Journal of Financial Economics 157 (2024) 103859

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/finec

When failure is an option: Fragile liquidity in over-the-counter markets ☆

Terrence Hendershott [a,*], Dan Li [b], Dmitry Livdan [a,e], Norman Schürhoff [c,d,e]

[a] Walter A. Haas School of Business, University of California, Berkeley, Berkeley, CA 94720, United States of America
[b] Board of Governors of the Federal Reserve System, Washington, DC, United States of America
[c] Faculty of Business and Economics, University of Lausanne, CH-1015, Lausanne, Switzerland
[d] Swiss Finance Institute, CH-8006, Zürich, Switzerland
[e] CEPR, London EC1V 0DX, UK

ARTICLE INFO

Dataset link: TCI Data (Original data)

JEL classification:
G12
G14
G24

Keywords:
Liquidity
Fragility
Cost of immediacy
Securities auctions
Over-the-counter markets
Structured products
Collateralized loan obligations

ABSTRACT

Markets can give false impressions of liquidity and stability if failed attempts to trade are ignored. For collateralized loan obligations, we quantify this bias by estimating the total cost of immediacy (TCI) which incorporates failure rates and failure costs. TCI is substantially higher than the observed cost, 0.3–3.8% versus 0.04–0.12% across credit-quality tranches because trade failures are frequent, failure costs are large, and failure costs and rates are correlated. TCI is almost double the realized gains from trade for low-rated tranches. Overall, auction-based over-the-counter markets become illiquid and fragile, especially during stressful periods for low-rated assets.

1. Introduction

The fundamental role of markets is to facilitate the exchange of goods, assets, and services in a frictionless and low-cost manner. If agents cannot trade or choose not to trade when costs are high relative to their outside options, then trading volume diminishes, welfare declines, and the cost of immediacy rises. Assessing the quality of market functioning is an empirical issue with a long tradition in both economics and finance starting with Demsetz (1968). In most markets, however, an unresolved problem is that only outcomes of completed trades are observed. The lack of data about failed trades biases observed market quality upward and trading cost downward as observed measures fail to account for the opportunity costs of not trading (Perold, 1988). This bias can give a false impression of liquidity and stability, especially in less transparent markets. These issues are particularly pertinent in financial markets where the 2007–2009 financial crisis and the 2020 pandemic have highlighted the importance of understanding and monitoring the fragility of liquidity. However, they apply to markets with search-and-match frictions like marriage markets, labor markets, housing markets, and international trade (Chade et al., 2017). Trade outcomes in these markets are strongly affected by the interaction between trading costs and participants' unobserved outside options.

We document frequent trade failures in the market for collateralized loan obligations (CLOs). CLOs are a primary source of funding for low-credit-quality corporate borrowers and have grown rapidly over the past decade. CLOs consist of repackaged portfolios of loans issued by low-rated, highly leveraged firms. CLOs have become popular due

☆ Nick Roussanov was the editor for this article. We thank Hank Bessembinder, Eric Budish, Thierry Foucault, Joel Hasbrouck, Burton Hollifield, Stacey Jacobsen, Pete Kyle, Albert Menkveld, Michael Schwert, Bart Yueshen, Haoxiang Zhu and participants at HEC Montreal, Southern Methodist University, University of Michigan, University of Rochester, 2020 INSEAD Finance Symposium, 2020 NBER Big Data Meeting, 2020 Microstructure Exchange, 2021 Search and Matching in Macro and Finance Virtual Seminar Series, and 2021 Western Finance Association meeting. An earlier draft of the paper was titled "True Cost of Immediacy". Hendershott gratefully acknowledges support from the Norwegian Finance Initiative, Norway. Schürhoff gratefully acknowledges research support from the Swiss Finance Institute, Switzerland and the Swiss National Science Foundation, Switzerland under SNF project #100018_192584, "Sustainable Financial Market Infrastructure: Towards Optimal Design and Regulation". The views of this paper do not represent the views of the Federal Reserve Board.

* Corresponding author.
E-mail addresses: hender@haas.berkeley.edu (T. Hendershott), dan.li@frb.gov (D. Li), livdan@haas.berkeley.edu (D. Livdan), norman.schuerhoff@unil.ch (N. Schürhoff).

https://doi.org/10.1016/j.jfineco.2024.103859
Received 14 February 2023; Received in revised form 26 April 2024; Accepted 29 April 2024
Available online 21 May 2024

to the reach for yield by asset managers in a low-interest-rate environment (Becker and Ivashina, 2015). CLOs hold roughly 60% of all leveraged loans outstanding and are important sources of funding for small and mid-sized businesses. Like many other structured products, CLOs trade in an OTC market that performed poorly during the 2007–2008 financial crisis (Dick-Nielsen et al., 2012; Friewald et al., 2012). We document that trade failures in CLOs occur more than 10% of the time and occur more often in lower-rated CLOs and stressful market conditions when failure rates exceed 50%. These frequent trade failures should raise concern among regulators and investors.

Centralized markets like stock exchanges have continuous, visible, and firm quotes, informing investors of the terms of trade in advance, thus limiting failed attempts to trade. Structured products and other financial securities trade in over-the-counter (OTC) markets, where there are typically no firm quotes and trade is infrequent,[1] and, therefore, trade failures are an option. To assess the liquidity and fragility of the CLO market, we estimate the opportunity cost of trade failure and incorporate it into the total cost of immediacy (TCI). We measure TCI by combining the direct intermediation cost, captured by the bid–ask spread, with the opportunity cost of trade failure. The cost of failing to trade is the expected difference between the price at which a buyer can purchase the asset minus the seller's expected payoff from continuing to hold the asset.

Auction and trade data from 2012–2020 allow us to infer the cost of failed attempts to trade and incorporate it into the cost of immediacy. Bids-wanted-in-competition (BWIC) is an auction-like email-based mechanism which existing holders predominantly use to sell their CLOs.[2] In a BWIC, an institutional seller contacts several securities dealers to solicit bids. Dealers decide whether to take the CLO into their inventory or search for buyers before responding. The investors then decide whether to sell to the dealer with the highest bid. While a large theoretical literature studies bilateral trading in OTC markets (Weill, 2020), little is known about auction-based trading in OTC financial markets. We study successful and failed BWICs using third-party collected auction data and trade reporting data.

Our estimates for TCI reveal substantial market fragility. The total costs of immediacy are significantly higher than the observed bid–ask spreads not only because trade failures are frequent but because failure costs are large and correlated with failure rates. Across our 2012–2020 sample period, the observed cost is 0.04/0.10/0.12% for senior/mezzanine/junior tranches while TCI is 0.31/1.32/3.79%. In the 2020 pandemic and other stressful periods when failure rates exceed 50% and failure costs exceed 10% for junior tranches, the observed cost of trading increases from 0.10% to 0.25% while TCI increases tenfold from about 1% to 10%. For high-rated tranches, TCI is low relative to sellers' gains from trade, which we define as the difference between the seller's reserve price and the highest bid. For low-rated tranches, TCI is almost double the gains from trade. When we estimate TCI for sellers with different reserve prices, TCI is more than 50% higher for sellers with low reserve prices than for sellers with the median reserve price. These findings suggest that sellers cannot trade when they have the greatest need to trade because there is insufficient liquidity even for the sellers with the lowest reserve prices, exacerbating market illiquidity and fragility.

The total costs of immediacy incorporate the observable cost of immediacy, often defined as the cost to an investor of simultaneously buying and selling an asset. In addition, if investors choose not to trade when those prices are not good, estimates of the cost of immediacy derived from successful transactions are downward biased. To capture the investors' opportunity costs when failing to trade, TCI must incorporate the value of continuing to hold the asset when a BWIC fails. We only observe the highest accepted bids, but not bids when auctions fail to trade. Using successful bids to estimate the sellers' reserve prices, which are measures of their outside options, TCI can be written as the expected bid–ask spread on observed trades plus the probability of BWIC failure times the cost of trade failure; where the cost of trade failure is the expected highest bid conditional on a trade minus the reserve price. The reserve price is the only component of TCI that is not directly observable in the data.

We estimate the reserve price using the theoretical link between the seller's unobserved outside option and observed quantiles of the distribution of accepted bids. Endogenous trade failures lead to a selection bias in observed accepted bids thus complicating the choice of the correct quantile for estimating sellers' reserve prices. We show that the selection bias must be accounted for by rotating the best bids' quantiles by an endogenous degree that depends on the amount of selection. We determine the endogenous rotation to recover sellers' reserve prices via the copula of errors in the outcome equation for best bids and errors in the selection equation. We call this the *optimal quantile rotation*. We provide the optimal rotation in closed form and show that it can be expressed as a function of the selection score and copula only.

Because TCI is sellers' loss relative to the frictionless sale price, our approach also provides measures of gains from trade. The realized seller's gains from trade are the difference between the frictionless sale price and the seller's reserve price minus TCI, which equals the difference between the expected best bid and the reserve price multiplied by the probability of trade. Over our sample period, for high-rated tranches, TCI is low relative to sellers' gains from trade (1.40/2.43% for senior/mezzanine tranches). For low-rated tranches, TCI is almost double the gains from trade (2.05% for junior tranches). In addition, TCI for sellers with a low reserve price (5.58%) is more than 50% higher than for sellers with a median reserve price (3.64%).

TCI is the bid–ask spread plus the failure rate times the failure cost. In the cross-section, for high-rated tranches, TCI moderately exceeds bid–ask spreads, but for low-rated tranches TCI substantially exceeds bid–ask spreads. In the time series, because failure costs and failure rates co-move with both increasing in stressful periods, this amplification causes TCI to spike significantly more than bid–ask spreads in stress periods (0.20/0.88/2.73% for senior/mezz/junior when dealer CDS spreads are low vs. 0.47/2.02/4.98% when dealer CDS spreads are high). TCI spiking so dramatically emphasizes how the selection bias in observed costs can cause them to dramatically underestimate illiquidity in crises when trade failures diminish realized gains from trade. These results demonstrate substantial heterogeneity in market functioning across investors, assets, and time.

To illustrate how TCI increases with the risk of the asset and spikes in market stress and how trade failures diminish gains from trade, Fig. 1 plots weekly average realized bid–ask spreads (dashed red lines) and TCI (solid blue lines) for different CLO credit tranches during the 2020 pandemic. The patterns are representative of other stress periods. Panel A shows that realized bid–ask spreads for senior tranche (AAA-rated) CLOs were only a few basis points (bps) at the start of 2020. Beginning in February 2020 spreads rose to about 10 bps. As the pandemic hit in March 2020, spreads widened to 20 bps. The realized bid–ask spreads for mezzanine tranches follow a similar pattern: spreads were about 5 bps at the beginning of the year, rose to 15 bps in February, and were 40 bps by the end of March 2020.[3] Interestingly, realized bid–ask spreads for below-investment-grade junior tranches rose more modestly in March. Panel A compares our TCI measure directly to realized bid–ask spreads. The magnitude of the increases in TCI dwarfs that of

[1] See Duffie et al. (2005), Lagos and Rocheteau (2007, 2009), Atkeson et al. (2015), Malamud and Rostek (2017), Hugonnier et al. (2020), and others.

[2] Around 40% of customer-sell-to-dealer trades in the sample of US CLOs can be mapped to a BWIC auction in our data, which suggests that at least 40% of investors sells are through auctions, instead of direct bilateral negotiations with a dealer. The dealer-sell-to-customer leg is typically done through search. Offers-wanted-in-competition (OWICs) are rare, likely reflecting the difficulty in pricing such securities for investors relative to dealers.

[3] Kargar et al. (2020) and O'Hara and Zhou (2020) study liquidity for corporate bonds through this time period. They find that bid–ask spreads widen to three to four times their pre-pandemic level, with the increase occurring in trades where the dealers act as principals.

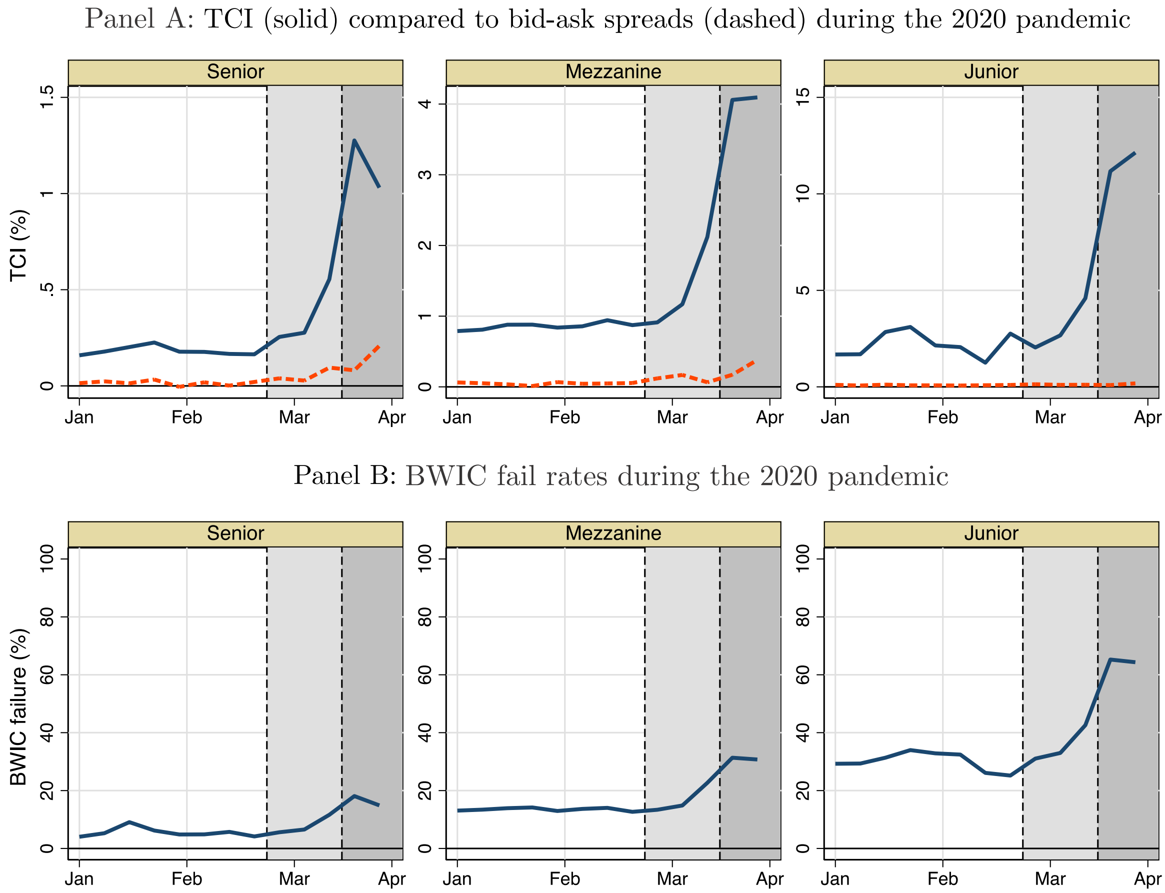

Fig. 1. TCI and BWIC fail rates during 2020 pandemic
The figure documents TCI (solid blue lines) compared to bid-ask spreads (dashed red lines) in the CLO market (Panel A) and weekly BWIC fail rates (Panel B) during the onset of the 2020 pandemic between January 2020 and March 2020, split by tranche into the senior tranche rated AAA, mezzanine tranche rated AA–BBB, and junior tranche rated BB–equity. The lightly shaded area indicates early stages of the pandemic between 3rd week of February and 2nd week of March (before the first US shelter-in-place orders), while the darker shaded area indicates later stages of the pandemic between 2nd and 4th weeks of March, i.e., after first US shelter-in-place orders have been implemented. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the spread. For the senior tranches, TCI exhibits similar dynamics as the spread until late February, but then TCI spikes to over 1%. For mezzanine tranches, TCI is above the spread at the start then increases to 4% by April. For the junior tranches, TCI starts around 2%–3%, then increases to over 10%. Overall, TCI suggests more substantial market deterioration with the onset of the pandemic. We show later in the paper that the dramatic increase in TCI is due to an increase in the BWIC failure rate as depicted in Panel B as well as an increase in the failure cost. Panel B shows that failure rates for senior (mezzanine) CLOs increased from less than 10% (20%) to over 20% (30%) during the onset of the 2020 pandemic, while the failure rate for junior rose from 30% in early 2020 and exceeded 60% by late March 2020.

Literature overview. Measuring the opportunity costs of not trading has been recognized as important in many markets (Perold (1988)). In equity markets, because institutional trading data from Plexus and Ancerno only have realized trading volume, limited progress has been made in measuring how the opportunity costs bias execution costs for institutions.[4] Our estimation of reserve prices from failed trades is most valuable in OTC markets where quotes and prices are not continuously observable. A number of approaches have been used to calculate transaction costs in OTC markets.[5] TCI extends these approaches by incorporating the costs of not trading.[6]

[4] Harris and Hasbrouck (1996) examine the opportunity costs for individual orders by comparing the performance of limit versus market orders on the New York Stock Exchange. Limit orders impose opportunity costs due to delay and failure to execute. Harris and Hasbrouck (1996) impute the cost of a failed limit order execution as being the cost if that limit order is converted to a market order. The Harris–Hasbrouck approach shows how a market price estimate could be used to estimate reserve prices. Bessembinder et al. (2009) extend Harris and Hasbrouck (1996) to examine opportunity costs for hidden limit orders. However, with the advent of high-frequency trading, limit orders are more likely to be replaced by repriced limit orders than converted into market orders. Hence, the imputed-fills approach is difficult to apply in recent periods.

[5] The simplest approach is to compare roughly contemporaneous buy and sell prices of the same security to impute a spread. Hong and Warga (2000) follow this approach by subtracting the average sale price from the average buy price each day when there is both a buy and a sell. While imputed costs are easy to calculate, infrequent trading limits the amount of usable data. To allow buy and sell trades to occur on different days, Harris and Piwowar (2006) and Bessembinder et al. (2006) calculate trading costs using regressions of the change in price between transactions on the change in the trade sign. Similar to our approach to measuring effective bid–ask spreads, Green et al. (2007) and Li and Schürhoff (2019) identify matching buys and sells to calculate round-trip trading costs.

[6] Bessembinder et al. (2018) find evidence consistent with a decrease in customers' ability to trade and an increase in the opportunity costs of failed trades. However, the paper does not directly measure either of these. Dick-Nielsen et al. (2012) and Friewald et al. (2012) show that corporate bond liquidity decreased during the 2007–2009 financial crisis. Our results show that bid–ask spreads underestimate the cost of immediacy during crises

We estimate the bias in transaction costs due to the opportunity cost of non-trading by modeling the selection in which trades occur. Another approach to estimating the cost of immediacy is to identify investors always needing to complete a trade regardless of how attractive a price they receive. This eliminates the selection problem as the failure rate is zero and the observed expected best bid conditional on trade equals the unconditional best bid. Ellul et al. (2011) study insurance companies' forced bond selling due to downgrades and regulatory constraints. Dick-Nielsen and Rossi (2019) use index exclusions to measure the cost of uninformed index trackers demanding immediacy. Indexers likely have the lowest reserve prices because of the cost of a tracking error, so they always trade. The approach using specific events enables unbiased measurement of the costs for certain investors. However, this approach cannot measure the opportunity costs for investors who choose not to trade because, by construction, the cost of immediacy is measured only for investors who always trade. Our approach can measure the TCI for investors with any reserve price, rather than for only investors with the lowest reserve price, and it shows that TCI varies dramatically across investors, assets, and time.

Auctions have been studied in many contexts, including financial markets (for example Hortaçsu et al., 2018). The empirical literature typically focuses on revenue to the seller/issuer and bidders' strategies and surplus. Because TCI is a fair-value benchmark minus seller revenue, we also analyze seller revenues. We further analyze when trade fails to occur to estimate sellers' reserve prices.[7] These results enable the calculation of overall seller gains from trade as well as the heterogeneity of TCI and gains across sellers with different reserve prices. Because sellers' reserve prices are not observed, our approach must account for unobserved heterogeneity and selection.[8]

Selection-adjusted quantiles have been used to estimate unconditional outcomes in a variety of settings.[9] For example, Arellano and Bonhomme (2017) determine a selection-adjusted rotation to recover a fixed, exogenous quantile of the unconditional wage distribution to adjust for endogenous unemployment. We move beyond the estimation conditional distribution of the outcome variable to estimate the unobserved unconditional distribution of the outside option. In our setting, we prove that the distribution of best bids can be used to recover the sellers' reserve prices. Our optimal quantile rotation is endogenously determined, observation specific, can be derived in closed form and can be applied in other markets where reserve prices, reservation values, or outside options are unobserved.

2. Background: CLOs and how they trade

2.1. What are CLOs?

CLO securities are asset-backed securities (ABS) where the underlying collateral is a pool of syndicated loans issued by indebted low-rated corporations. The issuers of CLO securities are special purpose vehicles or trusts that own the pool of collateral and finance the holding by issuing debt and equity claims that are of different seniority to the cash flow from the collateral. The collateral for CLOs is mostly leveraged loans that Standard & Poor's defines as senior secured bank loans rated BB+ or lower (i.e., below investment grade) or yielding at least 125 basis points above a benchmark interest rate (typically LIBOR or EURIBOR) and secured by a first or second lien.[10]

CLOs are complicated structures that consist of many debt tranches and a sliver of equity. The tranches are ranked highest to lowest in order of credit quality and cash flow priority and, thus, lowest to highest in order of riskiness. Although leveraged loans themselves are mostly rated below investment grade, a large fraction of CLO tranches are rated investment grade, benefiting from diversification and subordination of cash flows. CLOs aim to capture the excess spread between the portfolio of leveraged bank loans (assets) and the classes of CLO debt (liabilities), with the equity investors receiving any excess cash flows after the debt investors are paid in full. The popularity of CLOs has grown tremendously since the financial crisis, both in the US and in Europe.

Fig. 2 shows the amount of US CLO securities outstanding over the sample period 2012–2020, split by tranche into the senior tranche rated AAA, mezzanine tranche rated AA–BBB, and junior tranche rated BB–equity.[11] The market more than doubled since 2012. This trend partly reflects the growth of the leveraged loan market as issuers took advantage of investors' reaching-for-yield demand under a generally low-interest rate environment. The amount of US leveraged loans outstanding was just over one trillion dollars as of late 2019. CLOs currently hold about two-thirds of the leveraged loan universe, and they are expected to bear the brunt of losses in leveraged loans when the business cycle turns. The three grey shaded areas in Fig. 2 correspond to periods of credit market stress: the 2012 European debt crisis, the 2015–16 energy sector-related credit stress, and the 2020 COVID-19 pandemic.

Ownership of CLOs varies by tranche. The least risky, most senior tranches are mainly owned by insurance companies (which favor income-producing investments) as well as banks (which need high-quality capital to meet regulatory requirements), particularly foreign

thus suggesting that illiquidity may have been even worse than previously measured.

[7] Hendershott and Madhavan (2015) study corporate bond trading via traditional OTC search and request for quote (RFQ) auctions, which are similar to BWICs. Corporate-bond RFQs fail often and costs are higher for trades that occur after a failed auction. However, the paper does not estimate the costs for attempted, yet failed trades. The cost of failed auctions is higher in CLOs than corporate bonds for several reasons: Corporate-bond RFQs are shorter lived (5–10mins) than BWICs; corporate bonds trade more frequently; the number of corporate-bond dealers is higher; and a larger share of customer sells for CLOs are done through BWICs than via RFQs in corporate bonds. Riggs et al. (2020) study RFQs and bilateral trades in the index credit default swaps market. The vast majority of index CDS inquiries resulted in trades, so the cost of failure is less relevant in their setting. Allen and Wittwer (2023) examine trading using data on OTC and RFQ trading in Canadian government bonds without examining failures.

[8] Haile and Kitamura (2019) show that accounting for unobserved heterogeneity that affects bidders' valuations is important, especially in first-price auctions. To address this issue, the control function approach conditions on observables, such as variation in the number of bidders (Campo et al., 2003) or seller's reserve price if it is known (Roberts, 2014), to implicitly condition on the unobservables. Krasnokutskaya (2011) develops a multiple measurement approach showing how independent noisy measurements, e.g., data on multiple bids in the independent private value model, can be used to back out a latent variable. Cassola et al. (2013) study multi-unit ECB repo auctions, and Hortaçsu et al. (2018) study multi-unit Treasury auctions where they address unobserved heterogeneity through a resampling method that separately uses bids from each auction instead of pooling bidding data across auctions. However, accounting for observed heterogeneity in their setting requires slicing bidders into groups which lowers the number of observations available for resampling. In our setting, we pool data across auctions in a quantile regression setting and develop a moment-based approach using dealers' private inventory information as an instrumental variable to capture unobserved heterogeneity that affects sellers' reserves.

[9] Koenker and Bassett (1978) pioneered the quantile regression approach. Chernozhukov and Hansen (2005) develop an instrumental variables approach to quantile regression. Machado and Silva (2019), among others, show how quantiles can be estimated via moment restrictions. Gimenes and Guerre (2021) apply quantile regression in an auction setting.

[10] Several characteristics make leveraged loans particularly suitable for securitizations: They pay interest on a consistent monthly or quarterly basis; they have a reasonably active secondary market; they have a high recovery rate in the event of default historically; and they originate from a large, diversified group of issuers.

[11] The growth in the AAA tranche in 2012–2014 reflects a substantial amount of AA rated tranches being upgraded to AAA.

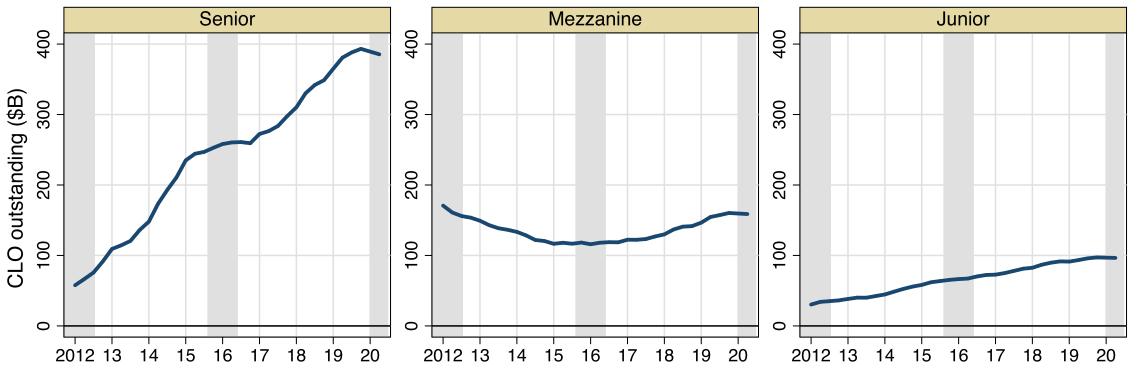

Fig. 2. CLO amount outstanding by credit risk
The figure plots quarterly CLO amount outstanding in $ billions over the sample period 2012–2020, split by tranche into the senior tranche rated AAA, mezzanine tranche rated AA–BBB, and junior tranche rated BB–equity. Grey areas indicate stress periods in the credit market: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic.

banks.[12] The equity tranche is the riskiest, offers potential upside and a degree of control, and appeals to a wider universe of investors. According to the BIS, banks hold more than 50% senior CLOs and none of the equity, which are held primarily by hedge funds, structured credit funds, and other asset managers, including CLO managers.[13]

2.2. How do CLOs trade?

Much like corporate bonds, municipal bonds, and other debt instruments, CLO securities are traded over the counter with dealers intermediating trade. Existing CLO investors sell their holdings through an auction where many dealers can simultaneously bid on the security. Such auctions are called bids-wanted-in-competition, or BWICs. Investors choose a dealer to run the BWIC. The dealer sends out the BWIC to a list of other dealers to solicit bids before the bidding deadline.[14] Typically, the organizing dealer decides which other dealers are contacted and collects bids by email. The dealer communicates with the seller about whether the seller accepts the highest bid or lets the BWIC fail. The email mechanism is more flexible than fully automated trading, so the exact process is not fully standardized.

