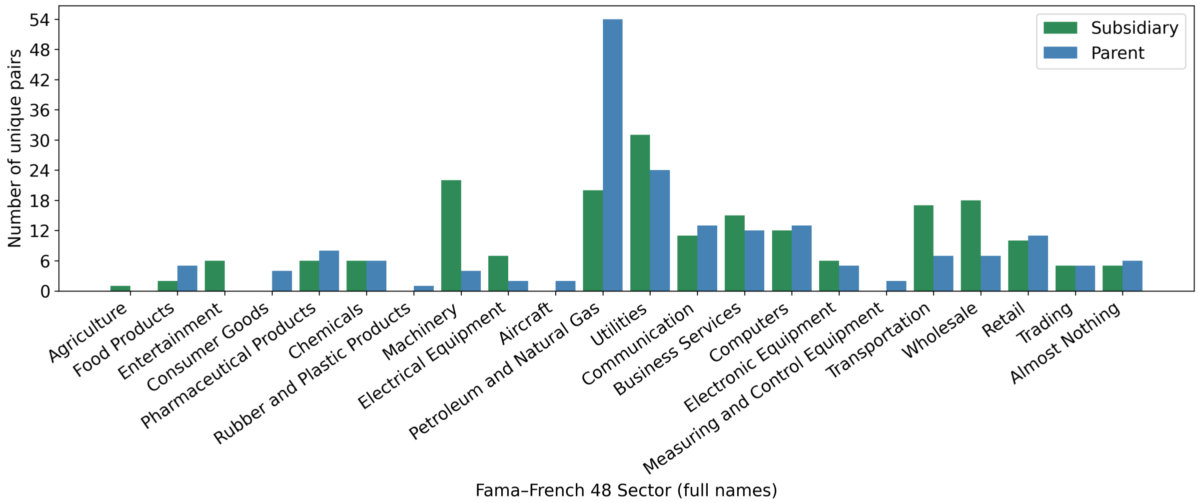

Figure 5: Parent–subsidiary pairs across Fama–French 48 sectors. Most utilities are oil and gas companies with some regulated operations.

C. Average Transmission Rates × characteristics

Table 6: Mean transmission rates estimated by Post-Double-Lasso for All (1999–2024), split by first pair start date.

	All	First<2002	First≥2002
Panel A: Subsidiary Pairs			
n	187	27	160
mean $\mathcal{T}$	0.381^{***}	0.262^{***}	0.401^{***}
SE of mean $\mathcal{T}$	(0.036)	(0.040)	(0.041)
z-stat	10.560	6.564	9.675
p-value	0.000	0.000	0.000
median $\mathcal{T}$	0.387	0.191	0.419
Avg R^2	0.684	0.565	0.704
Avg R^2 (selection on sub)	0.459	0.364	0.475
Panel B: Placebo Pairs			
n	187	27	160
mean $\mathcal{T}$	−0.022	0.066	−0.037
SE of mean $\mathcal{T}$	(0.044)	(0.069)	(0.050)
z-stat	−0.509	0.954	−0.746
p-value	0.611	0.349	0.457
median $\mathcal{T}$	0.023	−0.006	0.026
Avg R^2	0.628	0.515	0.647
Avg R^2 (selection on sub)	0.589	0.463	0.610

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 7: Mean transmission rates by excluding energy-related subsidiary sectors. Columns report statistics for the full sample and after successively excluding subsidiaries in Utilities, then Utilities+Oil, then Utilities+Oil+Transportation (FF sector short names). Sample sizes by column are $n = \{187, 156, 137, 120\}$.

	All	No Util	No Util+Oil	No Util+Oil+Trans
Panel A: Parent–Subsidiary Pairs				
n	187	156	137	120
mean $\mathcal{T}$	0.381^{***}	0.367^{***}	0.394^{***}	0.428^{***}
SE of mean $\mathcal{T}$	(0.036)	(0.032)	(0.034)	(0.035)
z-stat	10.560	11.421	11.617	12.311
p-value	0.000	0.000	0.000	0.000
median $\mathcal{T}$	0.387	0.397	0.413	0.435
Avg Adj. R^2	0.667	0.662	0.667	0.657
Avg R^2	0.684	0.680	0.684	0.674
Avg R^2 (sub)	0.459	0.456	0.457	0.452
Panel B: Parent–Placebo Pairs				
n	187	156	137	120
mean $\mathcal{T}$	−0.022	−0.003	0.018	0.023
SE of mean $\mathcal{T}$	(0.044)	(0.046)	(0.044)	(0.049)
z-stat	−0.509	−0.062	0.407	0.469
p-value	0.611	0.951	0.684	0.640
median $\mathcal{T}$	0.023	0.026	0.026	0.029
Avg Adj. R^2	0.607	0.593	0.590	0.574
Avg R^2	0.628	0.614	0.611	0.596
Avg R^2 (sub)	0.589	0.600	0.598	0.591

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Table 8: Regression of transmission rates on Weight only (placebo pairs). When common factors between parent and subsidiary returns are controlled poorly, a spurious negative correlation arises between portfolio weight and transmission rates. In $\tilde{r}_t^S = w_{t-1} r_t^S$, the same amount of omitted common factor inflates the estimated transmission rate when w_{t-1} is small. In the placebo regressions this negative correlation disappears with Post-Double-Lasso, providing additional support for this specification.

Variable	Simple	FF5+Sector	Post-Double-Lasso
Constant	3.084^{***}	0.955^{***}	-0.048
	(0.229)	(0.151)	(0.088)
Weight	-2.494^{***}	-0.689^{***}	0.051
	(0.265)	(0.154)	(0.090)
R^2	0.304	0.071	0.001
df	185	185	185

Notes: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. *HC3 heteroskedasticity-consistent standard errors in parentheses.*

D. Explaining the Cross-Sectional Variation in Transmission Rate

D.1. Placebo Regressions

As a sanity check, I report the same tables and plot for the placebo transmission rates. For the WLS regression, $\log(\widehat{\mathrm{UVR}}_i)$ has a coefficient of -0.013 with a standard error of 0.008, and for the OLS the coefficient is -0.050 with a standard error of 0.028. The relation is almost flat, but placebo pairs with small UVR tend to have slightly higher transmission rates. This is not surprising if we assume that a relatively homogeneous small positive bias remains in the transmission rates: it becomes more visible where the transmission rate is higher, that is, for small UVR.

