

A Missing Market: The Cost of Term Lending in Equity Markets

Spencer Andrews¹

First Draft: August 23, 2026. This version: September 14, 2026.

ABSTRACT

Short sellers borrow shares at fees that reset daily, and no term rate is quoted. Using U.S. equity lending data for 2007–2024, this paper measures the cost of supplying a fixed-rate term loan by rolling overnight borrows. I find that locking in the expected fee is nearly costless, because the current fee predicts it well. Most of the cost comes from the risk that fees move. This risk grows with the term and is concentrated in stocks with unforecastable fees. Loans in these stocks are also closed sooner. At six months, option prices imply a term borrowing rate of 71 basis points against a supply cost of 143. That rate covers the expected carry but only about a quarter of the risk. The shortfall is much smaller for a diversified provider. Cash settlement and pooling would reduce the cost of protection on the least predictable stocks by 86 percent.

JEL Classification: G12, G14, G23

Keywords: Securities lending, short selling, term structure, borrowing costs, limits to arbitrage

¹Villanova School of Business, Villanova University; spencer.andrews@villanova.edu. I thank Adam Reed and seminar participants for helpful comments, and a securities lending practitioner for a conversation that informed the institutional discussion and motivated several of the tests below. Code that reproduces every exhibit is available from the author. Tables and figures numbered IA.*n* appear in the Internet Appendix. I have no relevant financial interests to disclose.

1 Introduction

Every major funding market has a term structure. Repurchase agreements, foreign exchange swaps, and unsecured bank funding all quote rates from overnight out to one year, and market participants use these curves to price, hedge, and lay off financing risk. The equity lending market, which finances essentially all short selling, has no term structure. Nearly every equity loan is open: the lender can reset the fee daily, either side can terminate at any time, and no term rate is quoted at any horizon.²

The absence of a term market matters for short sellers. The literature has modeled the short-sale constraint as a level, a fee that must be paid and that lowers the return to shorting (Miller, 1977; Jones and Lamont, 2002; Duffie et al., 2002). A short seller who plans to hold a position for several months faces a second constraint, however, because she can observe today’s fee but cannot lock it in. If the fee rises sharply, she must pay the higher rate or close the position, and if her shares are recalled and cannot be replaced, she must buy them back. She therefore bears a risk in addition to a cost, and she bears it in the stocks where anomaly returns are most concentrated on the short side (Stambaugh et al., 2012). An arbitrageur who faces the risk of forced liquidation treats these two constraints very differently (Shleifer and Vishny, 1997; Gromb and Vayanos, 2002). In this paper, I measure this second constraint and ask why the contract that would remove it does not exist.

Answering this question requires both the price of a term loan and its cost of supply. Option prices reveal the price, since a synthetic short position embeds a borrowing rate over the life of the options, and several papers extract shorting-cost term structures in this way (Weitzner, 2023; Cocquemas, 2017; Muravyev et al., 2022). Option prices cannot reveal the cost of supply, however, and a product is written only if its price covers that cost. Nor can term loans be observed directly. Prime brokers have written them bilaterally, but no source reports their terms, and a former head of a securities lending division whom I interviewed could neither confirm nor deny that single-name term loans are written today.³

I reconstruct the cost of supply using daily Markit Securities Finance lending data on U.S. common stocks from July 2007 through December 2024. A provider who sells a fixed rate for a given term can deliver it only by borrowing the stock overnight, rolling the loan every day at the market fee, and buying the shares in on any day they cannot be borrowed. I price this strategy along every realized path at horizons from two trading days to one year. The cost has three components, or legs. The *expected carry* is the average fee the provider expects to pay, the *replication risk* is the risk that the realized fee differs from that expectation, and the *sourcing* leg is the expected cost of buying in shares that cannot be borrowed.

My main result is that the cost of term protection comes almost entirely from risk and not

²Industry estimates place global lendable supply in the tens of trillions of dollars and annual lending revenue near ten billion dollars (DataLend and ISLA market reports).

³Interview conducted in August 2026. Section 2.2 describes the arrangement when it is written, and Section 2.3 discusses what the absence of observed terms means for the analysis.

from the expected fee. The risk grows with the term, is concentrated in stocks whose fees are hard to forecast, and is only partly covered by the option-implied price of the closest substitute.

I begin by measuring the expected carry, and I find that it is close to today’s fee. Averaged across stocks, the realized cost of carrying a term differs from the fee on the day the loan is written by less than one basis point at horizons out to two weeks and by seven basis points at one year, compared with a mean fee of about 130 basis points (Table IA.2; Figure 1 plots the same wedge by fee decile). A forecasting model built from standard fee predictors performs worse out of sample than simply quoting today’s fee at horizons from one month to one year, because fees are highly persistent and today’s fee already contains most of what is forecastable.

Next, I measure replication risk, and I find that it grows with the term instead of averaging away. If daily fees were today’s fee plus independent noise, the dispersion of the term cost around today’s fee would shrink from 95 basis points at two days to 8 basis points at one year. Instead, it rises to 271 basis points (Figure 2). A one-year term borrow carries thirty-two times the risk of an equivalent sequence of independent overnight borrows, because a fee shock early in the term remains in the average at the end of it.

I then ask in which stocks term protection is expensive to supply. In regressions that control for a dozen characteristics observable at origination, including the fee level, fee volatility, utilization, size, and the short-selling risk measure of Engelberg et al. (2018), fee unpredictability predicts replication risk at every horizon (Table 3). The markup a provider needs to break even on 95% of six-month contracts, which I call the charge, rises from 15 basis points in the most predictable decile to 424 basis points in the least predictable decile (Table 2). Much of this spread reflects the fee level. Within fee deciles, the ratio of the charge to the origination-day fee still rises from 0.30 to 0.70. Once regressions compare stocks with similar fees in the same month, however, unpredictability is clearly distinct from the fee level only out to three months. The sourcing leg is concentrated in the same stocks: 13.2% of six-month terms in the least predictable decile contain a day on which shares could not be borrowed, compared with fewer than one in five hundred in the six most predictable deciles. Term protection is therefore costliest to supply where borrowers would most want it, a pattern that can keep a market thin even when protection is valuable (Froot, 2001).

To determine whether anyone would pay this cost, I compare it with the price of the closest traded substitute, the borrowing rate implied by put–call parity. Because a large borrowing fee acts like a large dividend yield, it gives American calls an early-exercise premium, and I correct the implied fee for this premium. At six months, the corrected implied rate is 71 basis points, compared with a full cost of supply of 143 basis points (Figure 3 and Table 5). The option price covers the expected carry but only about the lowest quarter of the replication-risk distribution at three and six months, and about the lowest two-fifths at one year. This statement requires no assumption about how much safety margin a provider needs. The small gap between the implied rate and expected carry also forecasts the realized cost of the term and the share of days on which shares cannot be sourced.

