Multimarket Contact and the Cost of Trading Corporate Bonds*

Marco Grotteria[1], Andreas Rapp[2], Ruslan Sverchkov[3], and Ram Yamarthy[2]

[1]London Business School and CEPR.

[2]Board of Governors of the Federal Reserve.

[3]Warwick Business School, University of Warwick.

August 14, 2026

Abstract

We show that multimarket contact (MMC)—the extent to which the same dealers interact across many securities—disciplines the pricing of corporate bonds. In a repeated-game model, where dealers alternate between supplying and demanding liquidity, a dealer that is tempted to widen spreads for a short-run gain is deterred by the prospect of facing wider spreads when next demanding liquidity. MMC pools this discipline across securities: a deviation in one security triggers retaliation in the others. The model predicts that higher MMC for one dealer lowers spreads for *all* dealers active in that bond, with the largest effects when adverse selection is high or trading is infrequent. Using TRACE data, we exploit within-bond-month variation across dealers and within-dealer-month variation across bonds. Higher MMC compresses interdealer spreads, with reductions passed through to customers. A shift–share instrument based on the sharp increase in investment-grade bond issuance during COVID confirms the result.

Keywords: Over-the-counter markets; multimarket contact; dealer networks; corporate bonds.

JEL: G12, G14, L13, D85, C73.

*For helpful comments, we thank Maxime Bonelli, Ana Babus, Joao Cocco, Jean-Edouard Colliard, Julian Franks, Kostas Koufopoulos, Albert Menkveld, Lakshmi Naaraayanan, Florian Nagler, Giorgio Ottonello, Tim McQuade, and Milena Wittwer. We also thank seminar and conference participants at 6th David Backus Memorial Conference, Bocconi University, Federal Reserve Board, and SKEMA Conference on the IO of Financial Markets. Corresponding authors: Grotteria (mgrotteria@london.edu), Rapp (andreas.c.rapp@frb.gov), Sverchkov (ruslan.sverchkov@wbs.ac.uk), and Yamarthy (ram.s.yamarthy@frb.gov). The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the research staff or the Board of Governors.

1 Introduction

U.S. corporate bonds trade in a highly concentrated dealer-intermediated market: the top 10 dealers account for over 70% of dealer-to-customer trading volume (O'Hara and Zhou, 2021).[1] Such concentration immediately raises classic market-power concerns. When a few intermediaries dominate market-making, standard oligopoly logic predicts supracompetitive prices—namely, wide bid–ask spreads (Green, 2007; Green, Hollifield, and Schürhoff, 2007; Foucault, Pagano, and Röell, 2013). And because the set of dealers remains largely the same over time, repeated interaction should make it easier to sustain even wider spreads through tacit collusion (Christie and Schultz, 1994, 1995; Dutta and Madhavan, 1997). Yet both interdealer and dealer–customer spreads appear to be driven primarily by dealer inventory and trading frictions rather than pricing power (Feldhütter and Poulsen, 2018; Bessembinder, Spatt, and Venkataraman, 2020). This raises a natural question: what limits dealers' ability to convert concentration into pricing power?

We show that dealer *multimarket contact* (MMC)—the extent to which the same dealers trade with each other across many bonds—disciplines pricing. Because the corporate bond market is fragmented across thousands of CUSIPs, dealers often trade with each other in the interdealer market to offset their positions and manage inventory risk. MMC links their pricing incentives across bonds: a dealer who widens her quote in one bond faces retaliation in every other bond she shares with the same counterparties. Holding total bilateral volume fixed, we show that incentives differ sharply depending on whether trading activity is concentrated in a single bond or spread across many heterogeneous bonds.

We formalize the mechanism in a dynamic model with adverse selection in which dealers alternate between supplying and demanding liquidity. As in Glosten and Milgrom (1985), trade is either information- or liquidity-driven, but only the latter generates gains from trade. While a dealer is tempted to widen quotes for a short-run gain, she is deterred from doing so by role reversals as then she would pay wider spreads when demanding liquidity herself.[2]

[1] As trading migrated from voice to electronic platforms, essentially the same institutions remained central to intermediation: nine of the ten largest voice-trading dealers rank among the top fifteen in electronic trading (O'Hara and Zhou, 2025). Electronification changed the venue but not the underlying dealer structure.

[2] In over-the-counter (OTC) corporate bond markets, liquidity demand flows through bilateral negotiations with a limited dealer set, so quote-widening is not instantly undercut by a public limit-order book (O'Hara and Zhou, 2025).

MMC pools interaction across bonds, so that a deviation in one bond triggers retaliation in all other bonds where the same dealers are active. Bond heterogeneity is central: in frequently traded, low-adverse-selection bonds, dealers trade with each other often enough that the within-bond threat of retaliation alone keeps spreads tight. MMC extends this discipline to illiquid or information-sensitive bonds, where dealers trade too rarely for the within-bond threat to work. Moreover, high-MMC dealers optimally tighten their quotes to all counterparties, and they do so especially aggressively to low-MMC dealers. This, in turn, induces low-MMC dealers to tighten their own quotes, so that MMC held by *some* dealers compresses spreads quoted by *all* intermediaries in that bond.

The model implies that sustainable transaction costs should be lower for dealers with greater MMC. We test this prediction using regulatory TRACE data, which identify dealers in each corporate bond trade. We measure MMC at the dealer–bond–month level, capturing dealer i's overlap with other dealers active in bond b across their other bonds. We construct our baseline MMC measure from *trailing* 12-month trading networks (months $t-12$ through $t-1$), so the overlap measure is predetermined relative to spreads in month t and not mechanically driven by contemporaneous trading in bond b. We then estimate dealer–bond–month regressions of transaction costs on MMC and show robustness to alternative network windows of 6 and 1 months.

Interpreting this relationship requires separating MMC from other determinants of transaction costs that move with participation. Dealers with broad franchises may both overlap more and quote more tightly; liquid bonds may attract more dealers, mechanically increasing overlap while exhibiting lower spreads. We address these possibilities by including bond–month fixed effects, absorbing time-varying bond-level determinants of transaction costs and participation, and dealer–month fixed effects, absorbing time-varying dealer-level determinants such as balance-sheet capacity, funding advantages, and technology. Identification comes from two within-month comparisons: across dealers trading the same bond in the same month and across bonds traded by the same dealer in the same month.

The data strongly support the theory. Within a given bond and month, dealers with higher MMC charge lower transaction costs. A one-standard-deviation increase in MMC is associated with a 4–6 basis point decline in spreads relative to an average of about 23

basis points—a 17–26% reduction. The relationship is stable when we add predetermined controls for dealer–bond activity and trading conditions, and it is robust across credit quality, maturity, and issuer sectors, as well as alternative constructions of MMC and transaction-cost measures and to alternative fixed-effect and clustering specifications.

The empirical results also align with comparative statics of the repeated-game model. Using an unsigned price impact measure—the slope from a regression of absolute price changes on trading volume (Hasbrouck, 2007; Foucault et al., 2013)—as a proxy for adverse selection, we find that MMC has a larger effect in bonds with greater price impact, where the temptation to widen spreads is strongest. Consistent with the threat of retaliation being weaker when dealers interact less often in a given bond, MMC also has a stronger effect in infrequently traded bonds. Finally, consistent with the model's pass-through prediction, we show that tighter interdealer pricing associated with MMC translates into lower dealer–customer transaction costs as well.

Finally, we complement the fixed-effects design with a shift–share instrumental-variables strategy (Bartik, 1991; Blanchard and Katz, 1992) that exploits the sharp and uneven issuance waves following the COVID-19 market stress in March 2020. These surges were concentrated in investment-grade segments, while high-yield issuance did not rise in tandem (Becker and Benmelech, 2025). If market shares are persistent, when a segment experiences a burst of new issuance, the set of securities in which dealers interact expands, mechanically increasing MMC for dealers with larger footprints in the affected segment. We interact segment-level abnormal issuance with each dealer's pre-COVID market share in that segment. Since the shares are measured over the decade preceding the pandemic and held fixed, all cross-sectional variation in the instrument reflects persistent specialization patterns rather than contemporaneous trading decisions. The estimation sample excludes bonds issued after 2019, ruling out mechanical links between the issuance shock and the pricing of bonds in the sample. The IV specification includes the same bond–month and dealer–month fixed effects as the baseline, so identification comes from comparing dealers trading the same bond in the same month who differ in how exposed they were to segment-level issuance shocks elsewhere in the market. Issuance exposure strongly predicts higher MMC in the first stage, and the second-stage estimates are negative but smaller than the fixed-effects OLS estimates

consistent with the hypothesis that the 2SLS identifies a local average treatment effect for dealer–bond pairs whose MMC responds to COVID-era issuance shocks.

Our results help reconcile two seemingly contradictory facts about corporate bond markets: intermediation is highly concentrated, yet observed spreads appear driven by inventory and trading frictions rather than pricing power. The reconciliation comes from a mechanism, grounded in repeated-game incentives, that reverses the standard comparative static of multimarket contact in this setting. We contribute to three strands of prior work.

The first is the literature on multimarket contact. The canonical view in industrial organization is that overlap across markets facilitates tacit coordination by linking punishments across markets (Bernheim and Whinston, 1990), and empirical work in airlines and other industries typically finds that greater overlap raises prices (Evans and Kessides, 1994; Parker and Röller, 1997; Ciliberto and Williams, 2014; Schmitt, 2018; Hatfield and Wallen, 2022). Our setting yields the opposite prediction because role reversal is intrinsic to OTC dealer markets (Duffie, Gârleanu, and Pedersen, 2005): the same intermediary alternates between supplying and demanding liquidity across repeated interactions. The relevant deviation is therefore not aggressive undercutting but opportunistically *widening* spreads when on the liquidity-supplying side of a trade. In this environment, greater cross-bond overlap can strengthen the threat of punishment and lower, rather than raise, sustainable intermediation costs.

The second strand concerns OTC market structure and corporate bond trading costs. A large literature emphasizes search frictions, limited transparency, and balance-sheet constraints as key determinants of spreads (Duffie et al., 2005; Duffie, Gârleanu, and Pedersen, 2007), and documents substantial dispersion in transaction costs across bonds and dealers (Bessembinder, Maxwell, and Venkataraman, 2006; Edwards, Harris, and Piwowar, 2007; Goldstein, Hotchkiss, and Sirri, 2006; Bao, Pan, and Wang, 2011; Friewald, Jankowitsch, and Subrahmanyam, 2012; Friewald and Nagler, 2019). More recent work highlights economies of scale in intermediation (Goldberg and Nozawa, 2021) and the role of dealer–dealer and dealer–customer relationships (Di Maggio, Kermani, and Song, 2017; Hendershott, Li, Livdan, and Schürhoff, 2020; O'Hara, Wang, and Zhou, 2018; Choi, Huh, and Seunghun Shin, 2024; O'Hara, Rapp, and Zhou, 2025). Our channel is complementary to these forces. We

provide model-based and empirical evidence that MMC is associated with tighter pricing even after controlling for bilateral trading activity.

The third strand is the growing literature on financial networks (Allen and Gale, 2000; Jackson, 2008; Elliott, Golub, and Jackson, 2014; Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2015) and trading networks in corporate bond intermediation (Di Maggio et al., 2017; Li and Schürhoff, 2019; Colliard, Foucault, and Hoffman, 2021). That literature shows that a dealer's network position—in particular centrality—helps explain spreads and liquidity. We depart from this work in two ways. First, existing network channels operate through mechanisms such as search, bargaining, or congestion, whereas our mechanism operates through repeated-game incentives shaped by the number of heterogeneous bonds dealers trade in common. Second, our multimarket contact measure varies across bonds within a dealer–month. Dealer–month fixed effects absorb a dealer's overall network position in a given month, so the identifying variation is distinct from centrality or other slow-moving network characteristics. The spread compression we document survives a stringent fixed-effects design and a shift–share instrumental variable, suggesting that cross-bond overlap captures a margin of dealer interaction not subsumed by standard network measures.

2 Theoretical Framework

We develop a repeated-game model of dealer cooperation under multimarket contact. We first show how repeated interaction can sustain cooperative pricing between two dealers in a single market. We then add a second market and show that MMC can make cooperation incentive-compatible in markets where repeated interaction alone is insufficient: the key for MMC to be relevant is bond heterogeneity. We next extend the framework along several dimensions including a setup with many dealers and many bonds, where we allow for heterogeneous market participation and counterparty-specific adverse selection. All proofs and the extensions are in the Online Appendix.

2.1 One Market and Two Dealers

Stage Game. There are $N = 2$ dealers trading in $K = 1$ market. The stage game follows Glosten and Milgrom (1985). One of the two dealers seeks liquidity either because she is liquidity-motivated (L) or information-motivated (I) and the other dealer acts as a market-maker (M) who provides liquidity. All trading in the model is bilateral (e.g., as in Duffie et al., 2005) and counterparties know the identity of each other when they trade. The market-maker's quotes to the liquidity-seeking counterparty are private and not observable to the rest of the market. This structure closely matches the US corporate bond market, where negotiations are bilateral and pre-trade transparency is limited. We assume all dealers are risk-neutral, which allows us to abstract from inventory risk considerations.

The dealers trade a risky asset: its common-value component v is equally likely to be positive or negative, and its magnitude $|v|$ is independently distributed according to a distribution $F_I(v)$ with support on $[0, 1]$. With probability $\alpha \in [0, 1]$, the liquidity-seeking dealer is information-motivated, in which case she privately knows the common-value component. With complementary probability $1 - \alpha$, she is liquidity-motivated, in which case she has a private-value component b for trading the asset (and knows nothing about the common-value component). The dealer's demand in this case can be driven by the need to offset customer trades due to balance sheet constraints or a dealer's private hedging needs. The private-value component is equally likely to be positive or negative, and its magnitude $|b|$ is also independently distributed according to a distribution $F_L(b)$ with support on $[0, 1]$. For ease of exposition, we assume both $F_I(v)$ and $F_L(b)$ are uniform.[3]

The market-making dealer knows neither v nor b and provides take-it-or-leave-it quotes for selling and buying the asset. Since $\mathbb{E}[v] = \mathbb{E}[b] = 0$, the problem of the market maker

[3]Our results do not depend on the uniform assumption. Under symmetric distributions satisfying standard regularity conditions, the payoff expressions and cooperation-capacity constraint have the same form, although s^*, s^B, and s^W are generally defined implicitly. For example, if $v \sim \mathcal{N}(0, \sigma_v^2)$ and $b \sim \mathcal{N}(0, \sigma_b^2)$, then $\Pi(s) = \frac{1-\alpha}{2}\sigma_b\phi(s/\sigma_b)$, where $\phi(\cdot)$ denotes the standard-normal density; this is the normal analogue of $\frac{1-\alpha}{2}\frac{1-s^2}{2}$. The framework can also accommodate richer forms of adverse selection. In particular, instead of varying the probability α that the liquidity-seeking dealer is informed, one can hold α fixed and let adverse selection vary through signal precision: if informed dealers observe $\sigma = v + \varepsilon$, with $\varepsilon \sim \mathcal{N}(0, 1/\tau)$, then the informed dealer's expected profit is $\omega(\tau)\phi(s/\omega(\tau)) - s\bar{\Phi}(s/\omega(\tau))$, where $\bar{\Phi}(\cdot) = 1 - \Phi(\cdot)$ is the standard-normal survival function and $\omega(\tau) = \sqrt{\tau\Sigma^2/(1 + \tau\Sigma)}$ is the standard deviation of $\mathbb{E}[v \mid \sigma]$. The repeated-game and MMC results are unchanged.

is symmetric and thus her quotes are symmetric around zero: the ask (bid) is s ($-s$). The information-motivated (liquidity-motivated) dealer chooses to trade if the magnitude of the common-value component (private-value component) exceeds s. Given the symmetry, our derivations throughout focus on the case of the market-maker selling the asset at the ask. In a stage game, the market maker's ex-ante profit is

$$\pi_M(s) = (1-\alpha)\int_s^1 s dF_L(b) - \alpha\int_s^1 (v-s)dF_I(v) = (1-\alpha)(1-s)s - \alpha\frac{(1-s)^2}{2}, \quad (1)$$

which is the sum of the expected gains when trading with a liquidity-motivated dealer and losses when trading with an information-motivated dealer. The ex-ante profit of the information-motivated dealer is

$$\pi_I(s) = \int_s^1 (v-s)dF_I(v) = \frac{(1-s)^2}{2}, \quad (2)$$

which is entirely made at the expense of the market-making dealer. In contrast, the profit of the liquidity-motivated dealer is

$$\pi_L(s) = \int_s^1 (b-s)dF_L(b) = \frac{1-s^2}{2} - (1-s)s = \frac{(1-s)^2}{2}, \quad (3)$$

namely, the realized private-value component net of the payments made to the market-making dealer.

Lemma 1 (Static monopoly and break-even spreads)**.** *Let s^* maximize $\pi_M(s)$ and let s^B satisfy $\pi_M(s^B) = 0$. Then $s^* = \frac{1}{2-\alpha}$ and $s^B = \frac{\alpha}{2-\alpha}$.*

If there is only one round of trade, the market-making dealer maximizes her current profits $\pi_M(s)$ and quotes $s^* > 0$. There are two reasons for the dealer to quote a strictly positive spread. First, the dealer uses the quote to protect herself from adverse selection by the information-motivated dealer. Second, the dealer's market power allows her to extract rents by raising the quote above the level at which she breaks even, $s^* > s^B$.

Dealer Cooperation in Repeated Game. In the repeated game, the dealers play the stage game ad infinitum and have a common discount factor $\delta \in (0,1)$. The values of the

common and private components are redrawn from the respective distributions each period. The dealer roles alternate between periods. For simplicity, we assume that a given dealer seeks liquidity in stage t with probability $\frac{1}{N}$.

We focus on symmetric equilibria in which both dealers make the same cooperative quote s whenever they act as market-makers. Deviations from cooperative quotes are observable because of the bilateral nature of trade and are punished by grim-trigger strategies (Fudenberg and Maskin, 1986): if a dealer ever deviates in period t, both dealers quote s^* whenever they make the market in all subsequent periods.[4]

When both dealers expect to make and receive quote s, a dealer's expected per-period profit (before her role is known) is

$$\Pi(s) = \frac{1}{2}\pi_M(s) + \frac{1}{2}\left(\alpha\pi_I(s) + (1-\alpha)\pi_L(s)\right) = \frac{(1-\alpha)}{2}\frac{1-s^2}{2}. \tag{4}$$

Expected profit is proportional to the surplus generated by liquidity-motivated trades, $\frac{1-s^2}{2}$, which is decreasing in s. The factor $\frac{1-\alpha}{2}$ is the unconditional frequency of liquidity-motivated trades. Hence, lower spreads or a higher probability of liquidity trades increase total expected surplus. However, despite the higher continuation value under lower spreads, a dealer acting as a market maker has an incentive to deviate in the current period by quoting a higher spread. In particular, the most profitable one-shot deviation is to quote the static monopoly spread s^*, which maximizes her current-period profits.

A quote $s < s^*$ is sustainable in equilibrium if and only if the continuation value from cooperation outweighs the one-shot gain from the deviation. Formally,

$$\pi_M(s) + \frac{\delta}{1-\delta}\Pi(s) \geq \pi_M(s^*) + \frac{\delta}{1-\delta}\Pi(s^*), \tag{5}$$

[4]Our choice of punishment is not critical for the results. Reversion to the static Nash equilibrium s^* is moderate relative to alternatives in the literature. Harsher punishments such as optimal penal codes (Abreu, 1988) or complete termination of the relationship (Hendershott et al., 2020) would only expand the set of sustainable cooperative quotes, making our characterization conservative. Under weaker punishments—such as finite-period reversion or proportional responses—the IC-feasible set shrinks, but all comparative statics we document in our framework are preserved: cooperation capacity still increases in δ and decreases in α, and multimarket contact still expands the feasible set whenever bond heterogeneity generates positive slack. Weaker punishments *raise* the relative importance of MMC, as each market's standalone cooperation capacity declines, *increasing* the marginal value of imported slack.

that is, the expected payoff from quoting s must exceed the payoff from deviating to s^*. Rearranging yields the incentive-compatibility condition

$$C(s) \equiv \frac{\delta}{1-\delta}\left(\Pi(s) - \Pi(s^*)\right) - \left(\pi_M(s^*) - \pi_M(s)\right) \geq 0, \tag{6}$$

where we define $C(s)$ as the *cooperation capacity* of the market at quote s. Intuitively, $C(s)$ measures the present value from cooperation in future periods net of the opportunity cost of not deviating in the current period.