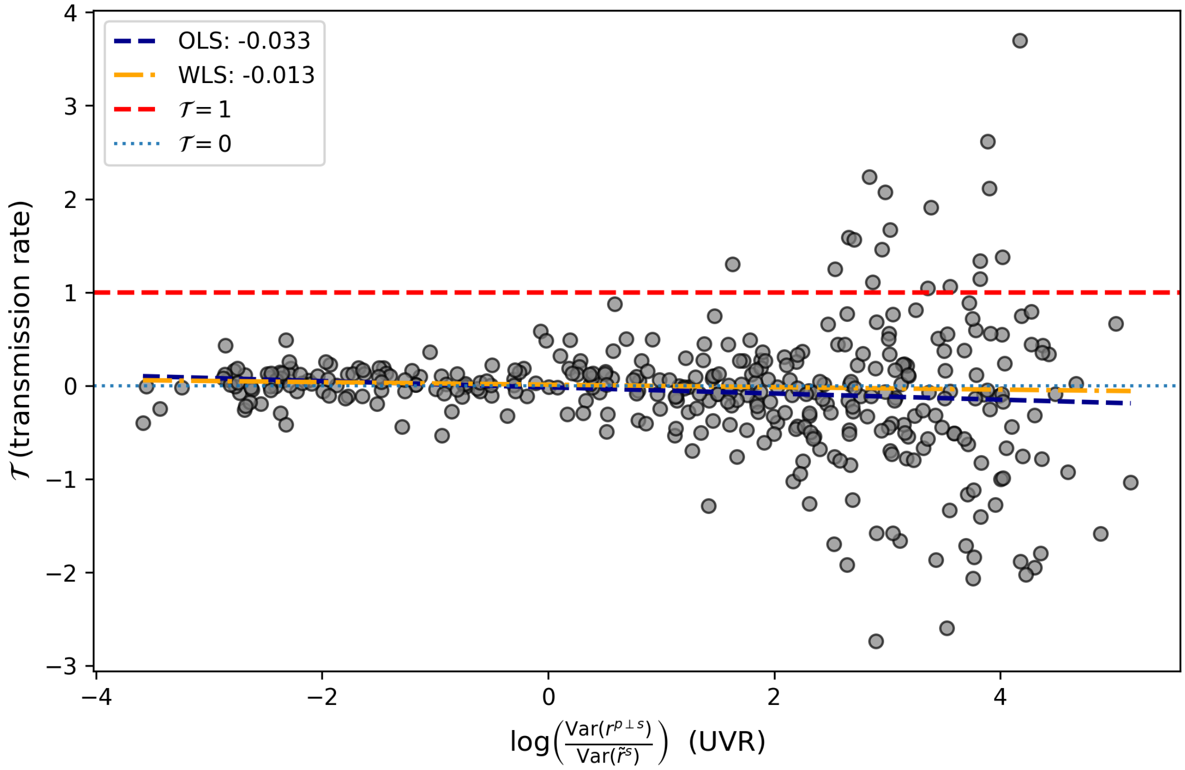

Figure 6: Cross-fitting scatter plots of $\mathcal{T}$ against $\log(UVR)$, where UVR denotes the Unspanned Variance Ratio.

D.2. Idiosyncratic and UVR Regressions

Table 9: Explaining Placebo Transmission Rates Using Weighted Least Squares (WLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\mathrm{UVR}_i}) + \beta_2 \textsc{Float}_i + \beta_3 \mathrm{MC}_i + \beta_4 \textsc{Turnover}_i + \beta_5 \textsc{Amihud}_i + \beta_6 \log(\textsc{Size}_i) + \beta_7 \textsc{Post2002}_i + \varepsilon_i.$$

i denotes a parent–placebo pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. $\log(\widehat{\mathrm{UVR}_i})$ measures the ratio of the parent's unspanned variance to the subsidiary's variance, capturing uncertainty about the parent's other assets. $\textsc{Float}_i$ is the share of the subsidiary's stock available for trading after deducting the parent's holdings. MC_i is a dummy equal to one when the parent controls more than 50% of the voting rights. $\textsc{Turnover}_i$ is the average daily trading volume relative to shares outstanding, and $\textsc{Amihud}_i$ is the average price impact of trading, computed as the absolute return divided by the dollar volume. $\log(\textsc{Size}_i)$ is the logarithm of the combined market value of the parent and subsidiary, and $\textsc{Post2002}_i$ equals one for observations after the internet bubble. Cross-fitting ensures that $\log(\widehat{\mathrm{UVR}_i})$ and $\hat{\mathcal{T}}_i$ are estimated on separate folds. Standard errors are clustered by pair (HC3), and WLS weights are the inverse squared standard errors of $\hat{\mathcal{T}}_i$.

	Parent–Placebo (WLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–Am, –P02
Constant	0.011	0.134**	0.128**	0.119**	0.108**
	(0.016)	(0.057)	(0.057)	(0.049)	(0.049)
log(UVR)	-0.013*	-0.025**	-0.022**	-0.024**	-0.021**
	(0.008)	(0.010)	(0.009)	(0.010)	(0.010)
Amihud	–	-0.012	-0.016	–	–
		(0.038)	(0.038)		
Turnover	–	-5.007**	-4.671**	-4.842**	-4.400**
		(2.211)	(2.168)	(2.169)	(2.097)
Float	–	-0.077	-0.060	-0.081	-0.062
		(0.054)	(0.044)	(0.050)	(0.043)
log(Size)	–	-0.015	-0.011	-0.012	-0.006
		(0.014)	(0.015)	(0.012)	(0.012)
MC	–	-0.041	-0.026	-0.039	-0.020
		(0.030)	(0.028)	(0.031)	(0.028)
Post-2002	–	0.033	–	0.036	–
		(0.040)		(0.036)	
R^2	0.012	0.046	0.043	0.045	0.042
df	372	366	367	367	368

Notes: Weighted least squares (WLS) estimates with weights $1/\mathrm{SE}(\hat{\mathcal{T}}_i)^2$. Standard errors are clustered by pair (HC3). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 10: Explaining Placebo Transmission Rates Using Ordinary Least Squares (OLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\text{UVR}}_i) + \beta_2 \text{Float}_i + \beta_3 \text{MC}_i + \beta_4 \text{Turnover}_i \\ + \beta_5 \text{Amihud}_i + \beta_6 \log(\text{Size}_i) + \beta_7 \text{Post2002}_i + \varepsilon_i.$$

i denotes a parent–placebo pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. $\log(\widehat{\text{UVR}}_i)$ measures the ratio of the parent's unspanned variance to the subsidiary's variance, capturing uncertainty about the parent's other assets. Float_i is the share of the subsidiary's stock available for trading after deducting the parent's holdings. MC_i is a dummy equal to one when the parent controls more than 50% of the voting rights. Turnover_i is the average daily trading volume relative to shares outstanding, and Amihud_i is the average price impact of trading, computed as the absolute return divided by the dollar volume. $\log(\text{Size}_i)$ is the logarithm of the combined market value of the parent and subsidiary, and Post2002_i equals one for observations after the internet bubble. Cross-fitting ensures that $\log(\widehat{\text{UVR}}_i)$ and $\hat{\mathcal{T}}_i$ are estimated on separate folds. Standard errors are clustered by pair (HC3).