BWIC auctions operate as first-price sealed bid auctions, in that the winning dealer pays her bid. After the auction is completed, only the seller and the winner know the transaction price. The only information disseminated back to the market and to losing dealers is the second highest bid in the auction known as the "cover". In the case of not enough bids or bids below the seller's reserve price, the auction is disclosed as "DNT (Did Not Trade)". Because the seller is unable to exit his/her position in the security, we classify such an outcome as "auction failure". Sellers' reserve prices are private and are not announced during or after the auction. As intermediaries, dealers mainly bid in auctions with the intention of resale. Dealers' profits are the difference between the resale price and their winning bid net of search cost and inventory holding cost.

CLOs are traded in a unique hybrid fashion where the dealer-buy leg is conducted through auction and the dealer-sell leg is conducted through bilateral search. Shorting is rare in this market, likely due to the idiosyncratic nature of CLOs. Therefore, dealer intermediation almost always starts with a dealer buy followed by a dealer sell, rather than the opposite. Many other types of structured products such as auto-ABS, credit card ABS, and CMBS/RMBS are traded in a similar hybrid fashion. Our analysis is likely applicable to those securities as well.

3. Data

Our main data sources combine BWIC auction data from Creditflux with regulatory transaction data from FINRA.

3.1. Auction data

Our BWIC auction data are from Creditflux, a commercial provider of granular data on CLOs and other structured products. For each auction, we have information about the CUSIP, the cover (second highest bid), and a flag that indicates if the auction failed. We limit our sample to the period from 2012 to March 2020, when TRACE reporting for structured products was available and when the number of auctions was large enough for our estimation. The sample includes the volatile period at the onset of the COVID-19 pandemic. We adopt a couple of filters for the sample, including requiring the CLO securities to be "US CLOs" denominated in USD. We keep only CUSIPs that appear in the TRACE master file. The majority of BWICs correspond to CUSIPs that are in the TRACE master file. The Appendix reports statistics on the cleaning filters that we have applied to the CLO BWIC data. After cleaning, we are left with a sample of 33,408 BWICs.

Additional information about the CLOs such as the amount outstanding and dates of issuance are from Bloomberg. Leveraged loan market spreads from JP Markets capture the CLO market condition. Information about the CLO dealer CDS spread is from Markit. Dealer CDS spread captures CLO dealers' funding conditions. There are about 20 active dealers in the CLO market. We use the average CDS spread of all primary dealers to proxy for the funding costs of dealers in the CLO market.

3.2. Trade data

We supplement the auction data with post-trade reporting data to infer the best bids (final transaction prices) in BWIC auctions as well as dealers' resale prices. The SEC started requiring dealers of structured

[12] The Volcker Rule prevents banking entities from holding ownership interests in "covered funds". Many CLOs, especially those issued prior to the publication of the final rule, have features that make them fall into the category of "covered funds". As a result, banking entities would not be permitted to hold an "ownership interest" in those CLOs. An exemption to the rule applies to CLOs that hold only loans. CLOs issued since 2014 (CLO 3.0) typically hold only loans and can be held by banks.

[13] See https://www.bis.org/publ/qtrpdf/r_qt1909w.htm. Cordell et al. (2023) examine CLO and manager performance. The majority of outstanding CLOs are actively managed. Managers of these CLOs play an important role in selecting loans for the portfolio, actively monitoring the quality of the collateral and trading in and out of loans.

[14] In our sample, there are roughly 20 dealers who regularly win BWICs. See Atkeson et al. (2015) for a theoretical analysis of dealer entry and exit in OTC markets.

products, including CLOs, to report trades in May 2011. Reported trades are not disseminated to the public. We match FINRA supervisory TRACE to the auction data to infer winning bids in auctions. To the extent that some residual clients sell directly to their relationship dealers, instead of conducting auctions, we observe those prices as well.

There are roughly 290K observations from TRACE for securities labeled as CDO/CLOs. Roughly 190K of these trades are for CUSIPs in the auction data. Roughly 40% of TRACE customer-to-dealer trades of CLOs can be matched to our auction data on a CUSIP-date basis. About 85% of our successful BWICs can be uniquely matched to a TRACE transaction that is customer-buy-from-dealer of the same CLO security on the same day. The Appendix reports statistics on the number of transactions that can be merged between TRACE and Creditflux.

CLOs trade very infrequently. A typical CLO in our sample only trades 11 times (roughly 4 dealer buys, 6 dealer sells, 1 interdealer) over the entire span of our sample. The trading frequency also varies by the seniority of the tranches, with a typical (median) AAA-rated tranche trading 15 times in the sample and a typical equity tranche trading 9 times.

Another notable observation is that relative to the municipal bond and corporate bond markets, inter-dealer trades are not as prevalent in the CLO market. Roughly speaking, there is only 1/3 of an interdealer trade for every dealer-buy trade. The comparable number is 1.7 for the municipal bond market in 2018, 1.8 for investment-grade corporate bonds, and 1.4 for high-yield corporate bonds in 2019. This is perhaps not surprising given the competitive nature of the customer-to-dealer leg of the transaction. In the empirical OTC market literature, a common way to measure transaction costs or liquidity of infrequently traded fixed-income securities is through effective bid–ask spread, or realized return for dealers, which calculates the difference between dealer buy prices and dealer sell prices. The effective bid–ask spread can be viewed as a measure of compensation for dealers' intermediation services.

Following Li and Schürhoff (2019), which builds on Green et al. (2007), we form dealer roundtrips by matching trades of the same CLO security that are dealer-buy-from-customer to the dealer-sell-to-customer trade by the same dealer. Roughly 46K roundtrips can be found for the sample of CLO securities in our auction data. Over 95% of these roundtrips have only one dealer between the selling customer and the buying customer (CDC roundtrips, as in Li and Schürhoff, 2019). Roughly 17K of these CDC roundtrips can be uniquely matched to a successful auction in our BWIC data, where the first leg of the roundtrip happens on the same day of the auction. These roundtrips correspond to about 60% of all successful auctions. Approximately half of the matched roundtrips occur on the same day, about 11% have an inventory time of one day; 16% of the roundtrips have an inventory time between one day and seven days; the rest of the roundtrips have an inventory time of up to 300 days. For effective bid–ask spread, we restrict the sample to roundtrips no more than 1 day apart, to limit the effect of mark-to-market price changes in the price difference between the buy leg and sell leg.

Table 1 provides summary statistics of key variables in our sample. We split the BWICs into auctions of senior, mezzanine, and junior CLO tranches. All of the Senior tranche CLOs are rated AAA by definition. Mezzanine tranche CLOs are roughly evenly split among AA, A, and BBB rated. Junior tranche CLOs are all below investment grade. The majority of the junior tranche CLOs are rated BB and only 16% of the junior tranche auctions are for the equity CLO tranche, underscoring how unlikely it is for an equity CLO tranche to trade in the secondary market. Depending on when the CLO was issued, each CLO is classified into vintage 1.0, 2.0, or 3.0. CLO 1.0s were issued before the 2007 financial crisis. CLO 2.0s were issued between 2008 and 2013, and are considered to have a robust structure compared to securities from earlier vintage. CLO 3.0s were issued after 2014, with most of them allowing only loans (no bonds) in the collateral pool, so as to be "Volcker" compliant. CLO 3.0s account for over half of our auctions across different tranches.

Because CLOs are held by institutional investors, they trade in large blocks. The median size of a trade (auction) in our sample is around $3 million. The median CLO security in our sample has a par outstanding value of $222/$24/$17 million for the senior/mezzanine/junior tranche.

Effective bid–ask spreads are small in our sample. Median bid–ask spreads are 2 bps/5 bps/9 bps of the CLO face value for the senior/mezzanine/junior tranche. These are a little smaller than the round-trip costs of corporate bonds with comparable ratings. This may be due to the fact that trade sizes for CLOs are larger than the average size of an institutional trade for a comparable corporate bond and the bid–ask spread tends to decline with trade size for fixed-income securities.

Table 1 also provides the average fraction of failed BWICs. On average, only 7% of the senior tranche CLO BWICs fail and 16% of BWICs in the mezzanine tranche, while the junior tranche BWICs fail 30% of the time.

3.3. CLO trading volume and bid–ask spreads

Fig. 3 reports the CLO trading volume (Panel A), CLO realized bid–ask spread (Panel B), and market and dealers' funding conditions (Panel C) over the 2012–2020 sample period. Results reported in Panels A and B are split by tranche. Grey areas indicate periods of market stress: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic. All data are quarterly.

Panel A of Fig. 3 shows CLO trading volume. CLO trading volume is consistent with the amount of CLO securities outstanding from Fig. 2. Senior and junior amounts outstanding are comparable to each other in 2012 when the mezzanine tranche was approximately twice the size of the senior tranche. Consequently, senior and junior CLO tranches had lower trading volume in 2012 than the mezzanine tranche. Trading volume seems to be a leading indicator for credit stress events for all CLO tranches. Trading volume tends to fall ahead of credit stresses and recovers by the end of a stress period.

Panel B of Fig. 3 plots the realized bid–ask spread. As discussed in Section 3, we measure bid–ask spreads on successful BWICs as a difference between the best bid submitted by dealers in the BWIC and the ask price that the CLO buyer pays to the winning dealer, as a percent of the par value of the trade. We restrict the sample to roundtrips where the dealer sell-leg and dealer buy-leg are no longer than one day apart. When the legs are further apart market movements start to contaminate the measurement of bid–ask spreads. Realized bid–ask spreads are low and equal to 5 bps/10 bps/12 bps of the CLO face value for senior/mezzanine/junior tranches. Spreads tend to fluctuate with market-wide credit risk.

Panel C of Fig. 3 shows the leveraged loan spread (left plot) and dealer CDS spread (right plot). The leveraged loan spread fluctuates between 3% to 5.3% during most of the sample period, and rises to almost 7% in March 2020, during the COVID-19 pandemic. Dealer CDS spreads follow the same time trend as the leveraged loan spread. The average dealer CDS spread was over 2% during the European crisis in 2012 and is mostly below 1% after 2013.

Appendix B provides a more formal analysis of bid–ask spreads in CLOs. We examine the determinants of effective bid–ask spreads using multivariate regressions. Table A.1 reports the parameter estimates. All results are split by CLO tranche. Overall, effective bid–ask spreads do not vary much within a quarter for CLOs with the same vintage and credit rating. When bid–ask spreads do vary within a quarter, the variation is mainly due to major economy-wide credit events affecting the riskiness of leveraged loans held by the CLOs. The next section investigates BWIC failures.

Table 1
Summary statistics.
The table reports summary statistics for BWICs in all CLOs over the sample period 2012–2020 for different sample splits. Panel A reports results for the senior tranche rated AAA, Panel B for the mezzanine tranche rated AA–BBB, and Panel C for the junior tranche rated BB–equity.

Variable	Mean	S.D.	5%	25%	50%	75%	95%
Panel A: Senior tranche (AAA, I = 9000 BWICs)							
CLO vintage: 1.0/2.0/3.0	15%/11%/74%						
CLO rating: AAA	100%						
BWIC fail rate	7%						
Par value of trade ($M)	2.77	4.16	0.25	1.00	2.80	8.00	27.25
CLO issue size ($M)	228.24	2.45	30.00	225.72	302.50	363.05	492.75
Bid–ask spread (%)	0.04	0.08	0.00	0.01	0.02	0.05	0.13
Panel B: Mezzanine tranche (AA–BBB, I = 12,955 BWICs)							
CLO vintage: 1.0/2.0/3.0	24%/16%/60%						
CLO rating: AA/A/BBB	27%/31%/42%						
BWIC fail rate	16%						
Par value of trade ($M)	2.48	2.69	0.49	1.25	2.75	5.00	11.00
CLO issue size ($M)	33.94	1.74	15.95	24.50	32.00	46.40	78.00
Bid–ask spread (%)	0.10	0.29	0.00	0.02	0.05	0.12	0.31
Panel C: Junior tranche (BB–equity, I = 11,453 BWICs)							
CLO vintage: 1.0/2.0/3.0	12%/29%/58%						
CLO rating: BB/B/Equity	70%/14%/16%						
BWIC fail rate	30%						
Par value of trade ($M)	2.72	2.32	0.60	1.70	3.00	5.00	9.80
CLO issue size ($M)	22.53	1.91	8.40	17.00	23.00	31.75	56.50
Bid–ask spread (%)	0.12	0.19	0.01	0.02	0.09	0.13	0.38

4. Trade failures

This section explores why failure rates on BWICs are large and how they behave prior to and during periods of market stress. The choice of explanatory variables, X, that we use in our analysis is important as they explain variation in CLO failure rates by observable characteristics and market conditions. Therefore, X includes the par value of trade (log of par), the CLO issue (log of the amount outstanding, vintage, and rating), the JPM leveraged loan spread, and dealer CDS spread to capture trade, market, and dealer conditions, respectively. Dealer CLO inventory is an additional characteristic included in X. Data on the dealer inventories are from the New York Fed Primary Dealer Statistics. The data is on the aggregate net position of all primary dealers in "Other Asset-Backed Securities" as these include dealers' CLO holdings, but exclude credit card, student loan, and automobile loan-backed ABS securities. Finally, we add credit rating splits for mezzanine and junior tranches.

The number of successful and the number of failed trades are natural liquidity measures. While the number of successful trades, or trading volume, has been extensively used in the literature, our paper is the first to perform a comprehensive study of failed trades. Trading volume alone may be a downward biased measure of illiquidity as it does not capture the lost trade value due to failure to strike a deal.

Panel A of Fig. 4 documents quarterly BWIC activity (solid line) and trading volume (dashed line) in the CLO market. Panel B of Fig. 4 plots the quarterly failure rate in percent, where the solid line is actual and the dashed line is predicted using probit regression (1). Each plot covers the 2012–2020 sample period, and both panels are split by tranche. Panel A shows that the number of BWICs increases over time across all tranches, consistent with the growth of the market. The number of BWICs and trades tends to decline during credit stress events and the BWIC failure rate tends to increase during stress periods. The average failure rate is different across tranches. It ranges from 0% to 15% for the senior tranche, 5%–30% for the mezzanine tranche, and 8%–60% for the junior tranche.

It is notable that the number of BWICs and failure rate seem to be leading indicators of market stress. This is particularly notable during the 2015–2016 stress episode when the number of BWICs dropped substantially early in the stress period and the failure rate rose sharply. This is possibly due to CLO sellers being slow to adjust their reserve prices in response to sharply declining demand. Sellers catch up with the declining demand by either reducing their reserve prices or/and reducing the supply by not even attempting to trade.

Table 2 compares BWIC fail rate using sample splits based on different explanatory variables. The Low/Medium/High columns indicate values of the varied characteristic in its bottom/middle/top tercile. Panels A/B/C report results for the senior/mezzanine/junior tranche, respectively.

For the senior tranche, larger trades fail more often with fail rates equal to 3%/7%/12% for Low/Medium/High trade sizes. The relation between BWIC fail rate and CLO issue size is weak—BWICs of smaller CLOs fail a bit more at 9%, while medium and large CLOs both fail about 7% of the time. CLO 3.0 BWICs fail slightly less than CLO 1.0 and 2.0 BWICs, i.e., 7% compared to 10%. The BWIC fail rate increases with the dealer CDS spread, 5%/8%/9% for Low/Medium/High CDS spreads. This suggests that dealers may bid higher when the cost of funding is lower. BWIC fail rates also increase with the dealer CLO inventory position, 5%/5%/11% for Low/Medium/High, as dealers are more reluctant to buy when inventories are larger.

For the mezzanine tranche, BWIC fail rates vary a lot with trade, market, and dealer conditions. Similar to the senior tranche, BWICs with a larger par amount fail more often. Fail rates increase with the riskiness of the tranches, with rates equal to 12%/16%/20% for AA/A/BBB ratings. BWICs of CLOs with a larger issue size fail less often. This is potentially due to the fact that larger issues may have a wider investor base, making it easier to find buyers. BWIC fail rates are higher when the leveraged loan spread is higher, when dealers' funding costs are higher, and when dealers' inventories are larger. The qualitative nature of all these relations remains the same for the junior tranche as Panel C demonstrates. In summary, unlike the effective bid–ask spread, BWIC fail rates vary significantly with trade, market, and dealer conditions.

How costly a failed BWIC is for the seller depends on whether the asset can be subsequently sold at a good price and whether the risk of failure can be hedged. Table 3 reports how often failed auctions are followed by another BWIC trade (Panel A) or an OTC trade (Panel B) over different time horizons: on the same day, within one day, and within 10 days of the failed auction. The probability of BWIC being

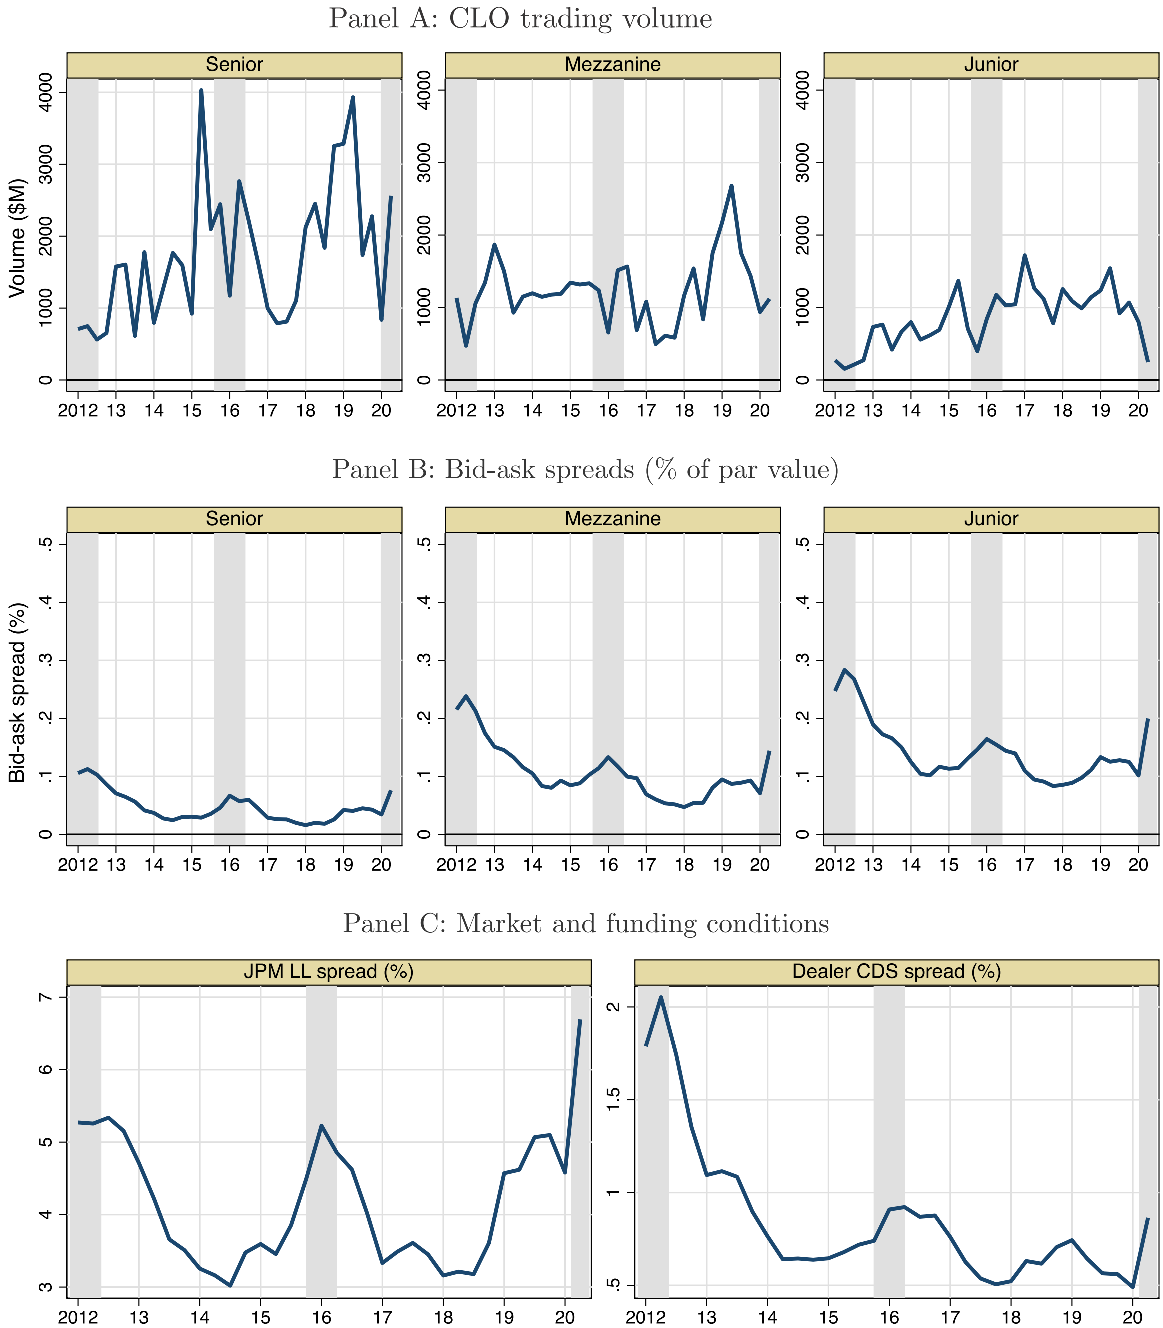

Fig. 3. CLO trading volume and bid–ask spreads
The figure documents quarterly trading volume (Panel A) and bid–ask spreads (Panel B) in the CLO market over the sample period 2012–2020, split by tranche into the senior tranche rated AAA, mezzanine tranche rated AA–BBB, and junior tranche rated BB–equity. Grey areas indicate periods of market stress: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic. The figure also documents market and funding conditions in the CLO market (Panel C) over the sample period 2012–2020. The left plot shows the evolution of the JPM leverage loan spread index and the right plot shows the evolution of the average CDS spread on all primary dealers active in the CLO market.

followed by BWIC is less than 0.1% on the same day and it is no greater than 1.2% over the next day for all CLO tranches. It increases with window size but even over 10 days does not exceed 3.9% on average. Across tranches, AAA has the highest chance of rerunning a BWIC, and Equity has the lowest. When we match on trade size, the repeat BWIC probability drops to less than 1%. Panel B shows that BWICs are more likely to be followed by an OTC trade than another BWIC.[15] Panel B of Fig. 4 shows failure rates spiking in times of stress. Table 2 shows that leveraged loan spreads and dealer CDS spreads are important determinants of BWIC fail rates. These suggest hedging possibilities in instruments correlated with these spreads. Finally, the reservation values we later estimate should incorporate such possible hedging as well as hedging of the risk of the underlying CLO.

We next examine the determinants of BWIC failure by estimating the following Probit regression (Φ is the normal distribution function):

[15] Hendershott and Madhavan (2015) find similar results following failed corporate bond auctions. They attribute the low likelihood of trade after a failed auction to dealers being unlikely to increase their bids after a failed auction.

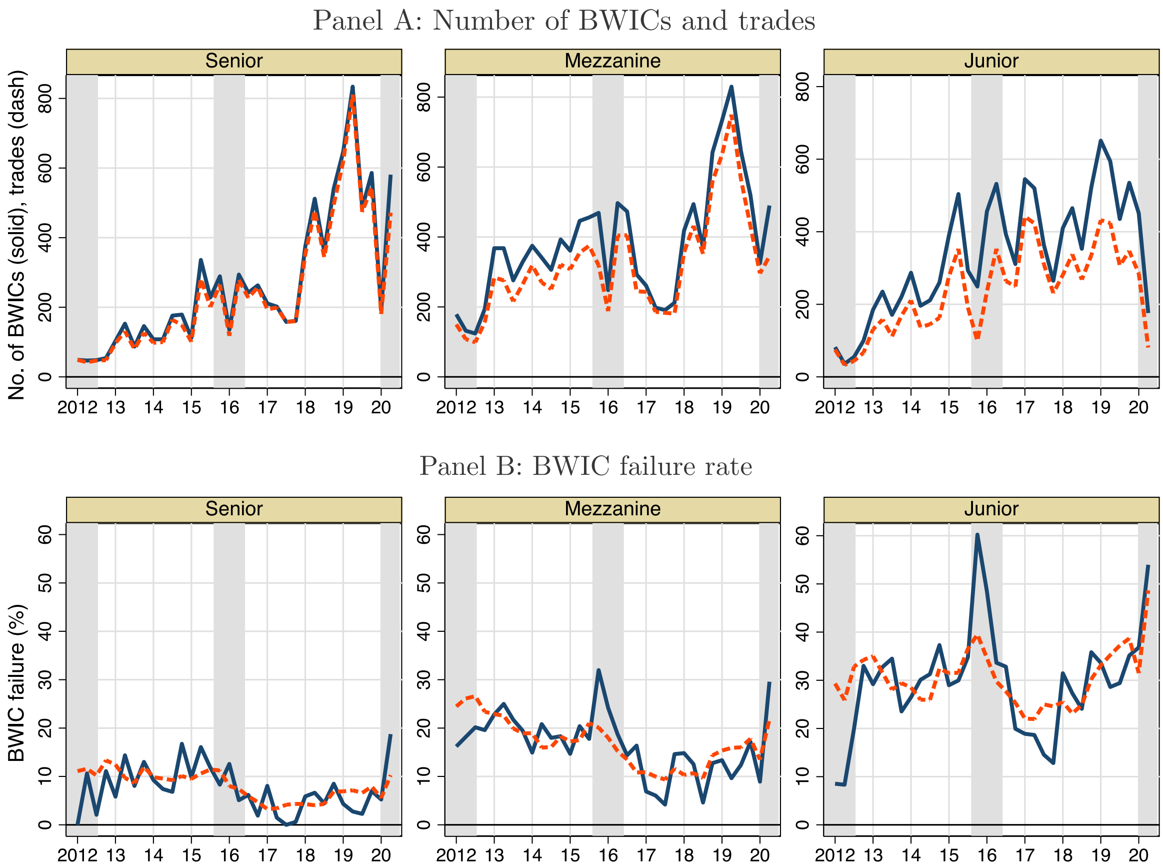

Fig. 4. BWIC activity and failure rates
The figure documents quarterly BWIC activity (Panel A) and BWIC failure rates (Panel B) in the CLO market over the sample period 2012–2020, split by tranche into the senior tranche rated AAA (left), mezzanine tranche rated AA–BBB (middle), and junior tranche rated BB–equity (right). In Panel B, the solid line is actual, and the dashed line is predicted. Grey areas indicate periods of market stress: 2012 (European debt), 2015–16, and 2020 (COVID-19 pandemic).