These results are subject to four qualifications. First, replication risk is almost entirely idiosyncratic, so a provider running a diversified book faces a six-month shortfall of about 14 to 17 basis points instead of 72, and a shortfall of that size is no larger than my calibration error. The sourcing leg does not shrink with the book, because it is priced at its expected cost, which averaging across names leaves unchanged. Second, price remains below cost when the risk leg is set at the median of its distribution, but the median shortfalls lie within the ± 20 basis point reach of the level calibration. Only the shortfalls at high quantiles are large relative to it. Third, short sellers can take positions through options instead. Where option markets are deepest, however, the option-implied fee rises by about 0.88 basis points for each basis point of borrowing fee, so this route tracks the cost it replaces. Fourth, the interest-rate environment raises the charge in the aggregate but explains little of the cross-section or of the price–cost gap (Section 4.5).

I also test whether the constraint affects short sellers’ behavior and prices, and the evidence is mixed. Short interest does not build more slowly in stocks with unpredictable fees. Among stocks that are already expensive to borrow, those with the least predictable fees earn 3.9 percentage points less over six months, but this result does not survive value weighting (Section 6.4). Loan tenure does respond, however. A one-standard-deviation increase in fee unpredictability shortens it by about four days, against a mean tenure of 74 days. A decomposition of the loan book shows that terminations rise where replication risk is higher, while new loans arrive at the usual rate. The aggregated data cannot identify whether borrowers or lenders close the loans.

Finally, I explore why the market is thin and how the contract could be changed. A diversified fund would pay roughly half a basis point to eliminate idiosyncratic fee risk under standard assumptions, and only tens of basis points under aggressive ones. A leveraged short seller facing forced cover, by contrast, would pay close to the full charge. Because a diversified provider can supply protection for a fraction of the single-name charge, the gains from trade between such a provider and a constrained short seller appear substantial. A key obstacle between them is the obligation to deliver shares in the stocks where delivery fails, because the cost of meeting that obligation, unlike fee risk, does not diversify. Settling in cash removes the sourcing leg, and pooling across stocks diversifies the risk. Together, these two changes reduce the six-month cost of protection in the least predictable decile from 441 basis points to 61 basis points, and shortening the term to one month reduces it to 29 basis points (Table 8 and Section 6.5). I offer two explanations for why even this cheaper contract is missing (Section 6.6). First, a cash-settled contract needs a rate to settle against, and none exists. Second, exclusives already pool lending portfolios, but they are sold as access to inventory and not as protection against changes in the fee. Loan-level reporting under SEC Rule 10c-1a, which begins in 2028 with public dissemination from 2029, will make both explanations testable.

This paper contributes to three literatures. First, it contributes to work that builds shorting-cost term structures from option prices (Weitzner, 2023; Cocquemas, 2017; Muravyev et al., 2022) by providing, to my knowledge, the first term structure of equity borrowing costs built from realized lending data. This term structure covers horizons below one month and includes a sourcing leg,

neither of which option prices can measure. Second, it contributes to the literature on short-sale constraints (Miller, 1977; Jones and Lamont, 2002; Duffie et al., 2002; Engelberg et al., 2018) by identifying and pricing a constraint that operates alongside the fee level, namely the inability to lock in the fee. Third, it contributes to the literature on thin or missing markets for risk transfer (Froot, 2001) by documenting a missing market in which both the cost of supply and the price of the closest substitute are separately measurable.

The remainder of the paper is structured as follows. Section 2 describes the institutional setting and the data, and Section 3 sets out the framework. Section 4 measures each leg of the cost of supply. Section 5 compares that cost with the option-implied price. Section 6 examines why the market is thin and prices four changes to the contract, and Section 7 concludes.

1.1 Related literature

The paper closest to mine is Muravyev et al. (2022), who compare the option-implied fee with the spot lending fee. I instead compare the implied fee with the cost of supplying a term loan, which speaks to whether the product is worth manufacturing. My results also bear on theirs. When the implied fee is extracted by European put–call parity, my price-minus-spot margin is negative, whereas they report a positive premium. Once I correct for the early-exercise premium on American calls, price minus spot is close to zero and well within the reach of the level calibration.

More broadly, this paper relates to work on term structures in other markets, including whether forward rates forecast spot rates (Fama and Bliss, 1987; Campbell and Shiller, 1991; Cochrane and Piazzesi, 2005), where along the curve compensation is concentrated (Lustig et al., 2019), and how to build a curve where none is quoted (van Binsbergen et al., 2012; Giglio et al., 2024; Dew-Becker et al., 2017). Special repo rates are the closest analogue in a financing market (Duffie, 1996; Jordan and Jordan, 1997), and work on limits to arbitrage in financing markets motivates the mechanism (Krishnamurthy, 2002; Vayanos and Weill, 2008; Du et al., 2018). Because my conclusion concerns a contract that is not written, I also draw on the theory of security design, which attributes missing markets to fixed costs, informational frictions, or the absence of a contractible settlement variable (Allen and Gale, 1994; Duffie and Rahi, 1995; Shiller, 1993). My evidence points to the last of these as one reason that even a redesigned, cash-settled contract is missing.

Finally, this paper relates to work on the determinants of lending fees (D’Avolio, 2002; Geczy et al., 2002; Cohen et al., 2007; Kolasinski et al., 2013; Blocher et al., 2013; Prado et al., 2016; Chague et al., 2017). I study the cost of insuring against fee variation instead of the level of fees. Andrews et al. (2026) document a strong common factor in fee levels, and the two results are complementary. A term provider nets out the fee level at origination, and the fee innovations she is left holding are almost entirely idiosyncratic. As a result, fee risk diversifies across a book. Because fees are persistent, however, it does not diversify over the life of a term.

2 Institutional Background and Data

Three institutional facts carry the argument. Loans are open, so the fee can be reset at will and the shares can be recalled. Term contracts are rare and leave no observable trace. And a borrower who loses a borrow and cannot replace it must buy the shares back in the market. The first fact creates a risk to price, the second means that the cost has to be reconstructed, and the third makes the last leg of the cost a purchase instead of a fee.