	Parent–Placebo (OLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–Am, –P02
Constant	0.018	-0.004	-0.090	-0.041	-0.080
	(0.024)	(0.202)	(0.213)	(0.185)	(0.192)
log(UVR)	-0.050*	-0.035	-0.042	-0.035	-0.043
	(0.028)	(0.028)	(0.027)	(0.028)	(0.027)
Float	–	0.166	0.112	0.170	0.109
		(0.195)	(0.184)	(0.196)	(0.185)
MC	–	0.179*	0.152	0.181*	0.150*
		(0.099)	(0.092)	(0.098)	(0.090)
Turnover	–	3.183	2.605	3.416	2.510
		(8.290)	(8.206)	(8.312)	(8.150)
Amihud	–	-0.040	0.012	–	–
		(0.090)	(0.067)		
log(Size)	–	-0.034	-0.039	-0.029	-0.040
		(0.048)	(0.047)	(0.044)	(0.041)
Post-2002	–	-0.173	–	-0.157	–
		(0.110)		(0.105)	
R^2	0.017	0.038	0.033	0.038	0.033
df	372	366	367	367	368

Notes: Standard errors clustered by pair (HC3). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 11: Explaining Transmission Rates Using Weighted Least Squares (WLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\text{UVR}}_{\text{idiosyn.},i}) + \beta_2 \text{F}\textsc{loat}_i + \beta_3 \text{MC}_i + \beta_4 \text{T}\textsc{urnover}_i$$
$$+ \beta_5 \text{A}\textsc{mihud}_i + \beta_6 \log(\text{S}\textsc{ize}_i) + \beta_7 \text{P}\textsc{ost}2002_i + \varepsilon_i.$$

i denotes a parent–subsidiary pair observed over one year. The main regressor is the log Idiosyncratic UVR,

$$\log(\widehat{\text{UVR}}_{\text{idiosyn.},i}) = \log\left(\frac{\text{Var}(e_i^O)}{\text{Var}(\tilde{e}_i^S)}\right),$$

where both e_i^O and $\tilde{e}_i^S$ are residualized with respect to the same set of controls, ensuring comparability. Cross-fitting ensures that both $\hat{\mathcal{T}}_i$ and $\log(\widehat{\text{UVR}}_{\text{idiosyn.},i})$ are estimated on separate folds. Standard errors are clustered by pair (HC3).

	Parent–Subsidiary (WLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–Am, –P02
Constant	0.380***	0.391***	0.446***	0.260**	0.225
	(0.025)	(0.134)	(0.155)	(0.126)	(0.193)
$\log(\text{UVR}_{\text{idiosyn.}})$	-0.118***	-0.118***	-0.100***	-0.119***	-0.095**
	(0.022)	(0.019)	(0.032)	(0.021)	(0.039)
Float	–	0.098	0.278	0.050	0.251
		(0.202)	(0.357)	(0.207)	(0.408)
MC	–	-0.052	0.091	-0.029	0.181
		(0.112)	(0.143)	(0.115)	(0.140)
Turnover	–	-5.117	-2.820	-4.416	-0.762
		(3.158)	(3.477)	(3.011)	(3.854)
Amihud	–	-0.167	-0.269**	–	–
		(0.104)	(0.136)		
log(Size)	–	-0.099***	-0.083***	-0.077***	-0.037
		(0.020)	(0.026)	(0.020)	(0.034)
Post-2002	–	0.357***	–	0.424***	–
		(0.126)		(0.120)	
R^2	0.223	0.575	0.471	0.544	0.378
df	372	366	367	367	368

Notes: Weighted least squares (WLS) estimates with weights $1/\text{SE}(\hat{\mathcal{T}}_i)^2$. Standard errors are clustered by pair (HC3). Significance levels: $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 12: Explaining Transmission Rates Using Ordinary Least Squares (OLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\text{UVR}}_{\text{idiosyn.},i}) + \beta_2\textsc{Float}_i + \beta_3\text{MC}_i + \beta_4\textsc{Turnover}_i \\ + \beta_5\textsc{Amihud}_i + \beta_6 \log(\textsc{Size}_i) + \beta_7\textsc{Post2002}_i + \varepsilon_i.$$

i denotes a parent–subsidiary pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. The main regressor is the log Idiosyncratic UVR, defined as

$$\log(\widehat{\text{UVR}}_{\text{idiosyn.},i}) = \log\left(\frac{\text{Var}(e_i^O)}{\text{Var}(\tilde{e}_i^S)}\right),$$

where e_i^O and $\tilde{e}_i^S$ are the residual returns of the parent and subsidiary after partialling out the same set of controls. This construction isolates the component of the parent's risk that is orthogonal to the subsidiary. Cross-fitting ensures that both $\hat{\mathcal{T}}_i$ and $\log(\widehat{\text{UVR}}_{\text{idiosyn.},i})$ are estimated on separate folds. Standard errors are clustered by pair (HC3).