Table 2
BWIC fail rate: Splits by dealer inventory, trade size, CLO type, and market state.
The table reports BWIC fail rates for different sample splits. Panel A reports results for the senior tranche rated AAA, Panel B for the mezzanine tranche rated AA–BBB, and Panel C for the junior tranche rated BB–equity. Low/Medium/High columns indicate values of the varied characteristics while other characteristics are held fixed at their medium value.

Split variable	BWIC fail rate split by		
	Low	Medium	High
Panel A: Senior tranche (AAA, I = 9000 BWICs)			
Par size of trade	0.03	0.07	0.12
CLO vintage (1.0/2.0/3.0)	0.10	0.09	0.06
CLO issue size	0.09	0.07	0.06
JPM LL spread	0.07	0.07	0.08
Dealer CDS spread	0.05	0.08	0.10
Dealer CLO inventory	0.05	0.05	0.11
Panel B: Mezzanine tranche (AA–BBB, I = 12,955 BWICs)			
Par size of trade	0.12	0.17	0.20
CLO vintage (1.0/2.0/3.0)	0.19	0.18	0.15
CLO rating (AA, A, BBB)	0.12	0.16	0.20
CLO issue size	0.19	0.16	0.14
JPM LL spread	0.13	0.17	0.19
Dealer CDS spread	0.12	0.17	0.20
Dealer CLO inventory	0.13	0.15	0.21
Panel C: Junior tranche (BB–equity, I = 11,453 BWICs)			
Par size of trade	0.26	0.32	0.35
CLO vintage (1.0/2.0/3.0)	0.25	0.32	0.31
CLO rating (BB, B, Equity)	0.27	0.36	0.39
CLO issue size	0.30	0.28	0.33
JPM LL spread	0.27	0.29	0.35
Dealer CDS spread	0.27	0.31	0.33
Dealer CLO inventory	0.25	0.34	0.32

Table 3
Probability of repeat BWICs.
The table reports the probability of a failed BWIC to be followed by another BWIC (Panel A) or OTC trade (Panel B) on the same day, within a day, or within 10 days of the auction.

Panel A: BWIC followed by BWIC						
	BWIC followed by BWIC; any size			BWIC followed by BWIC; matched size		
Seniority	Same day	$[t, t+1]$	$[t, t+10]$	Same day	$[t, t+1]$	$[t, t+10]$
All	0.000	0.005	0.039	0.000	0.001	0.008
AAA	0.000	0.012	0.090	0.000	0.000	0.012
AA	0.000	0.002	0.025	0.000	0.000	0.005
A	0.000	0.008	0.035	0.000	0.002	0.009
BBB	0.000	0.005	0.039	0.000	0.000	0.014
BB	0.000	0.005	0.044	0.000	0.001	0.009
B	0.000	0.000	0.012	0.000	0.000	0.000
EQUITY	0.000	0.001	0.009	0.000	0.000	0.001
Panel B: BWIC followed by OTC trade						
	BWIC followed by OTC; any size			BWIC followed by OTC; matched size		
Seniority	Same day	$[t, t+1]$	$[t, t+10]$	Same day	$[t, t+1]$	$[t, t+10]$
All	0.052	0.080	0.149	0.030	0.043	0.072
AAA	0.042	0.073	0.215	0.032	0.036	0.067
AA	0.032	0.073	0.144	0.022	0.054	0.100
A	0.044	0.087	0.129	0.038	0.049	0.064
BBB	0.063	0.083	0.149	0.034	0.046	0.086
BB	0.062	0.091	0.163	0.033	0.045	0.077
B	0.040	0.062	0.104	0.025	0.040	0.066
EQUITY	0.043	0.061	0.102	0.016	0.030	0.038

$$\Pr(\text{BWIC failure}_{it}) = \Phi(\alpha_t + \beta' X_{it} + \eta_{it}). \tag{1}$$

where α_t are quarter fixed effects. Table 4 reports parameter estimates for the determinants of BWIC failure, split by tranche. The baseline specification (Column 1) includes the trade and CLO characteristics only. We add credit rating dummies to the baseline specification of the mezzanine and junior tranches. Across specifications 2 to 5, we vary the set of explanatory variables by including market and dealer characteristics, and quarter fixed effects α_t, incrementally.

Panel A of Table 4 presents results for the senior tranche. The baseline specification (Column 1) is consistent with our univariate findings from sample splits (Table 2) that the fail rates increase with trade size and dealer CLO inventory position, and decrease with CLO issue size. Adding the leveraged loan spread (Column 2) slightly improves the pseudo-R^2 from 4.9% in the baseline specification to 5.5%. In addition, the coefficient on the leveraged loan spread is positive and statistically significant. The dealer CDS spread (Column 3) adds little in terms of the explanatory power to the baseline specification, although the coefficient on the dealer CDS spread is positive and both statistically and economically significant. Depending on whether quarter fixed effects are included (Columns 4 and 5), only one of the leveraged loan spreads and dealer CDS spreads stays significant and positive. Panel B of Table 4 shows results for the mezzanine tranche. The results are similar to those for the senior tranche. The main difference is that the regression coefficients on both the leveraged loan spread and dealer CDS spread are positive and statistically significant at 1% level (Column 4). However, once the quarter fixed effects are included (Column 5), only the dealer CDS spread remains positive and statistically significant at 1% level. Another noteworthy result is that fail rates increase monotonically with the credit rating. In fact, BWICs of A/BBB-rated CLOs are 14%/32% more likely to fail than BWICs of AA-rated CLOs (Column 5).

Panel C of Table 4 reports results for the junior tranche. Results are qualitatively similar to the ones reported for more senior tranches. Notably, the vintage effect is the most prominent in the junior and mezzanine tranches, with BWICs of CLO 2.0s/3.0s more likely to fail than BWICs of CLO 1.0s. This is likely because newer vintages have smaller senior tranches and relatively larger subordinated tranches to share potential credit loss.

These multivariate results confirm that there exists significant variability in BWIC failure rates with trade, market, and dealer conditions. In addition, the low pseudo–R^2s indicate that the variables used do not explain much of the variability observed in failure rates. This suggests that the total costs of immediacy that incorporate the expected losses from the failure to trade vary significantly across time, CLO types, trade characteristics, and other unobserved characteristics. The value lost to both the seller and the buyer due to the inability to trade is not directly observed in the data. Therefore, it is estimated from the data in the next section, taking into account that failure rates and failure costs comove with observables and unobservables.

5. Total cost of immediacy — definition and measurement

In this section, we define the total cost of immediacy (TCI) that accounts for both successful and failed trades and discuss its empirical implementation in the CLO market.

5.1. Definition

In a BWIC an owner wanting to sell an asset solicits bids $\{B_n\}_{n=1}^N$ from N dealers. The owner has a reserve price of R. While we treat R as a constant here, later, when we perform the estimation, we will use R as a source of heterogeneity across CLOs and their sellers. For instance, R may depend on the type of CLO and market conditions. More patient/impatient CLO sellers and sellers with higher/lower reputation are likely to set a higher/lower R, and more so during normal times than crisis periods when CLO investors may be forced sellers. Let $B \equiv B^{1:N} = \max\{B_n | n \in \{1, \ldots, N\}\}$ be the *best* bid.[16] The trade takes place only if the best bid, B, is at least the seller's reserve price, R:

$$\text{Owner/Seller's participation constraint: Trade} = 1 \Leftrightarrow B \geq R. \tag{2}$$

Let the probability that a trade fails to materialize be

$$\Pr(\text{Fail}) = \Pr(B < R) = 1 - \Pr(\text{Trade}).$$

[16] For instance, the CLO secondary market is organized as a sealed-bid first-price auction and, therefore, $B \equiv B^{1:N} = \max\{B_n | n \in \{1, \ldots, N\}\}$, where B_n is the bid price for the CLO by dealer n out of $n = 1, \ldots, N$ dealers participating in the BWIC (N varies across BWICs).

Table 4
BWIC fail rate.
The table reports the determinants of BWIC failure. Estimates are obtained from Probit regressions. Panel A reports results for the senior tranche rated AAA, Panel B for the mezzanine tranche rated AA–BBB, and Panel C for the junior tranche rated BB–equity. Standard errors are robust to heteroskedasticity.

	Probability of BWIC failure				
	(1)	(2)	(3)	(4)	(5)
Panel A: Senior tranche (AAA, I = 9000 BWICs)					
log(Par value of trade)	0.20 ***	0.20 ***	0.20 ***	0.21 ***	0.21 ***
CLO 2.0 vintage	0.10	0.11	0.13 *	0.10	0.10
CLO 3.0 vintage	0.15 **	0.10	0.21 ***	0.10	0.10
log(CLO issue size)	−0.07 ***	−0.07 ***	−0.07 ***	−0.07 ***	−0.08 ***
CLO issue size missing	−0.04	−0.05	−0.03	−0.05	−0.02
Dealer CLO inventory	0.37 ***	0.37 ***	0.36 ***	0.37 ***	0.43 **
JPM LL spread		0.09 ***		0.09 ***	−0.11 *
Dealer CDS spread			0.20 **	−0.07	1.02 **
Quarter FE	No	No	No	No	Yes
Pseudo-R^2	0.056	0.061	0.057	0.061	0.103
Panel B: Mezzanine tranche (AA–BBB, I = 12,955 BWICs)					
log(Par value of trade)	0.12 ***	0.12 ***	0.12 ***	0.12 ***	0.12 ***
CLO 2.0 vintage	0.02	0.05	0.08 *	0.06	0.05
CLO 3.0 vintage	0.05	0.03	0.15 ***	0.07	0.17 ***
log(CLO issue size)	−0.08 ***	−0.09 ***	−0.08 ***	−0.09 ***	−0.08 ***
CLO issue size missing	0.09	0.09	0.13 **	0.11 *	0.15 **
A rating	0.15 ***	0.15 ***	0.14 ***	0.14 ***	0.13 ***
BBB rating	0.31 ***	0.31 ***	0.31 ***	0.31 ***	0.31 ***
Dealer CLO inventory	0.31 ***	0.33 ***	0.29 ***	0.31 ***	0.08
JPM LL spread		0.11 ***		0.10 ***	−0.08 *
Dealer CDS spread			0.34 ***	0.13 **	0.84 ***
Quarter FE	No	No	No	No	Yes
Pseudo-R^2	0.025	0.032	0.028	0.032	0.052
Panel C: Junior tranche (BB–equity, I = 11,453 BWICs)					
log(Par value of trade)	0.14 ***	0.15 ***	0.14 ***	0.15 ***	0.15 ***
CLO 2.0 vintage	0.18 ***	0.20 ***	0.21 ***	0.18 ***	0.13 ***
CLO 3.0 vintage	0.33 ***	0.30 ***	0.39 ***	0.27 ***	0.21 ***
log(CLO issue size)	−0.06 **	−0.07 **	−0.06 **	−0.07 **	−0.07 **
CLO issue size missing	0.22 ***	0.20 ***	0.24 ***	0.19 ***	0.18 ***
B rating	0.22 ***	0.25 ***	0.23 ***	0.25 ***	0.28 ***
Equity rating	0.30 ***	0.33 ***	0.32 ***	0.34 ***	0.36 ***
Dealer CLO inventory	0.33 ***	0.32 ***	0.33 ***	0.33 ***	−0.21 *
JPM LL spread		0.16 ***		0.18 ***	−0.03
Dealer CDS spread			0.24 ***	−0.17 **	0.85 ***
Quarter FE	No	No	No	No	Yes
Pseudo-R^2	0.024	0.033	0.026	0.035	0.057

* Significance at the 10% level.
** Significance at the 5% level.
*** Significance at the 1% level.

The owner of the asset gets the expected sale price if the trade succeeds, $\mathbb{E}[B|\text{Trade}]$, and she gets the reserve price, R, if the trade fails and she keeps the asset. Therefore, the seller/owner's expected total payoff, Π, is equal to[17]

$$\mathbb{E}[\Pi] = \underbrace{(1-\Pr(\text{Fail})) * \mathbb{E}[B|\text{Trade}]}_{\text{Trade}} + \underbrace{\Pr(\text{Fail}) * R}_{\text{Outside option}}. \tag{3}$$

Let P be a *benchmark* price to measure the quality of the transaction. P can be, for instance, a fair value of the asset in a perfectly competitive auction, where $N \to \infty$. Alternatively, P can be the price the winning dealer, but not the owner, can immediately resell the asset for. We define the *total cost of immediacy* or TCI as the expected difference between what the asset's owner should be getting if P would be available to her, and what she expects to get, that is, Π, $\text{TCI} \equiv \mathbb{E}[P-\Pi]$. Combining this relation with (3) leads to the following definition of the TCI.

Definition 1. The total cost of immediacy, TCI, is equal to

$$\text{TCI} \equiv \mathbb{E}[P-\Pi] = \underbrace{\mathbb{E}[P] - \mathbb{E}[B|\text{Trade}]}_{\text{Effective spread}} + \underbrace{\Pr(\text{Fail})}_{\text{Fail rate}} * \underbrace{(\mathbb{E}[B|\text{Trade}] - R)}_{\text{Cost of trade failure}}. \tag{4}$$

Eq. (4) highlights three key components of the TCI. The first component is the spread between the benchmark and transaction prices conditional on trade occurring. In markets without firm quotes, it can be calculated from successful transactions where a dealer buys the asset at the bid price B and sells it at the ask price A, and it is often referred to as the effective bid–ask spread or round-trip markup. The bid–ask spread is widely used as the observed cost of immediacy in both academia and industry. However, using only the bid–ask spread to measure the cost of immediacy ignores that trades fail. The failure to trade is especially costly when the seller's need to sell is very high, which is often when the value of the seller's outside option is close

[17] If the owner strategically sets her reserve price above her outside option to increase bids, then she raises her revenues when she trades, but her outside option must be valued below R. This results in (3) overestimating the seller's payoff. This in turn results in a lower measured TCI. Hence, our TCI measure is biased downwards by using the reserve price to value the seller's outside option when the trade fails.

to zero. For instance, CLO investors are long-term investors and tend to sell their holdings when hit by a liquidity shock such as hitting capital constraints, portfolio rebalancing, and the need for cash. A seller unable to sell the CLO defaults to keeping it and, therefore, may suffer large losses due to CLO devaluation, legal fees, regulatory penalties, and credit rating downgrade. These losses have to be incorporated into the cost of immediacy. Therefore, expression (4) is more comprehensive than the bid–ask spread alone because it incorporates expected payoffs from all options available to the seller.

The expected opportunity costs of trade failure are captured by the product of two TCI components, the trade fail rate and the cost of trade failure. It follows from Eq. (4) that failure costs are equal to the expected best bid from a successful BWIC, net of the seller's reserve price:

$$\text{Failure Cost} = \mathbb{E}[B - R | B \geq R]. \tag{5}$$

TCI is equal to the effective spread when the fail rate is zero. In the case of the CLO secondary market, Pr(Fail) is the probability of a BWIC failure and, therefore, it is directly observable in the data.[18] Fig. 4 shows that the fail rate is rarely zero even for the senior CLO tranche which has a mean(median) fail rate equal to 8%(6%). This implies that even for the "safest" CLOs the TCI should be greater, and, much of the time significantly greater, than the bid–ask spread. Expected gains from trade conditional on trading, $\mathbb{E}[B|B \geq R] - R$, which we refer to as failure cost, is the third component of the TCI. It declines with R and increases in the expected sale price. Intuitively, the costliest failures are the ones where the seller has a poor outside option and is expected to sell the CLO for a high price. The cost of failure is not directly observable in the data and, therefore, needs to be estimated. In the next subsection, we discuss how to measure the TCI in the data.

While Table 3 shows the seller is unlikely to sell the CLO after a BWIC fails, it may occur in some cases. Our measure of failure costs in Eq. (5) assumes that the value of continuing to hold the asset is the reserve price. If attempting to resell the asset is better for the seller then true failure costs differ from Eq. (5). To examine whether our approach under or over-estimates failure costs relative to a dynamic setting with many attempts to sell we can recursively write the seller's expected payoff from all attempts to sell incorporating per-period holding costs h for continuing to hold the CLO:

$$\mathbb{E}[\Pi] = \underbrace{(1 - \Pr(\text{Fail})) * \mathbb{E}[B|\text{Trade}]}_{\text{Trade}} + \underbrace{\Pr(\text{Fail}) * (\mathbb{E}[\Pi] - h)}_{\text{Continuation value}}. \tag{6}$$

The reserve price must be at least as large as $\mathbb{E}[\Pi] - h$, otherwise the seller would be accepting bids below the continuation value and, hence, be better off with a higher reserve price. This implies that our failure costs are lower than the failure costs in a dynamic setting. Therefore, using the reserve price in Eq. (4) underestimates the dynamic TCI because using the reserve price correctly measures failure probability and underestimates failure costs.

5.2. Estimating seller's reserve price

To estimate the TCI in Eq. (4), we need to know the effective spreads, BWIC failure rates, and BWIC failure costs. The first two TCI components are either directly observable in the data or have readily available proxies. However, BWIC failure costs are not directly observable and thus have to be estimated. The main challenge is to impute sellers' reserves R_i. Consistent with a standard theory of reserve-price auctions, a reserve price can be viewed as a minimum bid the CLO seller accepts in a BWIC[19]:

$$R_i = \inf_{(B_i^1, B_i^2, \ldots, B_i^N)} \{B_i^{1:N} | \text{Trade}_i\}, \tag{7}$$

where the infimum is taken over all possible realizations of N bids $B_i^1, B_i^2, \ldots, B_i^N$ leading to a trade.[20] As Eq. (7) illustrates, R_i equals the lower support of the distribution of best bids given that the auction was successful and, hence, it is the 0-quantile of the best-bid distribution conditional on trade.

We next outline our identification strategy for sellers' reserve prices. While the individual R_i's are not directly observable, we can use identity (7) to learn about the determinants of reserve prices from the accepted best bids across BWICs as follows. We assume the dealers' best bids B and the sellers' reserve prices R vary across BWICs with observable and unobservable determinants which allows us to pool BWICs across sellers (here we drop the individual CLO sub-index i for the sake of clarity). We adopt a flexible specification for best bids and reserve prices by making no parametric assumptions other than the regularity conditions introduced below:

$$\begin{aligned} B &= \mu_B(X) + \sigma_B \varepsilon_B, \\ R &= \mu_R(Z) + \sigma_R \varepsilon_R, \end{aligned} \tag{8}$$

where $\mu_j(\cdot), j \in \{B, R\}$, are unspecified functions to be estimated in the data and $X = Z \cup W$. The dealers' bids B are determined by public information, captured by observable determinants Z, dealers' private information, captured by observable determinants W, and an unobservable random component $\sigma_B \varepsilon_B \in \mathbb{D}_B$ drawn from some distribution G_B with density g_B. We assume that reserve prices are determined by public information Z and an unobservable random component $\sigma_R \varepsilon_R \in \mathbb{D}_R$ drawn from some distribution G_R with density g_R. Such differences in sellers' reserve prices can be due to liquidity needs, risk capacity, hedging needs, and non-trading motives for running a BWIC, e.g., information gathering.[21]

The distributions G_B and G_R have variances (σ_B, σ_R) and a correlation coefficient $\omega \in (-1, 1)$. The special case $\omega = 0$ represents a private-value setting. Let $H(\cdot)$ be the distribution of the standardized error $\eta \equiv \frac{\sigma_R \varepsilon_R - \sigma_B \varepsilon_B}{\sqrt{\sigma_R^2 + \sigma_B^2 - 2\omega\sigma_R\sigma_B}}$ and $C(\cdot)$ be the copula of (ε_B, η). The score $S(X)$ is the standardized difference between expected best bid and reserve price:

$$S(X) = \frac{\mu_B(X) - \mu_R(Z)}{\sqrt{\sigma_R^2 + \sigma_B^2 - 2\omega\sigma_R\sigma_B}}. \tag{9}$$

Then the BWIC success probability in terms of $S(X)$ is equal to

$$\Pr(\text{Trade}) = \Pr(R \leq B) = \Pr(\eta \leq S(X)) = H(S(X)). \tag{10}$$

The low pseudo–R^2s in Table 4 suggest that the explanatory variables X do not explain all the variability observed in success/failure rates. Therefore, it is important to account in the copula $C(\cdot)$ for the correlation in the errors ε_B and η. Two parameters are of special interest in

[18] Note that we are unable to measure the frequency and costs for when investors want to sell because the market is sufficiently illiquid.

[19] Our CLO data supports the view that when auctions fail, the seller is unable to trade. Table 3 shows that the probability of a failed auction being followed by an OTC trade either on the same day or within a day of the auction is low.

[20] The seller's outside option is likely to be at or below her reserve price for two reasons. First, if the BWIC fails, bidders have no incentive to bid above the reserve price in any subsequent auction. Second, sellers have no incentive to set the reserve price below their outside option. Both of these cause our estimated TCI to be downward biased.

[21] The capacity for running BWICs to chase information in dealer's bids may be limited because the seller is not anonymous to the dealer running the auction.

the results that follow[22]:

$$\rho \equiv Corr(\varepsilon_B, \eta) = \frac{\omega\sigma_R - \sigma_B}{\sqrt{\sigma_R^2+\sigma_B^2-2\omega\sigma_R\sigma_B}}, \quad \frac{1}{\gamma} \equiv \rho - \frac{\omega\sigma_R}{\sqrt{\sigma_R^2+\sigma_B^2-2\omega\sigma_R\sigma_B}} = -\frac{\sigma_B}{\sqrt{\sigma_R^2+\sigma_B^2-2\omega\sigma_R\sigma_B}}. \tag{11}$$

We expect $\rho < 0$ so long as auctions fail predominantly because of variation in bidders' valuations or when ω is small or negative. In turn, when auctions fail predominantly because of variation in sellers' reserve prices and $\omega > 0$, ρ can become positive. The definitions of the parameters nest three special cases: The case $\rho = 0$ corresponds to no selection in which auctions fail or succeed. The case $\rho = -1$ corresponds to no unobserved heterogeneity across auctions due to variation in reserve prices. The case $\rho = \frac{1}{\gamma}$ captures the pure private-value auction setting.

Estimating R as a fixed quantile of the best-bid distribution introduces a selection bias. The participation constraint (2) and our model of best bids and reserve prices (8) yield that $\Pr(B < R|X, \text{Trade}) = 0$ equivalent to the following extra restriction

$$\mathbb{E}[\mathbb{1}(B < \mu_R(Z) + \sigma_R\varepsilon_R)|X, \text{Trade}] = 0. \tag{12}$$

When $\sigma_R = 0$ in (12), we immediately have that the determinants of reserve prices can be recovered from the lowest accepted bids, or the 0-quantile of the observed best bid distribution because

$$\mathbb{E}[\mathbb{1}(B < \mu_R(Z))|X, \text{Trade}] = 0 \text{ if } \sigma_R = 0. \tag{13}$$

However, when $\sigma_R \neq 0$ the seller's participation constraint, $R \leq B$, biases the 0-quantile because accepted and, hence, observed best bids are greater than the unobserved best bids when BWICs fail.[23]

5.3. Optimal quantile rotation approach to capture seller's reserve

The selection bias introduced when replacing (12) by (13) can be corrected by rotating the quantile at which the reserve price is recovered away from zero. To determine the optimal rotation τ^*, we now develop a novel optimal quantile rotation (OQR) approach. The observed distribution of best bids in BWICs contains information about the sellers' reserve prices. In the absence of selection bias, the lowest accepted bid, that is, the 0-quantile of the best-bid distribution, recovers the reserve prices. However, winning bids are subject to sample selection due to sellers' participation constraint (2).

The optimal quantile rotation $\tau^* \in (0, 1)$ can be found by shifting the percentile level as a function of the amount of selection by optimally rotating away from the 0-quantile according to

$$\mathbb{E}[\tau^* - \mathbb{1}(B < \mu_R(Z))|X, \text{Trade}] = 0, \tag{14}$$

thus immediately yielding $\tau^* \equiv \Pr(B - \mu_R(Z) < 0|X, \text{Trade}) = \Pr(\varepsilon_R < -\frac{1}{\sigma_R}(B - R)|X, \text{Trade})$.

Following the discussion in the prior section, we can model the sample selection via the bivariate cumulative distribution function, or *copula*, of the errors in the winning bids (i.e., observed outcome) and the participation constraint (i.e., selection equation). Specifically, the selection-correcting quantile rotation τ^* can be expressed via the copula of ε_B, the errors in the outcome Eq. (8) for best bids, and $\sigma_R\varepsilon_R - \sigma_B\varepsilon_B$, the errors in the selection Eq. (2): $R - B \leq 0$. The selection-adjusted quantile rotation to recover the expected reserve price from the observed best bids is given by the following result proved in Appendix C.

[22] Two limiting cases are of interest. In the limit when $\omega = 0$, we have $\gamma = \frac{1}{\rho}$. When $\omega = \frac{\sigma_B}{\sigma_R} \equiv 1/\lambda$, we have $\gamma = -\sqrt{(\frac{\sigma_R}{\sigma_B})^2 - 1}$ and $\rho = 0$. Simple algebra yields useful relations $\lambda\omega = 1 - \rho\gamma$ and $\lambda = \sqrt{\gamma^2 + 1 - 2\rho\gamma}$.

[23] Arellano and Bonhomme (2016) and Maasoumi and Wang (2019) provide side-by-side comparisons of existing approaches to correct for selection bias in quantile regressions. Our setting differs from theirs because we determine τ^* endogenously.

Proposition 1. *The quantile of the best bid distribution adjusted for selection at $\varepsilon_R = 0$ is observation specific and equal to $H(\gamma S(X))$ with score $S(X)$ defined in (9) and parameter γ given by (11).*

Proposition 1 can be used to consistently estimate μ_R as a function of the observables, as summarized by the following Lemma proved in Appendix C.