2.1 The open loan

Equity lending is an over-the-counter, bilateral market intermediated by agent lenders acting for beneficial owners (principally pension funds, insurers, sovereign funds, and mutual funds) and by prime brokers acting for borrowers (D’Avolio, 2002; Kaplan et al., 2013). Loans are collateralized, marked to market daily, and, in the overwhelming majority of cases, open: neither the fee nor the duration is fixed at initiation. The lender may re-rate the loan at its discretion, typically daily, and either side may terminate. Because there is no exchange and no central quotation, the borrower observes the rate it is charged but not the rate the market would charge, a feature that Duffie et al. (2002) model as search and that helps explain the cross-sectional dispersion in fees documented by Cohen et al. (2007) and Kolasinski et al. (2013).

A stock is *general collateral* when it is easy to borrow and lends at a low and stable fee, which describes most of the market. It is *special* when borrowing demand commands a materially higher fee, and it is *warm* in the band between. The open structure serves the lender. A beneficial owner may need a lent security back in order to sell it, vote it, or capture a corporate action, and an open loan also lets the lender re-price into rising demand. Both features are given up under a term contract.

Recall, fails, and buy-in. If a lender recalls a loan, the borrower must return the shares within the settlement window, and a borrower that cannot source a replacement borrow fails to deliver. Before 2008, failing carried a cost largely unrelated to the lending fee and functioned as an informal outside option for a short seller (Fotak et al., 2014). The SEC’s temporary Rule 204T in September 2008, made permanent as Rule 204 in July 2009, imposed a mandatory close-out requirement that substantially removed this option. Under the current regime, a borrower who loses the borrow and cannot replace it must buy the shares in the market. The cost of failing to source is therefore a market-price cost, and term protection is valuable to a borrower not only for the fixed fee but also for the guarantee of continued availability.

2.2 Why term contracts are rare

A term loan requires the lender to forgo re-rating and recall for the life of the contract. For a beneficial owner, this forecloses selling the position, complicates voting and corporate actions, and gives up the upside from re-pricing into demand. Two term-adjacent arrangements do trade.

Beneficial owners auction *exclusives*, which are guaranteed-fee arrangements that grant a single borrower access to a lending portfolio for a fixed period. Owners therefore do commit lending revenue forward when the commitment is written at the portfolio level, where fee risk largely diversifies (Section 4.2.4). Prime brokers have also written term loans bilaterally for large clients, but these contracts leave no trace in any consolidated source, and no quoted market exists in single-name term protection.

My interviewee described two features of term arrangements as they have been written in practice. First, the commitment is typically relational: the lender retains a recall right and is restrained from exercising it by the prospect of losing the counterparty's future business, so the rate is fixed while availability is guaranteed only by reputation. Second, terms run to three or six months and not beyond one year, since a longer commitment invites treatment as a securitized loan. Term equity loans were described as more common in the early-to-mid 1990s, when they fixed the rebate on cash collateral, and as written mainly on warm names that were neither general collateral nor hard to borrow. Section 4.5 asks whether my cost estimates recover that band.

One might object that a beneficial owner who already holds the shares rolls nothing and sources nothing. An owner who commits to a fixed rate for H days, however, forgoes what the open market would have paid her over the same span, so her profit and loss relative to staying open is the same commitment wedge the dealer bears, with the same distribution. Only the third leg changes, with sourcing replaced by the forgone right to recall.

2.2.1 How long loans actually last

In my sample, the average age of a stock's outstanding loans is 74 days, and across stock-months this average has a median of 58 days and a 90th percentile of 143 days, so the horizons I price, from two trading days to one year, bracket the observed distribution. This statistic sits alongside the trade-level evidence that short sellers turn over quickly (Boehmer et al., 2008; Diether et al., 2009). The two statistics describe different objects: theirs is the life of a typical short trade, and mine is the age of a typical outstanding loan, a statistic that over-represents long-lived positions because it comes from a snapshot of the book. A term provider inherits the book, so the loan statistic is the relevant one for supply. On the demand side, a borrower whose position lasts one month values a six-month guarantee less than the charges below suggest. Section 6.5 therefore also prices contracts with shorter tenors.

2.3 What I do not observe

I do not observe a quoted term rate, since no consolidated source reports what a prime broker charges to fix a borrowing cost, which names it will quote on, or at what spread. Nor do I observe whether a given contract fixes or caps the fee, or whether rate protection and recall protection are sold together. The quantities I estimate are therefore the cost of manufacturing term protection out of the open loan, not the price at which such protection changes hands. The two coincide only under competitive supply and risk-neutral pricing, and I assume neither. Read as a manufacturing

cost, the expected-carry and sourcing legs are a floor under any quoted rate, since a provider would also require compensation for capital, overhead, and counterparty exposure. The risk leg instead depends on a quoting convention, and Table 6 shows how the cost moves with it. Loan-level disclosure under Rule 10c-1a will make contracted terms publicly observable from 2029.

2.4 Data

I use the Markit Securities Finance daily equity lending file merged to CRSP, from July 2007 through December 2024, following the sample conventions of Saffi and Sigurdsson (2011). Table IA.1 defines every constructed quantity. Markit collects daily position-level data from agent lenders, custodians, and prime brokers, who report their outstanding loans, the fees on them, and their lendable inventory. The fee measure is the simple average fee on all outstanding loans (Markit’s SAF), in annualized basis points, and utilization is the value of shares on loan divided by the value available to lend. I restrict the sample to U.S. common stocks on the NYSE, AMEX, and NASDAQ, exclude prices below \$5, and winsorize fees at the 1st and 99th percentiles within month.

A provider rolling an overnight position pays the rate on new loans, so the average across outstanding loans is the wrong object in principle. In practice, Markit’s new-loan series is distributed as a coarse score over most of my sample, which rules it out.⁴ Because averaging smooths the series, my estimates of fee variation are conservative, and Section 4.2.1 shows that the central result does not depend on the smoothing. A charge rebuilt from unwinsorized fees has the same shape and roughly the same 28-fold spread (Table IA.3); I report winsorized figures for comparability with the rest of the panel.

The merged daily file contains 13.7 million stock-days, of which 10.0 million carry a valid fee. I compute term costs on a grid of 2, 3, 5, 10, 15, 21, 42, 63, 126, and 252 trading days using running sums, on a month-end origination grid, and I retain a term only where the full daily path is present and the calendar span is consistent with the number of trading days. This leaves 425,179 stock-month observations at the two-day horizon and 297,480 at the one-year horizon. Coverage grows from roughly 2,600 stocks per month early in the sample to over 3,200 by the end, so all cross-sectional sorts are formed within month, and the panel regressions absorb month effects through month fixed effects or finer month interactions. Table 1 reports summary statistics. The option-implied analysis in Section 5 rests on a smaller sample that covers 26% of lending-panel stock-months.