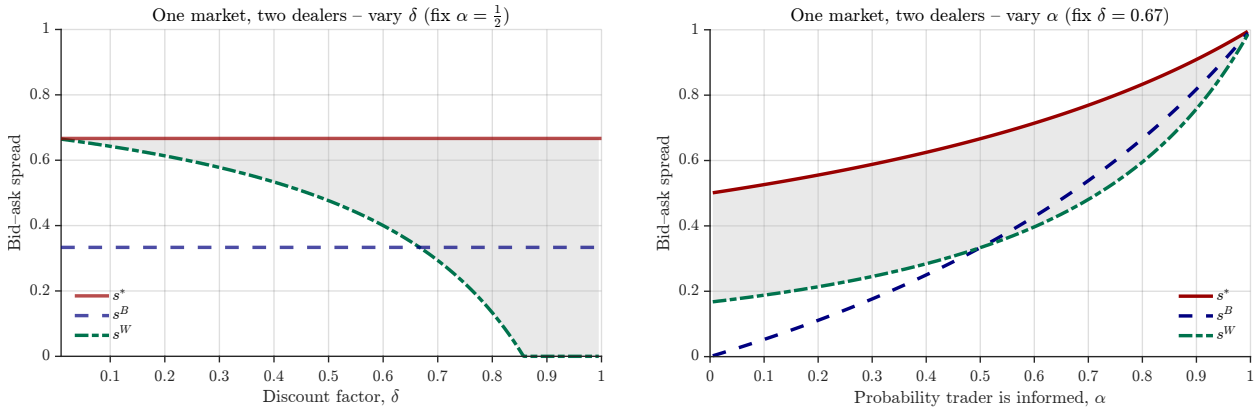
The cooperation capacity provides a natural way to characterize the set of sustainable cooperative quotes. A quote s is sustainable if $C(s) \geq 0$, that is, if the continuation value of cooperation is large enough to offset the temptation to deviate. We restrict our attention to non-negative spreads, which prevents arbitrage opportunities. The upper boundary of the set is always given by s^*. The lowest sustainable quote, which we call s^W, is the *welfare-maximizing cooperative quote*, that is, the quote that maximizes the joint surplus of the two dealers:[5]

$$s^W \;=\; \min\{\, s \in [0,1] \text{ s.t. } C(s) \geq 0 \,\}. \tag{7}$$

Proposition 1 (Cooperation capacity)**.** *The cooperation capacity $C(s) \geq 0$ for any $s \in [s^W, s^*]$, where s^W is weakly decreasing in the discount factor δ and weakly increasing in the probability of informed trading α.*

A higher discount factor δ increases the weight placed on future payoffs, thereby raising the value of continued cooperation. Since the gains from cooperation accrue over time, a larger δ amplifies the continuation value relative to the one-shot deviation gain. As a result, the incentive constraint is satisfied more easily, and a wider range of spreads can be sustained. In contrast, an increase in α reduces the continuation value of cooperation, as $\Pi(s)$ is proportional to $1-\alpha$. At the same time, higher adverse selection makes the cooperative outcome less attractive relative to the one-shot deviation s^*. Both effects tighten the incentive constraint and reduce $C(s)$, implying that only higher spreads can be sustained

[5]The welfare-maximizing cooperative quote s^W represents the lower bound of the set of sustainable spreads. We do not select s^W as the equilibrium outcome. Our main point is that it is individually optimal for each dealer to coordinate on some spread below the monopolistic level s^*: any quote $s \in [s^W, s^*]$ satisfies the incentive-compatibility constraint and is therefore sustainable. The range $[s^W, s^*]$ characterizes the set of cooperative equilibria.



(a) Varying δ for fixed $\alpha = \frac{1}{2}$.

(b) Varying α for fixed $\delta = \frac{2}{3}$.

Figure 1. The panels plot the static monopoly spread $s^*(\cdot)$, the competitive break-even spread $s^B(\cdot)$, and the welfare-maximizing cooperative spread $s^W(\cdot)$ that satisfies the incentive constraint. The shaded band shows the IC-feasible region $[\, s^W,\, s^* \,]$ where cooperative quotes are sustainable. Welfare maximization is subject to the constraint $s \geq 0$. Left: comparative statics in δ holding $\alpha = \dfrac{1}{2}$ fixed. Right: comparative statics in α holding $\delta = \dfrac{2}{3}$ fixed.

in equilibrium. Figure 1 illustrates the monopolistic spread s^*, the welfare-maximizing cooperative quote s^W, and the break-even spread s^B across different values of δ and α. The shaded area between s^W and s^* represents the set of quotes for which cooperation capacity is non-negative, $C(s) \geq 0$.

Corollary 1 (Cooperation and frequency of trade)**.** *Cooperation capacity $C(s)$ strictly increases in δ for any $s \in [0, s^*)$. Thus, it is easier to sustain cooperation for assets that trade more frequently.*

To capture differences in trading frequency in our model, we vary the length of the trading period relative to a fixed reference period. This implies that for more (less) frequently traded assets, the effective discount factor is higher (lower) than the discount factor δ associated with the reference period. For example, consider an asset that trades twice within the reference period in a high-frequency market h. For such a market, all future payoffs arrive twice as fast, implying an effective discount factor of $\delta_h = \sqrt{\delta} > \delta$ per h-period. Since $C(s)$ increases with the effective discount factor, more frequent trading raises the continuation value of cooperation and expands the set of sustainable spreads.

2.2 Two Markets and Two Dealers

We now add a second security, with a corresponding market, and consider $N = 2$ dealers that operate in both markets, $K = 2$. Each period t, the stage game is played *simultaneously* in both markets. In each market k, exactly one dealer is the market-maker and the other is the liquidity taker. The role assignments are independent across markets. Hence, in a given period a dealer may be the market-maker in neither market, in exactly one market, or in both markets.

We focus on symmetric trigger strategies across markets. Fix cooperative quotes (s^{k_1}, s^{k_2}). On the cooperative path, whenever a dealer is the market-maker in market k, she quotes s^k. A deviation in *any* market in any period triggers reversion to the static equilibrium in *both* markets from period $t+1$ onward, that is, market-makers quote $s^{*,k}$ forever in each market k. Because within-period payoffs are additively separable across markets, the best one-shot deviation is to quote the static monopoly spread in all markets where the dealer is currently a market-maker. Thus, it is sufficient to focus on the *tightest* incentive-compatibility constraint for cooperation in market k, which arises when the deviating dealer is a market-maker in *both* markets in the deviation period: if this constraint holds, cooperation is also sustainable under alternative role assignments across markets.

Consider a history on the cooperative path and a period in which a given dealer is the market-maker in market k. If she deviates in market k, she obtains the one-shot deviation payoff in market k today, and from $t+1$ onward receives the punishment continuation values in *both* markets. If she does not deviate, she retains cooperative continuation values in both markets. Using the fact that $\Pi^k(\cdot)$ is defined as the dealer's expected per-period payoff in market k, the tightest incentive-compatibility constraint for not deviating can be written as

$$\pi_M^k(s^k) + \frac{\delta_k}{1-\delta_k}\Pi^k(s^k) \;+\; C^\ell(s^\ell) \geq \pi_M^k(s^{*,k}) + \frac{\delta_k}{1-\delta_k}\Pi^k(s^{*,k}), \tag{8}$$

where $C^\ell(s^\ell)$ is cooperation capacity of the other market ℓ, defined as in (6). Defining $MMC^k \;\equiv\; C^\ell(s^\ell)$, the two-market incentive constraint becomes

$$C^k(s^k) + MMC^k \geq 0. \tag{9}$$

Proposition 2 (Cooperation under multimarket contact)**.** *If MMC^k is strictly positive, the set of sustainable quotes s^k in market k weakly expands.*

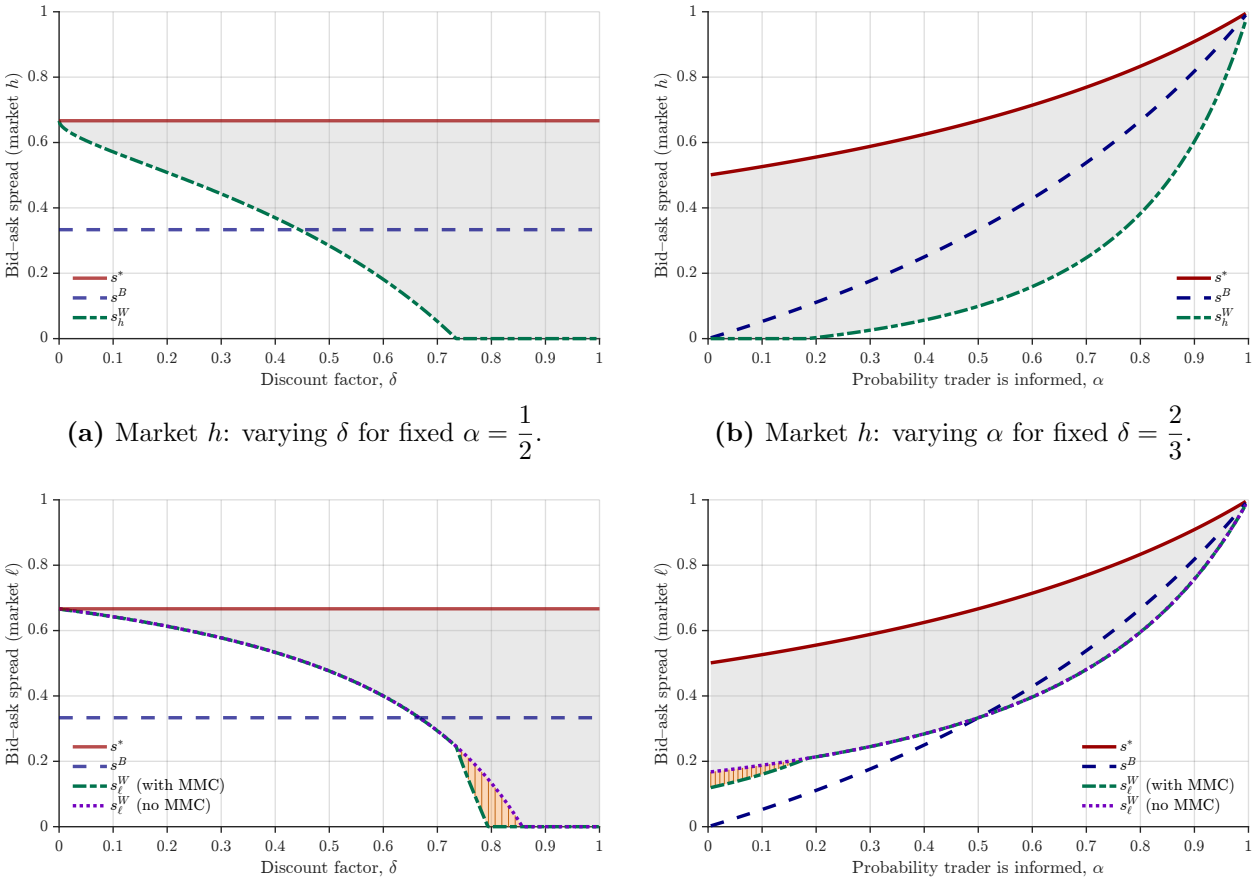
Intuitively, cooperation between dealers could improve when they operate in more than one market. In particular, if there is extra cooperation capacity in one market, measured by MMC, it can be transferred to extend cooperation in another market, that is, if $MMC^k > 0$ the constraint (9) is more slack compared to the one-market case. The extra capacity however arises only if the two markets differ as highlighted by the following corollary.

Corollary 2 (No improvement in cooperation in symmetric markets)**.** *When identical dealers trade with each other in identical markets, multimarket contact does not aid in sustaining cooperation outcomes.*

To see why, note that the incentive constraint (9) takes the form $C^k(s^k) + C^\ell(s^\ell) \geq 0$. If the two markets are identical and dealers cooperate on the same spread in each, this condition reduces to $2C^k(s^k) \geq 0$, which is the same constraint that applies when dealers operate in a single market alone. Thus, both constraints are simultaneously slack or simultaneously binding, and multimarket contact contributes nothing. Asymmetry across markets—in trading frequency, adverse selection, or both—is what allows slack cooperation capacity in one market to relax the binding constraint in another.

We illustrate Proposition 2 in the following two examples with heterogeneous markets.

Example 1 (Cooperation on the welfare-maximizing spread)**.** Suppose that the two markets differ only in the frequency of trade. The first security is traded every (reference) period, with the effective discount factor $\delta_\ell = \delta$, while the second security is traded twice as often, with the effective discount factor $\delta_h = \sqrt{\delta}$. Assume that dealers cooperate on the minimum sustainable spreads. The cooperation capacity in market h is higher than in market l for any given value of δ. By Proposition 2, if the cooperation capacity is positive at quote s^h, it expands the set of sustainable quotes s^l in market l. In this example, this happens when s^h reaches the lowest possible spread of 0. Figure 2 shows the sustainable quotes in both markets. The orange hatched regions illustrate how cooperation in market l expands when the cooperation capacity in market h is positive.



(a) Market h: varying δ for fixed $\alpha = \frac{1}{2}$.

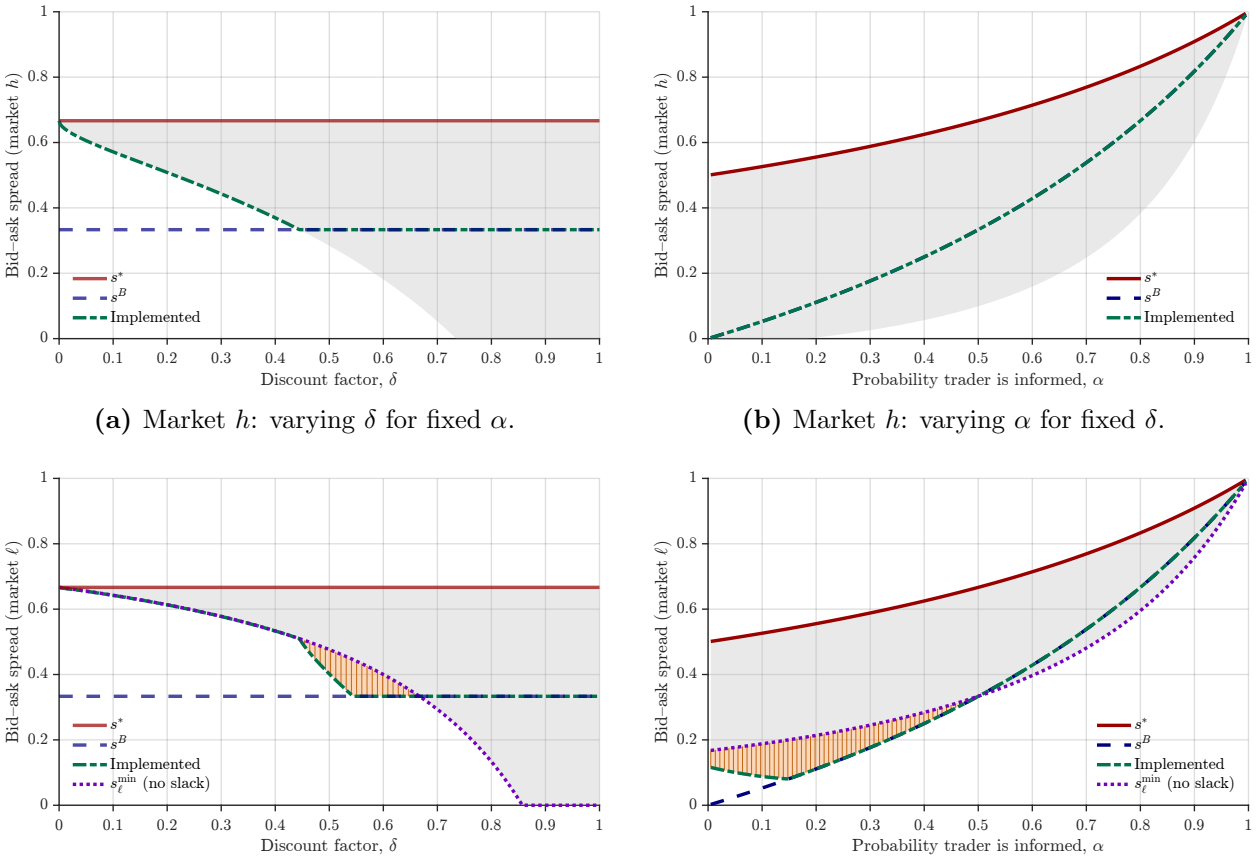
(b) Market h: varying α for fixed $\delta = \frac{2}{3}$.

(c) Market ℓ: varying δ for fixed $\alpha = \frac{1}{2}$.

(d) Market ℓ: varying α for fixed $\delta = \frac{2}{3}$.

Figure 2. Sustainable spreads in two markets (h and ℓ) under welfare-maximizing cooperation with multimarket contact. Top row: high-frequency market h; bottom row: low-frequency market ℓ with slack imported from h. Grey bands indicate the IC-feasible region in each market, while orange areas mark the additional feasible region in ℓ created by multimarket contact. Red solid lines plot the static monopoly spread s^*, blue dashed lines the competitive spread s^B, and green dash-dotted lines the welfare-maximizing cooperative spread s^W. Left column: comparative statics in δ holding $\alpha = \frac{1}{2}$ fixed. Right column: comparative statics in α holding $\delta = \frac{2}{3}$ fixed.

Example 2 (Cooperation on break-even spreads)**.** Consider the same two-market setting as in Example 1. Now suppose dealers coordinate on sustaining spreads $s \geq s^B$, where s^B is the break-even spread. In each market dealers sustain the break-even spread s^B whenever it satisfies the incentive constraint; otherwise, they cooperate on the lowest sustainable spread $s^{\min}$ specific to that market. As before, if the cooperation capacity is positive at quote s^h, it expands the set of sustainable quotes s^l in market l. In this example, this happens when s^h



(a) Market h: varying δ for fixed α.

(b) Market h: varying α for fixed δ.

(c) Market ℓ: varying δ for fixed α.

(d) Market ℓ: varying α for fixed δ.

Figure 3. Competitive-or-minimum rule with multimarket contact. Dealers quote the competitive spread s^B whenever it is incentive-compatible; otherwise they select the lowest sustainable spread (market h: $s_h^{\min}$; market ℓ: $s_\ell^{\min}$). Shaded grey regions indicate incentive-compatible spreads absent imported cooperation capacity. Orange hatched regions in the low-frequency market ℓ show the additional set of sustainable quotes created by slack imported from the high-frequency market h. Red solid lines denote the static monopoly spread s^*; blue dashed lines s^B; green dash–dot lines the implemented spread; and purple dotted lines (bottom panels) $s_\ell^{\min}$ without imported slack. The left column varies δ holding $\alpha = \frac{1}{2}$ fixed; the right column varies α holding $\delta = \frac{2}{3}$ fixed.

reaches the lowest possible spread of s^B. Figure 3 illustrates the sustainable quotes in both markets. The orange hatched regions illustrate how cooperation in market l expands when the cooperation capacity in market h is positive.

We conclude this section by characterizing how the marginal effect of multimarket contact varies with the characteristics of the market in which a dealer operates.