Lemma 1. *The optimal rotation to recover expected reserve prices μ_R from the observed distribution of best bids depends only on the score $S(X)$ and parameters (γ, ρ), and it is given by the copula*

$$\tau^* = \frac{C(H(\gamma S(X)), H(S(X)); \rho)}{H(S(X))}. \tag{15}$$

Expression (15) shows that the quantile rotation to adjust for selection and recover the reserve price is observation-specific. We can for some $\pi \in (0, 1)$ make the approximation

$$\tau^* \approx \frac{C(\pi \Pr(\text{Trade}), \Pr(\text{Trade}); \rho)}{\Pr(\text{Trade})}.$$

This means the selection adjustment to the 0-quantile in this case is known up to the coefficient ρ that depends on $\sigma_R \geq 0$.[24] This approximation illustrates that the quantile of the best bid distribution adjusted for selection in order to identify $\mu_R(Z)$ varies observation by observation and is equal to $H(\gamma S) \approx \pi \Pr(\text{Trade})$.[25]

Fig. 5 provides a visual illustration of the selection-corrected quantile of the best bids distribution, τ^*, containing sellers' reserve prices. In Fig. 5, the average reserve, μ_R, and the average best bid, μ_B, in (8) are normalized to zero, $\mu_R = \mu_B = 0$. We draw the best bid prices for 1,000 BWICs from the standard Normal distribution, $B \sim G_B = \mathcal{N}(0, 1)$ and use them in all three subplots. We draw respective reserve prices independently from B, $\omega = 0$, either from the standard Normal distribution, $R \sim G_R = \mathcal{N}(0, 1)$, in the left subplot, or from $R \sim G_R = \mathcal{N}(0, 0.25)$, in the middle subplot, or from $R \sim G_R = \mathcal{N}(0, 0.01)$, in the right subplot. In all subplots, the 45-degree dashed line, $B = R$, separates accepted, $B \geq R$, and rejected, $B < R$, best bids. Failed bids are shown with red rhombuses. Accepted bids shown with green triangles have $\varepsilon_R < 0$, and accepted bids shown with blue circles have $\varepsilon_R \geq 0$. τ^* is the mass of green triangular accepted bids relative to the total mass of accepted bids. The optimal rotation in the left subplot of Fig. 5, $\Pr(\varepsilon_R < -(B - R)|\text{Trade})$, is equal to 0.25 or 25%-quantile. The middle and right subplots of Fig. 5 illustrate how the size of the selection bias in observed bids is related to σ_R. The volatility σ_R is equal to 0.5 in the middle subplot, and it is equal to 0.1 in the right subplot. When there is less variation in reserve prices, there is less selection bias in the observed bids. Hence, the optimal rotation in the middle subplot of Fig. 5, $\Pr(\varepsilon_R < -2(B - R)|\text{Trade})$, is equal to 0.15 or 15%-quantile, and it is equal to $\Pr(\varepsilon_R < -10(B - R)|\text{Trade}) = 0.01$ or 1%-quantile in the right subplot.

Appendix D provides simulation evidence on the optimal quantile rotation to measure TCI. Table A.2 in Panel A (Panel B) provides simulation evidence when the unobserved variation from $(\varepsilon_B, \varepsilon_R)$ is small (large) compared to the observed variation in best bids and reserve prices.

[24] When $\sigma_R = 0$, $\rho = -1$ and $\tau^* = C(-\Pr(\text{Trade}), \Pr(\text{Trade}); -1)/\Pr(\text{Trade}) = 0$.

[25] Arellano and Bonhomme (2017) determine a selection-adjusted rotation to recover quantiles of the unconditional wage distribution. They fix a quantile $\tau \in (0, 1)$ and show that the rotation $\tau^\dagger(\tau) = C(\tau, \Pr(\text{Trade}); \rho)/\Pr(\text{Trade})$ adjusts for selection at the fixed, exogenous quantile τ. Two recent applications of Arellano and Bonhomme's (2017) quantile-copula approach are by Maasoumi and Wang (2019) to assess how selective participation of individuals in the labor market affects the gender wage gap, and by Bollinger et al. (2019) to account for the effects of the endogenous decision to respond to the administrative earnings survey on the income inequality estimates. The Arellano and Bonhomme (2017) rotation holds for any fixed τ and, hence, $\tau^\dagger(\tau)$ is observation independent and therefore does not allow for recovering reserve prices in our setting.

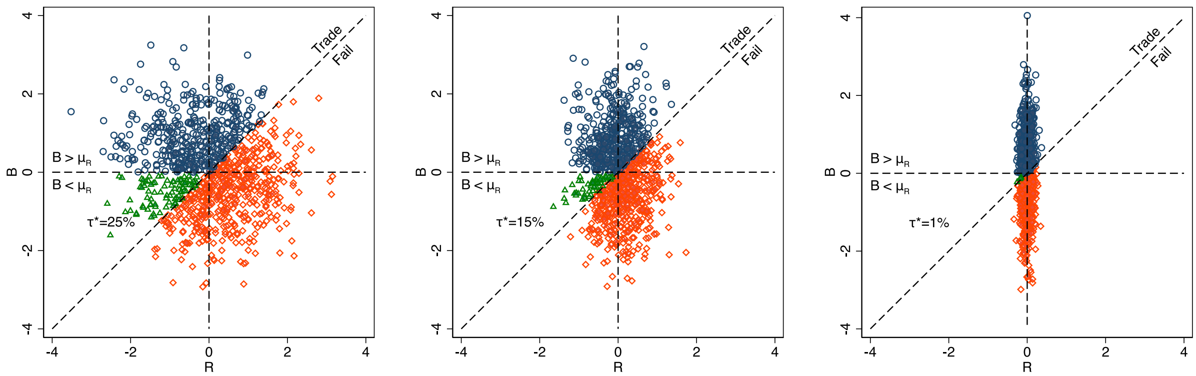

Fig. 5. Selection bias in quantiles of accepted bids
The figure illustrates the selection bias in observed accepted bids for different values of σ_R. The average reserve price, μ_R, and the bid price, μ_B, are both set to zero, $\mu_R = \mu_B = 0$, and are shown by the vertical and horizontal dashed lines, respectively. The standard Normal distribution is used to draw best bid prices, $B \sim G_B = \mathcal{N}(0,1)$, in both subplots for 1,000 BWICs. The standard Normal distribution is used to draw 1,000 reserve prices $R \sim G_R$ from (i) the standard Normal distribution, $G_R = \mathcal{N}(0,1)$, in the left subplot, (ii) from $G_R = \mathcal{N}(0,0.25)$ in the middle subplot, and (iii) from $G_R = \mathcal{N}(0,0.01)$ in the right subplot. B and R are drawn independently in all subplots. The 45-degree dashed line, $B = R$, separates accepted, $B \geq R$, and rejected, $B < R$, best bids. Green triangular accepted bids have $B < \mu_R$ ($\epsilon_B < 0$), and blue circular accepted bids have $B \geq \mu_R$ ($\epsilon_B \geq 0$). The mass of green triangular accepted bids relative to the total mass of accepted bids is τ^* which corrects for selection bias.

Next, we discuss the empirical implementation of the estimation procedure.

5.4. Empirical implementation of OQR

We estimate the OQR τ^* given by (15) for the CLO data in two stages. In the first stage, we obtain from (10) the conditional moment condition for the success of a BWIC depending on X:

$$\mathbb{E}[\mathbb{1}(\text{Trade}) - H(S(X))|X] = 0. \tag{16}$$

Condition (16) provides a set of unconditional moment conditions to estimate the determinants of the score $S(X)$, requiring information on both failed and successful BWICs. Under the linearity assumption, $S(X) = \alpha_S + \beta_S' X$, with observables $X = (Z, W)$. Standard arguments guarantee identification of (α_S, β_S). We estimate (α_S, β_S) by the Probit model (i.e., $H = \Phi$ is the Normal cdf).[26] The first-stage predictions $\hat{S}(X)$, given in the linear case by $\hat{S}(X) = \widehat{\alpha_S} + \widehat{\beta_S}' X$, and the predicted probability of BWIC success, $H(\hat{S}(X))$, are used in the second stage.

In the second stage, condition (15) provides conditional moments for identifying the OQR τ^*:

$$\mathbb{E}[\tau^* - \Pr(B \leq \mu_R(Z)|Z, \text{Trade})|X, \text{Trade}] = 0, \tag{17}$$

which can be estimated using GMM. Under the linearity assumption, $\mu_R(Z) = \alpha_R + \beta_R' Z$, the second-stage GMM condition takes the form

$$\mathbb{E}\left[\frac{C(H(\gamma\hat{S}(X)), H(\hat{S}(X)); \rho)}{H(\hat{S}(X))} - \mathbb{1}(B \leq \alpha_R + \beta_R' Z)\middle| X, \text{Trade}\right] = 0. \tag{18}$$

It is a function of $(\gamma, \rho, \alpha_R, \beta_R)$. The parametric assumptions required to implement copula (15) in (18) can be substantial. However, the statistical literature offers a number of convenient specifications, including the Gaussian, Frank, or Gumbel copulas, which provide a fair amount of flexibility. Alternatively, we can estimate (15) in (18) using flexible semi-parametric specifications. The first term in (18) is then known up to the deep parameters γ and ρ. Identification requires a shock affecting dealers' bids (demand) and through them the auction success probability, $H(S(X))$, without affecting sellers' reserve prices (supply). All model parameters combined are $(\alpha_S, \beta_S, \alpha_R, \beta_R, \gamma, \rho) \in \mathbb{R}^{2|Z|+|W|+4}$, where $|Z| \geq 0$ is the number of reserve price determinants Z and $|W| > 0$ is the number of determinants $W = X \backslash Z$ affecting bids only.

To illustrate identification, Appendix E describes in detail our GMM estimation procedure and use of aggregate dealer CLO inventories for the instrument, W, that is dealers' private information. We explain our choice of moments, demonstrate how condition (18) yields by conditioning down a set of moment conditions (E.3) that identify the parameters (α_S, β_S) and $(\alpha_R, \beta_R, \gamma, \rho)$ in two stages, provide intuition and formal analysis for why the moment conditions identify the parameters, discuss our inventory instrument, and consider three approaches to parametrize copula $C(\cdot)$ in (18). Based on these approaches, we employ six different model specifications for τ^*. Tables A.3–A.7 summarize the distribution of rotations τ^* across BWICs, the GMM estimates of the sellers' reserve prices, and specification tests.

The next section uses the $\mu_R(Z)$ estimates (α_R, β_R) obtained based on τ^* to calculate TCI.

6. TCI and the costs of trade failure

The estimates for the optimal quantile rotation τ^* from Table A.3 allow us to back out the sellers' expected reserves, calculate TCI across BWICs and time and quantify the cost of trade failure.

6.1. Mapping TCI to the data

The total cost of immediacy is given by (4) which comprises the expected spread, the probability of trade failure, and the cost of trade failure. The first term is the effective bid–ask spread in successful BWICs constructed and discussed in Section 3. The last two quantities map into the probability of BWIC failure and the cost of BWIC failure in our CLO data, respectively. The probability of failure is observed in the data. The two estimated components of TCI are the effective spread and the failure costs. The spread depends on the benchmark price and failure costs depend on the reserve price.

For the benchmark price, P, we use the ask price. This choice avoids modeling CLO fair value which is challenging given the infrequent nature of CLO trading. Using the ask price as the benchmark price incorporates dealer trading revenues on reselling the CLO to a buyer into TCI for the seller. Assuming the dealer does not resell the CLO at a loss, the effective spread overestimates the cost to the buyer. For senior tranches, failure rates are lower so the spread is more important for TCI. Failure rates for junior tranches are higher, making the effective spread less important for TCI.

[26] Under the linearity assumption, $\mathbb{E}[B - R|Z, W] = (\alpha_B - \alpha_R) + (\beta_B - \beta_R)'Z + \delta_B' W$ and, hence, $S(X) = \alpha_S + \beta_S' Z + \delta_S' W$ with $\alpha_S = \frac{\alpha_B - \alpha_R}{\sqrt{\sigma_R^2 + \sigma_B^2 - 2\omega\sigma_R\sigma_B}}$, $\beta_S = \frac{\beta_B - \beta_R}{\sqrt{\sigma_R^2 + \sigma_B^2 - 2\omega\sigma_R\sigma_B}}$, $\delta_S = \frac{\delta_B}{\sqrt{\sigma_R^2 + \sigma_B^2 - 2\omega\sigma_R\sigma_B}}$. From (10) and (16), for the Probit model the conditional moment condition is $\mathbb{E}[\mathbb{1}(\text{Trade}) - \Phi(\alpha_S + \beta_S' Z + \delta_S' W)|Z, W] = 0$.

We map all three quantities making up TCI onto a cross-section of both successful and failed BWICs as functions of observable characteristics $X = (Z, W)$. To extract TCI from the data conditional on these observables, denoted $\widehat{\mathrm{TCI}}(X)$, note that it is the sum of the expected bid–ask spread and the expected BWIC failure rate scaled by expected failure costs:

$$\widehat{\mathrm{TCI}}(X) = \widehat{\text{Bid–ask spread}}(X) + \widehat{\text{BWIC fail rate}}(X) * \widehat{\text{Failure cost}}(X), \tag{19}$$

where we condition on X and take expectation over all realizations of R as discussed in Appendix B. TCI averaged over BWICs is equal to

$$\begin{aligned}\widehat{\mathrm{TCI}} &\equiv \mathbb{E}_X[\widehat{\mathrm{TCI}}(X)] = \mathbb{E}_X[\widehat{\text{Bid–ask spread}}(X)] \\ &\quad + \mathbb{E}_X[\widehat{\text{BWIC fail rate}}(X) * \widehat{\text{Failure cost}}(X)] \\ &= \mathbb{E}_X[\widehat{\text{Bid–ask spread}}(X)] \\ &\quad + \mathbb{E}_X[\widehat{\text{BWIC fail rate}}(X)] * \mathbb{E}_X[\widehat{\text{Failure cost}}(X)] \\ &\qquad + \widehat{Cov}_X(\widehat{\text{BWIC fail rate}}(X), \widehat{\text{Fail cost}}(X)),\end{aligned} \tag{20}$$

where the expectation is taken over the values of X in all BWICs and $\widehat{Cov}_X(\cdot,\cdot)$ is the empirical covariance that arises because BWIC fail rates and failure costs are correlated across BWICs through X. Eq. (20) essentially calculates the average TCI over realizations of the observables X.

We construct the TCI for each BWIC $i = 1, \ldots, I$ and then average over them according to expression (20). In this expression, the relationship between bid–ask spreads and the explanatory variables, $X_i = (Z_i, W_i)$, can be estimated by OLS according to

$$\text{Bid–ask spread}_i = \alpha + \beta' X_i + \epsilon_i, \tag{21}$$

where the coefficients include quarter fixed effects, so that $\widehat{\text{Bid–ask spread}}(X_i) = \hat{\alpha} + \hat{\beta}' X_i$.[27] The bid–ask spread specification is estimated on successful BWICs as discussed in Appendix B and Table A.1 in the Appendix reports parameter estimates. The relationship between BWIC fail rate$(X_i) = 1 - \Pr(\text{Trade}_i | X_i)$ and the explanatory variables X_i is estimated in Section 4 using probit regression (1). Table 4 reports parameter estimates for the determinants of BWIC failure rates. The expected best bids conditional on trade can be estimated by OLS according to

$$B_i | X_i, \text{Trade}_i = \alpha_B + \beta_B' Z_i + \delta_B' W_i + \epsilon_{B,i}. \tag{22}$$

Taken together, the cost of BWIC failure conditional on X_i is measured by

$$\widehat{\text{Failure cost}}_i = (\alpha_B - \alpha_R) + (\beta_B - \beta_R)' Z_i + \delta_B' W_i, \;\; i = 1, \ldots, I, \tag{23}$$

with the coefficients $(\alpha_B, \beta_B, \delta_B)$ from OLS estimates of (22) and (α_R, β_R) from the GMM estimates reported in Tables A.4–A.6.

6.2. TCI across CLOs and over time

We first document $\widehat{\mathrm{TCI}}$ and the variation in $\widehat{\mathrm{TCI}}(X)$ across BWICs.

Average TCI and dispersion of TCI across BWICs. Table 5 reports estimates of $\widehat{\mathrm{TCI}}$ using the OQR approach across BWICs for the senior, mezzanine, and junior tranches in our base specification, the general Gaussian copula (E.4) (column (1) in Tables A.4–A.6). For each CLO tranche and each specification model, we report the mean and the standard deviation, followed by the 5%, 25%, 50%, 75%, and 95% quantiles. Standard errors reported in parenthesis are bootstrapped. TCI for the senior tranche ranges between 9 bps at the 5%-quantile and 78 bps at the 95%-quantile, with the mean and standard deviation equal to 31 bps and 23 bps, respectively. For the mezzanine tranche, these values increase to 1.32% and 0.87%. For the junior tranche, the average TCI equals 3.79% with a standard deviation of 2.94%.

Fig. 6 graphs the quarterly average total cost of immediacy in the CLO market constructed using specification (1) for the seller's reserve prices. In all subplots, the solid/dashed line depicts the mean/median value across BWICs and the dotted line depicts the bid–ask spreads in the CLO market. Both TCI and the bid–ask spread are shown as a percentage of the face value. The shaded area indicates the 5% to 95% range across BWICs. Results are split by tranche into the senior tranche rated AAA (left), mezzanine tranche rated AA–BBB (middle), and junior tranche rated BB–equity (right). The sample period is 2012–2020 and the vertical grey areas indicate periods of market stress: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic. TCI spikes in all of these stress periods.

TCI is the largest for the junior tranche, ranging between 1% and more than 10% and dwarfing bid–ask spread, which is around a few bps.[28] The mezzanine tranche has the second highest TCI ranging between 50 bps and 4% with the bid–ask spread in the range of 0 to 31 bps. The senior tranche has the lowest TCI among the three tranches ranging between 10 bps and 1%. TCI for the senior tranche is still much larger than the bid–ask spread ranging between 0 bps and 13 bps. TCI and bid–ask are only comparable during 2017Q4-2018Q1 for the senior tranche. Overall, Fig. 6 confirms that the total costs of immediacy are much larger than traditional liquidity measures such as the effective bid–ask spread.

Table A.8 reports the complete set of estimates for $\widehat{\mathrm{TCI}}$ using the OQR approach across BWICs for the senior (Panel A), mezzanine (Panel B), and junior (Panel C) tranches using all specification models (1)–(6) as in Tables A.4–A.6. For each CLO tranche and for each specification model we report the mean and the standard deviation, followed by the 5%, 25%, 50%, 75%, and 95% quantiles. In the best fitting model specifications, (1) and (5), TCI for the senior tranche ranges between 9 bps/8 bps at the 5%-quantile and 78 bps/62 bps at the 95%-quantile, with the mean and standard deviation equal to 31 bps/23 bps and 26 bps/19 bps, respectively. For the mezzanine tranche, these values increase to 1.32%/1.38% and 0.87%/0.95% in the best-fitting model specifications, (1) and (5). For the junior tranche, the average TCI is equal to 3.79%/3.53% with the standard deviation equal to 2.94%/2.72% in the best fitting model specifications, (1) and (5). Estimates for TCI are similar across all model specifications, ranging from 24 bps to 37 bps for the senior tranche, 1% to 2.8% for the mezzanine tranche, and 1.25% to 9.75% for the junior tranche, with the model specification (6) generally yielding the largest and least conservative estimates.

TCI represents a liquidity risk and Fig. 6 shows substantial time series variation in TCI. Failure rates show a similar pattern in Fig. 4 Panel B by spiking in times of stress. Panel B in Table 5 (as well as Panel B in Table 8) report a variance decomposition of TCI and confirm that TCI varies systematically over time. Leveraged loan spreads and dealer CDS spreads in Table 4 are significant determinants of BWIC fail rates. These results suggest hedging possibilities in instruments correlated with these spreads. Reservation values should incorporate such possible hedging as well as hedging of the risk of the underlying CLO.

[27] We do not include CLO fixed effects α_i because individual CUSIPs do not trade often enough.

[28] While the trading of CLOs has not been studied before, OTC markets in many important asset classes have been studied: foreign exchange, spot commodities, derivatives, and corporate and municipal bonds. For fixed-income securities Bessembinder et al. (2020) provide an extensive survey of the literature. Our bid–ask spread estimates for CLOs (5–10 bps) are somewhat smaller than institutional trading costs in other OTC markets. However, it is important to note that CLO BWICs last from a day to several days, complicating direct comparisons as the time dimension of immediacy likely differs. Moreover, the average trade size for CLO BWICs is around $5 million, larger than a typical institutional corporate bond trade.

Table 5
TCI using OQR.
The table reports in Panel A the estimates of TCI, $\widehat{TCI}$, across BWICs for senior tranches, mezzanine tranches, and junior tranches. The model uses the specification (1) in Tables A.4–A.6 which assumes that the optimal quantile rotation to back out seller's reserve price from best bids is given by a Gaussian copula, $\tau^* = \Phi(\gamma S, S; \rho)/\Phi(S)$. Standard errors reported in parenthesis are bootstrapped. The table reports in Panel B the decomposition of variation in TCI across BWICs for senior tranches, mezzanine tranches, and junior tranches.

	Senior tranche (AAA, I = 8683 BWICs)	Mezzanine tranche (AA–BBB, I = 12,341 BWICs)	Junior tranche (BB–equity, I = 10,794 BWICs)
Panel A: Distribution in $\widehat{TCI}$ (% of par) across BWICs			
$\widehat{TCI}$ mean	0.31	1.32	3.79
	(0.04)	(0.10)	(0.68)
$\widehat{TCI}$ S.D.	0.24	0.87	2.94
	(0.04)	(0.09)	(0.53)
Dispersion of $\widehat{TCI}$ across BWICs:			
5%	0.10	0.42	1.02
	(0.02)	(0.05)	(0.22)
25%	0.16	0.75	1.82
	(0.02)	(0.07)	(0.35)
50%	0.24	1.10	2.78
	(0.03)	(0.10)	(0.51)
75%	0.36	1.61	4.57
	(0.05)	(0.12)	(0.81)
95%	0.78	3.00	10.40
	(0.13)	(0.27)	(1.90)
Panel B: Variance decomposition (% of variation)			
Issuer	0.09	0.07	0.04
Time	0.84	0.82	0.54

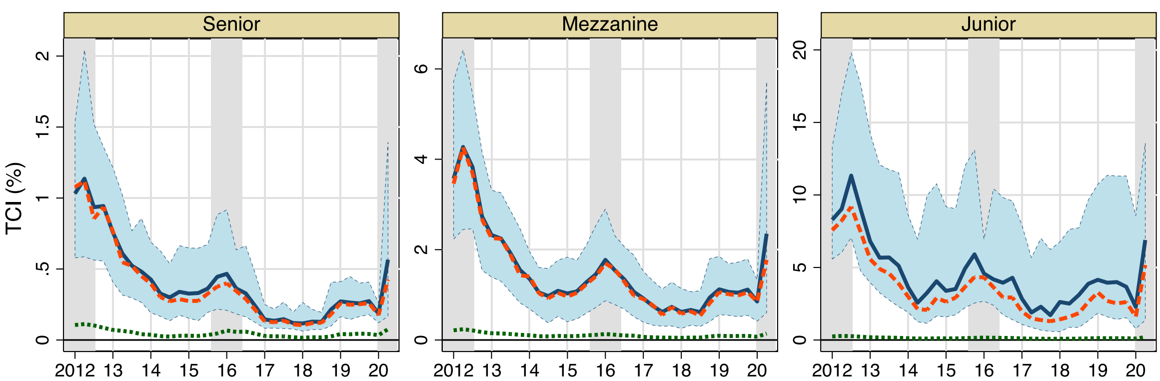

Fig. 6. Total cost of immediacy
The figure documents the quarterly total cost of immediacy in the CLO market using the parametric specification (1) for estimating the seller's reserve price. The solid line depicts mean $\widehat{TCI}$ over time. The dashed line corresponds to median $\widehat{TCI}$. The shaded area indicates the 5% to 95% range of $\widehat{TCI}$. The dotted line indicates the bid–ask spreads in the CLO market. The sample period is 2012–2020 with grey areas indicating periods of market stress: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic. Results are split by tranche into the senior tranche rated AAA (left), mezzanine tranche rated AA–BBB (middle), and junior tranche rated BB–equity (right). Fig. A.2 in the Internet Appendix documents the corresponding figure using the semiparametric specification (5) for estimating the seller's reserve price.

TCI determinants. Table 6 explores how TCI and its components vary with the trade, CLO, and market conditions – the explanatory variables X – using sample splits. Low/Medium/High columns indicate values of characteristics in its bottom/middle/top tercile. Panels A/B/C report results for the senior/mezzanine/junior tranche. Note that conditional on every characteristic in X, TCI increases with the riskiness of the tranches.

The results across tranches are generally consistent. Larger trades have higher TCI. This is because larger trades are harder to execute and they have a higher cost of failure. Older CLO vintages (CLO 1.0s and 2.0s) are more expensive to trade than the newer vintage (CLO 3.0) as the bid–ask spread, failure rate, and failure costs are all lower for CLO 3.0. TCI is smaller for CLOs with a larger issue size. CLO rating is the strongest cross-sectional determinant of TCI. In particular, the equity tranche has TCI as high as 9%. TCI is increasing with the size of dealers' inventories, that is, TCI is higher/lower when inventories are larger/smaller. TCI varies over time with leveraged loan market conditions and with dealer funding conditions. Higher spreads of the leveraged loans underlying CLOs lead to higher TCI. TCI increases also when dealer funding costs proxied for by the dealer CDS spread go up.

6.3. The cost of failing to trade

The opportunity cost of BWIC failure is defined in Eq. (5) with its empirical analog in expression (23). Failure costs are, in turn, the difference between the seller's reserve par discount and the expected dealer par discount. We can decompose the average TCI in terms of average expected failure cost, $\widehat{\text{Failure cost}}$, according to (20).

Table 7 reports the components of TCI using expression (20). The decomposition shows that the bulk of $\widehat{TCI}$ stems from a high chance of BWIC failure and large failure costs. Failure rates are 7% on average for a senior tranche, 16% for a mezzanine tranche, and 30% for a junior tranche. Failure costs as a fraction of par are 3.85% on average for a senior tranche, 7.00% for a mezzanine tranche, and 11.38% for a junior tranche. These failure costs are economically large as they are substantially larger than the leveraged loan spread in Fig. 3. By contrast, bid–ask spreads and empirical covariance contribute a small share to TCI. Bid–ask spread (failure rate–failure cost covariance) as a fraction of par is 0.04% (−0.02%) on average for senior, 0.10% (0.07%) for mezzanine, and 0.12% (0.23%) for junior tranches. The spread component is low for all tranches.

Table 6
TCI: Splits by trade size, CLO type, and market state.
The table reports TCI and its components for different sample splits. Panel A reports results for the senior tranche, Panel B for the mezzanine tranche, and Panel C for the junior tranche. The total costs of immediacy in a split are computed as the sample average over all BWICs falling in the respective tercile. The opportunity cost of BWIC failure is defined in Eq. (24).