⁴The value-weighted new-loan fee fields take integer values on a bounded scale with a median of zero. Where an underlying new-loan rate is reported, it is populated on roughly half of stock-days, with gaps scattered within stock, which precludes constructing a continuous replication path.

3 Framework: The Cost of Manufacturing Term Protection

A term equity loan is a fixed-for-floating swap on the borrowing fee. Writing $\bar{F}_{t,t+H}$ for the average daily fee actually paid over an H -day term, the fair quote is

$$K^{(H)} = \underbrace{E_t[\bar{F}_{t,t+H}]}_{\text{expected carry}} + \underbrace{\lambda \cdot \text{SD}_t[\bar{F}_{t,t+H}]}_{\text{replication risk}} + \underbrace{\text{Pr}(\text{fail}) \cdot c_{\text{buy-in}}}_{\text{sourcing failure}}, \quad (1)$$

and this section sets up the objects needed to measure each term.

3.1 The replication cost

Consider a provider who writes a term loan on stock i at date t for H trading days at a fixed rate $K_{i,t}^{(H)}$ and sources the shares by borrowing overnight and rolling. Let $F_{i,d}$ denote the annualized fee, in basis points, paid on day d . The realized cost of manufacturing the product is the day-weighted average fee over the term,

$$C_{i,t}^{(H)} = \frac{1}{H} \sum_{d=t+1}^{t+H} F_{i,d}, \quad (2)$$

and the provider's realized margin is $K_{i,t}^{(H)} - C_{i,t}^{(H)}$. Because the provider must quote before observing the path, the relevant benchmark is the overnight rate on the origination day. I therefore define the commitment wedge as

$$W_{i,t}^{(H)} = C_{i,t}^{(H)} - F_{i,t}, \quad (3)$$

which is the amount by which the realized cost of carrying the term exceeds the rate observable at the moment of commitment. $E_t[W]$ is the slope of the expected-fee term structure, and $\text{SD}_t[W]$ is the risk a provider bears if she quotes today's overnight rate for the whole term. Throughout the paper, "replication risk" refers to the dispersion of W , and "the charge" refers to the break-even markup defined in Section 4.2.2.

3.2 Rebate risk and fee risk

A term equity loan can fix either of two prices. Under cash collateral, the lender pays a *rebate* on the borrower's cash, and the borrow fee is the short-term interest rate less that rebate. A contract that fixes the rebate leaves the lender exposed to the level of interest rates, which is the most hedgeable exposure in financial markets. A contract that fixes the *fee* does not, because the fee is a spread net of the policy rate. This paper measures only the second exposure. Unlike many spreads over the short rate (Nagel, 2016; Drechsler et al., 2017), the commitment wedge is essentially unrelated to policy at the stock level, since rate variables explain a tenth of a percent of its variation, although they explain about half of the time-series variation in the aggregate charge (Section 4.5). There is no instrument correlated enough with a single stock's fee to hedge it, and no traded claim on it, so a provider must manufacture the term loan by rolling.

Practitioners describe the two exposures differently. When asked how a desk would manage the rate leg of a term commitment, my interviewee named an instrument immediately, but for the fee leg my interviewee had none to name and described the exposure as a risk one takes or declines. A provider writing term therefore warehouses the wedge, and the charge in Section 4.2.2 should be read as capital held against an exposure, not as the price of transferring it. The rate exposure of historical fixed-rebate contracts was also large: the standard deviation of the six-month change in the effective federal funds rate was 109 basis points across the 1990s, compared with 83 basis points over my sample.

3.3 Why the naive quote is the right benchmark

An alternative to quoting the overnight rate is to forecast $C^{(H)}$ with a model. An expanding-window panel regression of the realized term cost on stock fixed effects, the lagged fee, fee volatility, utilization, short interest, size, return volatility, turnover, and a lagged market-wide fee performs worse than quoting the current overnight rate at every horizon from one month to one year, with out-of-sample R^2 against that benchmark of -0.14 at one month and -0.43 at one year (Table IA.4). The model’s fitted values nonetheless correlate 0.96 with realized term costs at one month and 0.73 at one year. Fees are highly forecastable, but almost all of what is forecastable is already contained in today’s rate. For this reason, I measure risk relative to the overnight rate. Because the model’s fitted values track realized costs closely, the dispersion of its out-of-sample errors, which I use in Section 4.2.2 to sort stocks, captures what remains after a specification that explains most of the variation.

3.4 Diversification over the term

A provider needs to know whether the risk of manufacturing averages out over the life of the contract. The most favorable case is that deviations of daily fees from the origination-day rate are independent transitory noise. Under that null, the wedge is an average of H independent deviations, so its dispersion falls with the horizon:

$$\text{SD}\left[W^{(H)}\right] = \text{SD}\left[W^{(H_0)}\right] / \sqrt{H/H_0}. \quad (4)$$

Persistent fees break this relation, because a shock that arrives early in the term is still in the average at the end of it. I report the ratio of realized dispersion to this benchmark across horizons. The alternative null of i.i.d. fee *levels* (the $\rho = 0$ limit of the AR(1) in Section 4.2.1) implies a dispersion that is roughly flat in the horizon, so under either version of independence the dispersion does not rise.

4 What Term Protection Costs to Supply

This section measures the cost of the strategy described above: selling a fixed rate and delivering it by borrowing the stock overnight, rolling the position every day of the term, and buying the shares in on days they cannot be borrowed.

4.1 The expected-carry leg

Figure 1 plots the mean commitment wedge against the horizon, with one line per decile of the overnight fee at origination, scaled by the spot fee so that deciles with very different fee levels are comparable. Two features stand out. First, the pooled curve, which the figure does not plot and Table IA.2 reports in basis points, is nearly flat. Averaged across stocks, the realized cost of carrying a term differs from the origination-day rate by +0.05 basis points at two days and by -7.2 basis points at one year, against a mean fee of roughly 130 basis points, so a provider who quoted today’s rate would be approximately right on average at every horizon.