Proposition 3 (The marginal effect of MMC)**.** *The marginal increase in the set of sustain-*

able quotes s^k in market k due to the marginal increase in MMC^k weakly decreases with δ and weakly increases with α.

Ceteris paribus, MMC has the greatest effect when cooperative pricing is most difficult to sustain—either because the temptation to deviate is strong or because the present value of future interaction is small. When α is high, adverse selection is severe, so the incentive constraint is tight and MMC matters more. When δ is low, the effective discount factor is low, so the continuation value of cooperation is smaller and MMC again matters more.

2.3 Extensions

We develop three extensions. First, we generalize the framework to many dealers and many markets and highlight that MMC benefits all dealers involved in trade, beyond bilateral relationships. Second, we show that interdealer cooperation gains pass through to customer prices. Third, we demonstrate that dealers can cooperate along a second dimension—restraining their informed trading—which further tightens sustainable spreads. Below we summarize these extensions with the full analysis presented in Online Appendix.

2.3.1 Many Markets and Many Dealers

We generalize the model to $N > 2$ dealers in $K > 2$ markets. The full analysis is presented in Online Appendix B. Dealers may participate in different subsets of markets and face counterparty-specific adverse selection. We take the network structure and the different dealers' MMC as exogenous and focus on a single market k. As in the baseline, trading is bilateral and we assume dealers can detect deviations and condition future play on them.

Spreads are a matrix of directed quotes. Market-maker i quotes s_{ij}^k to liquidity taker j, where each liquidity taker carries a distinct informed-trading probability α_j^k. To characterize the existence and uniqueness of the solution, we consider the spread matrix that maximizes joint surplus. These spreads solve a constrained optimization in which each dealer's incentive-compatibility constraint includes cooperation capacity imported from all other markets via $MMC_i^k = \sum_{\ell \neq k} C_i^\ell$, exactly as in the two-dealer-two-market case.

Given that spreads are specific to each liquidity-taker, the deviation gain in each dealer's

IC constraint involves selecting the counterparty with respect to whom the deviation gain is maximized. To obtain a tractable formulation, we introduce an auxiliary variables d_i^k and *selector constraints* $d_i^k \geq \pi_M^k(s_j^{*,k}; \alpha_j^k) - \pi_M^k(s_{ij}^k; \alpha_j^k)$ for each $j \neq i$, with associated multipliers $\eta_{ij} \geq 0$. The multiplier η_{ij} is positive only when counterparty j attains the maximum deviation gap for maker-maker i.

Proposition 4 (Optimal counterparty-specific spreads)**.** *Under uniform trading probabilities, the welfare-maximizing spread quoted by maker-maker i to liquidity taker j is*

$$s_{ij}^k = \frac{\rho_k p(\lambda_i - \lambda_j) + \eta_{ij}}{p[(1-\alpha_j^k) + \rho_k(\lambda_i(2-\alpha_j^k) - \lambda_j)] + \eta_{ij}(2-\alpha_j^k)}, \tag{10}$$

where $\lambda_i \geq 0$ is the shadow price of dealer i's IC constraint, $\eta_{ij} \geq 0$ is a selector multiplier satisfying $\sum_{j \neq i} \eta_{ij} = \lambda_i$, and $\rho_k = \delta_k/(1-\delta_k)$.

The bilateral spread depends on the liquidity taker's adverse selection α_j^k, but also on the shadow prices of both parties' IC constraints. A dealer with ample multimarket contact has a low λ_i and quotes tight. She quotes especially tight to counterparties whose own λ_j is high—that is, to dealers who are more constrained in maintaining cooperation. The spread schedule thus redistributes surplus across dealers with heterogeneous outside options.

Proposition 5 (Within-maker-maker equalization under common adverse selection)**.** *When adverse selection is common across liquidity takers ($\alpha_j^k = \alpha^k$ for all j), each maker-maker optimally quotes the same spread to all counterparties.*

Corollary 3 (Spread ordering)**.** *At any interior solution:* (i) *s_{ij}^k is strictly increasing in α_j^k and in λ_i, and strictly decreasing in λ_j.* (ii) *Under common adverse selection, more constrained dealers quote wider.*

Example 3a (Common adverse selection)**.** Three dealers $\{A, B, C\}$ with $\alpha^k = 0.30$, $\delta = 2/3$, and $p = 1/6$. We vary MMC_A while fixing $MMC_B = 0.008$ and $MMC_C = 0$. Because adverse selection is common, Proposition 5 implies that each dealer quotes the same spread to all counterparties, collapsing the six directed spreads to three. As MMC_A rises dealer A's spread falls most steeply; dealer C's falls most slowly. Figure 4 plots the results.

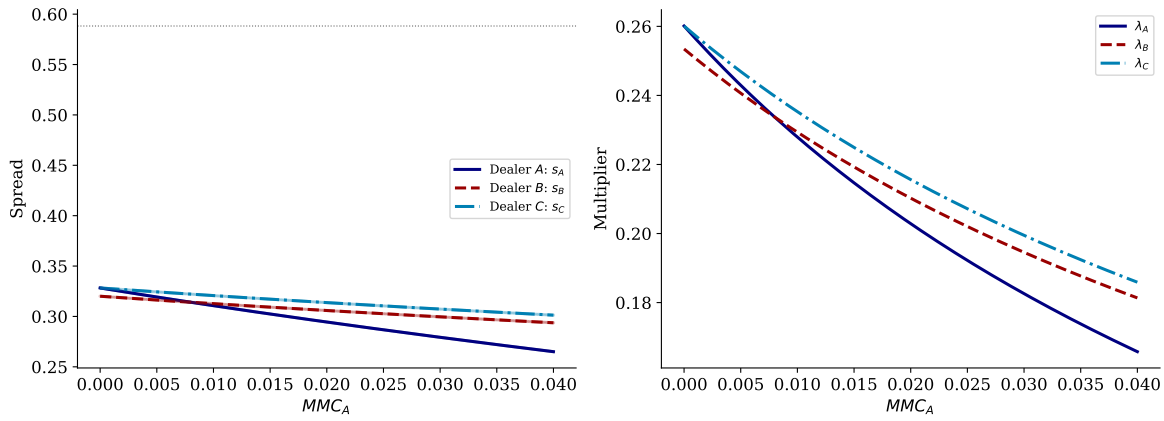


Figure 4. Common adverse selection ($\alpha = 0.30$, $\delta = 2/3$, $p = 1/6$). Left: welfare-maximizing spreads. Right: Lagrange multipliers. Parameters: $MMC_B = 0.008$, $MMC_C = 0$.

Example 3b (Counterparty-specific spreads)**.** We consider three dealers with heterogeneous adverse selection: $\alpha_A = 0.25$, $\alpha_B = 0.30$, $\alpha_C = 0.45$, so dealer C is the most informed (e.g., a hedge-fund-affiliated dealer). We vary MMC_A, fix $MMC_B = 0.008$, $MMC_C = 0$, with $\delta = 2/3$. Figure 5 plots the six optimal spreads as functions of MMC_A. Three patterns emerge.

First, each dealer quotes wider to more-informed liquidity takers: $s_{AB} < s_{AC}$, $s_{BA} < s_{BC}$, and $s_{CA} < s_{CB}$ throughout. The within-dealer differences (Figure 6, left panel) are driven entirely by the liquidity takers' adverse selection, since trading probabilities are identical across all pairs. Second, dealer C, the most constrained, quotes the widest spreads to every counterparty, reflecting the largest λ_C. Third, as MMC_A rises, A offers progressively tighter quotes, with the largest reductions accruing to the spreads C faces as a liquidity taker (s_{AC}, s_{BC})—high-MMC dealers effectively subsidize low-MMC dealers through the price schedule.

2.3.2 Pass-Through to Customers

Here, we show how the benefits of dealer cooperation on interdealer pricing pass through to end investors even under the extreme scenario of dealers behaving as monopolists with respect to customers. To allow for customer trades, as in Colliard et al. (2021), we assume that each customer is connected with a single dealer who provides an endogenous private

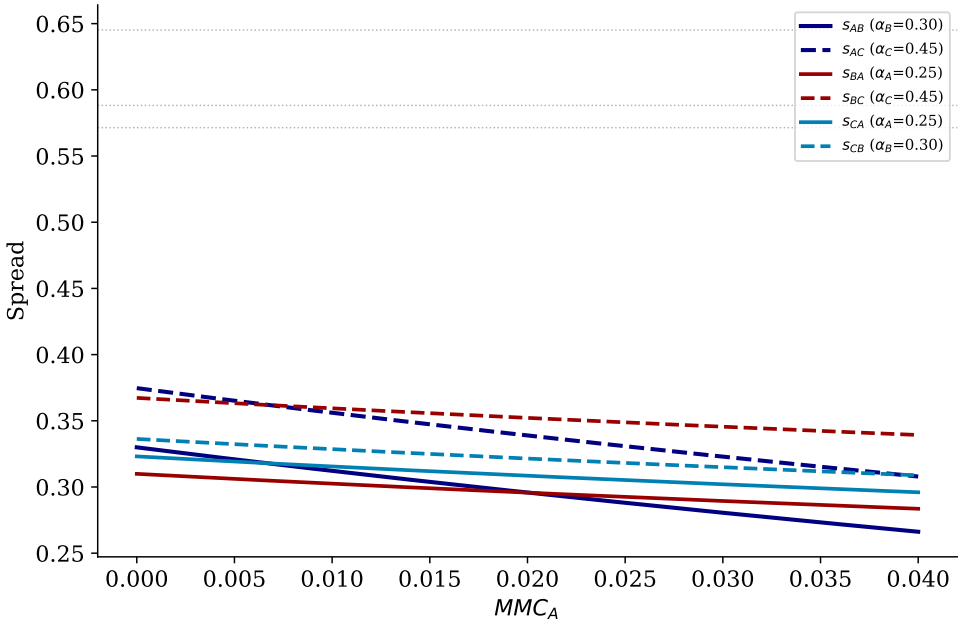


Figure 5. Counterparty-specific spreads as functions of MMC_A. $\alpha_A = 0.25$, $\alpha_B = 0.30$, $\alpha_C = 0.45$, $\delta = 2/3$, $MMC_B = 0.008$, $MMC_C = 0$, $p = 1/6$. Dotted lines mark liquidity taker-specific monopoly spreads.

quotes to the investor.[6]

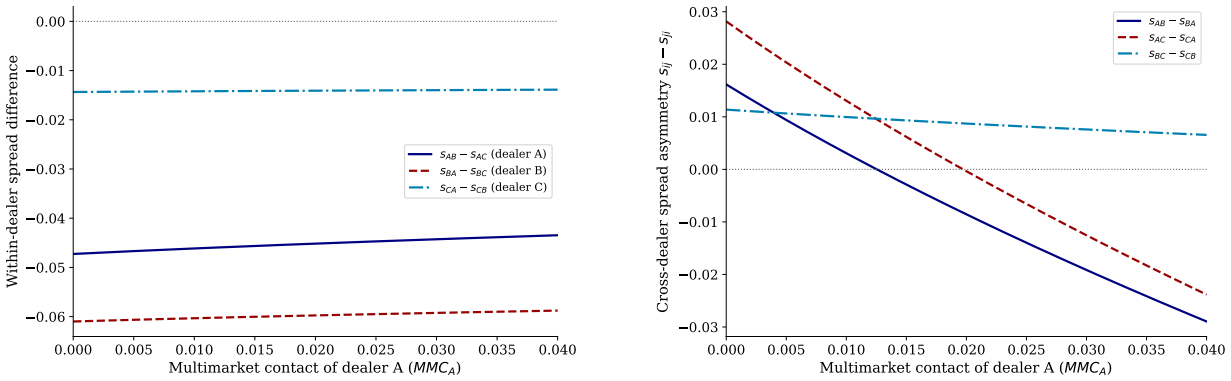
As with dealers, customer order flow can be either information-driven or liquidity-driven: with probability α_c, the customer is informed about the common-value component, while with probability $1 - \alpha_c$, the customer trades for liquidity reasons and has an idiosyncratic private value. We assume that customer trades are immediately offset in the interdealer market that follows bilateral market structure of the baseline model, so the dealer's marginal cost of serving customer demand is given by the prevailing interdealer quote s. Given this quote, the dealer optimally sets the customer quote as

$$s_c(s) = \frac{1+s}{2},$$

so customer prices inherit changes in interdealer pricing through a simple markup rule.

Adjusting for customer-driven dealer demand in the interdealer market, the characteri-

[6]The dealer thus acts as a monopolist when trading with the customer. In practice, a large fraction of investors trades with a single dealer (see, e.g., Hendershott et al., 2020). Allowing for dealer competition or customer-dealer relationships would only strengthen the pass-through to the customers.



(a) Within-dealer spread differences. **(b)** Cross-dealer asymmetry $s_{ij}^k - s_{ji}^k$.

Figure 6. Left: within-dealer differences arise from counterparty-specific adverse selection. Right: cross-dealer asymmetry arises from heterogeneous IC constraints.

zation of the interdealer quote s is otherwise analogous to the baseline environment: MMC relaxes the incentive constraints sustaining dealer cooperation and lowers interdealer quotes in equilibrium; see Online Appendix C.[7] In general, because customer quotes are increasing in s, the lower interdealer quotes induced by MMC translate directly into lower customer prices. Dealer cooperation therefore benefits not only interdealer counterparties, but also ultimate investors.

2.3.3 Cooperation on Information Exploitation

We further extend the model by allowing dealers to cooperate along the second dimension: as in Seppi (1990), dealers may agree to limit their informed trading at the expense of others. This form of informed restraint reduces adverse-selection losses incurred by dealers providing liquidity and can therefore support even tighter spreads.

To model this, we assume that, when privately informed, a dealer promises to trade only if her informational advantage exceeds an endogenous threshold $s_I \geq s$, rather than exploiting all profitable trades at the prevailing spread s. This promise is supported by a second incentive constraint added to the baseline problem. Raising s_I relaxes the market-maker's constraint by making informed trading less aggressive, but it tightens the informed

[7]The set of interdealer quotes sustainable through repeated-game incentives can be smaller than in the baseline, because the cancellation of adverse-selection terms within the dealer pair is only partial when a customer makes the trade. This can make MMC relatively more important.

dealer's constraint by increasing the gains from deviating and trading on weaker signals. This additional cooperation margin weakly enlarges the feasible set of sustainable outcomes and allows lower spreads than in the baseline model.

With multimarket contact, these effects propagate across markets. A market's cooperation capacity now reflects both spread discipline and informed restraint, so cooperation in one market can support cooperation in another through both channels. In this sense, the two dimensions of cooperation are complementary: restraining information exploitation can sustain tighter spreads locally and can also increase the cooperation capacity transmitted across markets. We present the full analysis in Online Appendix D.

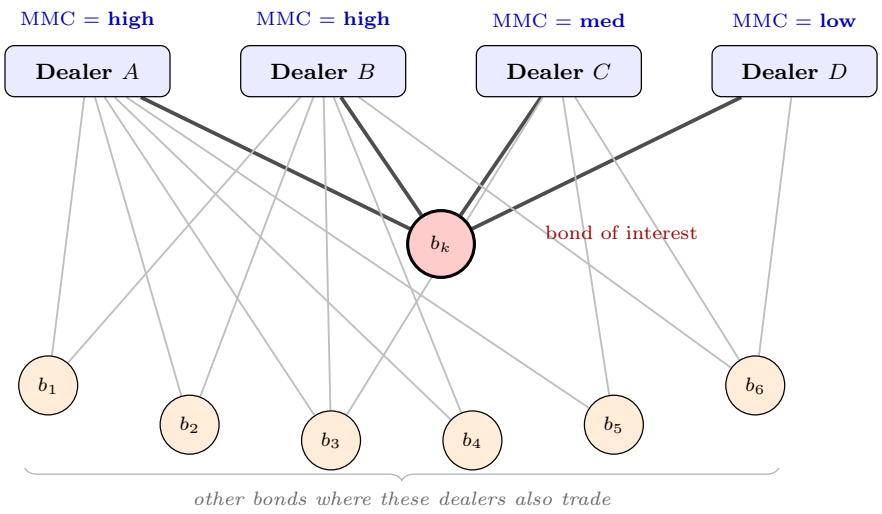
3 Data

We use transaction-level data from the regulatory version of TRACE maintained by FINRA. The dataset covers all U.S. corporate bond trades from 2003 through 2023 and reports detailed information on each transaction, including the date and time, terms, price, trade direction, and dealer identity (but not the identity of individual customers). Dealer identifiers allow us to differentiate whether a trade occurs between two dealers or between a dealer and a customer, to construct our measures of MMC, and to analyze how dealers' cross-bond overlap shapes spreads in both interdealer and dealer–customer trading. We follow standard practice in the literature to clean the data (for example, Dick-Nielsen and Rossi, 2018), removing trade duplicates, cancellations, reversals, and erroneous entries.

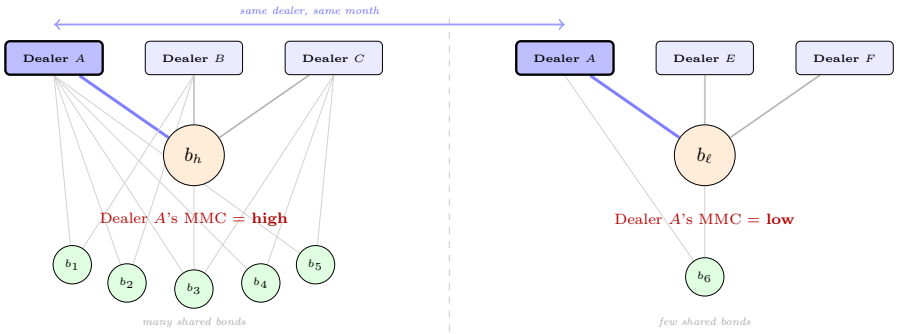
We complement the transaction data with bond characteristics from the Mergent Fixed Income Securities Database (FISD). FISD provides detailed issue- and issuer-level attributes, including coupon rates, maturity dates, credit ratings from major agencies, call and put provisions, and issuance amounts. We limit our sample to U.S. dollar-denominated, fixed-coupon or zero-coupon corporate debentures and medium-term notes issued by firms domiciled in the United States.[8] The FISD data enable us to control for key dimensions of bond heterogeneity that affect liquidity and transaction costs, such as credit quality and remaining time

[8]These bonds are identified by the FISD bond type fields CDEB, CMTN, CMTZ, CZ, and USBN. We further require bonds to be nonputtable, not privately placed, not issued under Rule 144A, not convertible or preferred, and senior unsecured.

to maturity.[9]



(a) Cross-dealer variation within a bond.



(b) Cross-bond variation within a dealer.

Figure 7. Multimarket contact: identifying variation. Panel (a): all four dealers trade bond b_k, but they differ in how many other bonds they share with each other. Dealers A and B overlap in many bonds and thus have high MMC in b_k; Dealer D overlaps in few and has low MMC. Identification exploits this within-bond–month variation across dealers. Panel (b): the same dealer (A), in the same month, faces counterparties with whom she shares many other bonds in b_h (high MMC) but few shared bonds in b_ℓ (low MMC). Identification exploits this within-dealer–month variation across bonds.

Construction of MMC Measures: Our measure of MMC captures how much a dealer trading bond k overlaps, in the cross-section of bonds, with the dealers that also trade bond

[9]To capture credit quality, we map rating notches to numeric values throughout a bond's life. When multiple rating agencies rate a bond, we assign the lower of two ratings or the median of three ratings. Bonds not rated by any agency are assigned to an unrated category.

k. The construction follows Evans and Kessides (1994), whose measure is estimated at the market–time level. The relevant unit in our setting is finer. In OTC bond markets, pricing depends not only on the bond being traded, but also on dealer-specific conditions, including inventory imbalances, client flow, balance-sheet capacity, and internal risk limits. A coarser measure of MMC would average out dealer-level variation that matters for pricing. For that reason, our baseline measure is defined at the dealer–bond–month level. We report bond–month results only as a robustness check and to assess whether bonds with greater overall MMC exhibit lower spreads.