Sample split	$\widehat{\text{TCI}}$ (% of par)		
	Low	Medium	High
Panel A: Senior tranche (AAA, I = 8683 BWICs)			
Par size of trade	0.19	0.31	0.42
CLO vintage (1.0/2.0/3.0)	0.59	0.31	0.25
log(CLO issue size)	0.40	0.27	0.24
JPM LL spread	0.20	0.32	0.41
Dealer CDS spread	0.20	0.26	0.47
Dealer CLO inventory	0.19	0.30	0.43
Panel B: Mezzanine tranche (AA–BBB, I = 12,341 BWICs)			
Par size of trade	1.08	1.37	1.52
CLO vintage (1.0/2.0/3.0)	1.87	1.58	1.02
CLO rating (AA, A, BBB)	0.82	1.22	1.71
log(CLO issue size)	1.58	1.30	1.05
JPM LL spread	0.91	1.30	1.76
Dealer CDS spread	0.88	1.09	2.02
Dealer CLO inventory	1.00	1.19	1.78
Panel C: Junior tranche (BB–equity, I = 10,794 BWICs)			
Par size of trade	3.10	3.51	4.89
CLO vintage (1.0/2.0/3.0)	4.97	4.37	3.25
CLO rating (AA, A, BBB)	2.64	3.45	8.89
log(CLO issue size)	2.99	2.87	5.70
JPM LL spread	3.00	3.79	4.59
Dealer CDS spread	2.73	3.72	4.98
Dealer CLO inventory	2.89	3.93	4.59

Table 7
Decomposition of average TCI.
The table reports the total cost of immediacy TCI and its components estimated from the data using specification (1). TCI is decomposed as the sample average over all BWICs $i = 1, \ldots, I$ in the tranche using expression (20). The table reports sample averages across all BWICs during the sample period from January 2012 to March 2020, split by CLO tranche. Standard deviations are reported in parenthesis below the mean. The table reports in Panel B the decomposition of variation in TCI across BWICs for senior tranches, mezzanine tranches, and junior tranches.

	Senior	Mezzanine	Junior
TCI (% of par)	0.31	1.32	3.79
Bid–ask spread (% of par)	0.04	0.10	0.12
Fail rate	0.07	0.16	0.30
Failure cost (% of par)	3.85	7.00	11.38
Best dealer bid (% discount of par)	−0.66	−3.24	−11.87
Seller's reserve (% discount of par)	−4.52	−10.25	−23.25
Cov(Fail rate, Failure cost)	−0.02	0.07	0.23

The size of the spread, which is the dealer's round-trip trading revenues, is informative regarding how our TCI estimates depend on our choice of benchmark price. For example, if we assume that the dealer's trading revenues are equally split between the dealer's purchase and subsequent resale, then the spread would be half as large. If the dealer gains on the initial customer sale are double the round-trip revenues, then the dealer loss on resale would be equal to the spread. In these two cases, the over and underestimation of the spread due to our choice of benchmark prices have an economically small impact on TCI in all tranches.

Given the importance of failure costs, Fig. 7 plots average failure costs and their components over time at a quarterly frequency. Panel A of the figure documents sellers' reserve prices constructed using the parametric specification (1) (solid and dashed line) and dealers' expected best bids (dotted line). Both sellers' reserve prices and dealers' expected best bids are expressed as a percentage of the face value. Panel B of the figure shows the BWIC failure cost (solid and dashed line) expressed in terms of par value in the CLO market at a quarterly frequency. In both panels, the solid/dashed line depicts the mean/median value across BWICs. The shaded area indicates the 5% to 95% range across BWICs.

Both sellers' reserve prices and dealers' expected bids tend to decline (rise when measured in par discounts) during credit stress periods and increase during regular (expansion) times.[29] The junior tranche shows the largest price swings, both in dealer bids and reserve prices, while the senior tranche shows the smallest price swings. Reserve prices fall by approximately 3%/10%/15% of the CLO face value between 2015Q3 and 2016Q1 for the senior/mezzanine/junior tranche. Expected bids declined by approximately 1%/5%/10% during the same time period for the senior/mezzanine/junior tranche. Both transaction and reserve prices recovered sharply after 2016Q1 and reached their highest values in 2017Q3. During this time the senior tranche was traded at approximately 0.5% above its face value, while the sellers' reserve price was equal to the face value. A noteworthy observation from Panel A is that sellers' reserve prices can be as low as 90% of the CLO face value for the senior CLO tranche. The reserve prices fall

[29] The reserve prices for the senior tranche are much lower and more volatile pre-2016Q1 than they are post-2016Q1, except for the COVID-19 pandemic period. This is largely due to the market being dominated by CLO 1.0s and 2.0s pre-2016Q1 and by CLO 3.0s post-2016Q1.

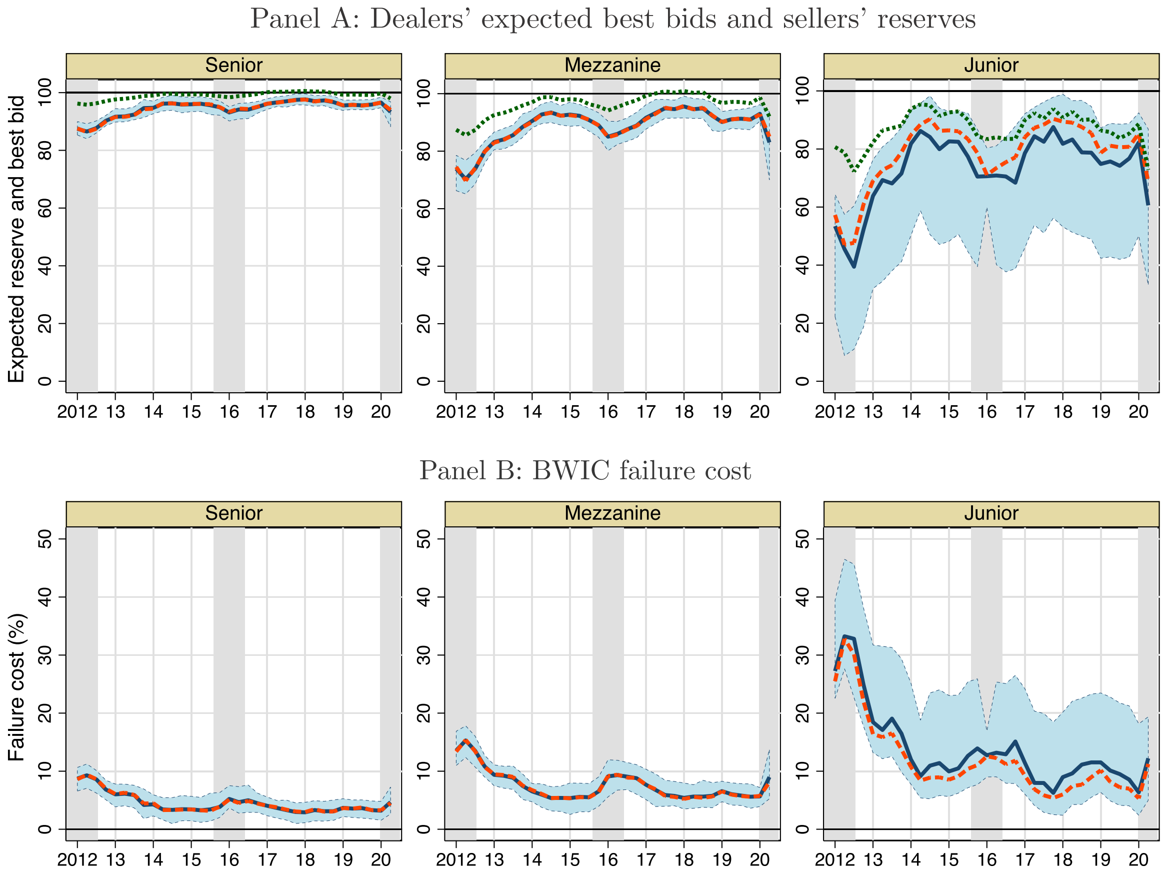

Fig. 7. Dealers' expected bids and sellers' reserves and BWIC failure cost
Panel A of the figure documents sellers' reserves constructed using the parametric specification (1) (solid and dashed line) and dealers' expected best bids (dotted line) expressed in terms of par value in the CLO market at a quarterly frequency. Panel B of the figure shows the BWIC failure cost (solid and dashed line) expressed in terms of par value in the CLO market at a quarterly frequency. In both panels, the solid/dashed line depicts the mean/median value across BWICs at a quarterly frequency. The shaded area indicates the 5% to 95% range across BWICs. The sample period is 2012–2020 with grey areas indicating periods of market stress: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic. In both panels, results are split by tranche into the senior tranche rated AAA (left), mezzanine tranche rated AA–BBB (middle), and junior tranche rated BB–equity (right). Fig. A.2 in the Internet Appendix documents the corresponding figures using the semiparametric specification (5) for estimating the seller's reserve price.

as low as 85% and 65% of the CLO face value for the mezzanine and junior tranches, respectively.

Most of the features from the plots in Panel A translate to BWIC failure costs depicted in Panel B of Fig. 7, which is the difference between the expected best dealer bids and seller reserve prices. Failure costs are, on average, larger and more volatile during 2012Q1 to 2016Q1 than they are during later sample periods. Failure costs are the largest during credit stress times, while they are the lowest during regular times for all tranches. For instance, the costs increased by approximately 1%/8%/10% between 2015Q3 and 2016Q1 for the senior/mezzanine/junior tranche. The junior tranche has the highest failure costs, which fluctuate between 5% and 40% most of the time. The senior tranche has the lowest failure costs and stays below 5% most of the time. Thus far we have discussed our empirical TCI measure. In the next section, we further discuss what TCI represents economically by calculating TCI for sellers with different reserve prices.

6.4. A simplified approach for applied research: Constrained quantile rotations

The OQR approach allows for the optimal quantile rotation to be observation-dependent. However, the approach can be computationally demanding. As a simpler alternative to the OQR approach, we can assume that the rotation is the same across auctions, which we term constrained quantile rotation (CQR).[30] We can then estimate the seller's reserve price $\mu_R(Z)$ using a non-parametric quantile regression on realized best-bids b_i at a fixed quantile $\tau^\star$. While the choice of an ideal quantile $\tau^\star$ requires some heuristics, our estimation using OQR offers some guidance on the range of such quantiles.

The procedure to determine TCI now involves the following steps: We again construct the TCI for each BWIC $i = 1, \ldots, I$ and then average over them according to expression (20). The only difference is that the determinants of sellers' reserve prices, $(\alpha_{\tau^\star}, \beta_{\tau^\star})$, are estimated by a standard quantile regression. The $\tau^\star$-quantile of the best bids regressed on Z_i yields: $\mu_R(Z_i) = \alpha_{\tau^\star} + \beta'_{\tau^\star} Z_i$.[31]

From now on the procedure evolves as before. The relationship between bid–ask spreads and the explanatory variables, $X_i = (Z_i, W_i)$, is estimated by OLS according to (21). The relationship between BWIC fail rate$(X_i) = 1 - \Pr(\text{Trade}_i|X_i)$ and X_i is estimated in Section 4 using probit regression (1). The expected best-bids conditional on characteristics X_i and trade success are estimated using linear regression

[30] In a setting where there exists no unobserved variation in reserve prices across BWICs, $\sigma_R = 0$, the approximation is correct and $\tau^* = 0$.

[31] In general, the τ%-quantile of a random variable B with distribution $G_B(b)$ is given by $q_B(\tau) = G_B^{-1}(\tau) = \inf\{b : G_B(b) \geq \tau\}$. We then postulate that the τth conditional quantile function can be parametrized as $q_{B|Z}(\tau) = \alpha_\tau + \beta'_\tau Z$. For the conditional distribution function of best-bids, $G_B(b|Z)$, α_τ and β_τ can be estimated from the pooled sample by solving $(\alpha_\tau, \beta_\tau) = \arg\min_{\alpha,\beta} \sum_{i=1}^{I} \xi_\tau(b_i - \alpha - \beta' Z_i)$, where $\xi_\tau(\cdot)$ is the tilted absolute value function at quantile τ.

Table 8
TCI using CQR.
The table reports in Panel A the estimates of TCI, $\widehat{\text{TCI}}$, across BWICs for senior tranches, mezzanine tranches, and junior tranches. The model uses the constrained quantile rotations to determine BWIC failure costs. The rotations used are $\tau^\star = 0.5\%$ for senior, 3% for mezzanine and 12% for junior tranches. The models use the same specifications as in Tables A.4–A.6. Standard errors reported in parenthesis are bootstrapped. The table reports in Panel B the decomposition of variation in TCI across BWICs for senior tranches, mezzanine tranches, and junior tranches.

	Senior tranche (AAA, I = 8683 BWICs)	Mezzanine tranche (AA–BBB, I = 12,341 BWICs)	Junior tranche (BB–equity, I = 10,794 BWICs)
Panel A: Distribution of $\widehat{\text{TCI}}$ (% of par) across BWICs			
$\widehat{\text{TCI}}$ mean	0.29	1.25	3.35
	(0.02)	(0.11)	(0.17)
$\widehat{\text{TCI}}$ S.D.	0.36	1.09	3.01
	(0.04)	(0.18)	(0.28)
Dispersion of $\widehat{\text{TCI}}$ across BWICs:			
5%	0.04	0.19	0.62
	(0.00)	(0.02)	(0.08)
25%	0.07	0.42	1.38
	(0.01)	(0.03)	(0.08)
50%	0.14	0.89	2.25
	(0.01)	(0.07)	(0.11)
75%	0.35	1.72	4.04
	(0.03)	(0.17)	(0.29)
95%	1.09	3.32	10.05
	(0.11)	(0.52)	(0.83)
Panel B: Variance decomposition (% of variation)			
Issuer	0.06	0.06	0.03
Time	0.81	0.89	0.52

(22). However, in this simplified approach the cost of BWIC failure conditional on X_i is measured by the difference between the expected best bid and (approximate) reserve price:

$$\widehat{\text{Failure cost}}(X_i) = (\alpha_B - \alpha_{\tau^\star}) + (\beta_B - \beta_{\tau^\star})' Z_i + \delta_B' W_i, \; i = 1, \dots, I, \quad (24)$$

instead of relations (23) used in OQR.

Table 8 reports estimates for TCI based on this CQR approach. We choose $\tau^\star = 0.5\%$ for the senior tranche, 3% for the mezzanine tranche, and 12% for the junior tranche, and the motivation for these choices is in the next section. Larger values of $\tau^\star$ lead to lower and thus more conservative TCI estimates for any given tranche. Standard errors reported in parenthesis are bootstrapped.

Table 8 shows that the total cost of immediacy is substantially higher than the observed cost for successful BWICs because trade failures are frequent and failure costs are sizeable. TCI is the smallest for the senior tranche and largest for the junior tranche. At the mean across BWICs, TCI is 29 bps/1.25%/3.35% for senior/mezzanine/junior tranche. Compared to bid–ask spreads (see Table 1 and Fig. 3), costs arising from trade failures make up the majority of the total cost of immediacy. The standard deviation and dispersion of TCI across BWICs is large, suggesting that TCI ranges between 4 bps/19 bps/62 bps and 1.09%/3.32%/10.05% for senior/mezzanine/junior tranche, each at 5% and, respectively, 95% of the distribution across BWICs.

The CQR approach is straightforward to implement, but the rotation needed to correct for selection, $\tau^\star$, is assumed to be known. Estimating CQR for various $\tau^\star$ can help researchers explore the sensitivity of their estimates to the degree of selection bias. Because $\tau^\star$ is higher for lower-rated CLOs, this sensitivity analysis may be more important for riskier assets.

7. TCI across sellers and realized gains from trade

The estimates for the model primitives allow us to conduct a counterfactual analysis. In the following, we compute the average TCI across all sellers and TCI for specific sellers and compare TCI to sellers' realized gains from trade (SGT). This allows us to study whether the cost of immediacy varies with how desperate or patient is the seller and whether dealers provide the same level of liquidity irrespective of the seller's type.

The TCI averaged over BWICs in expression (20) is intuitive but not necessarily representative of all sellers. Consider a seller with reserve price $R(r) \equiv \mu_R + \sigma_R r$, $r \in \mathbb{D}_r$, where $\mathbb{D}_r$ is a set of all reserve prices. The proof of Lemma 1 shows that $R(r)$ can be estimated from the optimal quantile of best bids, $\tau^*(r)$, equal to

$$\tau^*(r) = \frac{C(G_B(\gamma S + \lambda r), H(S))}{H(S)}. \quad (25)$$

The following result, proved in the Appendix, maps the seller's type (that is, her reserve price $R(r)$) into her individual TCI.

Proposition 2. *Estimated TCI for a seller with a reserve price $R(r)$ is equal to*

$$TCI(r) = \mathbb{E}[P - B|Trade, R(r)] + \mathcal{H}(r)\left(\int_{\mathbb{D}_b} b g_B(\frac{b-\mu_B}{\sigma_B}) C_{11}(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r))\frac{db}{\sigma_B} - R(r)\right), \quad (26)$$

where $\mathcal{H}(r)$ is the BWIC failure hazard rate given by (C.5) and $\mathbb{D}_b$ is the set of all bid prices.

From (26), TCI for the average/median seller, TCI(0), is equal to

$$\text{TCI}(0) = \mathbb{E}[P - B|\text{Trade}, \mu_R] + \left(\frac{1}{C_1(\frac{1}{2}, H(S))} - 1\right)\left(\int_{\mathbb{D}_b} b g_B(\frac{b-\mu_B}{\sigma_B}) C_{11}(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(0))\frac{db}{\sigma_B} - \mu_R\right). \quad (27)$$

The average TCI across all sellers is equal to

$$\mathbb{E}_R[\text{TCI}] = \mathbb{E}_R[\mathbb{E}[P-B|\text{Trade}, R]] + \underbrace{\int_{\mathbb{D}_r} \mathcal{H}(r)\left(\int_{\mathbb{D}_b} b g_B(\frac{b-\mu_B}{\sigma_B}) C_{11}(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r))\frac{db}{\sigma_B} - R(r)\right) dG_R(r)}_{\mathbb{E}_R[\Pr(\text{Fail}|R)\cdot(\mathbb{E}[B|\text{Trade},R]-R)]}$$
$$= \widehat{\text{TCI}} + \mathbb{E}_R[\mathbb{E}[B|\text{Trade}, R]] - \mathbb{E}_R[B|\text{Trade}] - Cov_R(\Pr(\text{Fail}|R), R). \quad (28)$$

Expression (28) illustrates that $\mathbb{E}_R$[TCI] is generally different from TCI(0), and, while intuitive and easy to compute, $\widehat{\text{TCI}}$ in (20) does not coincide with $\mathbb{E}_R$[TCI] because $\mathbb{E}_R[\mathbb{E}[B|\text{Trade}, R]] \geq \mathbb{E}_R[B|\text{Trade}]$ and $Cov_R(\Pr(\text{Fail}|R), R) \geq 0$. We can, however, compute upper and lower bounds on the average TCI across all sellers irrespective of any parametric assumptions. By the Cauchy–Schwarz inequality, $Cov_R(\Pr(\text{Fail}|R), R)$ is bounded by the standard deviation of the fail rate with respect to variation in R, $\text{SD}_R(\Pr(\text{Fail}|R))\cdot\sigma_R$ thus leading to the following Lemma.

Table 9
Average TCI and variation across sellers.
The table documents $\widehat{\text{TCI}}$ given by (20), TCI for the median seller given by (27), TCI(0), average TCI given by (28), $\mathbb{E}_R$[TCI], and the upper and lower bounds on average TCI, $\overline{\text{TCI}}/\underline{\text{TCI}}$, when using the estimated parameters from model (1) with a Gaussian copula and calibrating the remaining parameters to the specification $P = \sigma_B(\alpha+\beta)$, $B_i = \sigma_B(\alpha+\varepsilon_{B,i})$, $R_i = \sigma_R\epsilon_{R,i}$, $i = 1,\ldots,I$.

Panel A: Parameter estimate/calibration

	Senior	Mezzanine	Junior	Source
ρ	−0.82	−0.83	−0.82	GMM estimate from Table A.7
$1/\gamma$	−0.69	−0.76	−0.99	GMM estimate from Table A.7
λ	0.85	0.74	0.60	from $\lambda = \sqrt{\gamma^2+1-2\rho\gamma}$
ω	−0.22	−0.12	0.29	from $\omega = (1-\rho\gamma)/\lambda$
σ_B	0.65	1.72	2.83	matched to residual variance in observed bid–ask spreads
σ_R	0.56	1.27	1.70	from $\sigma_R = \lambda\sigma_B$
α	2.10	1.30	0.53	matched to fail rate
β	0.44	0.76	1.47	matched to $\widehat{\text{TCI}}$

Panel B: TCI for different sellers (% of par)

	Senior	Mezzanine	Junior
$\overline{\text{TCI}}$	0.32	1.38	4.00
$\underline{\text{TCI}}$	0.24	1.09	3.60
$\widehat{\text{TCI}}$	0.31	1.32	3.79
$\mathbb{E}_R$[TCI]	0.25	1.11	3.61
TCI(r) for specific sellers:			
Low-reserve seller, TCI(−2)	0.01	0.87	5.58
TCI(−1)	0.13	1.11	4.66
Median seller, TCI(0)	0.28	1.25	3.64
TCI(1)	0.38	1.16	2.55
High-reserve seller, TCI(2)	0.32	0.63	1.31

Panel C: Gains from trade for different sellers (% of par)

	Senior	Mezzanine	Junior
$\mathbb{E}_R$[SGT]	1.40	2.43	2.05
SGT(r) for specific sellers:			
Low-reserve seller, SGT(−2)	2.76	5.14	3.45
SGT(−1)	2.06	3.71	2.68
Median seller, SGT(0)	1.37	2.31	1.97
SGT(1)	0.72	1.12	1.42
High-reserve seller, SGT(2)	0.25	0.38	0.92

Lemma 2. *The average TCI across all sellers is bounded by*

$$\overline{TCI} \geq \mathbb{E}_R[TCI] \geq \underline{TCI}, \tag{29}$$

where the upper and lower bounds are given by

$$\begin{aligned}\overline{TCI} &= \widehat{TCI} + \mathbb{E}_R[\mathbb{E}[B|Trade, R]] - \mathbb{E}_R[B|Trade],\\ \underline{TCI} &= \widehat{TCI} - SD_R(\Pr(Fail|R)) \cdot \sigma_R,\end{aligned} \tag{30}$$

with $\widehat{TCI}$ defined in (20).

Having calculated expressions for TCI of different sellers and upper and lower bounds, we use calibration of the remaining parameters to compare the average TCI, $\widehat{\text{TCI}}$, given by (20), to the average TCI across all sellers, $\mathbb{E}_R$[TCI], given by (28), and the upper and lower bounds $\overline{\text{TCI}}/\underline{\text{TCI}}$ given by (30). We also compute TCI for the median seller given by (27), and we use (26) to calculate TCI for sellers with reserve prices 1 and 2 standard deviations above and below the mean. We use model 1 with a Gaussian copula $\tau^* = \Phi(\gamma S, S; \rho)/\Phi(S)$ to calculate TCI(r) and $\mathbb{E}_R$[TCI]. Table A.7 provides estimates of ρ and γ. λ and ω are given by expressions $\lambda = \sqrt{\gamma^2+1-2\rho\gamma}$ and $\omega = (1-\rho\gamma)/\lambda$, respectively. Using the estimated correlation and ratio of the variance of the bids and reserve prices, σ_B is matched so that the empirical variance in the noise of the observed bid–ask spread after controlling for observables X_i matches the model-implied variance in the noise of the bid–ask spread conditional on trade, $P - B_i|\text{Trade}_i$. σ_R is determined by $\sigma_R = \lambda\sigma_B$. α is calibrated to match the BWIC fail rate from Table 1, and β is calibrated to the estimated TCI from Column 1 of Table A.8.

Panel A of Table 9 reports the parameters from the calibration for the senior, mezzanine, and junior tranches. The estimates for ρ and γ from Table A.7 translate into estimates for the non-normalized correlation between the bids and reserve prices, ω. The estimated ω is negative for the senior and mezzanine tranches. This could arise if bidding in the auction is costly, possibly due to dealers' efforts in locating the buyer prior to the auction. In this case, dealers' bids and expected surplus when winning must account for incurring these costs. If dealers are aware of which sellers are more likely to have higher reserve prices, then dealers will bid less aggressively in auctions run by those sellers, causing $\omega < 0$. In the lower-rated tranches, there is more scope for bidders and dealers to have common information about the value of the CLO that is not observable to the econometrician. This naturally makes the correlation between bids and reserve prices more positive. Thus, for the junior tranche $\omega > 0$.

The various TCI measures in Panel B show that for all tranches the following ordering holds $\overline{\text{TCI}} > \widehat{\text{TCI}} > \text{TCI}(0) > \mathbb{E}_R[\text{TCI}] > \underline{\text{TCI}}$, indicating the estimated total costs of immediacy, $\widehat{\text{TCI}}$, is an upper bound on the TCI of the median seller, TCI(0), which, in turn, is an upper bound on the TCI across all sellers, $\mathbb{E}_R$[TCI], under the Gaussian copula.

TCI across sellers. Using the calibrated parameters from Panel A of Table 9, Panel B reports the variation in the total cost of immediacy across sellers by computing TCI(r) for different values of r. Panel B reports TCI(r) for sellers with $r = -2, -1, 0, 1, 2$. Sellers with $r = -2$ correspond to the most impatient sellers who must sell and will accept almost any bid. Their TCI will in the limit approach their bid–ask spread. At the other end of the spectrum, sellers with $r = 2$ correspond to the most patient sellers who only accept high prices. Sellers who run BWICs intending to gain information about valuations without transacting also fall into the category of sellers with the highest r.

Panel B shows that for the mezzanine and senior tranches TCI(r) is hump-shaped in sellers' reserve prices due to $\omega < 0$. TCI for sellers with low reserve prices are equal to 0.87% and 0.01%, respectively, and thus substantially lower than for the median seller, with 1.25%

and 0.28%. For low reserve prices, the probability of failure is close to zero. As reserve prices rise, the probability of failure remains low, but the cost of failure can actually increase due to $\omega < 0$. Hence, TCI increases for low reserve prices in the senior and mezzanine tranches. Once reserve prices are high enough TCI begins to fall due to the declining cost of failure. For the junior tranche, TCI for sellers with low reserve prices, TCI(−2), is 5.58%, which is 50% greater than TCI for the median seller, TCI(0), which is 3.64%. TCI is monotonically declining in reserve prices, consistent with the intuition from Eqs. (3)–(4). As reserve prices increase the cost of failure goes to zero.

Seller's gains from trade. The cost of immediacy is sellers' loss relative to the frictionless sale price. The realized gains from trade for the seller (SGT) are the difference between the frictionless sale price and the seller's reserve price minus TCI, which equals the difference between the expected best bid and the reserve price multiplied by the probability of trade[32]:

$$SGT(r) = \Pr(\text{Trade}|R(r)) \cdot (\mathbb{E}[B|\text{Trade}, R(r)] - R(r)). \tag{31}$$

Similar to TCI in Panel B, Panel C of Table 9 reports the average seller gains from trade as well as for sellers with a median reserve price and reserve prices 1 and 2 standard deviations above and below the median. For all tranches, as expected, TCI declines with reserve price. This is because both the expected gains from trade and the probability of trade decline with the reserve price. Gains from trade are higher for the mezzanine tranche than for the senior tranche. Gains from trade are higher for mezzanine than junior because TCI is so high for the junior tranche. For senior and mezzanine the average gains from trade are much higher than TCI suggesting that sellers capture much of their surplus. However, for the junior tranche, TCI is greater than SGT, implying the market is illiquid and the frictions consume more than half of the surplus.