Second, the conditional curve is strongly mean-reverting, and the reversion grows with the horizon. Among the cheapest decile of stocks (a mean fee of 21 basis points), the expected term cost exceeds the spot rate by 4.5 basis points at two days and by 21.5 basis points at one year, so expected carry over a year is roughly double the spot fee. Among the most expensive decile (a mean fee of 931 basis points), the expected term cost falls short of spot by 3.3 basis points at two days and by 199 basis points at one year, a decline of 21%. To my knowledge, this is the first equity lending fee term structure constructed from realized lending data, and its shape agrees with the option-implied estimates of Cocquemas (2017), which are downward-sloping for hard-to-borrow names.⁵

4.2 The replication-risk leg

4.2.1 The risk grows with the term

Committing for a longer term makes the average cost riskier. The dispersion of the commitment wedge rises monotonically with the horizon, from 94.6 basis points at two days to 270.9 basis points at twelve months, whereas the independence benchmark in equation (4) would have taken it to 8.4 basis points (Figure 2 and Table IA.5). Expressed as a ratio to that benchmark, replication risk is 1.27 times its diversified level at three days, 5.04 times at one month, and 32.1 times at one year. There is no flat region. Dispersion already rises from 94.6 to 98.0 basis points between two and three days, whereas under independence it would fall, so fees are persistent even at the daily frequency.⁶

⁵Sorting within month on the spot fee selects, at the extremes, stocks whose fees are transitorily high or low, so part of the reversion in the extreme deciles is induced by the sort. The decline in the top decile is far too large to be sorting noise alone, but the decile-level expected carries should be read as conditional means given the sort, which is also how Section 5.2 uses them.

⁶The within-stock version of the calculation gives similar ratios (29.0 at one year), so the result is not driven by cross-sectional differences in fee levels, and the one-year ratio is 43.3, 35.0, and 30.8 over 2007–2012, 2013–2018, and

Fee persistence implies this pattern. Under an AR(1) with daily autocorrelation ρ , the ratio has a closed form in which the stationary variance γ_0 cancels,

$$\frac{\text{Var}[W^{(H)}]}{\gamma_0} = \frac{1}{H^2} \left[H + 2 \sum_{k=1}^{H-1} (H-k) \rho^k \right] + 1 - \frac{2}{H} \frac{\rho(1-\rho^H)}{1-\rho}, \quad (5)$$

so each stock’s implied curve depends only on its own persistence. Estimating ρ stock by stock and averaging the implied curves gives a one-year ratio of 28.95, against a realized within-stock value of 29.00 (Table IA.7). The growth of replication risk with the term is therefore a consequence of fee persistence. The realized measurement adds the magnitude, 271 basis points at one year, in the units against which a provider would have to hold capital.

The AR(1) parameters are not, however, stable enough to price a stock’s risk from its distant history. When γ_0 and ρ are estimated on 2007–2015 and used to predict each stock’s six-month dispersion, the predictions span two orders of magnitude, from 4.5 basis points in the lowest decile to 420 basis points in the highest, while realized charges on 2016–2024 across those same deciles are flat, between 20 and 41 basis points. Stocks do differ durably in how forecastable their fees are, since the trailing measures of Section 4.2.2, re-estimated as the sample rolls forward, sort charges from 15 to 424 basis points. A structural parameter estimated once and carried forward does not price this variation, so borrowing-cost risk has to be estimated from recent data.

4.2.2 The charge, and which stocks are expensive to supply

The fact that risk grows with the term does not yet say which stocks are expensive to write. Table 2 reports the break-even charge by decile of fee unpredictability, which I measure as the trailing twenty-four-month dispersion of out-of-sample errors from the fee forecasting model. The charge is the empirical 95th percentile of $W^{(H)}$ within each cell, that is, the markup over the origination-day rate at which a provider breaks even on 95% of contracts. This is a value-at-risk quoting convention, not an estimate of an equilibrium price of risk, and it replaces the parameter λ in equation (1). The convention fits how the exposure is carried: because the fee leg cannot be hedged (Section 3.2), a provider must hold capital against it, and a quantile measures the capital needed to remain solvent against an unhedgeable position.⁷

The charge rises monotonically across deciles at both horizons in Table 2, and at every other horizon I compute (untabulated). At six months, it runs from 15.0 basis points in the lowest

2019–2024 (Table IA.5). The complete-path requirement drops 30.5% of stock-months with a two-day path, and the dropped names carry higher fees (176 against 112 basis points), so long-horizon levels are likely understated. On the 278,355 stock-months whose path is complete at every horizon the one-year ratio is higher, at 45.9, so attrition, if anything, flattens the profile. Because SAF averages across outstanding loans, smoothing attenuates short-horizon variation and could inflate a ratio anchored at two days. Re-basing the benchmark at twenty-one days still leaves one-year dispersion at 8.0 times its benchmark, and the divergence is largest among the stocks whose fees move most often (Table IA.6).

⁷I use the empirical percentile rather than 1.645 SD[W] because the distribution of W is fat-tailed, so the two differ substantially. A quantile of W also embeds the conditional mean of W , so in cells where fees mean-revert strongly, the charge is net of the expected decline in the fee.

unpredictability decile to 423.6 basis points in the highest, and at one month from 8.8 to 189.7 basis points. Resampling months in blocks of six puts the ratio between the extreme six-month charges at 28.2, with a 95% interval of [19.3, 42.1]. The tail behaves the same way: the probability that the realized cost exceeds the origination-day rate by more than 100 basis points rises from 0.11% to 13.45% across the six-month deciles.

Much of this spread reflects the fee level with which unpredictability is correlated. When the deciles are re-formed within fee decile $\times$ month, the charge still rises monotonically, but from 20.6 to 125.0 basis points, a sixfold spread (Table IA.8). A fee decile is also wide, so re-sorting within it does not hold the fee fixed: across the re-sorted cells, the mean origination-day fee still rises from 68.6 to 179.1 basis points. The cleanest form of the result is therefore the ratio of the charge to the fee, which rises from 0.30 in the most predictable re-sorted decile to 0.70 in the least predictable one.⁸

Across the unconditional deciles, the six-month charge is between 47% and 105% of the origination-day fee. The ratio peaks in the eighth and ninth deciles rather than the tenth, because the least predictable names carry an enormous charge but an even larger fee. The same ordering binds an owner who supplies term protection from her own holdings. The expected gain she forgoes from re-pricing into rising fees, $E[(W^{(H)})^+]$, rises at six months from 4.0 basis points in the most predictable decile to 73.5 basis points in the least predictable decile. The same stocks are therefore expensive to supply, in the same order, whether a dealer or an owner writes the protection.

4.2.3 Unforecastability versus the fee level

A lender deciding which names it could quote needs to know what predicts this cost. Table 3 regresses the absolute replication error $|W^{(H)}|$ on lagged fee unpredictability, controlling for the fee level, total fee volatility, size, return volatility, utilization, turnover, the Markit cost-of-borrow score, average loan tenure, the number of active brokers, lender concentration, the Engelberg et al. (2018) short-selling risk measure, and idiosyncratic volatility. The coefficient on unpredictability is positive and significant at every horizon, with t -statistics declining from 7.27 at two days to 3.31 at one year. Because the regression controls for total fee volatility, unpredictability is not simply a restatement of how much fees move.