At a given point in time, dealers differ in the extent to which they overlap with competitors across bonds. These differences generate two useful forms of variation. First, as shown in Figure 7 (a), among dealers trading the *same bond in the same month*, some dealers face higher MMC than others. Second, as shown in Figure 7 (b), for the *same dealer in the same month*, different bonds expose that dealer to different sets of overlapping competitors. Both sources of variation operate within a month, allowing us to hold time-varying bond-level conditions fixed as well as time-varying dealer-level conditions fixed.

We begin with a measure that captures the breadth of overlap between dealers in the cross-section of bonds. Formally, let there be K bonds and N dealers at month t. For each t, we construct a trailing 12-month dealer–bond activity matrix $D_t = (D_{i,k,t-12:t-1})$, where $D_{i,k,t-12:t-1}$ is an indicator equal to one if dealer i was among the top five dealers in bond k in at least one month during the previous 12 months, and zero otherwise. Using this matrix, define the $N \times N$ overlap matrix $A_t = (a_{ij,t})$ by

$$a_{ij,t} = \sum_{k=1}^{K} D_{i,k,t-12:t-1} D_{j,k,t-12:t-1}. \tag{11}$$

Thus, $a_{ij,t}$ counts the number of bonds in which dealers i and j both appear over the trailing 12-month window. Two dealers who specialize in entirely different segments of the market will have a small $a_{ij,t}$, while two dealers who consistently trade across many of the same bonds will have a large one.

Using this matrix, we define the baseline month-t MMC measure for dealer i trading

bond k as

$$AvgDealerOverlap_{i,k,t} = \frac{1}{NumDeal_{k,t} - 1} \sum_{j \neq i} a_{ij,t} \, D_{j,k,t-12:t-1}, \tag{12}$$

where $NumDeal_{k,t}$ is the number of distinct dealers that appear in bond k over the trailing 12-month window, and $D_{j,k,t-12:t-1}$ equals one if dealer j also appears in bond k at least once during that window. The measure is equal to 0 if no dealer other than i trades bond k. Therefore, MMC for dealer i in bond k is dealer i's average pairwise overlap in the universe of bonds with the other dealers trading the bond k.

The baseline MMC measure treats all dealers equally, regardless of how important is a dealer for trading a given bond. In practice, however, dealers account for a greatly different shares of intermediation activity for the same bond. A dealer who accounts for a large share of intermediation activity in a bond may have stronger incentives—and greater capacity—to lower spreads in that bond than a dealer with a trivial presence. To capture this intensity of participation, we construct a revenue-weighted variant of the MMC measure.

The starting point is a measure of each dealer's importance within a bond's revenue distribution. For dealer i and bond k, let $Rev_{i,k,t-12:t-1}$ denote the dealer's total revenue in the bond over the trailing 12 months.[10] We then define dealer i's share of total revenues in bond k over this trailing window as

$$\widetilde{R}_{i,k,t-12:t-1} \;=\; \frac{Rev_{i,k,t-12:t-1}}{\sum_{\ell=1}^{N} Rev_{\ell,k,t-12:t-1}}. \tag{13}$$

This share is bounded between zero and one and sums to one across all N dealers active in the bond. When a dealer accounts for a large fraction of total revenues in a bond, its revenue share is correspondingly high.

Using these revenue shares, we construct an $N \times N$ revenue-weighted overlap matrix

[10] Revenue is computed as trading volume multiplied by our transaction-cost measure. If the product of volume and transaction cost is negative for a given trade, we set the revenue contribution to zero.

$\widetilde{R}_t = (r_{ij,t})$ with entries

$$r_{ij,t} \;=\; \sum_{m=1}^{K} \widetilde{R}_{i,m,t-12:t-1}\, \widetilde{R}_{j,m,t-12:t-1}. \tag{14}$$

This object is the natural revenue-weighted analog of the unweighted overlap matrix described earlier. It is large when two dealers account for relatively large shares of trading activity in the same bonds. Two dealers who are both minor participants in many of the same bonds will produce a small $r_{ij,t}$, while two dealers who are both major intermediaries in the same bonds will produce a large one.

The corresponding revenue-weighted month-t MMC for dealer i trading bond k is

$$AvgDealerOverlapWeighted_{i,k,t} \;=\; \frac{1}{NumDeal_{k,t} - 1} \sum_{j \neq i} r_{ij,t}\, D_{j,k,t-12:t-1}. \tag{15}$$

This measure captures both the breadth of overlap—the number of bonds in which two dealers interact—and the extent to which the overlapping bonds are ones in which both dealers are important intermediaries. In this sense, the revenue-weighted measure provides a richer characterization of MMC than the unweighted version, because it distinguishes between settings where dealers have large stakes in the same bonds and settings where their co-presence is marginal.

To implement these measures in a tractable manner, we make two design choices. First, we focus on the cross-section of bonds that is most informative for capturing economically meaningful MMC. In each month we consider the top 2,000 traded bonds, which consistently account for more than 75 percent of market-wide trading volume. Limiting the bond universe reduces computational complexity, which is important given that the MMC calculation requires evaluating all dealer pairs across all bonds. Second, as mentioned, we use a 12-month look-back window to define dealer participation in each bond. This window smooths short-term noise in dealers' trading activity and more reliably identifies the dealers that consistently intermediate a given bond. As we show in our robustness analysis, our findings are not sensitive to shortening the look-back window to 6 or 1 months.

Figure 8 displays the monthly distribution of the MMC measures over time. The top

Table 1. Summary Statistics

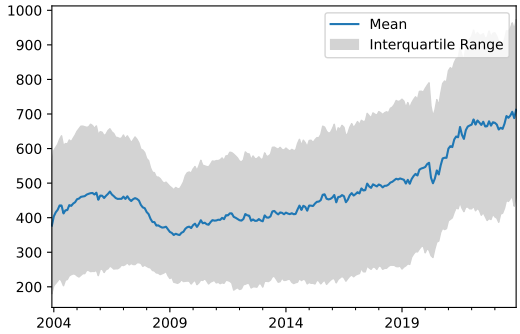
	Count	Average	Std	10pct	25pct	50pct	75pct	90pct
MMC								
Avg. Dealer Overlap	5,167,986	484.92	276.88	110.00	251.13	488.05	687.64	854.46
Avg. Dealer Overlap, Weighted	5,167,986	2.02	1.90	0.10	0.36	1.48	3.27	4.69
Spreads (bps.)								
Any–ID Cost (bps.)	1,213,660	23.23	59.85	-14.61	0.11	8.50	32.48	81.55
Any–Any Cost (bps.)	1,752,797	26.70	60.93	-8.28	0.89	10.19	35.17	85.65
CD–ID Cost (bps.)	830,982	33.63	65.90	-12.11	1.20	15.17	47.53	108.68
ID–CD Cost (bps.)	716,924	25.65	67.86	-24.29	-2.33	9.80	40.00	102.59

Notes: This table reports summary statistics for the MMC and transaction cost measures used in the analysis. *Avg. Dealer Overlap* and *Avg. Dealer Overlap, Weighted* correspond to the baseline and revenue-weighted MMC measures constructed at the dealer–bond–month level. Transaction costs are measured in basis points using the spread definition in Equation (16). The *Any–Any* and *Any–ID* measures compare trade prices, respectively, to benchmarks constructed from all counterparties or only interdealer trades. The *CD–ID* and *ID–CD* measures compare customer–dealer (interdealer) trades to interdealer (customer–dealer) benchmarks and are used in robustness analyses. All statistics are computed across the full sample period.

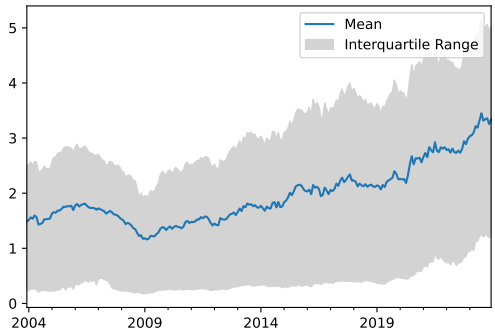
panel shows the baseline measure, $AvgDealerOverlap$, which exhibits a slight upward trend in the average number of bonds that two dealers trade in common. The figure also reveals substantial cross-sectional heterogeneity, which our empirical strategy exploits. The bottom panel displays $AvgDealerOverlapWeighted$, which exhibits a similar trend.

Table 1 reports summary statistics. The baseline measure has an average and median of about 485 and 488, respectively. Because this measure is expressed in units of bonds, these magnitudes imply that dealers associated with the same bond typically also overlap with one another in several hundred other bonds. The revenue-weighted measure has an average of 2 and a median of 1.5. While this measure is not naturally interpreted as a bond count, higher values indicate that the overlap is concentrated in bonds in which the relevant dealers account for larger shares of dealer revenue. Taken together with the observations from Figure 8, these statistics indicate that MMC among dealers is highly heterogeneous across dealers and bonds, and has been increasing over time.

Construction of Transaction Cost Measures: We measure transaction costs using a realized markup relative to a dealer-specific inventory-sourcing benchmark. The measure is defined at the dealer–bond–trade level. Building on Hendershott and Madhavan (2015), for transaction r in bond k executed by dealer i at time τ, we construct our transaction cost



(a) Baseline average pairwise dealer–bond overlap.



(b) Revenue-weighted average dealer overlap.

Figure 8. The panels plot the monthly average and interquartile range of the MMC measures over time. Panel (a) reports the baseline MMC measure, *AvgDealerOverlap*, which captures the average number of bonds that a dealer trades in common with other dealers active in a given bond–month. Panel (b) reports the revenue-weighted MMC measure, *AvgDealerOverlap, Weighted*, which places greater weight on dealer pairs that jointly account for larger shares of dealer revenue within the same bonds. Both measures display a modest upward trend over the sample period, together with substantial cross-sectional dispersion.

measure as follows:

$$Spread_{i,k,r,\tau} = 10{,}000 \times Q \times \ln\left(\frac{TradePrice_{i,k,r,\tau}}{BenchmarkPrice_{i,k,\tau}}\right), \tag{16}$$

where $TradePrice$ and $BenchmarkPrice$ are each computed from either customer–dealer, interdealer, or all transactions. We set $Q = 1$ when the $TradePrice$ reflects a dealer sale and $Q = -1$ for a dealer purchase. This normalization makes spreads directly interpretable as dealer i's effective markup, regardless of whether the counterparty is a customer or another dealer.[11]

For each transaction, the benchmark price is constructed using dealer i's own past transactions in the same bond on the opposite side of the market within a 24-hour look-back window—that is, benchmark trades must occur before the relevant transaction and no more than 24 hours earlier. When multiple benchmark trades are available, we use a volume-weighted average within the window. This construction delivers a dealer-specific proxy

[11]The measure is intended to capture the full execution markup faced by the dealer relative to its own benchmark price: not only explicit charges such as commissions or fees, but also implicit trading costs such as the bid–ask spread and price impact.

for the dealer's contemporaneous marginal valuation (or effective cost of sourcing inventory), smooths idiosyncratic noise in any single trade, and—because it relies only on past transactions—avoids look-ahead bias. A potential concern is that a 24-hour window is more likely to yield a benchmark for relatively liquid bonds, so our baseline sample may tilt toward more liquid securities. In robustness tests, we extend the benchmark window to five business days and show that results are robust to the choice of look-back window.

Beyond the look-back window, a key measurement choice concerns which trades enter the benchmark price. The transaction-cost measure that maps most directly to the model is what we label the *Any-to-ID (any-to-interdealer)* spread. This measure compares any transaction price—whether the counterparty is another dealer or a non-dealer customer—to dealer i's recent *interdealer* benchmark price in the same bond. The mapping is natural given that the model centers on interdealer trading relationships and predicts that MMC compresses interdealer spreads and customer spreads.

In the Appendix, we consider three alternative spread measures to verify that our results do not depend on a particular benchmark design. First, we construct an *Any-to-Any* spread that benchmarks each transaction against recent prices from *all* counterparties, including both dealers and customers, thereby capturing a broader notion of execution costs. This broader benchmark is motivated by recent evidence that non-dealer institutions have become an increasingly important source of liquidity in corporate bond markets (O'Hara and Zhou, 2025; O'Hara et al., 2025). Second, we construct a *customer-dealer-to-interdealer* spread that benchmarks customer–dealer transaction prices against dealers' recent interdealer benchmark prices only. Third, we construct an *interdealer-to-customer-dealer* spread that reverses the trade–benchmark pairing, to verify that our results do not depend on which side of the market provides the reference price.

Our transaction-cost measures are at the trade (date–timestamp) level. To match the monthly frequency of the MMC variables, we aggregate these high-frequency measures to the bond–dealer–month level, computing both volume-weighted and equally-weighted monthly averages. Figure 9 shows the resulting monthly distributions of the volume-weighted *any-to-interdealer* measure. The series is strongly countercyclical, with pronounced spikes during the Global Financial Crisis and the COVID-19 episode. Across bonds, dealers, and time, the

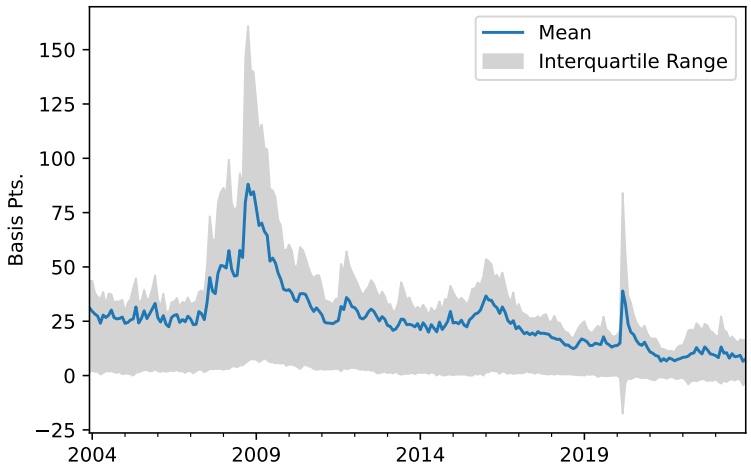


Figure 9. This figure plots the monthly distribution of the baseline transaction-cost measure. The *Any-ID* spread compares each trade price to the dealer's recent interdealer benchmark. The series exhibits strong countercyclicality, with pronounced increases during the Global Financial Crisis and the COVID-19 crisis. Summary statistics appear in Table 1.

average spread is 23 basis points. For the three directional variants used in robustness, average *any-to-any*, *customer-dealer-to-interdealer*, and *interdealer-to-customer-dealer* spreads are 27, 34, and 26 basis points, respectively. Summary statistics for all four measures are reported in Table 1.

4 Empirical Results

To relate MMC to transaction costs, we employ a fixed-effects design that exploits cross-sectional variation within months. For each dealer–bond–month triple (i, k, t), we estimate

$$Spread_{i,k,t} = \mu_{k,t} + \mu_{i,t} + \beta\, MMC_{i,k,t} + \gamma\, X_{i,k,t} + \varepsilon_{i,k,t}, \tag{17}$$

where $MMC_{i,k,t}$ is either $AvgDealerOverlap$ or $AvgDealerOverlapWeighted$, standardized to have average zero and unit standard deviation within each month, and $X_{i,k,t}$ is a vector of additional controls. The parameter of interest is β, the response of dealer-level transaction costs to variation in MMC.

The bond–month fixed effects $\mu_{k,t}$ absorb all time-varying bond conditions—credit fun-

damentals, interest-rate exposure, and shifts in demand—common to all dealers quoting that bond in month t. The dealer–month fixed effects $\mu_{i,t}$ absorb time-varying dealer conditions—balance-sheet capacity, risk appetite, and overall market activity—common to all bonds handled by dealer i in month t. With these fixed effects, β is estimated from within-month variation: comparisons across dealers within a bond–month, and comparisons across bonds within a dealer–month. This design isolates the pricing differences associated with cross-market overlap, holding fixed bond-level and dealer-level conditions at the monthly frequency.

Following Di Maggio et al. (2017), we show that our results are robust to including additional controls. The first is the log of lagged dealer–bond trading volume, measured separately for each dealer in each bond. Trading activity is known to be strongly negatively related to transaction costs: more active dealer–bond pairs tend to exhibit tighter spreads as search frictions are lower and inventory risks are more easily offset (Bessembinder et al., 2006; Edwards et al., 2007). Lagging volume restricts the control to information available prior to the pricing decision in month t, providing a cleaner comparison across dealers and bonds within a month. Including this measure also allows us to interpret the coefficient on MMC as the effect above and beyond trading volume, which is useful as the two may be positively correlated.

We also control for the number of comparable bonds to capture a dealer's scope for exerting market power across similar securities within a given month. The variable is defined as the number of bonds that dealer i trades in month t that fall into the same rating–maturity–sector group as bond k.[12] A dealer active in many comparable bonds may have greater flexibility to adjust pricing across alternative securities within the same category, potentially attenuating the disciplining effect of MMC in any given bond.

Table 2 reports the baseline estimates. Across all specifications, we find a negative and statistically significant relationship between MMC and transaction costs. The result holds for both $AvgDealerOverlap$ and $AvgDealerOverlapWeighted$. The estimates remain stable when we add controls for lagged dealer–bond trading volume and the number of comparable bonds. The magnitudes are economically large. A one-standard-deviation increase in MMC

[12]We form nine rating buckets (AAA through C and a group for unrated bonds), four maturity buckets (1–3 years, 4–7 years, 8–15 years, and above 15 years remaining time to maturity), and three sector groups (industrial, financial, and utility).

Table 2. Bond–Dealer Transaction Costs and MMC

	(1)	(2)	(3)	(4)
Avg. Dealer Overlap	-5.90***	-5.31***		
	(2.05)	(1.95)		
Avg. Dealer Overlap, Weighted			-4.54***	-3.94***
			(1.51)	(1.45)
Bond-Time FE	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓
Additional Controls		✓		✓
Obs	1,121,663	1,120,822	1,121,663	1,120,822
$Adj.R^2$	0.284	0.285	0.284	0.285

Notes: This table reports baseline estimates of the relationship between MMC and bond–dealer transaction costs. The dependent variable is the *Any-ID* spread, which compares each trade price to the dealer's recent interdealer benchmark. MMC is measured using either the unweighted average dealer overlap (*Avg. Dealer Overlap*) or the revenue-weighted measure (*Avg. Dealer Overlap, Weighted*), standardized each month to have average zero and variance one. Specifications include bond–month and dealer–month fixed effects. Additional controls include lagged dealer–bond trading volume and the number of comparable bonds traded each month. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

is associated with a reduction in transaction costs of roughly 4 to 6 basis points, depending on the MMC measure. Relative to an average spread of about 23 basis points, these estimates imply a decline of roughly 17 to 26 percent. Dealers with greater cross-bond overlap charge systematically tighter spreads, consistent with the theoretical prediction that MMC constrains dealers' ability to extract rents.

We now verify that the baseline relationship between MMC and spreads is robust along several dimensions. The robustness exercises fall into three groups: alternative specifications, alternative constructions of the MMC and transaction cost measures, and subsample analyses. We summarize the key findings here; full results appear in the Online Appendix.

Online Appendix Table E.1 reports two specification checks. In columns (1) and (2), replacing our baseline two-way clustering with triple clustering by dealer, bond, and month reduces the standard errors, leaving the statistical significance of the MMC coefficients unchanged. In columns (3) and (4), we add dealer–bond fixed effects, which absorb all time-invariant heterogeneity in the dealer–bond match and thus provide a more demanding specification. The estimated effects are smaller in this case—falling from roughly 5 basis points to about 3 basis points per one-standard-deviation increase in MMC—but remain statistically and economically significant.