	Parent–Subsidiary (OLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–Am, –P02
Constant	0.390***	0.126	0.146	0.061	0.078
	(0.024)	(0.156)	(0.157)	(0.142)	(0.144)
$\log(\text{UVR}_{\text{idiosyn.}})$	-0.053**	-0.045**	-0.043**	-0.044**	-0.042**
	(0.020)	(0.021)	(0.020)	(0.021)	(0.020)
Float	–	0.244	0.256	0.252	0.278*
		(0.163)	(0.158)	(0.162)	(0.157)
MC	–	0.191**	0.197**	0.194**	0.207***
		(0.080)	(0.076)	(0.080)	(0.075)
Turnover	–	7.571	7.690	7.977	8.345
		(5.196)	(5.141)	(5.184)	(5.104)
Amihud	–	-0.071	-0.082	–	–
		(0.076)	(0.073)		
log(Size)	–	-0.042	-0.041	-0.033	-0.028
		(0.030)	(0.030)	(0.027)	(0.027)
Post-2002	–	0.038	–	0.066	–
		(0.073)		(0.067)	
R^2	0.031	0.084	0.083	0.082	0.081
df	372	366	367	367	368

Notes: Standard errors clustered by pair (HC3). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 13: Explaining Placebo Transmission Rates Using Weighted Least Squares (WLS) and Cross-Fitting. The regression is

$$\begin{aligned}\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\mathrm{UVR}}_{\mathrm{idiosyn.},i}) + \beta_2\textsc{Float}_i + \beta_3\mathrm{MC}_i + \beta_4\textsc{Turnover}_i \\ + \beta_5\textsc{Amihud}_i + \beta_6 \log(\textsc{Size}_i) + \beta_7\textsc{Post2002}_i + \varepsilon_i.\end{aligned}$$

i denotes a parent–placebo pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. The main regressor is the log Idiosyncratic UVR, defined as

$$\log(\widehat{\mathrm{UVR}}_{\mathrm{idiosyn.},i}) = \log\left(\frac{\mathrm{Var}(e_i^o)}{\mathrm{Var}(\tilde{e}_i^s)}\right),$$

where e_i^o and $\tilde{e}_i^{PL}$ are the residual returns of the parent and placebo after partialling out the same set of controls. This construction isolates the component of the parent's risk that is orthogonal to the placebo. Cross-fitting ensures that both $\hat{\mathcal{T}}_i$ and $\log(\widehat{\mathrm{UVR}}_{\mathrm{idiosyn.},i})$ are estimated on separate folds. Standard errors are clustered by pair (HC3).

	Parent–Placebo (WLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–Am, –P02
Constant	0.018	0.134**	0.122*	0.109*	0.092*
	(0.013)	(0.068)	(0.066)	(0.059)	(0.054)
$\log(\mathrm{UVR}_{\mathrm{idiosyn.}})$	-0.022*	-0.039*	-0.029	-0.037*	-0.026
	(0.013)	(0.022)	(0.018)	(0.021)	(0.018)
Float	–	-0.074	-0.043	-0.079	-0.047
		(0.060)	(0.057)	(0.055)	(0.054)
MC	–	-0.036	-0.008	-0.032	-0.001
		(0.034)	(0.030)	(0.035)	(0.029)
Turnover	–	-4.887**	-4.182*	-4.580**	-3.788*
		(2.429)	(2.260)	(2.309)	(2.097)
Amihud	–	-0.021	-0.025	–	–
		(0.060)	(0.059)		
log(Size)	–	-0.020	-0.011	-0.014	-0.003
		(0.019)	(0.019)	(0.019)	(0.016)
Post-2002	–	0.069	–	0.072*	–
		(0.045)		(0.043)	
R^2	0.020	0.054	0.045	0.052	0.042
df	372	366	367	367	368

Notes: Weighted least squares (WLS) estimates with weights $1/\mathrm{SE}(\hat{\mathcal{T}}_i)^2$. Standard errors are clustered by pair (HC3). Significance levels: $^*p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

Table 14: Explaining Placebo Transmission Rates Using Ordinary Least Squares (OLS) and Cross-Fitting. The regression is

$$\hat{\mathcal{T}}_i = \beta_0 + \beta_1 \log(\widehat{\mathrm{UVR}}_{\mathrm{idiosyn.},i}) + \beta_2 \textsc{Float}_i + \beta_3 \mathrm{MC}_i + \beta_4 \textsc{Turnover}_i + \beta_5 \textsc{Amihud}_i + \beta_6 \log(\textsc{Size}_i) + \beta_7 \textsc{Post2002}_i + \varepsilon_i.$$

i denotes a parent–placebo pair observed over one year. The dependent variable is the cross-fitted transmission rate $\hat{\mathcal{T}}_i$, estimated via Post-Double-Lasso. The main regressor is the log Idiosyncratic UVR, defined as

$$\log(\widehat{\mathrm{UVR}}_{\mathrm{idiosyn.},i}) = \log\left(\frac{\mathrm{Var}(e_i^O)}{\mathrm{Var}(\tilde{e}_i^{PL})}\right),$$

where e_i^O and $\tilde{e}_i^{PL}$ are the residual returns of the parent and placebo after partialling out the same set of controls. This construction isolates the component of the parent's risk that is orthogonal to the placebo. Cross-fitting ensures that both $\hat{\mathcal{T}}_i$ and $\log(\widehat{\mathrm{UVR}}_{\mathrm{idiosyn.},i})$ are estimated on separate folds. Standard errors are clustered by pair (HC3).

	Parent–Placebo (OLS)				
Variable	Baseline	Full	No Post-2002	No Amihud	–Am, –P02
Constant	-0.002	-0.009	-0.100	-0.043	-0.085
	(0.022)	(0.203)	(0.215)	(0.188)	(0.195)
$\log(\mathrm{UVR}_{\mathrm{idiosyn.}})$	-0.044	-0.029	-0.035	-0.029	-0.036
	(0.027)	(0.027)	(0.026)	(0.027)	(0.026)
Float	–	0.158	0.101	0.162	0.096
		(0.197)	(0.185)	(0.197)	(0.186)
MC	–	0.176*	0.146	0.177*	0.144
		(0.099)	(0.092)	(0.099)	(0.090)
Turnover	–	3.395	2.827	3.605	2.685
		(8.325)	(8.256)	(8.350)	(8.204)
Amihud	–	-0.037	0.018	–	–
		(0.088)	(0.065)		
log(Size)	–	-0.035	-0.040	-0.031	-0.043
		(0.048)	(0.047)	(0.044)	(0.041)
Post-2002	–	-0.180*	–	-0.165	–
		(0.108)		(0.103)	
R^2	0.016	0.037	0.031	0.037	0.031
df	372	366	367	367	368

Notes: Standard errors clustered by pair (HC3). Significance levels: $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$.