8. Conclusion

The role of markets is to facilitate trade. Centralized markets like stock exchanges have continuous, visible, and firm quotes that inform investors of the terms of trade in advance and limit failed attempts to trade. Structured products and other financial securities trade in over-the-counter (OTC) markets without such quotes, making trade failures a more likely option. We measure the total cost of immediacy (TCI) in OTC markets by combining the direct intermediation cost, captured by the bid–ask spread, with the opportunity cost of trade failure. TCI is the expected bid–ask spread on observed trades plus the probability of trade failure times the cost of trade failure, which is the expected highest bid conditional on a trade minus the seller's reserve price. We construct a methodology to estimate the seller's reserve price using quantiles of the distribution of best bids adjusted for the selection bias in successful trades due to higher bids being more likely to be accepted.

Using data from 2012–2020 on bids-wanted-in-competition (BWIC) auctions for collateralized loan obligations (CLO), we document that trade failures are frequent; they occur more than 10% of the time, and occur more often in lower-rated CLOs and stressful market conditions when failure rates exceed 50%. Our estimates for TCI reveal significant illiquidity and fragility in auction-based OTC markets. TCI is substantially higher than the effective bid–ask spread from successful BWICs because trade failures are frequent, failure costs are large, and failure rates and costs are correlated.

In the 2020 pandemic and other stressful periods when failure costs exceed 10% for junior tranches, TCI increases tenfold from about 1% to 10%. For high-rated tranches, TCI is low relative to sellers' gains from trade, while for low-rated tranches TCI is almost double the gains from trade. When we estimate TCI for sellers with different reserve prices, TCI is more than 50% higher for sellers with low reserve prices than sellers with the median reserve price. Hence, when the need to trade is the most urgent and the cost of continuing to hold the asset is the highest, even sellers with the lowest reserve prices fail to trade because the OTC market provides insufficient liquidity, exacerbating fragility.

Our TCI results suggest that traditional trading cost measures underestimate true costs when failure rates are high and trading frequency is low. In corporate bonds failure rates in auctions/RFQs are close to 30% (Hendershott and Madhavan, 2015). Therefore, TCI may also be important in this setting. Auctions in more liquid OTC markets such as index CDS (Riggs et al., 2020) and Canadian government bonds (Allen and Wittwer, 2023) are much less likely to fail, suggesting that TCI is less relevant in those contexts. TCI in centralized markets may be easier to measure as subsequent prices can proxy for the value of holding the asset. Overall, TCI is likely most important in more risky and illiquid assets and in times of market stress.

Our approach can be extended to non-financial decentralized market settings to estimate outside option(s) of market participants. Natural settings include reservation wages in labor literature[33] and the utility of staying single for potential spouses in the marriage market.[34]

CRediT authorship contribution statement

Terrence Hendershott: Writing – review & editing, Writing – original draft, Validation, Resources, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Dan Li:** Writing – review & editing, Writing – original draft, Validation, Supervision, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Dmitry Livdan:** Writing – review & editing, Writing – original draft, Validation, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization. **Norman Schürhoff:** Writing – review & editing, Writing – original draft, Validation, Software, Resources, Project administration, Methodology, Investigation, Formal analysis, Data curation, Conceptualization.

Declaration of competing interest

Terrence Hendershott provides expert witness services to a variety of clients. He has taught a course for a financial institution that engages in liquidity provision and high-frequency trading activity. He gratefully acknowledges support from the Norwegian Finance Initiative.

Dan Li has no conflict of interest to disclose.

Dmitry Livdan has no conflict of interest to disclose.

Norman Schürhoff gratefully acknowledges research support from the Swiss Finance Institute and the Swiss National Science Foundation under SNF project #100018_192584, "Sustainable Financial Market Infrastructure: Towards Optimal Design and Regulation" and has no conflict of interest to disclose.

Data availability

TCI Data (Original data) (Mendeley Data)

.

[32] Calculating the full gains from trade for the round trip transaction where the dealer both buys and sells requires information about the eventual end buyer. Without such information, we focus on the gains from trade for the seller. Calculating the frictionless gains from trade for the seller requires endogenizing the benchmark price, which is beyond the scope of this paper.

[33] See Devine and Kiefer (1991) for the review of evidence on reservation wages.

[34] See Chiappori (2020) for the review of literature on marriage markets.

Appendix A. Data filters

The following table reports the steps to clean and filter the CLO BWIC data and the number of BWICs remaining after each step.

Step	Number of BWICs remaining
All auctions between 1.1.2012 and 3.31.2020	58,853
Keep only USD CLO	43,701
Keep only TRACE-eligible CUSIPs	34,138
Keep only if seniority/rating information is not missing	33,408

The following table reports summary statistics on merging TRACE and Creditflux (CF). We report for each dealer-to-client, interdealer, and client-to-dealer trade out of the total how many are in TRACE only, how many are in both TRACE and CF, and how high is the match rate. In the table, the statistic on 'In TRACE and CF' includes auctions with DNT flags or cover information that are nonstandard and cannot be translated into a decimal price.

	In TRACE only	In TRACE and CF	Total	Matched %
Dealer-to-client	88,573	13,093	101,666	13%
Inter-dealer	18,742	3112	21,854	14%
Client-to-dealer	43,194	27,890	71,084	39%

Appendix B. Evidence on bid–ask spreads in CLOs

This section examines the determinants of effective bid–ask spread using multivariate analysis. We estimate the following model for bid–ask spread that corresponds to a successful BWIC for CLO i at time t:

$$\text{Bid–ask spread}_{it} = \alpha_t + \beta' X_{it} + \epsilon_{it}, \tag{B.1}$$

where α_t are quarter fixed effects. The explanatory variables X include characteristics of the trade (log of par value), of the CLO (log of amount outstanding, vintage, CLO rating), as well as market and dealer conditions (JPM leveraged loan spread and dealer CDS spread). We do not include CLO fixed effects α_i because individual CUSIPs do not trade often enough.

Table A.1 reports parameter estimates for the determinants of the effective bid–ask spread on successful BWICs. All results are split by CLO tranche. The baseline specification (Column 1) includes only trade and CLO characteristics. We add credit rating dummies to the baseline specification in the case of the mezzanine and junior tranches. Across specifications 2 to 5, we vary the set of explanatory variables by adding market and dealer characteristics, and, finally, quarter fixed effects α_t. This approach allows us to trace the explanatory power of the incremental variables.

Panel A of Table A.1 presents results for the senior tranche. The vintage effect is robust to specifications, with newer vintages significantly cheaper to trade. The R^2 for the specification in Column (1) is relatively low at 1.4%, suggesting limited explanatory power from characteristics of the trade or the CLO. Adding market condition variables such as leveraged loan spread and dealer CDS spread (Columns 2 and 3) improves R^2 noticeably. The bid–ask spread is significantly positively related to both leveraged loan spread and dealer CDS spread. Therefore, when market conditions or dealer health worsens, bid–ask spread increases. When quarter-fixed effects are included together with market condition variables (Column 6), the leveraged loan market spread continues to be significant.

Panel B of Table A.1 shows results for the mezzanine tranche. The results are similar to the ones reported in Panel A for the senior tranche. Unique to this tranche, trade size is negatively significant across all specifications—larger trades are associated with smaller bid–ask spreads. Results reported in Column 5 indicate that once quarter fixed effects are included, only trade size retains explanatory power. The highest R^2 in Panel B is equal to 7.2% (Column 5). Panel C of Table A.1 reports results for the junior tranche. Notably, credit rating plays a more prominent role. When quarter fixed effects are included, only rating matters. Perhaps unsurprisingly, equity tranches have the highest bid–ask spread.

Overall, effective bid–ask spreads do not vary much within a quarter for CLOs with the same vintage and credit rating. When bid–ask spreads do vary within a quarter, the variation is mainly due to major economy-wide credit events affecting the riskiness of leveraged loans held by the CLOs.

Appendix C. Proofs

Proof of Proposition 1. The BWIC success probability is

$$\begin{aligned}\Pr(\text{Trade}) &= \Pr(R \le B) = \Pr(\sigma_R \varepsilon_R - \sigma_B \varepsilon_B \le \mu_B(X) - \mu_R(Z)) = \\ &= \Pr(\eta \le S(X)) = H(S(X)),\end{aligned} \tag{C.1}$$

thus completing the proof. □

Proof of Lemma 1. For $r \in \mathbb{D}_r$, we have $B \le \mu_R(Z) + \sigma_R r \Leftrightarrow \sigma_B \varepsilon_B \le -(\mu_B(X) - \mu_R(Z)) + \sigma_R r \Leftrightarrow \varepsilon_B \le \gamma S(X) + \lambda r$ with $\lambda = \frac{\sigma_R}{\sigma_B}$. Define

$$\begin{aligned}\tau^*(r) &\equiv \Pr(B \le \mu_R(Z) + \sigma_R r | \text{Trade}) \\ &= \Pr(\sigma_B \varepsilon_B \le -(\mu_B(X) - \mu_R(Z) - \sigma_R r) | \text{Trade}) \\ &= \frac{\Pr(\varepsilon_B \le \gamma S(X) + \lambda r, \eta \le S(X))}{\Pr(\text{Trade})} \\ &= \frac{C(H(\gamma S + \lambda r), H(S); \rho)}{H(S)}.\end{aligned} \tag{C.2}$$

By definition, $\tau^* = \tau^*(0)$. □

Proof of Proposition 2. Consider a seller with a noise component of the reserve price equal to $\sigma_R \varepsilon_R = \sigma_R r$, with $r \in \mathbb{D}_r$. We observe in the data for $R(r) \equiv \mu_R + \sigma_R r$,

$$\Pr(\text{Trade}) = \int_{\mathbb{D}_r} \Pr(\text{Trade} | R(r)) dG_R(r). \tag{C.3}$$

Using the properties of the copula $C(\cdot)$ capturing the joint relation of (ε_R, η) in the expression (C.3), the BWIC success probability conditional on $R(r)$ is equal to

$$\Pr(\text{Trade} | R(r)) = C_1(G_R(r), H(S)). \tag{C.4}$$

Note that the conditional hazard rate of BWIC failure, $\mathcal{H}(r)$, depends on the seller's reserve and equals

$$\mathcal{H}(r) \equiv \frac{\Pr(\text{Fail} | R(r))}{1 - \Pr(\text{Fail} | R(r))} = \frac{1}{C_1(G_R(r), H(S))} - 1. \tag{C.5}$$

We now determine the conditional expected best bids for a buyer $b \in \mathbb{D}_b$. We have for an event $B \le b$ equivalent to $\varepsilon_B \le \frac{b - \mu_B}{\sigma_B}$ the

Table A.1
Effective bid–ask spread on successful BWICs.
The table reports the determinants of the effective bid–ask spread on successful BWICs. The sample is restricted to CLO round-trip trades with one day or less in dealer inventory. Standard errors are robust to heteroskedasticity.

	Bid–ask spread (% of par value)				
	(1)	(2)	(3)	(4)	(5)
Panel A: Senior tranche (AAA, I = 2234 successful BWICs)					
Dealer CLO inventory	0.00	−0.00	−0.00	−0.00	0.01
log(Par value of trade)	−0.00	−0.00 *	−0.00 *	−0.00 **	−0.00 **
CLO 2.0 vintage	−0.01 **	−0.01 **	−0.00	−0.01	−0.01 *
CLO 3.0 vintage	−0.02 ***	−0.02 ***	−0.00	−0.02 ***	−0.01 **
log(CLO issue size)	−0.00	−0.00	−0.00	−0.00	−0.01
CLO issue size missing	0.03 **	0.03 **	0.03 **	0.03 **	0.04 ***
JPM LL spread		0.02 ***		0.01 ***	0.02 **
Dealer CDS spread			0.08 ***	0.04 ***	0.01
Quarter FE	No	No	No	No	Yes
R^2	0.014	0.068	0.046	0.073	0.105
Panel B: Mezzanine tranche (AA–BBB, I = 3370 successful BWICs)					
Dealer CLO inventory	0.02 **	0.03 ***	0.02 **	0.02 **	−0.03
log(Par value of trade)	−0.01 **	−0.01 ***	−0.01 ***	−0.01 ***	−0.01 ***
CLO 2.0 vintage	−0.02	−0.01	−0.00	−0.00	−0.01
CLO 3.0 vintage	−0.04 ***	−0.04 ***	−0.01	−0.02 **	0.01
log(CLO issue size)	−0.00	−0.00	−0.00	−0.00	−0.00
CLO issue size missing	−0.03 ***	−0.03 ***	−0.03 **	−0.03 ***	−0.02
A rating	0.01	0.00	0.01	0.00	0.00
BBB rating	0.01	0.01	0.01	0.01	0.01
JPM LL spread		0.03 ***		0.02 ***	0.02
Dealer CDS spread			0.12 ***	0.08 ***	0.12
Quarter FE	No	No	No	No	Yes
R^2	0.016	0.036	0.035	0.043	0.072
Panel C: Junior tranche (BB–equity, I = 2853 successful BWICs)					
Dealer CLO inventory	0.02 **	0.01	0.02 *	0.01	−0.01
log(Par value of trade)	0.00	0.00	−0.00	0.00	−0.00
CLO 2.0 vintage	−0.00	−0.00	0.01	0.01	−0.00
CLO 3.0 vintage	−0.02 **	−0.03 ***	0.00	−0.01	0.01
log(CLO issue size)	−0.00	−0.00	−0.00	−0.00	−0.00
CLO issue size missing	0.01	0.00	0.01	0.01	0.03 *
B rating	0.03 **	0.03 **	0.03 *	0.03 **	0.02
Equity rating	0.05 ***	0.05 ***	0.05 ***	0.05 ***	0.06 ***
JPM LL spread		0.03 ***		0.02 ***	0.01
Dealer CDS spread			0.13 ***	0.09 ***	0.05
Quarter FE	No	No	No	No	Yes
R^2	0.016	0.033	0.035	0.041	0.104

* Significance at the 10% level.
** Significance at the 5% level.
*** Significance at the 1% level.

following expression:

$$\Pr(B \le b|\text{Trade}, R(r)) = \frac{\Pr(B \le b, \text{Trade}|R(r))}{\Pr(\text{Trade}|R(r))} = \frac{\Pr(\varepsilon_B \le \frac{b-\mu_B}{\sigma_B}, \eta \le \mathcal{S}|R(r))}{\Pr(\eta \le \mathcal{S}|R(r))} = \frac{C_1(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r))}{C_1(G_R(r), H(\mathcal{S}))} - 1. \tag{C.6}$$

Using properties of the conditional expectation and expression (C.6), the best bid conditional on trade and reserve price $R(r)$ in (C.8) equals

$$\mathbb{E}[B|\text{Trade}, R(r)] = \int_{\mathbb{D}_b} b \frac{\partial}{\partial b} \Pr(B \le b|\text{Trade}, R(r)) db = \int_{\mathbb{D}_b} b \frac{\frac{\partial}{\partial b} C_1(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r))}{C_1(G_R(r), H(\mathcal{S}))} db = \frac{1}{\sigma_B} \int_{\mathbb{D}_b} b g_B(\frac{b-\mu_B}{\sigma_B}) \frac{C_{11}(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r))}{C_1(G_R(r), H(\mathcal{S}))} db. \tag{C.7}$$

Hence,

$$\begin{aligned}\Pr(\text{Fail}|R(r))\mathbb{E}[B|\text{Trade}, R(r)] &= \Pr(\text{Fail}|R(r)) \times \int_{\mathbb{D}_b} b \frac{\partial}{\partial b} \Pr(B \le b, \text{Trade}|R(r)) db \\ &= \frac{\Pr(\text{Fail}|R(r))}{C_1(G_R(r), H(\mathcal{S}))} \times \int_{\mathbb{D}_b} b g_B(\frac{b-\mu_B}{\sigma_B}) \cdot C_{11}(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r)) \frac{db}{\sigma_B} \\ &= \mathcal{H}(r) \times \int_{\mathbb{D}_b} b g_B(\frac{b-\mu_B}{\sigma_B}) \cdot C_{11}(G_B(\frac{b-\mu_B}{\sigma_B}), G_R(r)) \frac{db}{\sigma_B}.\end{aligned}$$

Expression (26) follows using the expressions for conditional failure rate (C.4) and best bid (C.7).

To derive expression (28), note that we observe in the data the expected best bids in successful BWICs, which are given by

$$\mathbb{E}_R[B|\text{Trade}] = \mathbb{E}_R[\frac{\Pr(\text{Trade}|R)}{\Pr(\text{Trade})} \mathbb{E}[B|\text{Trade}, R]]$$

Table A.2
Simulation evidence on optimal rotation τ^*.
The table documents the optimal rotation τ^* across BWICs $i = 1, \dots, I$ in simulated data from the specification $B_i = \alpha + z_i + w_i + \sigma_B \epsilon_{B,i}$, $R_i = z_i + \sigma_R \epsilon_{R,i}$. BWIC i is successful if and only if $B_i \geq R_i$. We simulate $I = 100{,}000$ BWICs. In the specification, the variable $z \sim \mathcal{N}(0, \sigma_z^2)$ is a determinant of best bid B and reserve price R. The variable $w \sim \mathcal{N}(0, \sigma_w^2)$ is a determinant of best bid B only. The errors $\epsilon \sim \mathcal{N}(0,1)$ are iid noise. We normalize $\sigma_B = 1$. Across rows we vary the selection bias through $\lambda = \frac{\mathrm{SD}(\sigma_R * \epsilon_R)}{\mathrm{SD}(\epsilon_B)}$. Across panels we vary the expected surplus through $\alpha = \mathbb{E}[B] - \mathbb{E}[R]$. In Panel A, we set $\sigma_z = \sigma_w = 1$ to simulate a small ratio of unobserved to the observed variation. In Panel B, we set $\sigma_z = \sigma_w = 1/10$ to simulate a large ratio of unobserved to the observed variation.

Panel A: Optimal rotation τ^* with small ratio of unobserved to observed variation

			Distribution of τ^* across BWICs				
λ	Implied ρ	Pr(Trade)	Mean	S.D.	5%	50%	95%
High expected surplus $\alpha = 2$							
0.1	−1.00	0.92	0.01	0.01	0.00	0.00	0.03
0.5	−0.89	0.91	0.02	0.04	0.00	0.01	0.10
1	−0.71	0.87	0.04	0.07	0.00	0.01	0.17
2	−0.45	0.79	0.05	0.10	0.00	0.01	0.25
10	−0.10	0.58	0.08	0.13	0.00	0.02	0.36
Low expected surplus $\alpha = 1$							
0.1	−1.00	0.76	0.01	0.01	0.00	0.01	0.04
0.5	−0.89	0.75	0.07	0.09	0.00	0.04	0.25
1	−0.71	0.72	0.12	0.14	0.00	0.07	0.43
2	−0.45	0.66	0.17	0.19	0.00	0.10	0.58
10	−0.10	0.54	0.23	0.23	0.00	0.15	0.72
No expected surplus $\alpha = 0$							
0.1	−1.00	0.50	0.03	0.02	0.00	0.02	0.07
0.5	−0.89	0.50	0.18	0.14	0.01	0.15	0.47
1	−0.71	0.50	0.30	0.22	0.02	0.25	0.73
2	−0.45	0.50	0.39	0.27	0.03	0.36	0.87
10	−0.10	0.50	0.49	0.29	0.05	0.48	0.94

Panel B: Optimal rotation τ^* with large ratio of unobserved to observed variation

			Distribution of τ^* across BWICs				
λ	Implied ρ	Pr(Trade)	Mean	S.D.	5%	50%	95%
High expected surplus $\alpha = 2$							
0.1	−1.00	0.98	0.00	0.00	0.00	0.00	0.00
0.5	−0.89	0.96	0.01	0.00	0.00	0.01	0.01
1	−0.71	0.92	0.01	0.00	0.01	0.01	0.01
2	−0.45	0.81	0.02	0.00	0.01	0.02	0.03
10	−0.10	0.58	0.06	0.01	0.04	0.06	0.08
Low expected surplus $\alpha = 1$							
0.1	−1.00	0.84	0.00	0.00	0.00	0.00	0.00
0.5	−0.89	0.81	0.03	0.00	0.02	0.03	0.04
1	−0.71	0.76	0.07	0.01	0.05	0.07	0.09
2	−0.45	0.67	0.10	0.02	0.07	0.10	0.13
10	−0.10	0.54	0.20	0.03	0.16	0.20	0.25
No expected surplus $\alpha = 0$							
0.1	−1.00	0.50	0.01	0.00	0.01	0.01	0.01
0.5	−0.89	0.50	0.12	0.01	0.10	0.12	0.13
1	−0.71	0.50	0.22	0.02	0.19	0.22	0.26
2	−0.45	0.50	0.34	0.03	0.29	0.34	0.39
10	−0.10	0.50	0.47	0.04	0.41	0.47	0.53

$$= \int_{\mathbb{D}_r} \frac{\Pr(\text{Trade}|R(r))}{\Pr(\text{Trade})} \mathbb{E}[B|\text{Trade}, R(r)] dG_R(r). \tag{C.8}$$

where $\mathbb{E}_R[\cdot]$ is the expectation over all possible reserve prices $r \in \mathbb{D}_r$. Using relation (C.8) we have

$$\begin{aligned} \mathbb{E}_R[\Pr(\text{Trade}|R) \cdot \mathbb{E}[B|\text{Trade}, R]] &= \Pr(\text{Trade}) \cdot \mathbb{E}_R[B|\text{Trade}], \\ \mathbb{E}_R[\Pr(\text{Fail}|R) \cdot R] &= \Pr(\text{Fail})\mu_R + Cov_R(\Pr(\text{Fail}|R), R). \end{aligned} \tag{C.9}$$

The expected total payoff across all sellers is equal to

$$\begin{aligned} \mathbb{E}_R[\mathbb{E}[\Pi]] &= \mathbb{E}_R[\Pr(\text{Trade}|R) \cdot \mathbb{E}[B|\text{Trade}, R]] + \mathbb{E}_R[\Pr(\text{Fail}|R) \cdot R] \\ &= \mathbb{E}_R[B|\text{Trade}] - \Pr(\text{Fail}) \cdot (\mathbb{E}_R[B|\text{Trade}] - \mu_R) \\ &+ Cov_R(\Pr(\text{Fail}|R), R). \end{aligned} \tag{C.10}$$

The average TCI across all sellers (28) hence equals

$$\begin{aligned} \mathbb{E}_R[\text{TCI}] &= \mathbb{E}_R[\mathbb{E}[A|\text{Trade}, R] - \mathbb{E}[\Pi]] \\ &= \mathbb{E}_R[\mathbb{E}[A|\text{Trade}, R]] - \mathbb{E}_R[\Pr(\text{Trade}|R) \cdot \mathbb{E}[B|\text{Trade}, R]] \\ &- \mathbb{E}_R[\Pr(\text{Fail}|R) \cdot R] \\ &= \mathbb{E}_R[\mathbb{E}[P - B|\text{Trade}, R]] \\ &+ \mathbb{E}_R[\Pr(\text{Fail}|R) \cdot (\mathbb{E}[B|\text{Trade}, R] - R)] \\ &= \mathbb{E}_R[\mathbb{E}[P - B|\text{Trade}, R]] + \Pr(\text{Fail}) \cdot (\mathbb{E}_R[B|\text{Trade}] - \mu_R) \\ &\quad + \mathbb{E}_R[\mathbb{E}[B|\text{Trade}, R]] - \mathbb{E}_R[B|\text{Trade}] \\ &- Cov_R(\Pr(\text{Fail}|R), R) \\ &= \widehat{\text{TCI}} + \mathbb{E}_R[\mathbb{E}[B|\text{Trade}, R]] - \mathbb{E}_R[B|\text{Trade}] \\ &- Cov_R(\Pr(\text{Fail}|R), R). \quad \square \end{aligned} \tag{C.11}$$

Table A.3
GMM estimates of τ^*.
The table reports GMM estimates of τ^* for senior (Panel A), mezzanine (Panel B), and junior (Panel C) tranches. Models (1)–(6) use the same determinants X_i across BWICs i as in all previous tables to capture observed heterogeneity in the reserve price $R_i = \mu_R(Z_i) + \sigma_R \varepsilon_{R,i}$. The models assume a linear form $\mu_R(Z_i) = \alpha_R + \beta_R' Z_i$. The models differ in the assumptions about the unobserved heterogeneity in reserve price R_i. Models (1)–(3) assume the optimal rotation to recover expected reserve prices conditional on the observables Z_i is given by a Gaussian copula, $\tau^* = \Phi(\gamma S, S; \rho)/\Phi(S)$. Models (4)–(5) assume the optimal rotation is given by a semiparametric copula, $\tau^* = C(S)/H(S)$, where the copula is approximated by a polynomial in the failure probability $f = 1 - H(S)$ of chosen degree L, $C(S) = 1 - \sum_{l=1}^{L} \kappa_l f^l$, and $H(S)$ is the predicted probability of trade from a Probit model. Model (6) assumes the optimal rotation is given by a fixed τ^*. Tables A.4/A.5/A.6 report all other relevant details of the estimation.