We also verify that the baseline relationship is robust across subsamples. Online Ap-

pendix Table E.2 reports estimates separately for investment-grade and high-yield bonds. In both segments, MMC enters with a negative and statistically significant coefficient. The effect is larger for high-yield bonds—roughly 5 basis points per standard deviation of MMC, compared with 3 basis points for investment-grade bonds. This pattern is consistent with the model's prediction that MMC has the largest marginal effect in bonds that exhibit greater information asymmetry and trade less frequently (Proposition 3).

Online Appendix Table E.2 also splits the sample into bonds with below-five-year and above-five-year remaining maturities. The MMC coefficient is negative and economically significant in both groups. A one-standard-deviation increase in MMC reduces transaction costs by roughly 6–7 basis points for longer-maturity bonds and by 3–4 basis points for shorter-maturity bonds. Online Appendix Table E.3 extends these results across issuer sectors and shows similarly stable estimates. Taken together, the subsample analyses show that the relationship between transaction costs and MMC is present across diverse segments of the corporate bond market and displays a stable economic magnitude across investment-grade versus high-yield bonds, short- versus long-maturity bonds, and different issuer sectors.

We next examine whether our findings are sensitive to alternative constructions of the MMC or transaction cost measures. Online Appendix Table E.4 columns (1) and (2) show that the estimates remain robust when the MMC look-back window is shortened from twelve months to six or even one month, allowing MMC to vary over shorter horizons and capture more transitory patterns in dealer participation. In column (3), we consider an alternative MMC measure that scales dealer overlap by the number of bonds a dealer trades, rather than by the number of dealers active in each bond, thereby incorporating heterogeneity in dealer scope. The resulting estimates are essentially unchanged.

Online Appendix Table E.4 columns (4) to (7) further show that our results are robust to alternative constructions of the benchmark price and the monthly aggregation of transaction costs. Our baseline benchmark uses dealer i's own opposite-side trades within the prior 24 hours. Because bonds that trade frequently are more likely to generate such benchmarks, the baseline sample may tilt toward more liquid securities. To assess sensitivity of the results, we relax the timing restriction and construct an alternative benchmark based on the dealer's average purchase price over the prior five business days. The resulting estimates are nearly

identical to the baseline. The results are also insensitive to how spreads are aggregated within a month: recomputing dealer–bond–month spreads using equally-weighted averages (rather than volume-weighted averages) yields MMC coefficients that are essentially unchanged.

Table F.1 shows that the relationship is not driven by which trades enter the spread construction. We begin with an *Any–Any* spread that benchmarks each trade to recent prices from *all* counterparties (dealers and customers). We then turn to two benchmark-swapping measures: *CD–ID*, which benchmarks customer–dealer trades to recent interdealer prices, and *ID–CD*, which benchmarks interdealer trades to recent customer–dealer prices. Across all three constructions, the MMC coefficient is negative, if anything larger in magnitude, and statistically significant, indicating that higher MMC is associated with lower spreads regardless of whether benchmarks are drawn from dealer trades, customer trades, or the pooled trading environment.

We extend all the robustness and subsample exercises mentioned above to these three spread constructions. Online Appendix Tables F.2–F.6 show that the magnitude, sign, and significance of the MMC coefficients remain stable across alternative benchmark definitions. These results indicate that the relationships identified so far are not sensitive to the construction of either the MMC measure or the transaction-cost measure and emerges consistently across a wide range of specifications.

Finally, we examine whether the baseline relationship between MMC and transaction costs remains once we account for bilateral trading relationships, which Di Maggio et al. (2017) show are an important determinant of pricing in OTC bond markets. To do so, we aggregate the data to the dealer-pair–bond–month level and regress dealer-to-dealer transaction costs on our MMC measures, including bond-by-month and dealer-pair-by-month fixed effects. This specification absorbs all time-varying conditions common to a given bond in a given month, as well as all time-varying conditions specific to the relationship between the two dealers. Identification therefore comes from comparing, within a bond and month, dealer pairs that differ in the extent of their overlap across other bonds. Table E.5 shows that dealer pairs with greater MMC transact at significantly lower costs. The result holds for both the unweighted and revenue-weighted MMC measures.

We now turn to the cross-sectional predictions of the model. The theory delivers two

testable implications about when MMC should matter most for pricing. The first concerns the severity of adverse selection. The second concerns the frequency of trading. We examine each in turn using the subsample splits reported in Table 3.

To test the model's prediction about adverse selection, we need a proxy for the information content of trading activity in each bond. The natural candidate in high-frequency microstructure data is Kyle (1985)'s λ, estimated from a regression of price changes on signed order flow. In corporate bond markets, however, reliably signing interdealer trades as buyer- or seller-initiated is difficult. As Hasbrouck (2007) and Foucault et al. (2013) emphasize, when signed order flow is unavailable or unreliable, one can instead exploit the relationship between unsigned returns and unsigned volume.[13]

We estimate the following specification for each bond k in each month:

$$|\Delta p_{k,\tau}| = \alpha_k + \lambda_k |Vol_{k,\tau}| + \eta_{k,\tau}, \tag{18}$$

where τ indexes trades (or intraday intervals), $|\Delta p_{k,\tau}|$ is the absolute value of the price change, and $|Vol_{k,\tau}|$ is the total trading volume. A large $\hat{\lambda}_k$ indicates that trading activity moves prices more in bond k, consistent with a greater likelihood that order flow reflects private information rather than liquidity. We then classify bonds based on whether their estimate in a given month is above or below the cross-sectional median of $\hat{\lambda}_k$ in that month.

For above-median $\hat{\lambda}$ bonds, a one-standard-deviation increase in MMC lowers transaction costs by 5 to 6 basis points, compared with 2 to 3 basis points for below-median $\hat{\lambda}$ bonds. The difference is notable. For instance, when we use the revenue-weighted measure, the difference between the two subsamples is statistically significant at the 1 percent level, confirming that the disciplining effect of MMC is stronger in bonds where adverse selection is more severe.

We also split the sample by trading frequency, defining bonds with above-median trading frequency per month as frequently traded. For these bonds, a one-standard-deviation increase in MMC is associated with a reduction in transaction costs by 4 to 5 basis points. For less active bonds, the corresponding effect is 5 to 7 basis points. Using the revenue-weighted

[13] Foucault et al. (2013) note that while trading volume and order imbalance are distinct concepts, they are likely to be correlated so that the slope from a regression of absolute price changes on trading volume can be interpreted as a measure of the price change associated with one additional unit of trading activity. Conceptually this measure is related to Amihud (2002) illiquidity ratio.

Table 3. Bond–Dealer Transaction Costs and MMC: Heterogeneity by Adverse Selection and Trading Frequency

	(1)	(2)	(3)	(4)
	Kyle's λ		Trade Frequency	
Avg. Dealer Overlap	-3.48**		-7.46***	
	(1.58)		(1.79)	
Avg. Dealer Overlap, Weighted		-1.81		-5.02***
		(1.17)		(1.29)
Above-med. λ × Avg. Dealer Overlap	-2.37			
	(2.08)			
Above-med. λ × Avg. Dealer Overlap, Weighted		-3.01***		
		(1.07)		
Above-med. freq. × Avg. Dealer Overlap			2.46	
			(1.55)	
Above-med. freq. × Avg. Dealer Overlap, Weighted				1.30*
				(0.77)
Bond-Time FE Interacted	✓	✓	✓	✓
Dealer-Time FE Interacted	✓	✓	✓	✓
Additional Controls Interacted	✓	✓	✓	✓
Obs	1,104,192	1,104,192	1,114,619	1,114,619
$Adj.R^2$	0.296	0.296	0.288	0.288

Notes: This table reports subsample estimates of the MMC–transaction-cost relationship. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have average zero and variance one. Columns (1)–(2) split bonds by below- versus above-median $\hat{\lambda}$, an unsigned price impact proxy estimated from the regression $|\Delta p_{k,\tau}| = \alpha_k + \hat{\lambda}_k |Vol_{k,\tau}| + \eta_{k,\tau}$, which serves as a measure of the severity of adverse selection (Hasbrouck, 2007; Foucault et al., 2013). Columns (3)–(4) split bonds by below- versus above-median monthly trading frequency. All specifications include bond–time and dealer–time fixed effects and the baseline controls. Standard errors are two-way clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

measure, the difference between the two subsamples is statistically significant at the 10 percent level,[14] providing evidence that the incremental discipline of MMC is correspondingly less valuable in more frequently traded bonds.

5 Shift–Share Instrument

We complement our fixed-effects design with a shift–share instrument (Bartik, 1991; Blanchard and Katz, 1992). The design exploits the sharp and uneven surge in corporate bond issuance after the COVID-19 shock in March 2020. Issuance rose especially strongly in investment-grade segments, while high-yield issuance did not increase in parallel (Becker and Benmelech, 2025). When issuance rises in a given segment, dealers can trade across a broader set of securities in that segment. This mechanically increases the scope for overlap in dealers' bond portfolios and, hence, MMC. Dealers with larger pre-2020 footprints in the segments hit by these issuance waves should therefore experience larger increases in MMC.

We define bond-market segments by the interaction of rating class (investment-grade or high-yield) and issuer's industry (Mergent FISD's two-digit industry code). Let g index segments and let $Iss_{g,t}$ denote the number of bonds issued in segment g in month t. For each month in the nine-month period after March 2020, we measure the segment-level shock as abnormal issuance relative to the same calendar month over the prior ten years:

$$AbIss_{g,t} = Iss_{g,t} - \frac{1}{10}\sum_{\ell=1}^{10} Iss_{g,t-12\ell}. \tag{19}$$

The cross-sectional component of the instrument comes from predetermined differences in dealers' historical market shares across segments estimated using the sample from 2010 to 2019. Dealer activity is persistent across segments: dealers are active in an average of 8.50 industry-rating segments per quarter, remain active within a segment for an average of 25 (out of 40 possible) quarters, and their market shares exhibit an autocorrelation coefficient of 0.80.[15] This persistence makes pre-COVID segment market shares a natural measure of dealers' exposure to segment-level issuance shocks during the COVID episode.

[14] Similarly, the interaction using the non-revenue-weighted measure has a p-value of 0.112.

[15] Table E.6 reports additional evidence on the persistence of dealer market shares within segments.

We compute dealers' exposure to issuance in a given segment in two separate ways. Our baseline measure is dealer i's pre-period market share in segment g over the 2010–2019 period:

$$Share_{i,g}^{(1)} = \frac{\sum_{\tau \in \mathcal{T}^{pre}} Vol_{i,g,\tau}}{\sum_j \sum_{\tau \in \mathcal{T}^{pre}} Vol_{j,g,\tau}}, \qquad \mathcal{T}^{pre} = \{2010, \ldots, 2019\}, \tag{20}$$

where $Vol_{i,g,\tau}$ is dealer i's trading volume in segment g in month τ. This measure captures dealer i's historical footprint in segment g relative to other dealers active in that segment. As a robustness, we construct a month-of-year version that compares only observations from the same calendar month:

$$Share_{i,g,m(t)}^{(2)} = \frac{\sum_{\tau \in \mathcal{T}^{pre}:\, m(\tau)=m(t)} Vol_{i,g,\tau}}{\sum_j \sum_{\tau \in \mathcal{T}^{pre}:\, m(\tau)=m(t)} Vol_{j,g,\tau}}, \qquad \mathcal{T}^{pre} = \{2010, \ldots, 2019\}, \tag{21}$$

where $m(t) \in \{1, \ldots, 12\}$ denotes the calendar month. This alternative measure removes potential seasonality in segment-level trading activity.

For dealer–bond–month observation (i, k, t), let $g(k)$ denote the segment of bond k. The shift–share instrument is

$$Z_{i,k,t}^{(s)} = Share_{i,g(k)}^{(s)} \times AbIss_{g(k),t}, \qquad s \in \{1, 2\}. \tag{22}$$

Thus, $Z_{i,k,t}^{(s)}$ is high when bond k belongs to a segment experiencing unusually strong issuance in month t and dealer i historically accounted for a larger share of intermediation in that segment.

We estimate two-stage least squares on the sample of bonds issued before 2020. Restricting the sample in this way ensures that the bonds whose transaction costs we measure are not themselves created by the COVID-era issuance wave. Given the instrument is constructed using issuance counts, the MMC measure that maps most directly is $AvgDealerOverlap_{i,k,t}$. Therefore, the first stage is

$$AvgDealerOverlap_{i,k,t} = \mu_{k,t} + \mu_{i,t} + \pi Z_{i,k,t} + \Gamma' X_{i,k,t} + u_{i,k,t}, \tag{23}$$

where $\mu_{k,t}$ are bond–month fixed effects, $\mu_{i,t}$ are dealer–month fixed effects, and $X_{i,k,t}$ is the

baseline set of controls. The second stage is

$$Spread_{i,k,t} = \mu_{k,t} + \mu_{i,t} + \beta \widehat{AvgDealerOverlap}_{i,k,t} + \Gamma' X_{i,k,t} + \varepsilon_{i,k,t}. \tag{24}$$

The coefficient of interest is β, the effect of exogenous variation in MMC on dealer-level transaction costs. Bond–month fixed effects absorb all time-varying bond-specific conditions, including changes in liquidity and fundamentals. Dealer–month fixed effects absorb time-varying dealer-specific conditions, such as funding capacity, balance-sheet pressure, and risk appetite. With bond–month and dealer–month fixed effects, identification comes from residual variation in exposure to issuance shocks across dealer–bond pairs after removing bond–month and dealer–month components.

As shown by Goldsmith-Pinkham, Sorkin, and Swift (2020), the exclusion restriction requires that conditional on the fixed effects and controls, pre-period dealer shares are uncorrelated with innovations in transaction costs:

$$E[\varepsilon_{i,k,t} \mid Share^{(s)}_{i,g(k)}, \mu_{k,t}, \mu_{i,t}, X_{i,k,t}] = 0. \tag{25}$$

This condition concerns *changes* in spreads, not levels: the identifying assumption is only that dealers' shares do not forecast spread innovations through channels other than MMC.

Several features of the design support the choice of this instrument. First, the exposure shares are measured over 2010–2019 and fixed before the COVID issuance wave, so they reflect persistent dealer specialization rather than trading responses to COVID. Second, because we exclude bonds issued in 2020 or later, the bonds in the estimation sample are not mechanically linked to the COVID-era issuance surge. Third, conditional on our choice of high-dimensional fixed effects, violating the exclusion restriction would require an omitted factor that varies across dealers within a bond–month and across bonds within a dealer–month, is correlated with decade-old specialization patterns interacted with segment-level issuance shocks, and affects spreads through a channel other than MMC.[16]

Table 4 reports the IV estimates. The first-stage relationship, reported in Panel B, is

[16]The design also fits the exogenous-shocks logic of Borusyak, Hull, and Jaravel (2021), since the abnormal issuance shocks are driven by issuers' financing decisions rather than by dealer behavior.

Table 4. Bond–Dealer Transaction Costs and Multimarket Contact — IV Regressions

	I. 2SLS second stage estimates Any-ID Cost	
	(1)	(2)
Instrumented Avg. Dealer Overlap	-0.785**	-0.828**
	(0.34)	(0.37)
Bond-Time FE	✓	✓
Dealer-Time FE	✓	✓
Additional Controls	✓	✓
	II. 2SLS first stage estimates MMC Basic Score	
	(1)	(2)
$Share^{(1)}_{i,g(k),m(t)} \times AbIss_{g(k),t}$	4.38***	
	(1.52)	
$Share^{(2)}_{i,g(k),m(t)} \times AbIss_{g(k),t}$		4.35***
		(1.62)
Bond-Time FE	✓	✓
Dealer-Time FE	✓	✓
Additional Controls	✓	✓
Obs	33,656	33,656
$Adj.R^2$	0.983	0.983

Notes: This table reports 2SLS estimates that instrument dealer–bond–month MMC with a shift–share exposure to abnormal COVID-era issuance waves. The shift is segment-month abnormal issuance, $AbIss_{g,t}$, defined as the number of bonds issued in segment g, rating × industry, at time t minus the average issuance in the same calendar month over the prior ten years. The shares are dealers' pre-period, 2010–2019, activity weights in each segment: $Share^{(1)}_{i,g}$ uses overall pre-period activity, while $Share^{(2)}_{i,g,m(t)}$ conditions on calendar month to purge seasonality. The instrument for observation (i,k,t) is $Z^{(s)}_{i,k,t} = Share^{(s)} \times AbIss_{g(k),t}$. Panel A reports second-stage estimates for $Spread^{Any\text{-}ID}_{i,k,t}$ using the fitted $AvgDealerOverlap$; Panel B reports first-stage estimates for $AvgDealerOverlap$. All specifications include bond–month and dealer–month fixed effects and the baseline controls. The sample is restricted to bonds issued before 2020. Standard errors are two-way clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

strong. Both the Kleibergen and Paap (2006) and Cragg and Donald (1993) F-statistics yield p-values well below the 1 percent significance level. Issuance exposure increases MMC with a coefficient of approximately 4.3 under either share definition, confirming that issuance waves raise MMC for dealers with larger pre-period footprints in the affected segments. Panel A reports the second-stage estimates. The 2SLS coefficients on instrumented $AvgDealerOverlap$ are approximately -0.8, confirming the central prediction: exogenous increases in MMC reduce transaction costs. The negative sign is robust across both instrument definitions and is statistically significant. The IV estimates are smaller than the corresponding OLS co-

efficients. One interpretation is that 2SLS identifies a local average treatment effect for dealer–bond pairs whose MMC responds to COVID-era issuance shocks. These compliers are concentrated in the segments with the largest issuance surges, which are investment-grade segments, where the marginal effect of additional overlap may be smaller than the average effect identified in OLS.

6 Conclusion

This paper studies how multimarket contact among dealers shapes transaction costs in corporate bond markets. In a dealer-intermediated OTC market, the relevant unit of interaction is not a single bond but the broader set of securities that dealers trade in common. A dealer that widens her quote in one bond earns a short-run gain, but risks retaliation in every other bond where she trades with the same counterparties. This cross-bond linkage strengthens incentives to maintain tight pricing, especially in illiquid or information-sensitive bonds. The resulting surplus is redistributed through the equilibrium price schedule across all dealers active in the bond, helping sustain cooperation market-wide beyond bilateral relationships.

The empirical evidence is consistent with these predictions. Within a given bond and month, dealers with greater cross-bond overlap charge lower spreads, and the magnitude of this association is larger in information-sensitive and infrequently traded bonds. An instrumental-variables strategy based on COVID-era issuance waves confirms that exogenous variation in multimarket contact reduces transaction costs.

Our results highlight a distinction that standard measures of market structure obscure: in decentralized markets, concentration and competitive conduct can diverge. A market intermediated by a small number of dealers may exhibit limited pricing power, provided those dealers interact across a sufficiently rich set of heterogeneous securities. This finding has implications for regulatory design. Regulations that fragment dealer activity across bonds—capital rules that raise the cost of warehousing specific securities (e.g., SEC Rule 15c3-1), liquidity requirements that privilege a narrow subset of investment-grade debt (the Liquidity Coverage Ratio, 12 CFR Part 249, especially the treatment of Level 2B corporate debt securities), or proprietary-trading restrictions that discourage broad inventory holding

(the Volcker Rule, 12 CFR Part 248)—may weaken the repeated-interaction mechanism that disciplines markups, even as they pursue other legitimate objectives. This point adds to existing concerns about fragmentation in OTC markets, which have emphasized the loss of netting efficiencies and higher counterparty exposures when activity is separated across clearing arrangements (Duffie and Zhu, 2011). Regulators should therefore consider not only whether a market is concentrated, but also whether the pattern of dealer participation sustains the cross-market linkages that restrain pricing power.

References

Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica 56*(2), 383–396.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2015). Systemic risk and stability in financial networks. *American Economic Review 105*(2), 564–608.