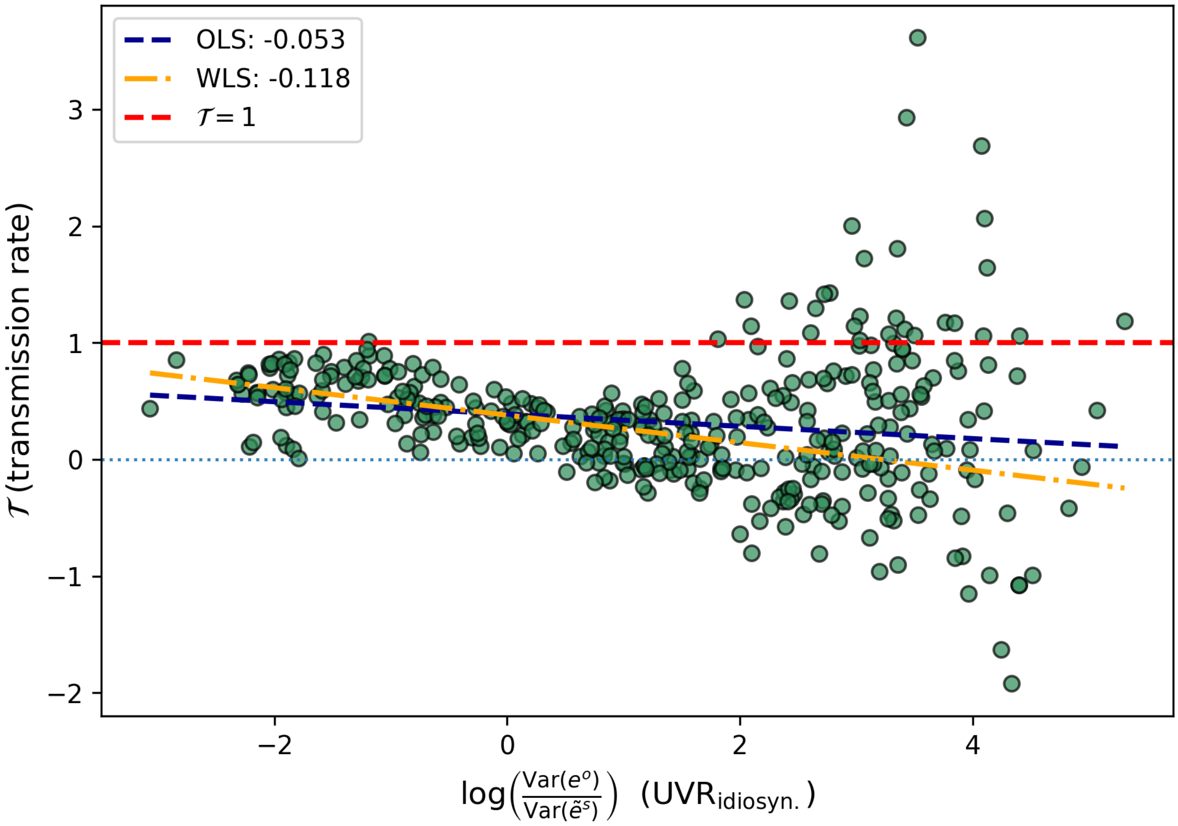

Figure 7: Cross-fitting scatter plots of $\mathcal{T}$ against $\log(UVR)$, where UVR denotes the Unspanned Variance Ratio.

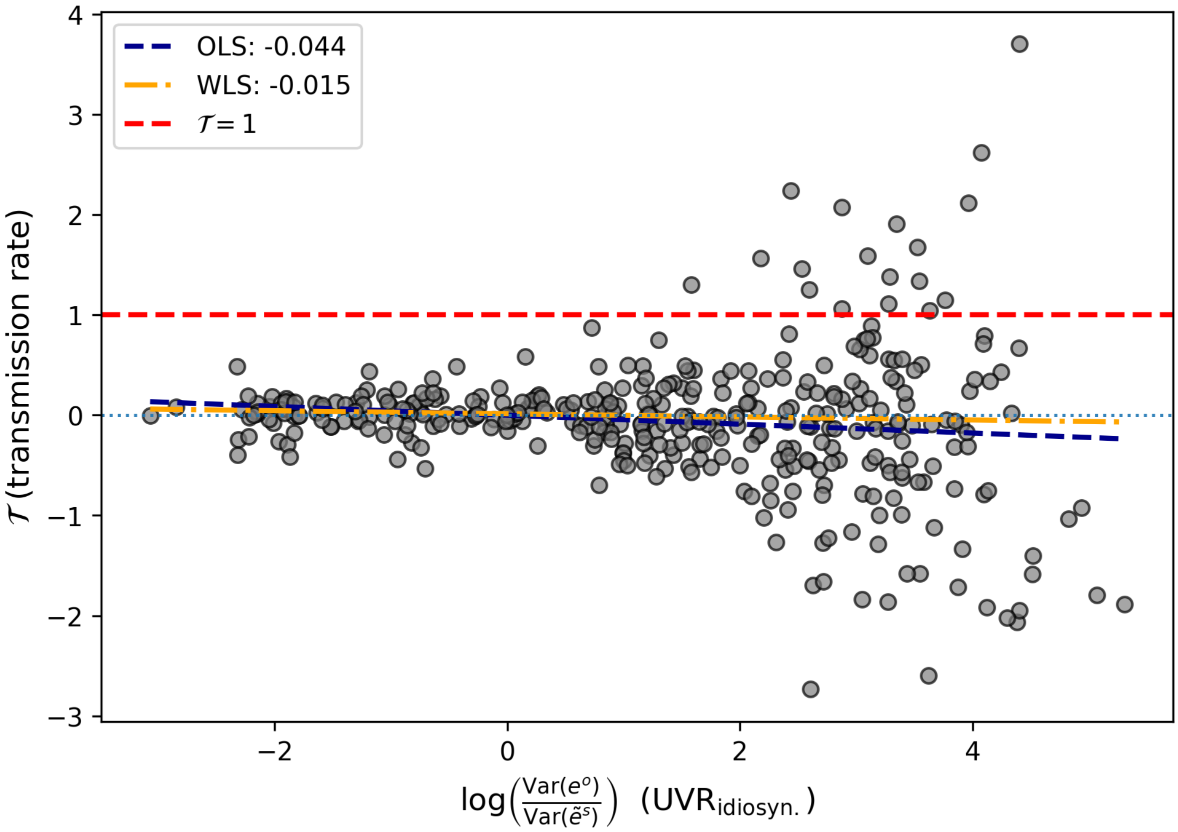

Figure 8: Cross-fitting scatter plots of $\mathcal{T}$ against $\log(UVR)$, where UVR denotes the Unspanned Variance Ratio.

E. An Information-Based Interpretation of the Limits-to-Arbitrage Mechanism

This appendix provides an information-based interpretation of the relation between uncertainty surrounding the parent's other assets and the transmission rate. In the baseline model, arbitrage is limited because a position in the parent–subsidiary spread exposes arbitrageurs to the non-hedgeable payoff risk associated with the parent's other assets. An analogous force arises when the value of these assets is difficult to infer.