A linear control is a weak control for the fee level, because the fee distribution is highly skewed (a mean of 133 basis points against a median of 30), and the natural objection is that unpredictability is simply a proxy for being hard to borrow. The middle panel of Table 3 therefore absorbs fee-decile $\times$ month fixed effects. The coefficient roughly halves and remains significant out to three months ($t = 3.68$ at one month and 2.39 at three months) and is marginal beyond that (a coefficient of 2.41 with $t = 1.93$ at six months, and 3.48 with $t = 1.97$ at one year). Tightening the control further to fee ventiles and fee percentiles, which leave residual fee spreads of only 1.41-fold

⁸An expanding-window charge that uses only information available at origination is somewhat smaller, covers 90–94% of subsequent contracts, and preserves the ordering across deciles (Table IA.9, Panel C). A twelve-month window for the unpredictability measure gives the same profile (Table IA.8), and so does a charge built from unwinsorized fees (Table IA.3).

and 1.12-fold, reduces the six-month coefficient to 1.11 ($t = 1.04$) and 0.81 ($t = 0.96$; Table IA.10). In regressions, then, unpredictability is clearly distinct from the fee level at short horizons, while at six months the design cannot separate the two.

The regression measures how much a continuous regressor adds once fee-percentile-by-month indicators, a hundred in each month, are in place. That is a narrower question than how the charge varies when the fee is held fixed. When stocks are sorted into terciles of unpredictability within fee decile and month, the six-month charge rises from the most to the least forecastable tercile in all ten fee deciles, by factors between 1.4 and 4.3. The gradient is weakest in the most expensive decile, where the fee level does most of the work. Among cheaper names, whether the fee is liable to move largely determines whether a term commitment is cheap or costly. The unpredictability sort also orders the charge better than a sort on the fee itself. A fee sort gives a top-to-bottom six-month charge ratio of 11.3, against 28.2, and it is not monotone, because general-collateral names occasionally spike from very low fees. For these reasons, I keep the unpredictability partition and, at long horizons, rely on the charge-to-fee ratio rather than on the regression coefficient.

Two further points concern the measure. First, the sorting variable is built from the errors of a model that loses to the naive quote, but the model is used only to rank stocks by the size of their fee surprises, which it can do even though it cannot beat the current fee on average. A model-free alternative, the trailing mean absolute daily fee change, also predicts replication risk across stocks at every horizon, so the result reflects the fees and not only the forecasting equation. Second, the measure ranks stocks but does not time them. With stock fixed effects (bottom panel of Table 3), the coefficient is positive at two days ($t = 2.11$) and significantly negative from three months on, reaching -5.52 at one year ($t = -3.42$). The measure therefore has a large and durable between-stock component and a mean-reverting within-stock component. It is useful for deciding which names a desk could quote on, and I make no claim about when to quote on a given name.

4.2.4 Does the risk diversify across names?

A provider typically writes a book of term loans, so I next ask whether replication risk diversifies across names. I find that it does. Regressing each stock-month’s commitment wedge on that month’s cross-sectional mean wedge yields an R^2 of roughly 0.006 at every horizon. This does not conflict with the strong comovement in fee levels documented by Andrews et al. (2026), because the wedge is measured net of the origination-day fee, which removes the common level, and the fee changes that remain over the term turn out to be almost entirely idiosyncratic. At the six-month horizon, the single-name standard deviation of 245 basis points falls to 42 basis points for a book of fifty names and to about 25 basis points in the large-book limit.

Two implications follow, and they point in opposite directions. The charges in Table 2 are single-name quantities, so a provider running a diversified book requires roughly one-sixth to one-tenth of the compensation per position. Against that, diversification across names does not undo the growth of risk over the horizon, and it does not help with the sourcing leg, which is concentrated in the same names. The undiversifiable residual also does not appear to be priced as

a systematic factor, since it is unrelated to the intermediary capital ratio of He et al. (2017), the Chicago Fed leverage subindex, and the OFR financial stress and funding indices.

4.3 The sourcing leg

Averaging across names reduces the cost of fee variation, but it does nothing about the possibility that shares cannot be obtained at any fee. A provider who has written a term loan and cannot source the shares on a given day must buy in, at a price that is worst when the borrow is tightest. An implied fee extracted from put–call parity does not price whether shares were borrowable, but the daily lending panel can measure it. I flag a stock-day as a sourcing failure when utilization exceeds 95% or the Markit cost-of-borrow score is in its top two categories, and I count such days inside each term window.

Unconditionally, the leg is small: 1.1% of days inside a one-month term are failure days, and 1.6% of one-month terms contain at least one. The leg is, however, extremely concentrated. In the bottom six unpredictability deciles, the mean failure share is at most 0.02%, and fewer than one six-month term in five hundred contains a failure day. In the top decile, the failure share is 5.14%, and 13.2% of six-month terms contain at least one day on which the shares could not be sourced (Table 2). The sourcing leg and the replication-risk leg therefore load on the same stocks, and the sourcing leg is the more concentrated of the two: the charge rises roughly 28-fold across deciles, while the share of failure days rises more than 200-fold.⁹

4.3.1 Benchmarking the flag against realized fails

Because the flag carries a whole leg of the pricing equation, I check it against SEC fails-to-deliver data, which play no part in its construction. Nearly every stock-month contains some fail, so I apply a materiality screen modeled on the Reg SHO threshold condition: fails of at least 10,000 shares and at least 0.5% of shares outstanding. Across 324,481 stock-months from July 2009 to December 2024, unflagged stock-months show a material fail 1.5% of the time, whereas flagged stock-months show one 20.9% of the time. In the same window, among the 409,537 stock-months with an unpredictability measure, realized fails rise across unpredictability deciles from 0.29% to 4.32% of stock-months. The flag is thus a strong conditional signal but a weak classifier, with a sensitivity of about 27% and a precision of about 21%.¹⁰ A stronger test is also available. If the risk I measure is real rather than an artifact of the fee series, stocks whose replication cost has been more variable should subsequently prove harder to source. Regressing the materiality indicator on

⁹The thresholds do not drive the result. Under looser (utilization above 90% or a score of at least 8) and tighter (utilization above 98% or a score of 10) definitions, the top-decile failure probability is 16.7% and 7.4%, against 11.8% under the baseline flag in the daily reconstruction, while the bottom deciles stay below one percent (Table IA.3).