Allen, F. and D. Gale (2000). Financial contagion. *Journal of Political Economy 108*(1), 1–33.

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

Bao, J., J. Pan, and J. Wang (2011). The illiquidity of corporate bonds. *The Journal of Finance 66*(3), 911–946.

Bartik, T. J. (1991). *Who Benefits from State and Local Economic Development Policies?* W.E. Upjohn Institute.

Becker, B. and E. Benmelech (2025). The resilience of the U.S. corporate bond market during financial crises. *International Journal of Central Banking 21*(4), 405–463. International Journal of Central Banking, Volume 21, Issue 4.

Bernheim, B. D. and M. D. Whinston (1990). Multimarket contact and collusive behavior. *The RAND Journal of Economics 21*(1), 1–26.

Bessembinder, H., W. Maxwell, and K. Venkataraman (2006). Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics 82*(2), 251–288.

Bessembinder, H., C. Spatt, and K. Venkataraman (2020). A survey of the microstructure of fixed-income markets. *Journal of Financial and Quantitative Analysis 55*(1), 1–45.

Blanchard, O. and L. Katz (1992). Regional evolutions. *Brookings Papers on Economic Activity 23*(1), 1–76.

Borusyak, K., P. Hull, and X. Jaravel (2021). Quasi-experimental shift-share research designs. *The Review of Economic Studies 89*(1), 181–213.

Choi, J., Y. Huh, and S. Seunghun Shin (2024). Customer liquidity provision: Implications for corporate bond transaction costs. *Management Science 70*(1), 187–206.

Christie, W. G. and P. H. Schultz (1994). Why do nasdaq market makers avoid odd-eighth quotes? *The Journal of Finance 49*(5), 1813–1840.

Christie, W. G. and P. H. Schultz (1995). Policy watch: Did nasdaq market makers implicitly collude? *The Journal of Economic Perspectives 9*(3), 199–208.

Ciliberto, F. and J. W. Williams (2014). Does multimarket contact facilitate tacit collusion? inference on conduct parameters in the airline industry. *The RAND Journal of Economics 45*(4), 764–791.

Colliard, J.-E., T. Foucault, and P. Hoffman (2021). Inventory management, dealers' connections, and prices in over-the-counter markets. *The Journal of Finance 76*(5), 2199–2247.

Cragg, J. G. and S. Donald (1993). Testing identifiability and specification in instrumental variable models. *Econometric Theory 9*(2), 222–240.

Di Maggio, M., A. Kermani, and Z. Song (2017). The value of trading relations in turbulent times. *Journal of Financial Economics 124*(2), 266–284.

Dick-Nielsen, J. and M. Rossi (2018). The cost of immediacy for corporate bonds. *The Review of Financial Studies 32*(1), 1–41.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2005). Over-the-counter markets. *Econometrica 73*(6), 1815–1847.

Duffie, D., N. Gârleanu, and L. H. Pedersen (2007). Valuation in over-the-counter markets. *The Review of Financial Studies 20*(6), 1865–1900.

Duffie, D. and H. Zhu (2011). Does a central clearing counterparty reduce counterparty risk? *Review of Asset Pricing Studies 1*(1), 74–95.

Dutta, P. K. and A. Madhavan (1997). Competition and collusion in dealer markets. *The Journal of Finance 52*(1), 245–276.

Edwards, A. K., L. E. Harris, and M. S. Piwowar (2007). Corporate bond market transaction costs and transparency. *The Journal of Finance 62*(3), 1421–1451.

Elliott, M., B. Golub, and M. O. Jackson (2014). Financial networks and contagion. *American Economic Review 104*(10), 3115–53.

Evans, W. N. and I. N. Kessides (1994). Living by the "golden rule": Multimarket contact in the us airline industry. *The Quarterly Journal of Economics 109*(2), 341–366.

Feldhütter, P. and T. Poulsen (2018). What determines bid-ask spreads in over-the-counter markets? SSRN Working Paper.

Foucault, T., M. Pagano, and A. Röell (2013). *Market Liquidity: Theory, Evidence, and Policy*. New York: Oxford University Press.

Friewald, N., R. Jankowitsch, and M. G. Subrahmanyam (2012). Illiquidity or credit deterioration: A study of liquidity in the us corporate bond market during financial crises. *Journal of Financial Economics 105*(1), 18–36.

Friewald, N. and F. Nagler (2019). Over-the-counter market frictions and yield spread changes. *The Journal of Finance 74*(6), 3217–3257.

Fudenberg, D. and E. Maskin (1986). The folk theorem in repeated games with discounting or with incomplete information. *Econometrica 54*(3), 533–554.

Glosten, L. R. and P. R. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics 14*(1), 71–100.

Goldberg, J. and Y. Nozawa (2021). Liquidity supply in the corporate bond market. *The Journal of Finance 76*(2), 755–796.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020). Bartik instruments: What, when, why, and how. *American Economic Review 110*(8), 2586–2624.

Goldstein, M. A., E. S. Hotchkiss, and E. R. Sirri (2006). Transparency and liquidity: A controlled experiment on corporate bonds. *The Review of Financial Studies 20*(2), 235–273.

Green, R. C. (2007). Presidential address: Issuers, underwriter syndicates, and aftermarket transparency. *The Journal of Finance 62*(4), 1529–1550.

Green, R. C., B. Hollifield, and N. Schürhoff (2007). Financial intermediation and the costs of trading in an opaque market. *The Review of Financial Studies 20*(2), 275–314.

Hasbrouck, J. (2007). *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. New York: Oxford University Press.

Hatfield, J. W. and J. Wallen (2022). Many markets make good neighbors: Multimarket contact and deposit banking. Working Paper.

Hendershott, T., D. Li, D. Livdan, and N. Schürhoff (2020). Relationship trading in over-the-counter markets. *The Journal of Finance 75*(2), 683–734.

Hendershott, T. and A. Madhavan (2015). Click or call? auction versus search in the over-the-counter market. *The Journal of Finance 70*(1), 419–447.

Jackson, M. O. (2008). *Social and Economic Networks*. Princeton, NJ: Princeton University Press.

Kleibergen, F. and R. Paap (2006). Generalized reduced rank tests using the singular value decomposition. *Journal of Econometrics 133*(1), 97–126.

Kyle, A. (1985). Continuous auctions and insider trading. *Econometrica 53*(6), 1315–35.

Li, D. and N. Schürhoff (2019). Dealer networks. *The Journal of Finance 74*(1), 91–144.

O'Hara, M. and X. A. Zhou (2021). The electronic evolution of corporate bond dealers. *Journal of Financial Economics 140*(2), 368–390.

O'Hara, M. and X. A. Zhou (2025). U.S. corporate bond markets: Bigger and (maybe) better? *Journal of Economic Perspectives 39*(2), 215–34.

O'Hara, M., A. C. Rapp, and X. A. Zhou (2025). Bond market resiliency: The role of insurers. *The Review of Financial Studies*, hhaf097.

O'Hara, M., Y. Wang, and X. A. Zhou (2018). The execution quality of corporate bonds. *Journal of Financial Economics 130*(2), 308–326.

Parker, P. M. and L.-H. Röller (1997). Collusive conduct in duopolies: Multimarket contact and cross-ownership in the mobile telephone industry. *The RAND Journal of Economics 28*(2), 304–322.

Schmitt, M. (2018). Multimarket contact in the hospital industry. *American Economic Journal: Economic Policy 10*(3), 361–387.

Seppi, D. J. (1990). Equilibrium block trading and asymmetric information. *the Journal of Finance 45*(1), 73–94.

Online Appendix

Contents

A. Proofs of baseline model 45

B. Model extension: Many Markets and Many Dealers 48

C. Model extension: Pass–Through to Customers 51

D. Model extension: Cooperation on Information Exploitation 53

E. Robustness Checks 62

F. CD–ID and ID–CD: Baseline Results, Subsample Analysis, and Robustness 69

A Proofs

This appendix presents the proofs of the results presented in the main text.

Proof of Lemma 1. The profit of the market-making dealer $\pi_M(s)$ is given by (1). The direct computation confirms that $s^* = \arg\max_{s\geq 0} \pi_M(s)$ and $\pi_M(s^B) = 0$. □

Proof of Proposition 1. The cooperation capacity is given by (6). Substituting, $\pi_M(s)$ and $\Pi(s)$ from (1) and (4), respectively, obtain

$$C(s) = -\frac{1}{2}\left((2-\alpha) + (1-\alpha)\frac{\delta}{2(1-\delta)}\right)s^2 + s + \left(-\frac{\alpha}{2} + \frac{\delta(1-\alpha)}{4(1-\delta)(2-\alpha)^2} - \frac{(1-\alpha)^2}{2(2-\alpha)}\right), \tag{A.1}$$

which is a quadratic function in s. Thus, the equation $C(s) = 0$ has at most two roots, s^* and s^W, with $C(s) > 0$ for all $s \in (s^W, s^*)$.

By the implicit function theorem,

$$\frac{ds^W}{d\delta} = -\frac{\frac{\partial C}{\partial \delta}}{\frac{\partial C}{\partial s^W}} < 0 \tag{A.2}$$

because $\frac{\partial C}{\partial \delta} = \frac{1}{(1-\delta)^2}(\Pi(s) - \Pi(s^*)) > 0$ and $\frac{\partial C}{\partial s^W} > 0$ at $s = s^W$. Similarly,

$$\frac{ds^W}{d\alpha} = -\frac{\frac{\partial C}{\partial \alpha}}{\frac{\partial C}{\partial s^W}} > 0 \tag{A.3}$$

because $\frac{\partial C}{\partial \alpha} = -\frac{4-2\alpha-\delta(4-3\alpha)-(2-\alpha)^3(2-\delta)(s^W)^2}{4(2-\alpha)^3(1-\delta)} < 0$ and $\frac{\partial C}{\partial s^W} > 0$ at $s = s^W$. □

Proof of Corollary 1. The cooperation capacity $C(s)$ is given by (6). Because $\Pi(s) - \Pi(s^*) > 0$ for any $s \in [0, s^*)$, it follows that $C(s)$ strictly increases with δ for such s. □

Proof of Proposition 2. By the proof of Proposition 1, $C^k(s^k)$ in (9) is a quadratic function. Thus, the smallest s^k solving $F(s^k, MMC^k) \equiv C^k(s^k) + MMC^k = 0$, decreases with MMC^k. □

Proof of Corollary 2. Taken separately, the incentive constraint in each of the two markets is given by (8). If the two markets are identical, these two incentive constraints are identical. Thus, when dealers cooperate on the same spreads in both markets, $s^{k_1} = s^{k_2}$, either the two constraints are simultaneously slack, in which case the cooperation capacity in both markets is positive but the welfare can be improved by lowering both spreads, possibly up to $s^{k_1} = s^{k_2} = 0$, or the two constraints are simultaneously binding, in which case the cooperation capacity in both markets is zero. □

Proof of Proposition 3. Consider $F(s^k, MMC^k)$ defined in the proof of Proposition 2. By the implicit function theorem,

$$\frac{ds^k}{dMMC^k} = -\frac{\frac{\partial F}{\partial MMC^k}}{\frac{\partial F}{\partial s^k}} = \frac{1}{\left((2-\alpha) + (1-\alpha)\frac{\delta}{2(1-\delta)}\right) s^k - 1} < 0 \tag{A.4}$$

because $\frac{\partial F}{\partial s^k} = 1 - \left((2-\alpha) + (1-\alpha)\frac{\delta}{2(1-\delta)}\right) s^k > 0$ for the smallest s^k solving $F(s^k, MMC^k) =$

0. Define $D \equiv \left(\left((2-\alpha)+(1-\alpha)\frac{\delta}{2(1-\delta)}\right)s^k\right)^2$. Now compute

$$\frac{dD}{d\delta} = \frac{(1-\alpha)}{(1-\delta)^2}\left(MMC^k + \frac{\delta(1-\alpha)}{2(1-\delta)(2-\alpha)^2}\right) > 0, \tag{A.5}$$

which means that $\left|\frac{ds^k}{dMMC^k}\right|$ decreases with δ. Similarly,

$$\frac{dD}{d\alpha} = -\left(\frac{MMC^k(2-\delta)}{(1-\delta)} + \frac{\delta^2(1-\alpha)}{2(1-\delta)^2(2-\alpha)^3}\right) < 0, \tag{A.6}$$

which means that $\left|\frac{ds^k}{dMMC^k}\right|$ increases with α. □

B Extension: Many Markets and Many Dealers

This appendix states and proves all results from Section 2.3.1 for $N > 2$ dealers in $K > 2$ markets under uniform trading probabilities.

In each market k and each period, each of the N dealers seeks liquidity with probability $1/N$, and conditional on seeking liquidity she contacts each of the remaining $N-1$ dealers with equal probability:

$$p_{ij}^k = p \equiv \frac{1}{N(N-1)} \qquad \text{for all } i \neq j, \qquad p_{ii}^k \equiv 0. \tag{B.1}$$

Each dealer's total probability of being the maker (or the taker) in a given period is $(N-1)p = 1/N$.

Let s_{ij}^k denote the spread that dealer i quotes to dealer j in market k, and let α_j^k denote the probability that taker j is informed in market k. Dealer i's expected per-period profit in market k is

$$\Pi_i^k(S^k) \equiv p \sum_{j \neq i} [\pi_M^k(s_{ij}^k; \alpha_j^k) + T^k(s_{ji}^k)], \tag{B.2}$$

where $\pi_M^k(s; \alpha_j^k) = (1-\alpha_j^k)(1-s)s - \frac{\alpha_j^k}{2}(1-s)^2$ is maker profit when facing taker j, $T^k(s) = \frac{(1-s)^2}{2}$ is the taker payoff, and $S^k = (s_{ij}^k)$ with $s_{ii}^k \equiv 0$. The static monopoly spread when facing taker j is $s_j^{*,k} = 1/(2-\alpha_j^k)$.

Write $\rho_k := \delta_k/(1-\delta_k)$. Dealer i's cooperation capacity in market k is

$$C_i^k(S^k) \equiv \rho_k p \sum_{j \neq i} [\pi_M^k(s_{ij}^k; \alpha_j^k) - \pi_M^k(s_j^{*,k}; \alpha_j^k) + T^k(s_{ji}^k) - T^k(s_j^{*,k})] - d_i^k, \tag{B.3}$$

where $d_i^k \equiv \max_{j \neq i}(\pi_M^k(s_j^{*,k}; \alpha_j^k) - \pi_M^k(s_{ij}^k; \alpha_j^k))$ is the deviation gain. The IC constraint is

$$C_i^k(S^k) + MMC_i^k \geq 0, \tag{B.4}$$

with $MMC_i^k = \sum_{\ell \neq k} C_i^\ell$. Total per-period welfare equals $\mathcal{W}^k(S^k) = p \sum_{i \neq j} G^k(s_{ij}^k; \alpha_j^k)$,

where $G^k(s;\alpha_j^k) = \pi_M^k(s;\alpha_j^k) + T^k(s)$ is strictly decreasing in s. The planner's problem is

$$S^{W,k} = \arg\max_{S^k \geq 0} \ \mathcal{W}^k(S^k) \qquad \text{s.t.} \quad C_i^k(S^k) + MMC_i^k \geq 0 \ \ \forall\, i \in \mathcal{N}^k. \tag{B.5}$$

The max in the deviation gain is non-smooth. We introduce auxiliary variables d_i^k and selector constraints $d_i^k \geq g_{ij}^k \equiv \pi_M^k(s_j^{*,k};\alpha_j^k) - \pi_M^k(s_{ij}^k;\alpha_j^k)$ for each $j \neq i$. Let $\lambda_i \geq 0$ be the multiplier on dealer i's IC constraint and $\eta_{ij} \geq 0$ the multiplier on the selector constraint linking d_i^k to counterparty j. The Lagrangian is

$$\mathcal{L} = p \sum_{i \neq j} G^k(s_{ij}^k;\alpha_j^k) + \sum_i \lambda_i [C_i^k - d_i^k + MMC_i^k] + \sum_i \sum_{j \neq i} \eta_{ij}[d_i^k - g_{ij}^k]. \tag{B.6}$$

B.1 Proofs

Lemma 2 (Diagonal Hessian of the Lagrangian)**.** *For each ordered pair (i,j) with $i \neq j$, the variable s_{ij}^k enters the Lagrangian only through terms that depend on s_{ij}^k alone, so $\frac{\partial^2 \mathcal{L}}{\partial s_{ij}^k \, \partial s_{i'j'}^k} = 0$ for $(i,j) \neq (i',j')$.*

Proof. The spread s_{ij}^k appears in exactly four places: (i) the welfare term $p\, G^k(s_{ij}^k;\alpha_j^k)$; (ii) dealer i's continuation value through $\rho_k p\, [\pi_M^k(s_{ij}^k;\alpha_j^k) - \pi_M^k(s_j^{*,k};\alpha_j^k)]$; (iii) dealer j's continuation value through $\rho_k p\, [T^k(s_{ij}^k) - T^k(s^{*,k})]$ (using $p_{ji} = p_{ij} = p$); (iv) the selector constraint $d_i^k \geq g_{ij}^k$, through $\eta_{ij}\pi_M^k(s_{ij}^k;\alpha_j^k)$. No other spread appears jointly with s_{ij}^k, so all cross-partials vanish. □

Proof of Proposition 4. Stationarity with respect to d_i^k gives $\sum_{j \neq i} \eta_{ij} = \lambda_i$. Stationarity with respect to s_{ij}^k gives

$$p\, G^{k\prime}(s_{ij}^k;\alpha_j^k) + (\rho_k p\, \lambda_i + \eta_{ij})\, \pi_M^{k\prime}(s_{ij}^k;\alpha_j^k) + \rho_k p\, \lambda_j\, T^{k\prime}(s_{ij}^k) = 0. \tag{B.7}$$

Substituting

$$G^{k\prime}(s;\alpha_j^k) = -(1-\alpha_j^k)s, \qquad \pi_M^{k\prime}(s;\alpha_j^k) = 1 - (2-\alpha_j^k)s, \qquad T^{k\prime}(s) = -(1-s),$$

and collecting terms:

$$[\rho_k p(\lambda_i - \lambda_j) + \eta_{ij}] - [p(1-\alpha_j^k) + (\rho_k p\,\lambda_i + \eta_{ij})(2-\alpha_j^k) - \rho_k p\,\lambda_j]\, s_{ij}^k = 0,$$

which gives (10). When $\eta_{ij} = 0$ the factor p cancels and the spread depends only on $(\lambda_i, \lambda_j, \alpha_j^k)$. When $\eta_{ij} = \lambda_i$ (unique pinning to counterparty j), direct substitution yields a second special case. If multiple counterparties simultaneously attain the maximum deviation gap, η_{ij} splits λ_i across them. □

Proof of Proposition 5. Suppose $\alpha_j^k = \alpha^k$ for all j and $s_{ij}^k \neq s_{ij'}^k$ for some $j, j' \neq i$. Replace all of i's spreads with their average $\bar{s}_i$.

Welfare. $G^k(\cdot;\alpha^k)$ is strictly concave ($G^{k\prime\prime} = -(1-\alpha^k) < 0$), so Jensen's inequality gives a strict increase.

Maker i's IC. The continuation-value maker terms rise by concavity of π_M^k, and the deviation gain d_i^k weakly falls because the minimum maker profit rises. Both effects relax i's constraint.

Other dealers' IC. Taker surplus T^k is convex, so some taker terms rise and others fall. But the strict welfare and maker-IC gains dominate this mean-preserving redistribution, contradicting optimality of unequal spreads.