The parent's other assets are not separately traded. As a result, arbitrageurs, like the econometrician, cannot directly observe their market value. To assess whether the subsidiary stake is appropriately reflected in the parent's price, arbitrageurs must infer the value of the parent's other assets from the prices of the parent and the subsidiary. When these prices provide only noisy information about this value, arbitrageurs trade less aggressively and the transmission rate remains below one.

To isolate this mechanism, I consider a simplified version of the model in which the price of the subsidiary stake is exogenous to arbitrage demand. This assumption makes the information-based mechanism transparent and yields a closed-form expression for the transmission rate.

E.1. Environment

There are two dates, $t = 0, 1$, and the gross risk-free rate is normalized to one. There are three assets, indexed by $k \in \{a, b, c\}$. Assets a and b are publicly traded, while asset c is not directly traded. Asset a is a claim on the combined payoff of assets b and c:

$$\theta_{a,1} = \theta_{b,1} + \theta_{c,1}. \tag{51}$$

In the empirical application, asset a represents the parent, asset b represents the parent's stake in the publicly traded subsidiary, and asset c represents the parent's remaining assets net of liabilities.

Terminal payoffs satisfy

$$\begin{aligned} \theta_{b,1} &= \theta_{b,0} + \eta_{b,0}, \\ \theta_{c,1} &= \theta_{c,0} + \eta_{c,0}, \end{aligned} \tag{52}$$

where

$$\eta_{b,0} \sim \mathcal{N}\left(0, \tau_{\eta,b}^{-1}\right), \qquad \eta_{c,0} \sim \mathcal{N}\left(0, \tau_{\eta,c}^{-1}\right), \tag{53}$$

and

$$\eta_{b,0} \perp \eta_{c,0}. \tag{54}$$

It follows from equation (51) that

$$\eta_{a,0} = \eta_{b,0} + \eta_{c,0}. \tag{55}$$

There is a continuum of dedicated investors in asset a, indexed by $i \in I_a$, and a continuum of dedicated investors in asset b, indexed by $i \in I_b$. Each continuum has unit mass. Investors in I_a trade only asset a, while investors in I_b trade only asset b. A third continuum of arbitrageurs trades the parent–subsidiary spread.

E.2. Dedicated Investors

Dedicated investors have CARA utility. Each investor $i \in I_a$ observes a private signal about the payoff innovation of asset a:

$$s_{a,0}^i = \eta_{a,0} + \varepsilon_{a,0}^i, \qquad \varepsilon_{a,0}^i \sim \mathcal{N}\left(0, \tau_{s,a}^{-1}\right). \tag{56}$$

The investor holds a prior belief

$$\eta_{a,0} \sim \mathcal{N}\left(\bar{\eta}_{a,0}^i, \tau_{\pi,a}^{-1}\right), \tag{57}$$

where

$$\bar{\eta}_{a,0}^i = n_{a,0} + u_{a,0}^i. \tag{58}$$

The common component

$$n_{a,0} \sim \mathcal{N}\left(0, \tau_{n,a}^{-1}\right) \tag{59}$$

captures aggregate sentiment in the market for the parent, while $u_{a,0}^i$ is an idiosyncratic prior component that averages to zero across investors.

Investors in asset b are defined symmetrically. Each investor $i \in I_b$ observes

$$s_{b,0}^i = \eta_{b,0} + \varepsilon_{b,0}^i, \qquad \varepsilon_{b,0}^i \sim \mathcal{N}\left(0, \tau_{s,b}^{-1}\right), \tag{60}$$

and holds a prior belief

$$\eta_{b,0} \sim \mathcal{N}\left(\bar{\eta}_{b,0}^i, \tau_{\pi,b}^{-1}\right), \tag{61}$$

where

$$\bar{\eta}_{b,0}^i = n_{b,0} + u_{b,0}^i, \qquad n_{b,0} \sim \mathcal{N}\left(0, \tau_{n,b}^{-1}\right). \tag{62}$$

All primitive shocks are mutually independent except for the payoff relation in equation (51).

Under CARA-normality, aggregate dedicated-investor demands are linear:

$$Q_{a,0} = S_a \left[\theta_{a,0} + \kappa_a \eta_{a,0} + (1-\kappa_a) n_{a,0} - p_{a,0}\right], \tag{63}$$

and

$$Q_{b,0} = S_b \left[\theta_{b,0} + \kappa_b \eta_{b,0} + (1-\kappa_b) n_{b,0} - p_{b,0}\right], \tag{64}$$

where

$$\kappa_j \equiv \frac{\tau_{s,j}}{\tau_{s,j} + \tau_{\pi,j}}, \qquad S_j \equiv \frac{\tau_{s,j} + \tau_{\pi,j}}{\gamma_j}, \qquad j \in \{a, b\}. \tag{65}$$

The parameter κ_j is the weight that dedicated investors place on their private signals, while S_j measures the sensitivity of their demand to expected net payoffs.

Abstracting from constant supply terms, define the price signals generated by dedicated investors as

$$\begin{aligned} x_{a,0} &\equiv \kappa_a \left(\eta_{b,0} + \eta_{c,0}\right) + (1-\kappa_a) n_{a,0}, \\ x_{b,0} &\equiv \kappa_b \eta_{b,0} + (1-\kappa_b) n_{b,0}. \end{aligned} \tag{66}$$

The parent-market signal $x_{a,0}$ reflects both the shared payoff innovation $\eta_{b,0}$ and the innovation in the parent's other assets $\eta_{c,0}$. The subsidiary-market signal $x_{b,0}$ reflects the shared payoff innovation, but is also affected by subsidiary-market sentiment.

E.3. Exogenous Price of the Subsidiary Stake

I assume that the price of asset b is exogenous to arbitrage demand. Up to a constant,

$$p_{b,0} = x_{b,0}. \tag{67}$$

This assumption can be interpreted as the market for the subsidiary being deep relative to the market for the parent, so that arbitrage trades have a negligible effect on the subsidiary's price. Arbitrageurs can nevertheless use $p_{b,0}$ as a signal about the shared payoff component.

In the absence of arbitrage, the parent price is similarly given, up to a constant, by

$$p^n_{a,0} = x_{a,0}, \tag{68}$$

where the superscript n denotes the no-arbitrage or segmented-market benchmark.