¹⁰Requiring the materiality condition on five consecutive settlement days, as the SEC’s threshold list does, leaves 0.22% of stock-months as threshold months. The flag catches 60% of these, compared with 27% of one-day material fails, so it identifies the persistent failures that defeat a term commitment. In the top unpredictability decile, the flag fires on 13.2% of six-month terms, against about 2% that contain a five-day threshold episode and about 17% that contain any material fail, so the probability component of the sourcing leg is an upper bound only against the stricter standard, while the buy-in cost is a lower bound (Section 4.3.2).

the trailing dispersion of the six-month commitment wedge, with month fixed effects and the full control set, gives a positive coefficient ($t = 2.10$), and the count of fail days gives $t = 11.51$. The left-hand variable comes from SEC filings and the right-hand variable from the lending panel, so this test supports the framework with data that the flag itself does not use.

4.3.2 What a buy-in costs

A provider forced to buy in pays the quoted spread and buys into whatever price move is underway, and sourcing fails precisely when the borrow is tight and the stock is under pressure. The quoted half-spread on failure days plus the excess five-day return relative to days when sourcing is available gives a buy-in cost of 133.6 basis points (Table IA.11). The comparison describes the states in which a buy-in occurs, not the causal cost of executing one, and a provider faces those states. It excludes market impact, so the estimate is a lower bound. Doubling it would double the sourcing leg without changing any ordering in the paper.¹¹

4.4 How the legs scale with the book

On a single position, the risk leg dominates. For the least predictable decile at six months, the charge is 424 basis points, while the sourcing leg, a 13.2% chance of at least one failure day times a buy-in cost of roughly 134 basis points, is about 18 basis points. The two legs scale differently with the size of the book (Table 4). Replication risk is idiosyncratic, so within the top decile the six-month charge drops to 73 basis points across fifty names and to 42 basis points in the large-book limit.¹² The sourcing leg does not scale that way, because it enters as an expected cost, the failure probability times the buy-in cost, which does not fall with the number of positions. It is also concentrated in a small group of names, so adding positions in easy-to-borrow stocks dilutes the average without touching the exposure. The sourcing leg rises from about 4% of the risk charge on a single position to about 42% at book scale. Fee risk remains the larger component at every book size, but the sourcing leg becomes material at scale, where a provider needs reliable access to inventory in the names where protection is most wanted, in addition to capital against fee variation.

4.5 Where the policy rate matters, and where it does not

An alternative explanation is that term lending is unattractive because of the macroeconomic environment rather than anything about a stock. To test it, I classify each origination month by the realized change in the effective federal funds rate over the following six months, as *flat*

¹¹Realized fail episodes in the SEC data give a similar average: across 18,835 episodes, the half-spread plus the market-adjusted price move is 152.7 basis points. Pairing the probability of a five-day threshold episode with its realized cost gives a top-decile sourcing leg of 20.1 basis points, and pricing every material episode at its own realized cost gives 26.2 basis points, compared with 17.6 basis points in my construction, so the leg as constructed is conservative.

¹²These book-level figures extrapolate the single-name tail. Drawing actual fifty-name books month by month gives somewhat higher 95th percentiles: 79 basis points against 73 for fifty names, and 61 basis points against 42 in the large-book limit (Table IA.9, Panel D), because adverse months are shared across a book's positions.

when the absolute change is below 25 basis points (123 months) and as *moving* when it is at least 50 basis points in either direction (63 months). Three findings emerge, and they do not all point the same way. First, the environment matters in the aggregate. The six-month charge pooled across all stocks is 33 basis points in flat months and 62 basis points in moving months (Table IA.12), and rate variables explain roughly half of the time-series variation in the aggregate charge (Table IA.13). Second, the cross-sectional ordering is unchanged, since the charge rises monotonically across unpredictability deciles in both regimes (Table IA.12, Panel A). Third, when the fee level is held fixed, the regime effect is large below 100 basis points, where the charge rises by a factor of about 1.8 to 2.8, and absent between 100 and 1,000 basis points (Panel B). At the stock level, month fixed effects explain 0.6% of the variation in the commitment wedge and rate variables explain 0.1%, and the same variables account for only about 6% of the monthly variation in the price–cost margin studied in Section 5.¹³

The same table indicates where the product is writable at all. Two conditions must hold at once: the charge must be small relative to the fee it fixes, or the borrower gains little by fixing it, and the shares must be sourceable, or the provider cannot deliver. In flat-rate months, protection costs 0.88 of the fee on a general-collateral name and 0.58 on a warm name between 25 and 100 basis points, compared with 1.50 and 1.28 in the two special buckets. The ratio falls to 0.30 in the most expensive bucket, but there sourcing fails on 61% of six-month terms. The writable band is therefore the warm middle, which matches my interviewee’s account that term loans were written on warm names and rarely on hard-to-borrow ones. In that band, the cost of term protection depends on the rate environment, whereas for special names with fees between 100 and 1,000 basis points it does not.

5 What the Market Pays

Outside a prime broker, no one observes the rate at which term protection is quoted, so I compare the cost computed above with a price from the options market. A short position can be assembled synthetically from a long put and a short call, and put–call parity then reveals the borrowing cost embedded in that position over the life of the options. That implied fee is a term rate, which makes it the natural counterpart to the cost of supply.

5.1 Extracting an implied fee

Reading a shorting cost out of option prices goes back to Ofek et al. (2004), who show that put–call parity deviations line up with short-sale constraints, and to Evans et al. (2009), who show that the option market maker’s own borrowing problem puts the cost there. For European options at a

¹³Sorting instead on terciles of trailing Treasury yield volatility, which a provider could observe at origination, gives the same answers. The aggregate time series also does not line up with the level of rates: the all-stock six-month charge rose from 21–33 basis points in 2019–2021 to 114–127 basis points in 2022–2023 but fell to 70 basis points in 2024, when the funds rate was at its sample high, and in the meme-stock year of 2021, when rates were near zero, it was already above its 2019 value.

common strike and expiry, $C - P = e^{-rT}(F - K)$, so the forward is $F = K + e^{rT}(C - P)$ and the implied borrowing fee is $r - \ln(F/S)/T$. I compute this from raw quotes, using matched call-put pairs at the strike nearest the forward on month-end origination dates from 2007 to 2024.¹⁴ Three screens do real work. I require at least seventy days to expiry, because $\ln(F/S)$ is divided by T and short-dated contracts amplify quote noise; I drop contracts with a dividend inside their life; and I require moneyness within ten percent.