With all spreads from maker i equal, all deviation gaps are equal, so $\eta_{ij} = \lambda_i/(N-1)$. Substituting into (10) and defining $A^k \equiv (N-1)\rho_k p + 1$:

$$s_i^k = \frac{A^k \lambda_i - \rho_k p \sum_{j\neq i} \lambda_j}{(N-1)p(1-\alpha^k) + A^k \lambda_i (2-\alpha^k) - \rho_k p \sum_{j\neq i} \lambda_j}. \tag{B.8}$$

□

Proof of Corollary 3. Write $s_{ij}^k = \mathcal{N}/\mathcal{D}$ from (10). The numerator $\mathcal{N} = \rho_k p(\lambda_i - \lambda_j) + \eta_{ij}$ is increasing in λ_i and η_{ij}, decreasing in λ_j. The denominator $\mathcal{D}$ is positive and increasing in all three. At any interior solution with $s_{ij}^k \in (0, s_j^{*,k})$, the quotient rule gives $\partial s_{ij}^k/\partial \lambda_i > 0$, $\partial s_{ij}^k/\partial \alpha_j^k > 0$, and $\partial s_{ij}^k/\partial \lambda_j < 0$. Part (ii) follows from the same argument applied to (B.8). □

C Extension: Pass–Through to Customers

We now add customer trades to the model and show that lower interdealer quotes resulting from higher MMC also benefit ultimate investors.

To incorporate customers, we assume for simplicity that each customer is connected to a single dealer. Whenever a customer needs to trade, she trades with her dealer at an endogenous price s_c quoted by the dealer. As with liquidity-seeking dealers, customer trades are either liquidity-motivated or information-motivated. With probability $\alpha_c \in [0,1]$, the customer is information-motivated, in which case she privately knows the common-value component. With complementary probability $1-\alpha_c$, the customer is liquidity-motivated, in which case she has a private-value component b_c for trading the asset (and knows nothing about the common-value component). The private-value component is equally likely to be positive or negative, and its magnitude is independently distributed according to a distribution $F_{LC}(b_c)$ with support on $[0,1]$. For simplicity, we assume $F_{LC}(b_c)$ is uniform.

We assume that all customer trades are immediately offset in the interdealer market. Thus, the liquidity-seeking dealer now trades in the interdealer market either because her trade is customer-driven (CD) or because she trades on her own account, in which case she is either (L) or (I) as in our baseline model. We assume that the liquidity-seeking dealer's trade is customer-driven with probability β and driven by her own interests with probability $1-\beta$. Finally, since the interdealer quote s is effectively a take-it-or-leave-it offer from the market-making dealer, the dealer providing liquidity to her customers takes s as given.

Customer Quote. Given the above assumptions, the profit of the customer's dealer, when offering quote s_c to the customer, is $\pi_{CD}(s_c) = (1-s_c)(s_c-s)$ and thus the dealer's optimal quote is

$$s_c(s) = \frac{1+s}{2}. \tag{C.1}$$

The dealer's quote s_c passes through the interdealer quote s to the customer with a markup.

Interdealer Quote. The characterization of the interdealer quote and how MMC affects it is similar to our baseline model. In the case of one market and two dealers, the profit of the market-making dealer is now the sum of profits from trade with a customer-driven liquidity

seeker (CD) and from trade with a dealer acting on her own interests, either (L) or (I):

$$\pi_M(s) = \beta\left((1-\alpha_c)(1-s_c(s))s - \alpha_c\int_{s_c}^{1}(v-s)dv\right)+(1-\beta)\left((1-\alpha)(1-s)s - \alpha\int_{s}^{1}(v-s)dv\right), \tag{C.2}$$

where the market-making dealer internalizes the impact of her quote on the customer's quote because the customer's dealer takes the interdealer quote s as given. Solving for the monopolistic quote of the market-making dealer we obtain

$$s^* = \frac{\beta\frac{1}{4}(2+\alpha_c)+(1-\beta)}{\beta\frac{1}{4}(4-\alpha_c)+(1-\beta)(2-\alpha)}. \tag{C.3}$$

The dealer's expected per-period profit when both quote s in the interdealer market is

$$\Pi(s) = \frac{1}{2}\pi_M(s) + \frac{1}{2}\left(\beta\pi_{CD}(s) + (1-\beta)[\alpha\pi_I(s) + (1-\alpha)\pi_L(s)]\right). \tag{C.4}$$

The incentive-compatibility constraint remains the same—a market-making dealer prefers to quote s rather than the static monopoly spread s^*, which maximizes (C.2), whenever

$$\pi_M(s) + \frac{\delta}{1-\delta}\Pi(s) \geq \pi_M(s^*) + \frac{\delta}{1-\delta}\Pi(s^*). \tag{C.5}$$

Relative to the baseline without customers, the customer's surplus extraction weakens the repeated-game incentive to cooperate: because informed gains and losses no longer fully cancel within the dealer pair when a customer makes the trade, the ratio of cooperation rents to deviation gains falls with the customer share β. As a result, when customer trades are at least as adversely selected as own trades, $\alpha_c \geq \alpha$, the set of interdealer quotes sustainable by repeated-game incentives strictly shrinks as β increases.

Finally, with multiple markets the incentive-compatibility constraint for dealer i in market k becomes

$$C_i^k(S^k) + MMC_i^k \geq 0. \tag{C.6}$$

As before, if MMC_i^k is strictly positive, it is easier to sustain a lower interdealer quote s_i^k for dealer i in market k. The customer's quote therefore also decreases with MMC, since it passes through the interdealer quote with a markup.

D Extension: Cooperation on Information Exploitation

In the baseline model, cooperation among dealers operates only through the market-making channel: dealers coordinate to quote lower spreads when providing liquidity. We now introduce a second dimension of cooperation: when privately informed, a dealer can promise to *restrain information exploitation* by trading less aggressively than her information would permit. This additional margin can sustain lower spreads than market-making cooperation alone.

We develop this extension in parallel to the baseline framework. We begin with one market and two dealers, derive the incentive constraints for two-dimensional cooperation, and characterize the welfare-maximizing cooperative outcome. We then extend the analysis to two markets and show how multimarket contact interacts with two-dimensional cooperation.

Throughout, we assume that $|v|$ and $|b|$ are independently distributed uniformly on $[0, 1]$ and that $\alpha \in (0, 1)$.

D.1 One Market and Two Dealers

Modified Stage Game. The stage game follows the baseline setup. The modification concerns the behavior of the informed dealer. In the baseline model, an information-motivated dealer trades whenever $|v| > s$, where s is the quoted spread. Under information-exploitation cooperation, the informed dealer instead commits to trade only when

$$|v| > s_I \qquad \text{for some } s_I \geq s,$$

thereby voluntarily foregoing profitable trades in the interval (s, s_I) and reducing the adverse-selection losses faced by the market-maker.

Under spread s and informed threshold s_I with $s \leq s_I$, the market-maker's expected profit is

$$\pi_M(s, s_I) = (1 - \alpha)(1 - s)s - \alpha \int_{\max(s, s_I)}^{1} (v - s)\, dv. \tag{D.1}$$

The informed trader only trades when $v > \max(s, s_I)$: if the market-maker's quote exceeds the informed trader's threshold, the informed trader will not trade at a loss.

For the cooperative outcome with $s \leq s_I$, we have $\max(s, s_I) = s_I$, so:

$$\pi_M(s, s_I) = (1-\alpha)(1-s)s - \alpha\left(\frac{1-s_I^2}{2} - s(1-s_I)\right), \tag{D.2}$$

the informed dealer's profit is

$$\pi_I(s, s_I) = \frac{1-s_I^2}{2} - s(1-s_I), \tag{D.3}$$

and the liquidity-motivated dealer's profit is $\pi_L(s) = (1-s)^2/2$ as before. For later use, note that

$$\frac{\partial \pi_M(s, s_I)}{\partial s_I} = \alpha(s_I - s) \geq 0 \qquad \text{and} \qquad \frac{\partial \pi_I(s, s_I)}{\partial s_I} = -(s_I - s) \leq 0. \tag{D.4}$$

Thus, for fixed s, increasing s_I weakly increases market-maker profits and weakly decreases informed-trader profits.

Lemma 3 (Static equilibrium with informed threshold)**.** *In the static game, the unique equilibrium has $s_I = s$: the informed dealer fully exploits her information. Anticipating this, the market-maker chooses $s^* = 1/(2-\alpha)$, as in the baseline.*

Proof. By (D.4), $\pi_I(s, s_I)$ is strictly decreasing in s_I for $s_I > s$. Hence the informed dealer maximizes current profit by choosing $s_I = s$. Substituting $s_I = s$ into (D.2) recovers the baseline profit $\pi_M(s, s) = (1-\alpha)(1-s)s - \alpha(1-s)^2/2$, which is maximized at $s^* = 1/(2-\alpha)$. □

Repeated Game with Two-Dimensional Cooperation. Dealers play the stage game infinitely often with common discount factor $\delta \in (0, 1)$. Each dealer has probability $\frac{1}{2}$ of being the liquidity-seeker per period. We focus on symmetric equilibria in which both dealers implement the same cooperative profile (s, s_I): when market-making, quote spread s; when informed, trade only if $|v| > s_I$.

Cooperation is supported by grim-trigger strategies. A deviation on *either* dimension—quoting a different spread when market-making, or trading in the interval (s, s_I) when

informed—triggers permanent reversion to the static equilibrium (s^*, s^*).

When both dealers adhere to (s, s_I) with $s \leq s_I$, the expected per-period profit of a dealer is

$$\Pi(s, s_I) = \frac{1}{2}\pi_M(s, s_I) + \frac{1}{2}\left[\alpha\pi_I(s, s_I) + (1-\alpha)\pi_L(s)\right] = (1-\alpha)\frac{1-s^2}{4}. \tag{D.5}$$

A key observation: $\Pi(s, s_I)$ depends on s but *not* on s_I. This is because total social surplus in each period comes only from liquidity trades—informed trades are pure transfers between dealers—so changing s_I redistributes contemporaneous profits between the market-maker and the informed trader without affecting total expected per-period payoff.

One-Shot Deviations. Fix a cooperative outcome (s, s_I) with $s \leq s_I$. We characterize the optimal deviations on each dimension.

Market-maker deviation. The market-maker can deviate to any spread $\tilde{s}$, while the counterparty still follows the cooperative threshold s_I. The optimal deviation spread is

$$\tilde{s}^D(s_I) = \begin{cases} s^* & \text{if } s_I < s^*, \\ \dfrac{1-\alpha s_I}{2(1-\alpha)} & \text{if } s_I \geq s^*. \end{cases} \tag{D.6}$$

When $s_I < s^*$, the optimal deviation is to quote s^*, which exceeds s_I, so the informed trader trades at threshold s^* and the deviation payoff equals the static monopoly profit: $D_M(s_I) = \pi_M(s^*, s^*)$. When $s_I \geq s^*$, the optimal deviation respects $\tilde{s} \leq s_I$, yielding $D_M(s_I) = \pi_M(\tilde{s}^D(s_I), s_I)$. Since only the case $s_I < s^*$ arises at the relevant cooperative outcomes (where $s < s^*$ and s_I is close to s), we focus on this case, in which $D_M(s_I) = \pi_M(s^*, s^*)$ is independent of s_I.

Informed-trader deviation. An informed dealer deviates by trading whenever $|v| > s$ rather than $|v| > s_I$. The deviation payoff is $\pi_I(s, s) = (1-s)^2/2$.

Lemma 4 (Informed-trader deviation gain)**.** *For $s \leq s_I$,*

$$\pi_I(s, s) - \pi_I(s, s_I) = \frac{(s_I - s)^2}{2}.$$

Proof. Direct computation using (D.3): $\pi_I(s,s) - \pi_I(s,s_I) = \frac{(1-s)^2}{2} - \frac{1-s_I^2}{2} + s(1-s_I) = \frac{s_I^2 - 2ss_I + s^2}{2} = \frac{(s_I - s)^2}{2}$. □

Cooperation Capacities. Define the cooperation capacities for the two dimensions:

$$C^{MM}(s,s_I) \equiv \frac{\delta}{1-\delta}\Big[\Pi(s) - \Pi(s^*)\Big] - \Big[D_M(s_I) - \pi_M(s,s_I)\Big], \tag{D.7}$$

$$C^{I}(s,s_I) \equiv \frac{\delta}{1-\delta}\Big[\Pi(s) - \Pi(s^*)\Big] - \frac{(s_I - s)^2}{2}, \tag{D.8}$$

where we use $\Pi(s) \equiv \Pi(s,s_I)$ to emphasize independence from s_I. The cooperative outcome (s,s_I) is sustainable if and only if

$$C^{MM}(s,s_I) \geq 0 \qquad \text{and} \qquad C^{I}(s,s_I) \geq 0. \tag{D.9}$$

Interaction between the Two Dimensions. The two cooperation dimensions interact in a natural way. For the relevant case $s_I < s^*$, the market-maker deviation payoff $D_M(s_I) = \pi_M(s^*,s^*)$ is independent of s_I, while $\pi_M(s,s_I)$ increases in s_I by (D.4). Hence,

$$\frac{\partial C^{MM}(s,s_I)}{\partial s_I} = \alpha(s_I - s) \geq 0. \tag{D.10}$$

At the same time, the informed-trader deviation gain $(s_I - s)^2/2$ increases in s_I, so

$$\frac{\partial C^{I}(s,s_I)}{\partial s_I} = -(s_I - s) \leq 0. \tag{D.11}$$

Thus, raising the informed threshold s_I *relaxes* the market-maker constraint (by reducing adverse selection) but *tightens* the informed-trader constraint (by increasing the temptation to deviate). The optimal s_I balances these two forces.

Welfare-Maximizing Cooperative Outcome. In the baseline (one-dimensional) problem, the welfare-maximizing spread is

$$\tilde{s}^W = \min\big\{s \in [0,1] : C^{MM}(s,s) \geq 0\big\}. \tag{D.12}$$

In the extended (two-dimensional) problem, the welfare-maximizing outcome is

$$(s^W, s_I^W) = \arg\min_{s,s_I} \left\{ s \in [0,1] : C^{MM}(s, s_I) \geq 0,\ C^I(s, s_I) \geq 0,\ s_I \geq s \right\}. \tag{D.13}$$

Since total surplus depends only on s and is decreasing in s, the welfare-maximizing outcome minimizes s subject to both incentive constraints.

Proposition 6 (Two-dimensional cooperation weakly expands the feasible set)**.** *For every (δ, α), $s^W \leq \tilde{s}^W$. Moreover, if the informed-trader constraint is slack at the baseline optimum, i.e., $C^I(\tilde{s}^W, \tilde{s}^W) > 0$, then $s^W < \tilde{s}^W$.*

Proof. Weak inequality. The baseline problem restricts $s_I = s$. We verify that $(\tilde{s}^W, \tilde{s}^W)$ satisfies both constraints: $C^{MM}(\tilde{s}^W, \tilde{s}^W) \geq 0$ by definition, and $C^I(\tilde{s}^W, \tilde{s}^W) = \frac{\delta}{1-\delta}[\Pi(\tilde{s}^W) - \Pi(s^*)] \geq 0$ since $\tilde{s}^W \leq s^*$ implies $\Pi(\tilde{s}^W) \geq \Pi(s^*)$. Hence every baseline-feasible outcome is feasible in the extended problem, so $s^W \leq \tilde{s}^W$.

Strict inequality. Suppose $C^I(\tilde{s}^W, \tilde{s}^W) > 0$, which implies $\tilde{s}^W < s^*$. Pick $s < \tilde{s}^W$ close to $\tilde{s}^W$. Since $\tilde{s}^W$ is the minimal baseline-feasible spread, $C^{MM}(s, s) < 0$. However, by (D.10), increasing s_I above s raises C^{MM}:

$$C^{MM}(s, s_I) = C^{MM}(s, s) + \alpha \frac{(s_I - s)^2}{2},$$

where the second term follows from integrating (D.10) and using $D_M(s_I) = \pi_M(s^*, s^*)$ for $s_I < s^*$. Choose $s_I > s$ so that $C^{MM}(s, s_I) = 0$, which is possible by continuity. For s and s_I sufficiently close to $\tilde{s}^W$, continuity and $C^I(\tilde{s}^W, \tilde{s}^W) > 0$ ensure $C^I(s, s_I) > 0$. Hence (s, s_I) is feasible with $s < \tilde{s}^W$. □

The intuition is as follows. At the baseline welfare-maximizing spread $\tilde{s}^W$, the market-maker constraint binds but the informed-trader constraint is generically slack. This slack can be "spent" by introducing a small amount of informed restraint ($s_I > \tilde{s}^W$), which relaxes the market-maker constraint enough to sustain a tighter spread.

D.2 Two Markets and Multimarket Contact

We now extend the two-dimensional cooperation framework to two markets, introducing multimarket contact. The setup follows Section 2: two dealers operate in markets k_1 and k_2, the stage game is played simultaneously in both markets each period, and dealer roles are assigned independently across markets.

Fix cooperative profiles (s^k, s_I^k) for each market $k \in \{k_1, k_2\}$. On the cooperative path, a market-maker in market k quotes s^k and an informed dealer in market k trades only when $|v| > s_I^k$. A deviation on *any* dimension in *any* market triggers reversion to the static equilibrium $(s^{*,k}, s^{*,k})$ in both markets from $t+1$ onward.

Incentive Constraints. Because within-period payoffs are additively separable across markets, the best one-shot deviation in a given period is to deviate on every available dimension in every market. The deviation gain in market m depends on the dealer's role:

- *Market-maker:* deviation gain $D_M^m(s_I^m) - \pi_M^m(s^m, s_I^m)$;
- *Informed trader:* deviation gain $(s_I^m - s^m)^2/2$;
- *Liquidity-motivated trader:* no deviation opportunity.

The no-deviation condition for a dealer in state (r^k, r^ℓ), where $r^m \in \{MM, I, L\}$ denotes the role in market m, is

$$\sum_{m \in \{k,\ell\}} \left[\frac{\delta^m}{1-\delta^m} \Big(\Pi^m(s^m) - \Pi^m(s^{*,m}) \Big) - G^m(r^m) \right] \geq 0, \tag{D.14}$$

where $G^m(MM) = D_M^m(s_I^m) - \pi_M^m(s^m, s_I^m)$, $G^m(I) = (s_I^m - s^m)^2/2$, and $G^m(L) = 0$.

Using the cooperation capacities (D.7)–(D.8), the condition for market m and role r^m can be written as

$$\frac{\delta^m}{1-\delta^m} \Big(\Pi^m(s^m) - \Pi^m(s^{*,m}) \Big) - G^m(r^m) = \begin{cases} C^{MM,m}(s^m, s_I^m) & \text{if } r^m = MM, \\ C^{I,m}(s^m, s_I^m) & \text{if } r^m = I, \\ \frac{\delta^m}{1-\delta^m} \left(\Pi^m(s^m) - \Pi^m(s^{*,m}) \right) & \text{if } r^m = L. \end{cases}$$

Since $C^{MM,m} \leq \frac{\delta^m}{1-\delta^m}(\Pi^m(s^m) - \Pi^m(s^{*,m}))$ and $C^{I,m} \leq \frac{\delta^m}{1-\delta^m}(\Pi^m(s^m) - \Pi^m(s^{*,m}))$, the tightest constraint arises when the cooperation-capacity term in each market is minimized. Define the *minimum cooperation capacity* of market m as

$$C^{\min,m}(s^m, s_I^m) \equiv \min\left\{C^{MM,m}(s^m, s_I^m),\ C^{I,m}(s^m, s_I^m)\right\}. \tag{D.15}$$

The following lemma characterizes the tightest constraint.