The corresponding transmission rate is

$$\mathcal{T}_n \equiv \frac{\operatorname{Cov}\left(x_{a,0}, x_{b,0}\right)}{\operatorname{Var}\left(x_{b,0}\right)}. \tag{69}$$

Define

$$V_a \equiv \operatorname{Var}\left(x_{a,0}\right), \qquad V_b \equiv \operatorname{Var}\left(x_{b,0}\right), \tag{70}$$

and

$$C_{ab} \equiv \operatorname{Cov}\left(x_{a,0}, x_{b,0}\right). \tag{71}$$

Using equation (66),

$$V_a = \kappa_a^2 \left(\tau_{\eta,b}^{-1} + \tau_{\eta,c}^{-1}\right) + (1-\kappa_a)^2 \tau_{n,a}^{-1}, \tag{72}$$

$$V_b = \kappa_b^2 \tau_{\eta,b}^{-1} + (1-\kappa_b)^2 \tau_{n,b}^{-1}, \tag{73}$$

and

$$C_{ab} = \kappa_a \kappa_b \tau_{\eta,b}^{-1}. \tag{74}$$

The no-arbitrage transmission rate is therefore

$$\mathcal{T}_n = \frac{\kappa_a \kappa_b \tau_{\eta,b}^{-1}}{\kappa_b^2 \tau_{\eta,b}^{-1} + (1-\kappa_b)^2 \tau_{n,b}^{-1}}. \tag{75}$$

The coefficient $\mathcal{T}_n$ need not equal one. Dedicated investors in the two markets may respond differently to the shared fundamental innovation, and the subsidiary price may contain sentiment that does not transmit to the parent in the absence of arbitrage.

E.4. Arbitrageurs' Inference Problem

Arbitrageurs trade the long-parent, short-subsidiary portfolio. Its terminal payoff is

$$\theta_{a,1} - \theta_{b,1} = \theta_{c,1}. \tag{76}$$

The profitability of the position therefore depends on the value of the parent's other assets. Because asset c is not separately traded, arbitrageurs do not observe this value directly. Instead, they infer the innovation $\eta_{c,0}$ from the price signals $x_{a,0}$ and $x_{b,0}$.

Because the model is Gaussian and linear, the posterior expectation is

$$m_0 \equiv \mathbb{E}\left[\eta_{c,0} \mid x_{a,0}, x_{b,0}\right] = \beta_a x_{a,0} + \beta_b x_{b,0}. \tag{77}$$

The projection coefficients are

$$\beta_a = \frac{\kappa_a \tau_{\eta,c}^{-1} V_b}{V_a V_b - C_{ab}^2}, \tag{78}$$

and

$$\beta_b = -\frac{\kappa_a \tau_{\eta,c}^{-1} C_{ab}}{V_a V_b - C_{ab}^2}. \tag{79}$$

The posterior variance of the innovation in the parent's other assets is

$$\operatorname{Var}\left(\eta_{c,0} \mid x_{a,0}, x_{b,0}\right) = \tau_{\eta,c}^{-1} - \frac{\kappa_a^2 \tau_{\eta,c}^{-2} V_b}{V_a V_b - C_{ab}^2}. \tag{80}$$

Let

$$\tau_{\eta,c|\mathcal{I}_{\text{arb}}} \equiv \left[\operatorname{Var}\left(\eta_{c,0} \mid x_{a,0}, x_{b,0}\right)\right]^{-1} \tag{81}$$

denote the posterior precision with which arbitrageurs estimate the value of the parent's other assets.

Arbitrageurs have CARA utility with risk aversion ξ. Their effective price sensitivity is

$$S_{\text{arb}} \equiv \frac{\tau_{\eta,c|\mathcal{I}_{\text{arb}}}}{\xi}. \tag{82}$$

Equation (82) combines two distinct limits to arbitrage. A higher ξ reduces arbitrageurs' willingness to bear risk. A lower posterior precision reduces their willingness to trade because they are less certain about the value of the non-traded component.

E.5. Arbitrage Demand and the Parent Price

Let $Q_{\text{arb},0}$ denote aggregate demand for the long-parent, short-subsidiary spread. Because the price of the subsidiary is exogenous, arbitrage demand affects only the price of the parent.

Up to constants and supply terms, market clearing in the parent implies

$$p_{a,0} = x_{a,0} + \frac{Q_{\text{arb},0}}{S_a}. \tag{83}$$

Arbitrageurs understand the price impact of their own demand. They therefore infer the dedicated-investor signal $x_{a,0}$ by subtracting their own price impact from the observed parent price:

$$x_{a,0} = p_{a,0} - \frac{Q_{\text{arb},0}}{S_a}. \tag{84}$$

The perceived value of the spread innovation is m_0, while its price, abstracting from constants, is $p_{a,0} - p_{b,0}$. Aggregate arbitrage demand is therefore

$$Q_{\text{arb},0} = S_{\text{arb}} \left[m_0 - (p_{a,0} - p_{b,0}) \right]. \tag{85}$$

Using equations (67) and (83),

$$p_{a,0} - p_{b,0} = x_{a,0} - x_{b,0} + \frac{Q_{\text{arb},0}}{S_a}. \tag{86}$$

Substituting this expression into equation (85) gives

$$Q_{\text{arb},0} = S_{\text{arb}} \left[m_0 - (x_{a,0} - x_{b,0}) - \frac{Q_{\text{arb},0}}{S_a} \right]. \tag{87}$$

Solving for aggregate arbitrage demand yields

$$Q_{\text{arb},0} = \frac{S_a S_{\text{arb}}}{S_a + S_{\text{arb}}} \left[m_0 - (x_{a,0} - x_{b,0}) \right]. \tag{88}$$

Define

$$\psi \equiv \frac{S_{\text{arb}}}{S_a + S_{\text{arb}}} = \frac{1}{1 + S_a/S_{\text{arb}}}. \tag{89}$$

The equilibrium parent price is then

$$p_{a,0} = x_{a,0} + \psi \left[m_0 - (x_{a,0} - x_{b,0}) \right]. \tag{90}$$

Equivalently,

$$p_{a,0} = (1 - \psi) x_{a,0} + \psi x_{b,0} + \psi m_0. \tag{91}$$

The coefficient ψ measures the strength of arbitrage. When $S_{\text{arb}} = 0$, arbitrageurs do not trade and $\psi = 0$. The parent price then equals the price implied by its dedicated investors. As S_{arb} increases, arbitrageurs place more weight on the price of the subsidiary and on their estimate of the parent's other assets.