	(1) Gaussian copula	(2)	(3)	(4) Semiparametric copula	(5)	(6)
	general	$\rho = 0$	$\omega = 0$	polynomial degree 1	polynomial degree 2	fixed τ^*
Panel A: Senior tranche (AAA, I = 9000 BWICs)						
Distribution of rotation τ^* across BWICs:						
Mean	0.0054	0.0160	0.0026	0.0015	0.0312	0.0041
S.D.	0.0077	0.0239	0.0015	0.0011	0.0151	0.0000
5%	0.0001	0.0003	0.0007	0.0003	0.0091	0.0041
50%	0.0026	0.0069	0.0023	0.0012	0.0300	0.0041
95%	0.0197	0.0609	0.0055	0.0036	0.0571	0.0041
Panel B: Mezzanine tranche (AA–BBB, I = 12,955 BWICs)						
Distribution of rotation τ^* across BWICs:						
Mean	0.0302	0.0603	0.0338	0.0197	0.0265	0.0056
S.D.	0.0194	0.0511	0.0262	0.0094	0.0159	0.0000
5%	0.0073	0.0076	0.0057	0.0077	0.0084	0.0056
50%	0.0260	0.0464	0.0272	0.0182	0.0230	0.0056
95%	0.0665	0.1585	0.0834	0.0368	0.0557	0.0056
Panel C: Junior tranche (BB–equity, I = 11,453 BWICs)						
Distribution of rotation τ^* across BWICs:						
Mean	0.1198	0.2919	0.0567	0.0415	0.1306	0.0097
S.D.	0.0358	0.0996	0.0205	0.0245	0.0413	0.0000
5%	0.0698	0.1520	0.0297	0.0178	0.0719	0.0097
50%	0.1165	0.2826	0.0542	0.0371	0.1270	0.0097
95%	0.1829	0.4696	0.0930	0.0793	0.2039	0.0097

Proof of (20). In the expression (19) for $\widehat{\mathrm{TCI}}(X)$, we condition on X and take expectation over all realizations of R as follows

$$\begin{aligned}
\widehat{\text{Bid–ask spread}}(X) &\equiv \mathbb{E}_R[\mathbb{E}[P - B|X, \text{Trade}, R]] \\
&= \int_{\mathbb{D}_R} \mathbb{E}[P - B|X, \text{Trade}, R(\varepsilon_R)] dG_R(\varepsilon_R), \\
\widehat{\text{BWIC fail rate}}(X) &\equiv \mathbb{E}_R[\Pr(\text{Fail}|X, R)] \\
&= \int_{\mathbb{D}_R} \Pr(\text{Fail}|X, R(\varepsilon_R)) dG_R(\varepsilon_R), \\
\widehat{\text{Failure cost}}(X) &\equiv \mathbb{E}_R[B|X, \text{Trade}] - \mu_R(Z).
\end{aligned} \tag{C.12}$$

In (C.12), the expected bid–ask spread is captured by (21), the expected BWIC failure rate is captured by (1), and the expected best bid is the observed price in successful BWICs:

$$\mathbb{E}_R[B|X, \text{Trade}] = \int_{\mathbb{D}_R} \frac{\Pr(\text{Trade}|X, R(\varepsilon_R))}{\Pr(\text{Trade}|X)} \mathbb{E}[B|X, \text{Trade}, R(\varepsilon_R)] dG_R(\varepsilon_R). \tag{C.13}$$

The expected best-bids conditional on the characteristics X and trade success that we require in expression (C.13) can be estimated by (22). *QED*

Appendix D. Simulation evidence on optimal quantile rotation

Here we provide simulation evidence on the optimal rotation τ^*. We simulate data on successful and failed BWICs. Best bids B_i and reserve prices R_i in BWIC $i = 1, \ldots, I$ are drawn randomly according to the following specification:

$$\begin{aligned}
B_i &= \alpha + z_i + w_i + \sigma_B \varepsilon_{B,i}, \\
R_i &= z_i + \sigma_R \varepsilon_{R,i}.
\end{aligned}$$

BWIC i is successful if and only if $B_i \geq R_i$. We simulate $I = 100{,}000$ BWICs. In the specification, the variable $z \sim \mathcal{N}(0, \sigma_z^2)$ is a determinant of best bid B and reserve price R. The variable $w \sim \mathcal{N}(0, \sigma_w^2)$ is a determinant of best bid B only. The errors $\varepsilon \sim \mathcal{N}(0, 1)$ are iid noise. Specifications with correlated errors yield similar results. We normalize $\sigma_B = 1$ and vary σ_R. Across rows we vary the selection bias through $\lambda = \frac{\mathrm{SD}(\sigma_R * \varepsilon_R)}{\mathrm{SD}(\varepsilon_B)}$ from 0.1 to 10. Across panels we vary the expected surplus through $\alpha = \mathbb{E}[B] - \mathbb{E}[R]$ between 0 and 2.

We estimate the model using the parametric specification (1). Table A.2, Panel A provides simulation evidence when the unobserved variation from $(\varepsilon_B, \varepsilon_R)$ is small compared to the observed variation in best bids and reserve prices, by setting $\sigma_z = \sigma_w = 1$. Table A.2, Panel B provides simulation evidence when the unobserved variation from $(\varepsilon_B, \varepsilon_R)$ is large compared to the observed variation in best bids and reserve prices, by setting $\sigma_z = \sigma_w = 1/10$. In the second column, we report the implied correlation ρ. In the third column, we report the probability of BWIC success. In the remaining columns, we report the distribution of optimal rotations τ^* across BWICs.

The implied correlation ρ varies between −1 for $\lambda = 0.1$ to −0.1 for $\lambda = 10$. The probability of BWIC success drops with λ and varies between 92% (98%) for $\alpha = 2$ and 50% for $\alpha = 0$ in Panel A. The simulations show the optimal rotations τ^* and its cross-sectional moments increase with the selection bias captured by λ. When $\lambda = 0.1$, the mean of τ^* across BWICs is no more than 3%. For $\lambda = 1$, the mean of τ^* across BWICs ranges between 4% and 30% in Table A.2 and between 1% and 22% in Panel B. The median of τ^* across BWICs is typically less than the mean because the right tail is skewed. Across sub-panels, the optimal rotations τ^* and its cross-sectional moments decrease with the expected surplus captured by α. When the trade probability exceeds 80%, the mean τ^* is less than 5%, and when the trade probability exceeds 90%, the mean τ^* drops below 3%. This range captures the trade probability observed in the data during normal times. Thus, we should expect τ^* to rarely exceed these numbers. When Pr(Trade) drops to 50%, however, τ^* can take a wide range of values. The observation-specific rotation τ^* captures this variability. Finally, comparing Panels A and B, the median τ^* decreases (increases) with the observed variation in best bids and reserve prices relative to the

Table A.4
GMM estimates of sellers' reserve prices for senior tranche.
The table reports the determinants of sellers' reserve prices across BWICs for senior tranches. Models (1)–(6) use the same determinants X_i across BWICs i as in all previous tables to capture observed heterogeneity in the reserve price $R_i = \mu_R(Z_i) + \sigma_R \varepsilon_{R,i}$. The models assume a linear form $\mu_R(Z_i) = \alpha_R + \beta_R' Z_i$. The models differ in the assumptions about the unobserved heterogeneity in reserve price R_i. Models (1)–(3) assume the optimal rotation to recover expected reserve prices conditional on the observables Z_i is given by a Gaussian copula, $\tau^* = \Phi(\gamma \mathcal{S}, \mathcal{S}; \rho)/\Phi(\mathcal{S})$. Models (4)–(5) assume the optimal rotation is given by a semiparametric copula, $\tau^* = C(\mathcal{S})/H(\mathcal{S})$, where the copula is approximated by a polynomial in the failure probability $f = 1 - H(\mathcal{S})$ of chosen degree L, $C(\mathcal{S}) = 1 - \sum_{l=1}^{L} \kappa_l f^l$, and $H(\mathcal{S})$ is the predicted probability of trade from a Probit model. Model (6) assumes the optimal rotation is given by a fixed τ^*. Panel A of Table A.3 reports τ^* estimates. Standard errors are robust to heteroskedasticity.

	(1) Gaussian copula	(2)	(3)	(4) Semiparametric copula	(5)	(6)
	general	$\rho = 0$	$\omega = 0$	polynomial degree 1	polynomial degree 2	fixed τ^*
Determinants of observed heterogeneity in reserve price R_i:						
log(Par value of trade)	0.66 ***	0.55 ***	0.13 *	0.22	0.59 ***	0.01
	(0.15)	(0.15)	(0.08)	(0.15)	(0.18)	(0.07)
CLO 2.0 vintage	2.04 ***	1.13 **	2.57 ***	2.75 ***	1.29 ***	2.41 *
	(0.14)	(0.52)	(0.23)	(0.68)	(0.33)	(1.34)
CLO 3.0 vintage	2.71 ***	1.81 ***	3.25 ***	2.89 ***	1.95 ***	3.17 **
	(0.23)	(0.51)	(0.31)	(0.58)	(0.28)	(1.54)
log(CLO issue size)	0.28 ***	0.32	0.75 ***	1.08 **	0.32 *	0.75
	(0.09)	(0.26)	(0.09)	(0.49)	(0.17)	(0.61)
CLO issue size missing	−0.88	−0.20	−1.09	−0.89	−0.37	−0.81
	(1.18)	(2.33)	(1.15)	(0.86)	(2.13)	(1.07)
JPM LL spread	−0.89 ***	−0.99 ***	−0.75 ***	−0.58 ***	−0.97 ***	−0.81 ***
	(0.18)	(0.22)	(0.09)	(0.10)	(0.03)	(0.10)
Dealer CDS spread	−4.70 **	−3.12 **	−8.42 ***	−8.63 ***	−3.67 ***	−8.23 ***
	(2.04)	(1.49)	(0.78)	(1.59)	(0.41)	(1.63)
Constant	−1.74 **	−0.87	−1.90 ***	−4.62	−0.93	−1.47
	(0.81)	(1.24)	(0.50)	(3.10)	(0.57)	(5.86)
Determinants of unobserved heterogeneity in reserve price R_i:						
γ	−1.44 ***	−1.60 ***				
	(0.22)	(0.15)				
$atanh(\rho)$	−1.16 **		−3.74 ***			
	(0.48)		(0.48)			
κ_1				0.02	−0.02	
				(0.01)	(0.04)	
κ_2					1.57 ***	
					(0.38)	
τ^*						0.00
						(0.00)
GMM statistics:						
No. of moments	10	10	10	10	10	10
No. of parameters	10	9	9	9	10	9
J	0.00	0.04	0.12	0.36	0.00	0.09
Q	1.43e−31	5.74e−06	1.68e−05	4.89e−05	5.39e−33	1.24e−05

* Significance at the 10% level.
** Significance at the 5% level.
*** Significance at the 1% level.

unobserved variation from $(\varepsilon_B, \varepsilon_R)$ for high (low) α. As one would expect, the relation reverses at the mean because the right tail is more skewed.

Appendix E. GMM estimation procedure

Choice of moments. Define the error term for each successful BWIC i by

$$\varepsilon_i = \frac{C(H(\gamma \mathcal{S}_i), H(\mathcal{S}_i); \rho)}{H(\mathcal{S}_i)} - \mathbb{1}(B_i \leq \alpha_R + \beta_R' Z_i), \tag{E.1}$$

where $\mathcal{S}_i$ is the score, B_i is the transaction price, and Z_i are the reserve price determinants if the BWIC is successful so that ε_i is missing otherwise. Then condition (18) yields by conditioning down the moment conditions

$$\mathbb{E}\left[\Xi(X_i) * \varepsilon_i \;|\text{Trade}_i\right] = 0, \tag{E.2}$$

for arbitrary functions $\Xi(\cdot)$ satisfying standard regularity conditions. We focus on functions Ξ/moments m that are a priori informative about the parameters we seek to estimate. Heuristically, a moment m is informative about an unknown parameter θ if that moment is sensitive to changes in the parameter and the sensitivity differs across parameters. Formally, local identification requires the Jacobian determinant, $\det(\partial m/\partial \theta)$, to be nonzero. For this condition to be satisfied, we require a suitable choice of moments (functions Ξ) to match, the instrumental variable W that is private information to dealers and affects their bidding, and functional form for copula C, which we explain now.

The score $\mathcal{S}_i$ and powers of the score are functions Ξ of X_i. They are informative about the parameters as $\mathcal{S}_i$ embodies X_i. In turn, X_i consists of two sets of observables, (Z_i, W_i). Based on this logic, we can deduce from (E.2) that the following moment conditions hold:

$$\begin{aligned} &\mathbb{E}\left[\varepsilon_i \;|\text{Trade}_i\right] = 0, \\ &\mathbb{E}\left[Z_i * \varepsilon_i \;|\text{Trade}_i\right] = 0, \\ &\mathbb{E}\left[\mathcal{S}_i * \varepsilon_i \;|\text{Trade}_i\right] = 0, \\ &\mathbb{E}\left[\mathcal{S}_i^2 * \varepsilon_i \;|\text{Trade}_i\right] = 0. \end{aligned} \tag{E.3}$$

Identification. It will be useful to provide some intuition and formal analysis for why the moment conditions (E.3) identify the parameters $(\alpha_R, \beta_R, \gamma, \rho)$. A sufficient condition for identification is a one-to-one mapping between the parameters and a set of data moments of the same dimension. For this to be the case, the model-implied moments must exhibit different sensitivity with respect to the parameters. Intuitively, the first moment in (E.3) identifies α_R, the second moment identifies β_R, the third moment identifies γ, and the fourth moment identifies

Table A.5
GMM estimates of sellers' reserve prices for mezzanine tranche.
The table reports the determinants of sellers' reserve prices across BWICs for mezzanine tranches. Models (1)–(6) use the same determinants X_i across BWICs i as in all previous tables to capture observed heterogeneity in the reserve price $R_i = \mu_R(Z_i) + \sigma_R \epsilon_{R,i}$. The models assume a linear form $\mu_R(Z_i) = \alpha_R + \beta_R' Z_i$. The models differ in the assumptions about the unobserved heterogeneity in reserve price R_i. Models (1)–(3) assume the optimal rotation to recover expected reserve prices conditional on the observables Z_i is given by a Gaussian copula, $\tau^* = \Phi(\gamma S, S; \rho)/\Phi(S)$. Models (4)–(5) assume the optimal rotation is given by a semiparametric copula, $\tau^* = C(S)/H(S)$, where the copula is approximated by a polynomial in the failure probability $f = 1 - H(S)$ of chosen degree L, $C(S) = 1 - \sum_{l=1}^{L} \kappa_l f^l$, and $H(S)$ is the predicted probability of trade from a Probit model. Model (6) assumes the optimal rotation is given by a fixed τ^*. Panel B of Table A.3 reports τ^* estimates. Standard errors are robust to heteroskedasticity.

	(1) Gaussian copula	(2)	(3)	(4) Semiparametric copula	(5)	(6)
	general	$\rho = 0$	$\omega = 0$	polynomial degree 1	polynomial degree 2	fixed τ^*
Determinants of observed heterogeneity in reserve price R_i:						
log(Par value of trade)	0.49	0.78 ***	0.79 **	0.30	0.39	−0.06
	(0.37)	(0.21)	(0.35)	(0.33)	(0.26)	(0.65)
CLO 2.0 vintage	0.85 *	1.26 **	1.00 *	0.73	0.82	0.21
	(0.44)	(0.50)	(0.59)	(1.78)	(0.68)	(1.20)
CLO 3.0 vintage	2.43 ***	2.15 ***	2.36 ***	2.59 ***	2.48 ***	3.21 ***
	(0.38)	(0.39)	(0.35)	(0.60)	(0.33)	(0.57)
log(CLO issue size)	−0.73 ***	−0.67 ***	−0.81 ***	−0.73 **	−0.69 **	−0.62
	(0.19)	(0.26)	(0.23)	(0.36)	(0.34)	(0.52)
CLO issue size missing	−0.87 **	0.59	0.40	−1.26 **	−0.90	−0.71
	(0.35)	(0.95)	(0.53)	(0.63)	(0.68)	(0.81)
A rating	−2.40 ***	−1.65 ***	−2.30 ***	−3.38 ***	−2.52 ***	−10.31 **
	(0.64)	(0.41)	(0.24)	(1.28)	(0.75)	(4.68)
BBB rating	−4.65 ***	−3.30 ***	−4.27 ***	−5.70 ***	−4.87 ***	−6.57 ***
	(0.56)	(0.43)	(0.28)	(0.79)	(0.55)	(0.57)
JPM LL spread	−2.30 ***	−1.94 ***	−2.31 ***	−2.30 ***	−2.30 ***	−2.27 ***
	(0.11)	(0.35)	(0.06)	(0.13)	(0.09)	(0.35)
Dealer CDS spread	−12.51 ***	−11.33 ***	−11.88 ***	−14.20 ***	−12.71 ***	−16.16 ***
	(1.40)	(0.58)	(0.31)	(0.44)	(0.65)	(3.38)
Constant	12.38 ***	10.39 ***	11.82 ***	13.34 ***	12.47 ***	12.96 ***
	(1.55)	(1.08)	(1.27)	(1.59)	(1.69)	(0.71)
Determinants of unobserved heterogeneity in reserve price R_i:						
γ	−1.32 ***	−1.68 ***				
	(0.15)	(0.06)				
$atanh(\rho)$	−1.18 ***		−0.79 ***			
	(0.29)		(0.06)			
κ_1				0.10 ***	0.09 *	
				(0.03)	(0.05)	
κ_2					0.24	
					(0.36)	
τ^*						0.01 **
						(0.00)
GMM statistics:						
No. of moments	12	12	12	12	12	12
No. of parameters	12	11	11	11	12	11
J	0.00	0.30	0.15	0.09	0.00	9.95
Q	9.76e−32	3.39e−05	1.70e−05	1.06e−05	4.72e−32	1.12e−03

* Significance at the 10% level.
** Significance at the 5% level.
*** Significance at the 1% level.

the correlation coefficient ρ. The last relation is more obvious once we recognize that $\mathbb{E}\left[S_i^2 \epsilon_i \,|\text{Trade}_i\right] = \mathbb{E}\left[S_i * (S_i \epsilon_i) \,|\text{Trade}_i\right]$ captures the covariance between the probability of trade, captured by the first term S, and the error term ϵ weighted by S, captured by the second term $S\epsilon$.

Fig. A.1 plots the moments in (E.3) as functions of the model parameters to check if this intuition holds in the model. We generate the figure by simulating BWICs as in Appendix D and then plot how the various moments m depend on each parameter, θ. The subplots in Fig. A.1 reveal that the model moments exhibit significant sensitivity to the model parameters. More importantly for identification, the sensitivities differ across parameters, such that one can find moments with $\det(\partial m/\partial \theta) \neq 0$. Moments 1 and 4 are highly sensitive to α_R while moment 2 is insensitive to α_R and moment 3 has intermediate sensitivity to α_R. By contrast, moment 2 is highly sensitive to β_R while moments 1, 3, and 4 are insensitive to β_R. Panel C shows that γ is identified from moment 3, $\mathbb{E}\left[S_i \epsilon_i \,|\text{Trade}_i\right] = 0$, as it is sensitive to γ while the other moments are not or at least less. Panel D shows that ρ is identified predominantly from moment 4, $\mathbb{E}\left[S_i^2 \epsilon_i \,|\text{Trade}_i\right] = 0$, as it is sensitive to ρ while in particular moment 3 is rather insensitive to ρ.

Choice of instrument. For identification of $(\alpha_R, \beta_R, \gamma, \rho)$ in the data, we require characteristics W that serve as instruments in relation (8). We use aggregate dealers' CLO inventory positions to instrument their demand. The exclusion restriction here is that dealers' private information affecting the BWIC success probability through a demand shock to dealers' bids is uncorrelated with the reserve price. Intuitively, as dealers' inventories increase, their demand declines, and so does their propensity to bid and the reduced number of bidders results in less competitive bids. Importantly, inventory statistics are collected by the New York Fed on Wednesdays and published with a one-week delay. Therefore, it is reasonable to assume that the seller does not know inventories at the time of the auction so they cannot directly affect the reserve price thus making dealers' inventories a natural choice for W.

Choice of copula(s). We consider three approaches to parametrize copula $C(\cdot)$ in (18).

1. Parametric copula. The most straightforward approach is to assume a copula form. For the Gaussian copula, the rotation in (18) is simply

Table A.6
GMM estimates of sellers' reserve prices for junior tranche.
The table reports the determinants of sellers' reserve prices across BWICs for junior tranches. Models (1)–(6) use the same determinants X_i across BWICs i as in all previous tables to capture observed heterogeneity in the reserve price $R_i = \mu_R(Z_i) + \sigma_R \epsilon_{R,i}$. The models assume a linear form $\mu_R(Z_i) = \alpha_R + \beta_R' Z_i$. The models differ in the assumptions about the unobserved heterogeneity in reserve price R_i. Models (1)–(3) assume the optimal rotation to recover expected reserve prices conditional on the observables Z_i is given by a Gaussian copula, $\tau^* = \Phi(\gamma \mathcal{S}, \mathcal{S}; \rho)/\Phi(\mathcal{S})$. Models (4)–(5) assume the optimal rotation is given by a semiparametric copula, $\tau^* = \mathcal{C}(\mathcal{S})/H(\mathcal{S})$, where the copula is approximated by a polynomial in the failure probability $f = 1 - H(\mathcal{S})$ of chosen degree L, $\mathcal{C}(\mathcal{S}) = 1 - \sum_{l=1}^{L} \kappa_l f^l$, and $H(\mathcal{S})$ is the predicted probability of trade from a Probit model. Model (6) assumes the optimal rotation is given by a fixed τ^*. Panel C of Table A.3 reports τ^* estimates. Standard errors are robust to heteroskedasticity.

	(1) Gaussian copula	(2)	(3)	(4) Semiparametric copula	(5)	(6)
	general	$\rho = 0$	$\omega = 0$	polynomial degree 1	polynomial degree 2	fixed τ^*
Determinants of observed heterogeneity in reserve price R_i:						
log(Par value of trade)	0.27	0.91 ***	0.73	0.52	0.35	−0.68
	(0.43)	(0.31)	(0.79)	(1.01)	(0.25)	(1.08)
CLO 2.0 vintage	−2.60 ***	−3.30 ***	−1.51	−1.31	−2.89 **	−1.07
	(0.98)	(1.16)	(1.56)	(6.46)	(1.18)	(5.23)
CLO 3.0 vintage	−0.12	−0.44	0.84	1.10	−0.32	0.61
	(1.22)	(1.09)	(2.34)	(4.43)	(1.20)	(3.75)
log(CLO issue size)	−1.23 **	−1.37 *	−1.41	−1.58	−1.10	−1.47 *
	(0.57)	(0.83)	(1.68)	(1.78)	(1.18)	(0.76)
CLO issue size missing	2.60	2.94 ***	3.72	3.61	2.41 ***	0.47
	(1.86)	(1.05)	(2.92)	(5.18)	(0.82)	(2.66)
B rating	−6.13 ***	−4.60 ***	−7.01 ***	−6.68 *	−5.78 ***	−10.19 ***
	(0.88)	(0.66)	(1.97)	(3.67)	(1.06)	(2.66)
Equity rating	−36.27 ***	−25.36 ***	−41.98 ***	−43.11 ***	−35.14 ***	−58.78 ***
	(2.90)	(3.49)	(3.21)	(1.01)	(0.98)	(2.30)
JPM LL spread	−5.32 ***	−4.06 ***	−6.24 ***	−6.31 ***	−5.07 ***	−9.02 ***
	(0.37)	(0.57)	(2.10)	(0.70)	(0.21)	(0.28)
Dealer CDS spread	−21.89 ***	−16.55 ***	−25.01 ***	−26.25 ***	−21.95 ***	−28.51 ***
	(2.46)	(1.63)	(9.22)	(7.02)	(1.54)	(7.75)
Constant	25.41 ***	22.05 ***	26.46 ***	26.23 ***	24.66 ***	32.88 ***
	(1.55)	(2.56)	(8.32)	(4.93)	(3.88)	(4.07)
Determinants of unobserved heterogeneity in reserve price R_i:						
γ	−1.01 ***	−1.06 ***				
	(0.09)	(0.12)				
$atanh(\rho)$	−1.16 ***		−1.84 ***			
	(0.16)		(0.35)			
κ_1				0.09 ***	0.43 ***	
				(0.02)	(0.08)	
κ_2					−0.43 ***	
					(0.13)	
τ^*						0.01 ***
						(0.00)
GMM statistics:						
No. of moments	12	12	12	12	12	12
No. of parameters	12	11	11	11	12	11
J	0.00	1.57	0.20	0.34	0.00	0.16
Q	5.28e−32	2.48e−04	3.08e−05	5.43e−05	4.48e−32	2.46e−05

* Significance at the 10% level.
** Significance at the 5% level.
*** Significance at the 1% level.

Table A.7
Specification tests.
The table reports specification tests for model (1) that assumes the optimal rotation to recover expected reserve prices conditional on the observables X_i is given by a Gaussian copula, $\tau^* = \Phi(\gamma \mathcal{S}, \mathcal{S}; \rho)/\Phi(\mathcal{S})$, where ρ and γ are the estimated structural parameters.

	Senior	Mezzanine	Junior
Point estimates:			
ρ	−0.82	−0.83	−0.82
$1/\gamma$	−0.69	−0.76	−0.99
(1) No selection test, $H_0 : \rho = 0$			
$\chi^2(1)$	27.78 ***	82.31 ***	250.66 ***
Prob > χ^2	0.0000	0.0000	0.0000
(2) No unobserved heterogeneity test, $H_0 : \rho = -1$			
$\chi^2(1)$	1.30	3.61 *	12.16 ***
Prob > χ^2	0.2542	0.0576	0.0005
(3) Pure private value test, $H_0 : \gamma * \rho = 1$			
$\chi^2(1)$	4.38 **	5.53 **	8.33 ***
Prob > χ^2	0.0364	0.0187	0.0039

* Significance at the 10% level.
** Significance at the 5% level.
*** Significance at the 1% level.

Table A.8
TCI across BWICs using OQR approach.
The table reports the estimates of TCI, $\widehat{\text{TCI}}$, across BWICs for senior tranches (Panel A), mezzanine tranches (Panel B), and junior tranches (Panel C). The models (1)–(6) use the same specifications as in Tables A.4–A.6. Models (1)–(3) assume the optimal rotation is given by a Gaussian copula, $\tau^* = \Phi(\gamma S, S; \rho)/\Phi(S)$. Models (4)–(5) assume the optimal rotation is given by a semiparametric copula, $\tau^* = C(S)/H(S)$. Model (6) assumes the optimal rotation is given by a fixed τ^*.

	(1) Gaussian copula	(2)	(3)	(4) Semiparametric copula	(5)	(6)
	general	$\rho = 0$	$\omega = 0$	polynomial degree 1	polynomial degree 2	fixed τ^*
Panel A: Senior tranche (AAA, I = 8683 BWICs)						
$\widehat{\text{TCI}}$ mean	0.31	0.24	0.34	0.37	0.26	0.34
$\widehat{\text{TCI}}$ S.D.	0.23	0.17	0.37	0.38	0.19	0.37
Dispersion of TCI across BWICs:						
5%	0.09	0.08	0.06	0.08	0.08	0.05
25%	0.16	0.13	0.13	0.15	0.14	0.12
50%	0.25	0.20	0.22	0.24	0.21	0.21
75%	0.36	0.28	0.41	0.45	0.30	0.40
95%	0.78	0.56	1.09	1.13	0.62	1.08
Panel B: Mezzanine tranche (AA–BBB, I = 12,341 BWICs)						
$\widehat{\text{TCI}}$ mean	1.32	0.99	1.28	1.55	1.38	2.80
$\widehat{\text{TCI}}$ S.D.	0.87	0.61	0.80	1.07	0.95	1.95
Dispersion of TCI across BWICs:						
5%	0.42	0.38	0.44	0.43	0.41	0.55
25%	0.74	0.61	0.76	0.84	0.76	1.45
50%	1.10	0.83	1.08	1.29	1.14	2.44
75%	1.61	1.17	1.53	1.91	1.70	3.55
95%	3.02	2.18	2.82	3.65	3.25	6.58
Panel C: Junior tranche (BB–equity, I = 10,794 BWICs)						
$\widehat{\text{TCI}}$ mean	3.79	1.25	5.45	6.11	3.53	9.75
$\widehat{\text{TCI}}$ S.D.	2.94	1.07	4.14	4.50	2.72	7.71
Dispersion of TCI across BWICs:						
5%	1.03	0.00	1.55	1.84	0.96	2.56
25%	1.82	0.61	2.71	3.13	1.70	4.68
50%	2.77	1.07	4.07	4.63	2.59	7.25
75%	4.55	1.70	6.49	7.27	4.29	11.58
95%	10.40	3.16	14.79	16.33	9.57	27.41

the ratio of binormal to normal CDF:

$$\tau^* = \frac{\Phi(\gamma \hat{S}, \hat{S}; \rho)}{\Phi(\hat{S})} = \frac{\Phi(\gamma \Phi^{-1}(\Pr(\text{Trade})), \Phi^{-1}(\Pr(\text{Trade})); \rho)}{\Pr(\text{Trade})}. \tag{E.4}$$

In the implementation of the moment conditions, we use powers of $\hat{S}$ that depend on W to identify γ and ρ. We also consider two special cases of our parametric specification. We estimate (E.4) under the assumption that the unobserved components of the bid and reserve prices are uncorrelated: $\omega = 0$, in which case $\gamma = 1/\rho$. Alternatively, we assume $\omega = \frac{\sigma_B}{\sigma_R} = 1/\lambda$, in which case $\gamma = -\sqrt{(\frac{\sigma_R}{\sigma_B})^2 - 1}$ and $\rho = 0$.