Lemma 5 (Tightest incentive constraint)**.** *The no-deviation condition* (D.14) *holds for all role assignments* (r^k, r^ℓ) *if and only if*

$$C^{\min,k}(s^k, s_I^k) + C^{\min,\ell}(s^\ell, s_I^\ell) \geq 0. \tag{D.16}$$

Proof. For each market m, the term in (D.14) is minimized at $C^{\min,m}(s^m, s_I^m)$. Since the overall constraint is the sum of two such terms and the role in each market is independent,

$$\min_{(r^k, r^\ell)} \left\{ \left[\frac{\delta_k}{1-\delta_k}(\Pi^k - \Pi^{*,k}) - G^k(r^k) \right] + \left[\frac{\delta_\ell}{1-\delta_\ell}(\Pi^\ell - \Pi^{*,\ell}) - G^\ell(r^\ell) \right] \right\} = C^{\min,k} + C^{\min,\ell},$$

where the equality uses $\min(a+c, a+d, b+c, b+d) = \min(a,b) + \min(c,d)$. □

Equivalently, defining $MMC^k \equiv C^{\min,\ell}(s^\ell, s_I^\ell)$ as the minimum cooperation capacity in the other market, the incentive constraints for market k can be written as

$$C^{MM,k}(s^k, s_I^k) + MMC^k \geq 0 \qquad \text{and} \qquad C^{I,k}(s^k, s_I^k) + MMC^k \geq 0. \tag{D.17}$$

Proposition 7 (MMC with two-dimensional cooperation)**.** *If* $MMC^k > 0$*, the set of sustainable cooperative outcomes* (s^k, s_I^k) *in market* k *weakly expands. In particular, the welfare-maximizing spread* $s^{W,k}$ *weakly decreases.*

Proof. Positive MMC^k relaxes both constraints in (D.17) relative to the single-market case. Thus, any (s^k, s_I^k) sustainable in the single-market case remains sustainable, and additional outcomes with lower s^k may become feasible. □

Multimarket contact interacts with two-dimensional cooperation through two channels.

First, as in the baseline, cooperation capacity from market ℓ can be transferred to sustain lower spreads in market k. Second, the cooperation capacity that market ℓ generates now reflects *both* dimensions: cooperation on informed restraint in market ℓ affects $C^{\min,\ell}$ and therefore the MMC available to market k.

Complementarity across Markets and Dimensions. In the single-market case, informed restraint allows tighter spreads by relaxing the market-maker constraint. MMC amplifies this mechanism: if informed restraint in market ℓ generates additional cooperation capacity (i.e., raises $C^{\min,\ell}$), this capacity can be transferred to market k via MMC^k.

Corollary 4 (Symmetric markets)**.** *When identical dealers trade in identical markets, multimarket contact does not aid in sustaining cooperation on either dimension. That is, the analysis reduces to the single-market case.*

Proof. With identical markets, the incentive constraint (D.16) becomes $2\,C^{\min,k}(s^k, s_I^k) \geq 0$, which is equivalent to the single-market constraint $C^{\min,k} \geq 0$. The argument is the same as in Corollary 2. □

Corollary 5 (Cross-dimension complementarity)**.** *The two cooperation dimensions are complementary: at a sustainable outcome (s, s_I) with $s_I < s^*$, a marginal increase in informed restraint (higher s_I) relaxes the market-maker incentive constraint by $\alpha(s_I - s)$ and tightens the informed-trader constraint by $(s_I - s)$. When one market's minimum cooperation capacity is determined by the market-maker constraint, informed restraint in that market increases $C^{\min}$ and thus expands cooperation in the other market via MMC.*

Proof. The effects on C^{MM} and C^I follow from (D.10)–(D.11). If $C^{MM,\ell} < C^{I,\ell}$ at the cooperative outcome in market ℓ, then $C^{\min,\ell} = C^{MM,\ell}$, and raising s_I^ℓ increases $C^{MM,\ell}$. This raises $C^{\min,\ell}$ (provided $C^{MM,\ell}$ remains the binding constraint), increasing $MMC^k = C^{\min,\ell}$ and expanding the feasible set in market k. □

Example 5 (MMC with two-dimensional cooperation)**.** Consider two markets that differ only in trading frequency, as in Example 1. The high-frequency market h has effective discount factor $\delta_h = \sqrt{\delta} > \delta_\ell = \delta$. Under one-dimensional cooperation (market-making only),

the high-frequency market generates cooperation capacity $C^h(s^h)$ that can be transferred to market ℓ via MMC. Under two-dimensional cooperation, dealers can additionally coordinate on informed restraint ($s_I^h > s^h$), which relaxes the market-maker constraint in market h and potentially raises $C^{\min,h}$, providing even more capacity to transfer to market ℓ. As a result, the welfare-maximizing spread in market ℓ is weakly lower under two-dimensional cooperation with MMC than under either mechanism alone.

Appendix E. Robustness Checks

Table E.1. MMC and Transaction Costs — Robustness on Clustering and Fixed Effects

	Any-ID Cost					
	Clustering		Fixed Effects			
	(1) Triple SE	(2) Triple SE	(3) Dealer-Bond FE	(4) Dealer-Bond FE	(6) Dealer-Issuer Industry FE	(6) Dealer-Issuer Industry FE
Avg. Dealer Overlap	-5.31***		-2.45**		-5.03***	
	(1.97)		(1.22)		(1.93)	
Avg. Dealer Overlap, Weighted		-3.94***		-2.70***		-3.83***
		(1.49)		(0.94)		(1.45)
Bond-Time FE	✓	✓	✓	✓	✓	
Dealer-Time FE	✓	✓	✓	✓	✓	
Dealer-Bond FE			✓	✓		
Dealer-Issuer Industry FE					✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓
Obs	1,120,822	1,120,822	1,019,210	1,019,210	1,119,900	1,119,900
Adj.R^2	0.285	0.285	0.351	0.351	0.290	0.290

Notes: This table reports robustness checks of the baseline MMC–spread relationship in Table 2. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. Columns (1) and (2) replicate the baseline specification but use triple clustering at the dealer, bond, and month levels instead of double clustering. Columns (3) and (4) add dealer–bond fixed effects to the baseline specification. Columns (5) and (6) add dealer–issuer industry fixed effects to the baseline specification. Besides these changes, the specifications include the same set of fixed effects and controls used in Table 2. Coefficients represent the change in transaction costs associated with a one-standard-deviation increase in MMC. Standard errors in columns (1) and (2) are clustered at the dealer, bond, and month levels. Standard errors in columns (3) through (6) are clustered following the baseline specification at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.2. MMC and Transaction Costs — By Rating and Maturity

	Investment-Grade vs. High-Yield				Maturity: ≤5Y vs. >5Y			
	(1) IG	(2) HY	(3) IG	(4) HY	(5) ≤5Y	(6) >5Y	(7) ≤5Y	(8) >5Y
Avg. Dealer Overlap	-3.71**	-5.12*			-3.88*	-7.04***		
	(1.63)	(3.02)			(2.16)	(2.06)		
Avg. Dealer Overlap, Weighted			-2.26**	-5.10**			-3.07	-5.78***
			(0.97)	(2.25)			(1.87)	(1.26)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓	✓
Obs	862,224	254,157	862,224	254,157	586,700	528,724	586,700	528,724
$Adj.R^2$	0.305	0.250	0.305	0.250	0.332	0.259	0.332	0.259

Notes: This table reports subsample estimates of the MMC–spread relationship for investment-grade and high-yield bonds, as well as for bonds with below-five-year and above-five-year remaining maturities. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Coefficients represent the change in transaction costs associated with a one-standard-deviation increase in MMC within each rating or maturity group. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.3. MMC and Transaction Costs — By Sector

	Any-ID Cost					
	(1) Industrials	(2) Finance	(3) Utility	(4) Industrials	(5) Finance	(6) Utility
Avg. Dealer Overlap	-4.68***	-6.64**	-3.50			
	(1.80)	(2.67)	(3.40)			
Avg. Dealer Overlap, Weighted				-3.64***	-4.38**	-2.09
				(1.22)	(2.01)	(2.49)
Bond-Time FE	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓
Obs	689,805	364,411	37,270	689,805	364,411	37,270
Adj.R^2	0.275	0.327	0.317	0.275	0.327	0.317

Notes: This table presents subsample estimates of the MMC–spread relationship across issuer industries. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. MMC coefficients represent the within-industry relationship between multimarket contact and transaction costs. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.4. Robustness to Alternative MMC and Transaction-Cost Constructions

	Alternative MMC constructions			Alternative cost constructions			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	6M lookback	1M lookback	over # bonds	Up to 1 Week	Up to 1 Week	Equal-Weighted	Equal-Weighted
Avg. Dealer Overlap	-3.68***	-1.38***	-2.63	-5.43***		-5.01**	
	(1.38)	(0.49)	(2.23)	(2.04)		(2.14)	
Avg. Dealer Overlap, Weighted					-4.00**		-4.40***
					(1.56)		(1.54)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓
Obs	907,788	311,340	1,120,822	1,253,658	1,253,658	1,120,822	1,120,822
Adj.R^2	0.287	0.278	0.285	0.201	0.201	0.307	0.307

Notes: This table reports robustness tests using alternative definitions of multimarket contact, MMC, including shorter lookback windows and a normalization based on the number of bonds traded by each dealer, and alternative constructions of transaction costs, including a shorter benchmark window, up to one week, and equal-weighted monthly spreads. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Coefficients measure the association between MMC and the *Any-to-ID* transaction-cost measure under each construction. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.5. MMC and Dealer-to-Dealer Transaction Costs

	(1)	(2)	(3)	(4)
Avg. Dealer Overlap	-4.09***	-3.51***		
	(0.68)	(0.90)		
Avg. Dealer Overlap, Weighted			-2.86***	-2.27***
			(0.43)	(0.59)
Bond-Time FE	✓	✓	✓	✓
Dealer 1-Dealer 2-Time FE	✓	✓	✓	✓
Additional Controls		✓		✓
Obs	4,186,892	2,087,684	4,186,892	2,087,684
$Adj.R^2$	0.280	0.303	0.280	0.303

Notes: The unit of observation is a dealer-pair–bond–month. Multimarket contact, MMC, is measured as either average dealer overlap or revenue-weighted average dealer overlap; both measures are standardized within month to have mean zero and standard deviation one. All specifications include bond-by-month and dealer-pair-by-month fixed effects. Columns (2) and (4) additionally control for lagged dealer–bond trading volume and the number of comparable bonds traded in the month. Standard errors are reported in parentheses and are two-way clustered by dealer-pair and bond-by-year-by-month. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table E.6. Persistence in Dealers' Market Shares by Industry–Rating Segment

	Market share (t)		
	(1)	(2)	(3)
	All Bonds	IG Bonds	HY Bonds
Market share (t-1)	0.80***	0.80***	0.80***
	(0.02)	(0.02)	(0.02)
Obs	273,982	273,982	273,982
$Adj.R^2$	0.648	0.648	0.648

Notes: This table presents pooled OLS estimates of market-share persistence based on an unbalanced panel of dealer–industry–rating-tier–quarter observations. The dependent variable is dealer market share in a given industry–rating-tier–quarter. Columns (1)–(3) report pooled OLS estimates for all bonds, investment-grade bonds only, and high-yield bonds only, respectively. Standard errors are clustered at the dealer–industry–rating-tier level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

F CD–ID and ID–CD: Baseline Results, Subsample Analysis and Robustness

This appendix presents subsample and robustness results for the CD–ID and ID–CD transaction-cost specifications. These specifications evaluate the MMC–spread relationship using alternative benchmark constructions—customer–dealer prices when the dealer purchases (CD–ID) and interdealer prices when the dealer sells (ID–CD). All regressions follow the baseline empirical design from the main text and include the full set of bond-time and dealer-time fixed effects.

Table F.1. Baseline MMC–Spread Estimates: Alternative Spread Constructions

	Any–Any Cost				CD–ID Cost				ID–CD Cost			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Avg. Dealer Overlap	-7.65***	-6.91***			-12.01***	-11.04***			-10.76***	-10.58***		
	(2.61)	(2.47)			(2.60)	(2.43)			(3.35)	(3.30)		
Avg. Dealer Overlap, Weighted			-5.46***	-4.83***			-8.42***	-7.48***			-8.55***	-8.47***
			(1.88)	(1.83)			(1.97)	(1.86)			(2.30)	(2.28)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls		✓		✓		✓		✓		✓		✓
Obs	1,671,779	1,670,403	1,671,779	1,670,403	730,646	730,164	730,646	730,164	613,787	613,340	613,787	613,340
$Adj.R^2$	0.336	0.337	0.336	0.337	0.354	0.355	0.354	0.355	0.297	0.297	0.297	0.298

Notes: This table reports estimates of the MMC–spread relationship using alternative benchmark constructions that swap the market used to form the reference price. *CD–ID* benchmarks customer–dealer transaction prices to recent interdealer prices, while *ID–CD* benchmarks interdealer transaction prices to recent customer–dealer prices. The *Any–Any* measure benchmarks each transaction to recent prices from all counterparties. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.2. MMC and Spreads by Rating Tier: Alternative Spread Constructions

	Any–Any Cost				CD–ID Cost				ID–CD Cost			
	(1) IG	(2) HY	(3) IG	(4) HY	(5) IG	(6) HY	(7) IG	(8) HY	(9) IG	(10) HY	(11) IG	(12) HY
Avg. Dealer Overlap	-6.45***	-2.68			-9.44***	-8.86***			-10.37***	-1.33		
	(1.96)	(2.51)			(2.20)	(2.69)			(2.62)	(3.39)		
Avg. Dealer Overlap, Weighted			-3.72***	-2.38			-6.20***	-6.37***			-7.34***	-1.56
			(1.30)	(1.85)			(1.51)	(1.71)			(1.78)	(2.40)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs	1,252,918	412,942	1,252,918	412,942	573,157	153,091	573,157	153,091	486,390	123,862	486,390	123,862
$Adj.R^2$	0.352	0.327	0.352	0.327	0.380	0.301	0.380	0.301	0.325	0.245	0.325	0.245

Notes: This table reports estimates of the MMC–spread relationship separately for investment-grade and high-yield bonds using alternative benchmark constructions: *Any–Any*, *CD–ID*, and *ID–CD*. Splitting by rating tier allows us to assess whether multimarket contact matters differently in segments with systematically different trading frictions, including baseline spreads, adverse selection, and dealer interaction frequency. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.3. MMC and Spread by Remaining Maturity: Alternative Spread Constructions

	Any–Any Cost				CD–ID Cost				ID–CD Cost			
	(1) Avg. Overlap ≤5Y	(2) Avg. Overlap >5Y	(3) Weighted Overlap ≤5Y	(4) Weighted Overlap >5Y	(5) Avg. Overlap ≤5Y	(6) Avg. Overlap >5Y	(7) Weighted Overlap ≤5Y	(8) Weighted Overlap >5Y	(9) Avg. Overlap ≤5Y	(10) Avg. Overlap >5Y	(11) Weighted Overlap ≤5Y	(12) Weighted Overlap >5Y
Avg. Dealer Overlap	-5.88**	-6.18***			-8.72***	-10.11***			-8.85***	-4.15		
	(2.60)	(2.04)			(2.30)	(2.50)			(2.84)	(3.17)		
Avg. Dealer Overlap, Weighted			-4.81**	-4.46***			-5.66***	-6.99***			-7.39***	-3.41**
			(2.16)	(1.46)			(2.05)	(1.42)			(2.21)	(1.69)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs	831,745	833,406	831,745	833,406	415,039	310,427	415,039	310,427	348,786	260,969	348,786	260,969
$Adj.R^2$	0.376	0.324	0.377	0.324	0.369	0.334	0.368	0.334	0.275	0.304	0.275	0.304

Notes: This table reports MMC–spread estimates separately for bonds with at most five years remaining maturity and for bonds with more than five years remaining maturity, using the *Any–Any*, *CD–ID*, and *ID–CD* spread constructions. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.4. MMC–Spread Estimates by Issuer Industry and Spread Construction

	Any–Any Cost						CD–ID Cost						ID–CD Cost					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)
	Ind.	Fin.	Util.	Ind.	Fin.	Util.	Ind.	Fin.	Util.	Ind.	Fin.	Util.	Ind.	Fin.	Util.	Ind.	Fin.	Util.
Avg. Dealer Overlap	-5.62**	-9.58***	-8.80***				-9.65***	-13.09***	-12.64***				-8.16**	-15.35***	0.15			
	(2.28)	(3.39)	(3.03)				(2.45)	(3.29)	(4.49)				(3.46)	(3.98)	(5.61)			
Avg. Dealer Overlap, Weighted				-4.01**	-6.25**	-4.00*				-6.10***	-9.23***	-8.40**				-6.11***	-12.32***	-2.33
				(1.63)	(2.45)	(2.18)				(1.68)	(2.50)	(3.55)				(2.11)	(2.89)	(5.05)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs	1,041,522	524,745	64,124	1,041,522	524,745	64,124	435,613	252,687	21,883	435,613	252,687	21,883	362,634	217,878	16,365	362,634	217,878	16,365
Adj.R^2	0.328	0.380	0.332	0.328	0.380	0.332	0.346	0.390	0.398	0.345	0.390	0.398	0.295	0.319	0.309	0.295	0.319	0.309

Notes: This table reports MMC estimates separately for industrial, financial, and utility sectors under three spread constructions: *Any–Any*, which benchmarks each trade to recent prices from all counterparties; *CD–ID*, which benchmarks customer–dealer trades to recent interdealer prices; and *ID–CD*, which benchmarks interdealer trades to recent customer–dealer prices. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.5. Alternative MMC Measures under Alternative Spread Constructions

	Any–Any Cost			CD–ID Cost			ID–CD Cost		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Avg. Dealer Overlap (6M lookback)	-4.82**			-7.81***			-8.67***		
	(1.95)			(1.87)			(2.68)		
Avg. Dealer Overlap (1M lookback)		-1.35**			-2.13***			-1.90*	
		(0.67)			(0.70)			(1.01)	
Avg. Dealer Overlap (over # bonds)			-7.79**			-9.69***			-14.81***
			(3.22)			(2.48)			(3.51)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs	1,390,193	546,631	1,670,403	583,938	189,296	730,164	495,739	169,134	613,340
$Adj.R^2$	0.335	0.318	0.337	0.356	0.347	0.355	0.291	0.250	0.298

Notes: This table reports robustness tests using alternative constructions of the multimarket-contact, MMC, measure, including shorter lookback windows, 6 months and 1 month, and an alternative normalization based on the number of bonds traded by the dealer. Estimates are shown for the *Any–Any*, *CD–ID*, and *ID–CD* spread measures. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table F.6. Transaction-Cost Robustness under Alternative Spread Measures

	Any–Any Cost				CD–ID Cost				ID–CD Cost			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	1-week window		Equal-weighted		1-week window		Equal-weighted		1-week window		Equal-weighted	
Avg. Dealer Overlap	-7.34***		-6.72**		-10.97***		-10.24***		-10.94***		-9.69***	
	(2.49)		(2.66)		(2.51)		(2.29)		(3.23)		(3.21)	
Avg. Dealer Overlap, Weighted		-5.08***		-5.18***		-6.89***		-7.63***		-8.23***		-7.54***
		(1.82)		(1.90)		(1.84)		(1.81)		(2.16)		(2.18)
Bond-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Dealer-Time FE	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Additional Controls	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Obs	1,705,497	1,705,497	1,670,403	1,670,403	885,675	885,675	730,164	730,164	791,961	791,961	613,340	613,340
$Adj.R^2$	0.300	0.300	0.355	0.355	0.278	0.278	0.396	0.396	0.225	0.225	0.317	0.317

Notes: This table reports robustness to alternative constructions of transaction costs under alternative spread definitions. We vary the benchmark window used to construct the reference price, extending the lookback to up to one week, and the monthly aggregation scheme, using equal-weighted rather than volume-weighted averages. MMC is measured using either the unweighted average dealer overlap, *Avg. Dealer Overlap*, or the revenue-weighted measure, *Avg. Dealer Overlap, Weighted*, standardized each month to have mean zero and variance one. All specifications include the full set of fixed effects and controls used in Table 2. Standard errors are clustered at the dealer and bond–year–month levels. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.