E.6. Transmission Rate

The equilibrium transmission rate is

$$\mathcal{T} \equiv \frac{\text{Cov}\left(p_{a,0}, p_{b,0}\right)}{\text{Var}\left(p_{b,0}\right)}. \tag{92}$$

Since $p_{b,0} = x_{b,0}$, equation (91) implies

$$\begin{aligned}\text{Cov}\left(p_{a,0}, p_{b,0}\right) = (1-\psi)\text{Cov}\left(x_{a,0}, x_{b,0}\right) \\ + \psi\text{Var}\left(x_{b,0}\right) + \psi\text{Cov}\left(m_0, x_{b,0}\right).\end{aligned} \tag{93}$$

Because m_0 is the conditional expectation of $\eta_{c,0}$ given $x_{a,0}$ and $x_{b,0}$,

$$\text{Cov}\left(m_0, x_{b,0}\right) = \text{Cov}\left(\eta_{c,0}, x_{b,0}\right). \tag{94}$$

The innovation $\eta_{c,0}$ is independent of both $\eta_{b,0}$ and $n_{b,0}$, and hence

$$\text{Cov}\left(\eta_{c,0}, x_{b,0}\right) = 0. \tag{95}$$

It follows that

$$\mathcal{T} = (1-\psi)\mathcal{T}_n + \psi. \tag{96}$$

Equivalently,

$\mathrm{T} = \mathrm{T}_n + (1-\mathcal{T}_n)\,\frac{S_{\text{arb}}}{S_a+S_{\text{arb}}}$.(97)

Substituting $S_{\text{arb}} = \tau_{\eta,c|\mathcal{I}_{\text{arb}}}/\xi$ gives

$$\mathrm{T} = \mathrm{T}_n + (1-\mathcal{T}_n)\,\frac{1}{1+\frac{S_a\xi}{\tau_{\eta,c|\mathcal{I}_{\text{arb}}}}}.$$

(98)

Proposition 1. *Suppose that the no-arbitrage transmission rate satisfies $\mathcal{T}_n < 1$. The equilibrium transmission rate satisfies*

$$\mathcal{T}_n \leq \mathcal{T} \leq 1. \tag{99}$$

It is increasing in arbitrageurs' posterior precision $\tau_{\eta,c|\mathcal{I}_{\text{arb}}}$ and decreasing in their risk aversion ξ.

Proof. Equation (96) expresses the equilibrium transmission rate as a convex combination of $\mathcal{T}_n$ and one, since $\psi \in [0,1]$. Therefore,

$$\mathcal{T}_n \leq \mathcal{T} \leq 1 \tag{100}$$

whenever $\mathcal{T}_n < 1$.

Moreover,

$$\psi = \frac{\tau_{\eta,c|\mathcal{I}_{\text{arb}}}}{S_a\xi + \tau_{\eta,c|\mathcal{I}_{\text{arb}}}}. \tag{101}$$

It follows that

$$\frac{\partial \psi}{\partial \tau_{\eta,c|\mathcal{I}_{\text{arb}}}} = \frac{S_a\xi}{\left(S_a\xi + \tau_{\eta,c|\mathcal{I}_{\text{arb}}}\right)^2} > 0, \tag{102}$$

and

$$\frac{\partial \psi}{\partial \xi} = -\frac{S_a\tau_{\eta,c|\mathcal{I}_{\text{arb}}}}{\left(S_a\xi + \tau_{\eta,c|\mathcal{I}_{\text{arb}}}\right)^2} < 0. \tag{103}$$

Since

$$\mathcal{T} = \mathcal{T}_n + (1 - \mathcal{T}_n)\psi \tag{104}$$

and $1 - \mathcal{T}_n > 0$, the result follows. $\square$

Proposition 1 shows that arbitrage moves the transmission rate toward the frictionless benchmark of one. The strength of this corrective force depends not only on arbitrageurs' risk-bearing capacity, but also on how precisely they can infer the value of the parent's other assets.

When the posterior precision $\tau_{\eta,c|\mathcal{I}_{\text{arb}}}$ is low, arbitrageurs are uncertain about the payoff of the spread. They consequently trade cautiously, and the transmission rate remains close to the segmented-market value $\mathcal{T}_n$. As prices become more informative about the value of the parent's other assets, posterior precision rises, arbitrage becomes more aggressive, and the transmission rate moves toward one.

E.7. Limiting Cases

No arbitrage. If arbitrageurs have no risk-bearing capacity, or if their information about the parent's other assets becomes arbitrarily imprecise, then

$$S_{\text{arb}} \to 0. \tag{105}$$

Consequently,

$$\psi \to 0, \tag{106}$$

and

$$\mathcal{T} \to \mathcal{T}_n. \tag{107}$$

Perfectly informed or risk-neutral arbitrageurs. If arbitrageurs infer the value of the parent's other assets with arbitrarily high precision, then

$$\tau_{\eta,c|\mathcal{I}_{\text{arb}}} \to \infty. \tag{108}$$

Alternatively, if arbitrageurs become risk-neutral, then

$$\xi \to 0. \tag{109}$$

In either case,

$$S_{\text{arb}} \to \infty, \qquad \psi \to 1, \tag{110}$$

and

$$\mathcal{T} \to 1. \tag{111}$$

Thus, informational uncertainty can limit arbitrage even when arbitrageurs have substantial risk-bearing capacity. Conversely, when arbitrageurs can infer the value of the non-traded component sufficiently precisely, they enforce the one-for-one transmission benchmark.

E.8. Relation to the Baseline Model

The baseline model and the information-based model emphasize different interpretations of uncertainty surrounding the parent's other assets.

In the baseline model, arbitrageurs know the expected value of the parent's other assets but remain exposed to their non-hedgeable payoff risk. Greater payoff uncertainty increases the risk of the parent–subsidiary spread and reduces arbitrage demand.

In the present model, arbitrageurs do not directly observe the value of the parent's other assets. Greater informational uncertainty lowers the posterior precision with which they estimate the spread payoff and similarly reduces arbitrage demand.

Despite this difference in interpretation, both mechanisms generate the same qualitative prediction. The transmission rate is close to one when uncertainty surrounding the parent's other assets is small and falls as this uncertainty increases. In both cases, uncertainty weakens the ability of arbitrageurs to enforce the pricing relation implied by the overlap in payoffs.

The absence of a separately observed market price for the parent's other assets therefore matters in two related ways. First, it prevents the econometrician from testing value additivity directly in price levels. Second, it may prevent arbitrageurs from directly identifying whether the subsidiary stake is appropriately reflected in the parent's value. The same feature that complicates empirical identification may consequently weaken the market force that would otherwise enforce the pricing restriction.