2. Semiparametric copula. An alternative approach to estimating (18) is to estimate $C(\cdot)$ semiparametrically. Here we can exploit the boundary conditions

$$\lim_{s\to\infty} H(s) = 1, \qquad \lim_{s\to\infty} C(s) = 1. \tag{E.5}$$

Condition (E.5) allows us to avoid estimating a constant term when we use appropriate basis functions. If we use the probability of trade failure, $f = 1 - H(s)$, as basis function and assume that the copula is well approximated by a polynomial in f of chosen degree L, we have

$$C(s) \approx 1 - \sum_{l=1}^{L} \kappa_l f^l, \tag{E.6}$$

where we have used that $\kappa_0 = 0$. In the implementation of the moment conditions, we again use powers of $\hat{S}$ depending on W to identify $\kappa_l, l = 1, \ldots, L$. We report results for $L = 1, 2$.

3. Constrained quantile rotation with endogenous $\tau^\star$. In a third approach, we estimate a fixed $\tau^\star$ endogenously based on GMM condition (17). This approach is equivalent to the CQR approach in Section 6.4, but we do not set $\tau^\star$ and instead estimate it together with the reserve price determinants.

We evaluate the performance of our OQR approaches outlined by (16)–(18) and (E.1)–(E.6) using simulations. We simulate data on successful and failed BWICs and then estimate the optimal rotation τ^* and the correlation ρ using the parametric specification (1).

Appendix D discusses the simulations in greater detail and the results are reported in Table A.2. The simulations show the optimal rotation τ^* increases with the selection bias captured by λ. When $\lambda = 0.1$, the mean of τ^* across BWICs is no more than 3%, while for $\lambda = 1$, the mean of τ^* across BWICs ranges between 1% and 30%. When the trade probability exceeds 80%, the mean τ^* is less than 5%, and when the trade probability exceeds 90%, the mean τ^* drops below 3%. This range captures the trade probability observed in the data during normal times. Thus, we should expect τ^* to rarely exceed these numbers.

GMM estimates for OQR. We estimate the optimal quantile rotation, τ^*, that depends on the parameters (γ, ρ) by GMM using (E.3). We use the same set of explanatory variables X_i as in Sections 3 and 4 to capture observed heterogeneity in the reserve price $R_i = \mu_R(Z_i) + \sigma_R \varepsilon_{R,i}$, where we assume a linear form for $\mu_R(Z_i) = \alpha_R + \beta_R' Z_i$. The explanatory variables include characteristics of the trade (size measured as log of par), characteristics of the CLO (log of CLO amount outstanding, CLO vintage, CLO rating), as well as market and dealer conditions (JPM leveraged loan spread and dealer CDS spread). We omit quarter fixed effects as leveraged loan spread and dealer CDS spread capture well the time series variation in the data and the estimation would otherwise be too high-dimensional. In addition, to account for differences in credit risks, we perform the estimation separately for senior, mezzanine, and junior CLO tranches.

We employ six different model specifications for τ^*. The model specifications differ in the assumptions about the unobserved heterogeneity in reserve price R_i. Models (1)–(3) assume the optimal rotation to recover expected reserve prices conditional on the observables Z_i

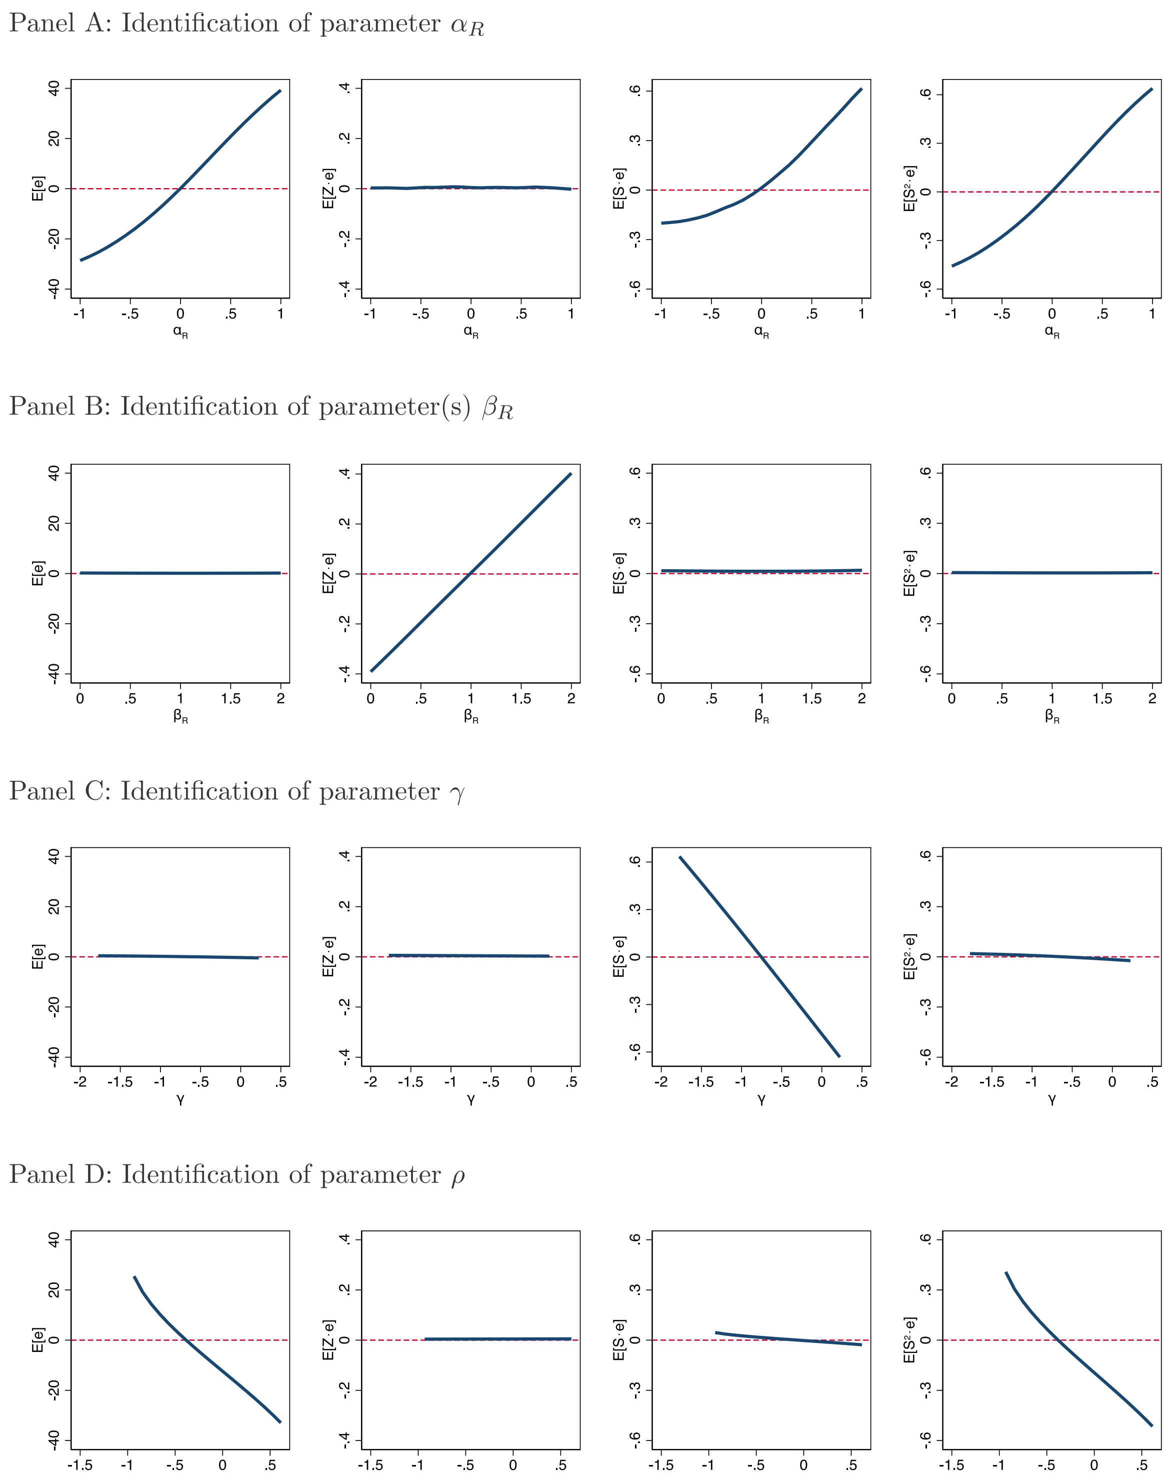

Fig. A.1. Identification of model parameters.

is given by a Gaussian copula, $\tau^* = \Phi(\gamma\mathcal{S},\mathcal{S};\rho)/\Phi(\mathcal{S})$. Models (4)–(5) assume the optimal rotation is given by a semiparametric copula, $\tau^* = C(\mathcal{S})/H(\mathcal{S})$, where the copula is approximated by a polynomial in the failure probability $f = 1 - H(\mathcal{S})$ of chosen degree L, $C(\mathcal{S}) = 1 - \sum_{l=1}^{L} \kappa_l f^l$, and $H(\mathcal{S})$ is the predicted probability of trade from a Probit model. Model (6) assumes the optimal rotation is given by a fixed τ^*.

Tables A.3–A.7 summarize the distribution of rotations τ^* across BWICs, the GMM estimates of the sellers' reserve prices, and specification tests. We discuss these findings here briefly. In Table A.3, we report the distribution of rotation τ^* across BWICs for senior (Panel A), mezzanine (Panel B), and junior (Panel C) CLO tranches. All selection-adjusted rotations are positive and show a significant variation across BWICs. The rotations τ^* are larger for lower-ranked tranches. For the

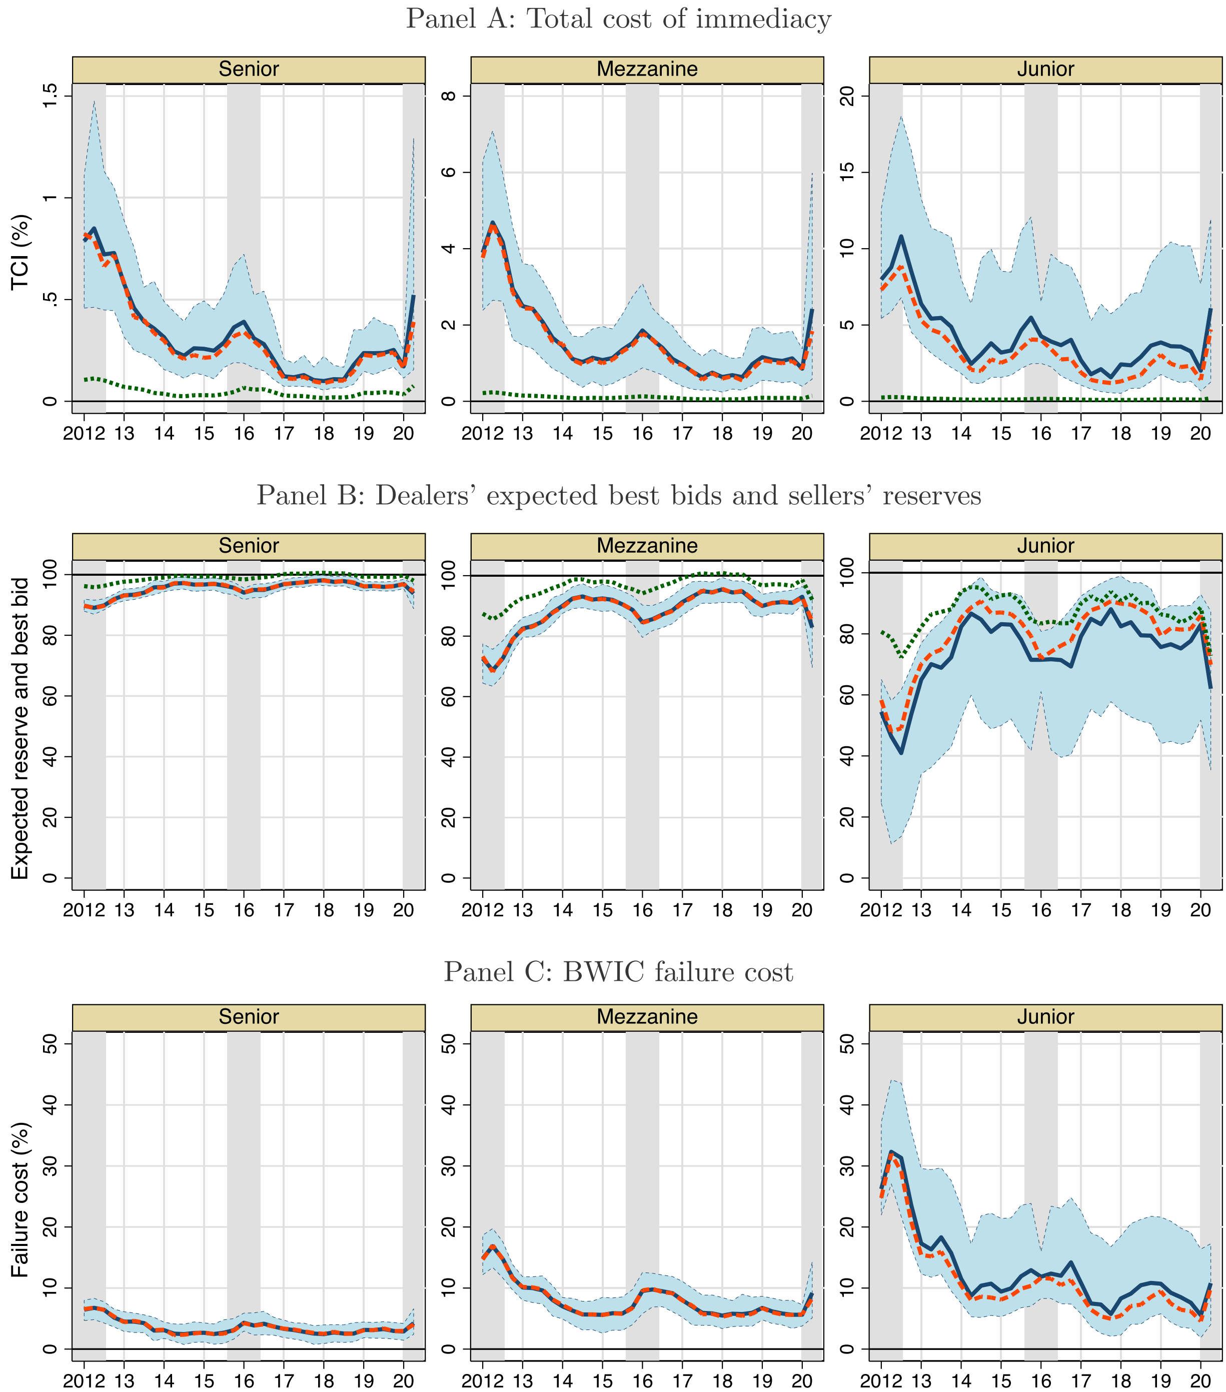

Fig. A.2. Semiparametric specification for TCI, sellers' reserves, and BWIC failure cost
Panel A of the figure documents the quarterly total cost of immediacy in the CLO market using the semiparametric specification (5) for estimating the seller's reserve price. Panel B of the figure documents sellers' reserves (solid and dashed lines) and dealers' expected best bids (dotted line) expressed in terms of par value in the CLO market at a quarterly frequency. Panel C of the figure shows the BWIC failure cost (solid and dashed line) expressed in terms of par value in the CLO market at a quarterly frequency. In all panels, the solid line depicts the mean value across BWICs at a quarterly frequency. The dashed line corresponds to the median across BWICs. The shaded area indicates the 5% to 95% range across BWICs. The dotted line in Panel A indicates the bid–ask spreads in the CLO market. The dotted line in Panel B indicates the dealers' expected best bids. The sample period is 2012–2020 with grey areas indicating periods of market stress: the 2012 European debt crisis, the 2015–16 credit stress, and the 2020 COVID-19 pandemic. Results are split by tranche into the senior tranche rated AAA (left), mezzanine tranche rated AA–BBB (middle), and junior tranche rated BB–equity (right).

senior tranche, τ^* is 54 bps with a standard deviation of 77 bps in specification (1). For specification (5), τ^* for the senior tranche is higher at 3.12% mean (1.51% S.D.). The mean τ^* for the mezzanine tranche is 3.02%/2.65% in specification (1)/(5), with a standard deviation equal to 1.94%/1.59%. The mean τ^* for the junior tranche is 11.98%/13.06% in specification (1)/(5), with a standard deviation 3.58%/4.13%. These numbers yield the fraction of the lowest best bids that need to be truncated to capture the sellers' expected reserve prices.

Tables A.4/A.5/A.6 report the GMM estimates of the sellers' reserve prices for the senior/ mezzanine/junior tranche and across six model specifications. Results presented in Tables A.4/A.5/A.6 rely on the distribution of rotation τ^* across BWICs reported in Table A.3 and thus these tables shares their layouts with Table A.3. Across panels, the tables report the determinants of observed heterogeneity in reserve prices, (α_R, β_R), the determinants of unobserved heterogeneity in reserve prices, $(\gamma, \rho, \kappa_l$, fixed $\tau^*)$, and GMM statistics. For the senior tranche, Table A.4 shows that sellers have higher reserve prices in BWICs with larger par value of trade and larger CLO issue size. Sellers have higher reserve prices in CLO 2.0 and CLO 3.0 vintages than in CLOs of 1.0 vintage. Higher leverage loan spreads and higher dealer CDS spreads lower sellers' reserve prices. Results for the mezzanine tranche in Table A.5 and the junior tranche in Table A.6 are generally similar to the senior tranche: The larger par value of trade, lower leverage loan spreads, and lower dealer CDS spreads increase the reserve price. As one would expect, sellers have lower reserve prices in CLOs with lower credit ratings. However, larger CLO issue size leads to lower reserve prices for the mezzanine and junior tranches. For all tranches, the general Gaussian copula (model 1) and the semiparametric copula with polynomial degree 2 (model 5) provide the best fit with the lowest Q value. For the Gaussian copula, the γ and ρ estimates are significant across all tranches, while the κ_l coefficients have varying significance.

Table A.7 reports specification tests for our main specification (1) (Gaussian copula $\tau^* = \Phi(\gamma S, S; \rho)/\Phi(S)$ where ρ and γ are the estimated structural parameters). The point estimates for ρ are −0.82, −0.83, and −0.82 for senior, mezzanine, and junior tranches, respectively. The point estimates for $1/\gamma$ are −0.69, −0.76, and −0.99. To understand what economic features are important to include, we perform three specification tests: (1) No selection test, $H_0 : \rho = 0$, which is rejected for all tranches with $\chi^2(1) = 27.78$, 82.31, 250.66; (2) No unobserved heterogeneity test, $H_0 : \rho = -1$, which is rejected for mezzanine at 5.76% and junior at 0.05%, with $\chi^2(1)$=1.30, 3.61, 12.16; for the senior tranche, the no unobserved heterogeneity cannot be rejected; (3) Pure private value test, $H_0 : \gamma * \rho = 1$, which is rejected at 3.64%, 1.87%, 0.39% with $\chi^2(1)$=4.38, 5.53, 8.33. The pure private value test is consistent with there being adverse selection risk, especially in the lower tranches.

References

Allen, J., Wittwer, M., 2023. Centralizing over-the-counter markets? J. Polit. Econ. 131, 3310–3351.
Arellano, M., Bonhomme, S., 2016. Sample selection in quantile regression: A survey. In: Koenker, Roger, Chernozhukov, Victor, He, Xuming, Peng, Limin (Eds.), Handbook of Quantile Regression. CRC, Boca Raton, FL.
Arellano, M., Bonhomme, S., 2017. Quantile selection models with an application to understanding changes in wage inequality. Econometrica 85, 1–28.
Atkeson, A., Eisfeldt, A., Weill, P.O., 2015. Entry and exit in OTC derivatives markets. Econometrica 83, 2231–2292.
Becker, B., Ivashina, V., 2015. Reaching for yield in the bond market. J. Finance 70, 1863–1902.
Bessembinder, H., Jacobsen, S., Maxwell, W., Venkataraman, K., 2018. Capital commitment and illiquidity in corporate bonds. J. Finance 73, 1615–1661.
Bessembinder, H., Maxwell, W., Venkataraman, K., 2006. Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. J. Financ. Econ. 82, 251–288.
Bessembinder, H., Panayides, M., Venkataraman, K., 2009. Hidden liquidity: an analysis of order exposure strategies in electronic stock markets. J. Financ. Econ. 94, 361–383.
Bessembinder, H., Spatt, C., Venkataraman, K., 2020. A survey of the microstructure of fixed-income markets. J. Financ. Quant. Anal. 55, 1–45.
Bollinger, C.R., Hirsch, B.T., Hokayem, C.M., Ziliak, J.P., 2019. Trouble in the tails? What we know about earnings nonresponse 30 years after Lillard, Smith, and Welch. J. Polit. Econ. 127, 2143–2185.
Campo, S., Perrigne, I., Vuong, Q., 2003. Asymmetry in first-price auctions with affiliated private values. J. Appl. Econometrics 18, 179–207.
Cassola, N., Hortaçsu, A., Kastl, J., 2013. The 2007 subprime market crisis in the EURO area through the lens of ECB repo auctions. Econometrica 81, 1309–1345.
Chade, H., Eeckhout, J., Smith, L., 2017. Sorting through search and matching models in economics. J. Econ. Lit. 55, 493–544.
Chernozhukov, V., Hansen, C., 2005. An IV model of quantile treatment effects. Econometrica 73, 245–262.
Chiappori, P.A., 2020. The theory and empirics of the marriage market. Annu. Rev. Econ. 12, 547–578.
Cordell, L., Roberts, M.R., Schwert, M., 2023. CLO performance. J. Finance 78, 1235–1278.
Demsetz, H., 1968. The cost of transacting. Q. J. Econ. 82, 33–53.
Devine, T.J., Kiefer, N.M., 1991. Empirical Labor Economics: The Search Approach. Oxford University Press, Oxford.
Dick-Nielsen, J., Feldhütter, P., Lando, D., 2012. Corporate bond liquidity before and after the onset of the subprime crisis. J. Financ. Econ. 103, 471–492.
Dick-Nielsen, J., Rossi, M., 2019. The cost of immediacy for corporate bonds. Rev. Financ. Stud. 32, 1–41.
Duffie, D., Garleanu, N., Pedersen, L., 2005. Over-the-counter markets. Econometrica 73, 1815–1847.
Ellul, A., Jotikasthira, C., Lundblad, C., 2011. Regulatory pressure and fire sales in the corporate bond market. J. Financ. Econ. 101, 596–620.
Friewald, N., Jankowitsch, R., Subrahmanyam, M.G., 2012. Illiquidity or credit deterioration: A study of liquidity in the U.S. corporate bond market during financial crises. J. Financ. Econ. 105, 18–36.
Gimenes, N., Guerre, E., 2021. Quantile regression methods for first-price auctions. J. Econometrics 226, 224–247.
Green, R.C., Hollifield, B., Schürhoff, N., 2007. Financial intermediation and the costs of trading in an opaque market. Rev. Financ. Stud. 20, 275–314.
Haile, P.A., Kitamura, Y., 2019. Unobserved heterogeneity in auctions. Econom. J. 22, 1–19.
Harris, L.E., Hasbrouck, J., 1996. Market vs. limit orders: The SuperDOT evidence on order submission strategy. J. Financ. Quant. Anal. 31, 213–231.
Harris, L.E., Piwowar, M.S., 2006. Municipal bond liquidity. J. Finance 61, 1330–1366.
Hendershott, T., Madhavan, A., 2015. Click or call? Auction versus search in the over-the-counter market. J. Finance 70, 419–444.
Hong, G., Warga, A., 2000. An empirical study of bond market transactions. Financ. Anal. J. 56, 32–46.
Hortaçsu, A., Kastl, J., Zhang, A., 2018. Bid shading and bidder surplus in the U.S. treasury auction system. Amer. Econ. Rev. 108, 147–169.
Hugonnier, J., Lester, B., Weill, P.O., 2020. Frictional intermediation in over-the-counter markets. Rev. Econ. Stud. 87, 1432–1469.
Kargar, M., Lester, B., Lindsay, D., Liu, S., Weill, P.O., Zuniga, D., 2020. Corporate bond liquidity during the COVID-19 crisis. Rev. Financ. Stud. 34, 5352–5401.
Koenker, R., Bassett, Jr., G., 1978. Regression quantiles. Econometrica 46, 33–50.
Krasnokutskaya, E., 2011. Identification and estimation of auction models with unobserved heterogeneity. Rev. Econ. Stud. 78, 293–327.
Lagos, R., Rocheteau, G., 2007. Search in asset markets: Market structure, liquidity, and welfare. Amer. Econ. Rev. 97, 198–202.
Lagos, R., Rocheteau, G., 2009. Liquidity in asset markets with search frictions. Econometrica 77, 403–426.
Li, D., Schürhoff, N., 2019. Dealer networks. J. Finance 74, 91–144.
Maasoumi, E., Wang, L., 2019. The gender gap between earnings distributions. J. Polit. Econ. 127, 2438–2504.
Machado, J.A.F., Silva, J.M.C.S., 2019. Quantiles via moments. J. Econometrics 213, 145–173.
Malamud, S., Rostek, M., 2017. Decentralized exchange. Amer. Econ. Rev. 107, 3320–3362.
O'Hara, M., Zhou, X., 2020. Anatomy of a liquidity crisis: Corporate bonds in the Covid-19 crisis. J. Financ. Econ. 142, 46–68.
Perold, A., 1988. The implementation shortfall: Paper versus reality. J. Portf. Manag. 114, 4–9.
Riggs, L., Onur, E., Reiffen, D., Zhu, H., 2020. Swap trading after dodd-frank: Evidence from index CDS. J. Financ. Econ. 137, 857–886.
Roberts, J.W., 2014. Unobserved heterogeneity and reserve prices in auctions. Rand J. Econ. 44, 712–732.
Weill, P.O., 2020. The search theory of OTC markets. Annu. Rev. Econ. 12, 747–773.