The level of the implied fee requires an anchor. The implied fee inherits the interest rate used in parity one for one, and the rate embedded in option prices sits above Treasury-based curves by roughly 35 basis points in the box-spread estimates of van Binsbergen et al. (2022), so an implied fee computed with the OptionMetrics zero curve is too low. I therefore anchor the level where the truth is approximately known. For each maturity, I shift the implied series so that, across the contracts in the cheapest fee decile formed within month, the implied fee equals their realized term cost: 29.6, 32.3, and 36.0 basis points at three, six, and twelve months. Because the shift is common to every stock at a given maturity, it absorbs any rate error that is constant across stocks, including one with its own term structure, and it leaves both the cross-section and the pass-through slope in Section 5.3 unchanged. It cannot absorb a rate error correlated with the cross-section, and the levels of the margins below depend on the accuracy of the anchor.¹⁵

Coverage is partial and selected. Of the lending-panel stock-months with a valid fee, 26% carry a six-month implied fee, and coverage falls to about 19% in the top fee decile, where the covered names have just over half the average fee of the uncovered names (580 versus 1,010 basis points) and substantially lower fee risk. The excluded names lack two-sided option quotes precisely because they are extreme, so the comparison below is estimated on the tamer part of the tail. Recomputing the cost side on the option-matched sample gives a charge close to its full-panel value at the extremes of the unpredictability distribution, which drive the comparison, although the option-matched sample understates the charge by as much as a half through the middle deciles. Reweighting the covered sample by the inverse of its estimated coverage propensity moves the margins only slightly (Table IA.14), so selection on observables does not manufacture the gap.¹⁶

¹⁴Raw quotes are essential. OptionMetrics' `forwardprice` field is a theoretical forward built from the zero curve and projected dividends, so it contains no borrowing information by construction: the fee it implies correlates -0.04 with the actual lending fee, compared with 0.69 for the fee implied by raw quotes.

¹⁵Shifting the anchor by ± 20 basis points moves the share of contracts on which the implied rate covers realized cost between 39% and 56% (Table IA.9, Panel A). Anchoring instead to the cheapest decile's overnight fee gives a single shift of about 20 basis points at every horizon, which understates the adjustment and overstates how quickly the shortfall widens. Two results survive any plausible anchor shift: the widening of the gap with maturity, and the top-decile shortfall of price against expected carry of 32.6 basis points, which exceeds the anchor's ± 20 basis point reach, although not by a wide margin. A maturity-matched box-spread rate would be the natural test of the level, but it cannot be run on these data, because my extraction retains only the strikes nearest the forward, and the box signal is smaller than the combined quote error on its four legs.

¹⁶Under inverse-propensity weighting, price minus expected carry moves from -3.3 to -5.9 basis points overall, price minus full cost, built from cell-level expected carry as the note to Table IA.14 describes, moves slightly toward zero, from -68.9 to -63.1 basis points, and the pass-through slope moves from 0.75 to 0.67 . Selection on unobservables is harder to rule out. Of the matched stock-months, 92.1% belong to names that had listed options before they first entered the top fee decile, and restricting to those names moves the six-month margin by about six basis points. Imputing the unmatched names' implied fees at the tenth and ninetieth percentiles of their fee-decile-month cell bounds the full-panel six-month margin between -235 and $+92$ basis points, against -88 basis points for the

5.1.1 American early exercise

Put–call parity is a European relation, while single-name equity options are American, so the extraction above is misspecified unless early exercise is worth nothing. Screening out contracts with a dividend inside their life removes the usual dividend channel, but it does not remove a second channel that runs through the borrowing fee. A stock that is expensive to borrow pays its holder the lending fee, so a large fee is economically a large dividend yield, which gives the American call a genuine early-exercise premium. An inflated call raises the extracted forward and lowers the implied fee, and it does so most where the fee is largest. The sensitivity is large: at six months, a call distorted by half a percent of spot moves the implied fee by about a hundred basis points. The American put also carries a premium, but it is driven by the interest rate rather than the borrow, so it is roughly common across stocks and the anchor absorbs it.

I therefore de-Americanize the quotes. For each contract, I price the American and European call and put on a binomial tree under a continuous yield equal to the current estimate of the borrow fee, subtract the two early-exercise premia from the observed mid quotes, re-apply parity, and iterate to a fixed point in the fee, using OptionMetrics’ volatilities. Table IA.15 reports the result. Across the six-month sample, the call premium rises almost eightfold from the ninth fee decile to the tenth, from 0.034 to 0.263 percent of spot, while the put premium falls, from 0.094 to 0.069. The call premium tracks the borrow fee and the put premium does not, as the two exercise boundaries imply. On the stock-month sample used in Table 5, the correction raises the implied fee by 43.7 basis points in the top six-month fee decile and changes it by between -0.5 and $+1.4$ basis points in deciles one through eight.

Unless noted otherwise, every number in the paper uses the corrected series. The correction moves price minus full cost, the comparison the paper rests on, by only about four basis points at six months. It matters more for quantities that depend on the implied fee alone. It cuts the top-decile shortfall of price against expected carry from 76 to 33 basis points, raises the break-even quantile of Section 5.2 by about five to seven points at three and six months, and brings price minus spot into line with Muravyev et al. (2022). At three and six months, each of these changes favors the options market. At twelve months, the correction lowers the break-even quantile and widens price minus full cost by about three basis points.¹⁷

5.2 Price against cost

Before interpreting the comparison, note that the option-implied rate is not a quote for a term loan, since no such quote exists. It is the price of the closest traded substitute, a synthetic short,

matched names when the charge is computed on the same full panel, so the price side should be read as a statement about optionable stocks.

¹⁷Two diagnostics that do not rely on the binomial model support the correction. Under European parity, the implied fee should not vary with the strike within a stock-month, but the uncorrected fee rises with the strike, as early exercise predicts (472 basis points per log point, $t = 14.51$); the corrected slope is 233 ($t = 7.07$), so the correction removes about half of the strike dependence overall and reduces it in every fee decile without reversing its sign. The corrected series also satisfies the American parity bound on the implied fee for more than 99% of contracts. Scaling the volatility input by 0.75 or 1.25 moves the top-decile correction by under three basis points.

and Section 5.4 shows that the two routes are close but imperfect substitutes. A term loan delivers something the synthetic does not, namely an actual borrow with availability committed for the term, so a borrower who needs the physical position would pay more for it. The implied rate is therefore a lower bound on what she would pay, and the gap reported below is an upper bound on the cost obstacle to writing a term loan. The bound is loose, but the gap is large: at six months, a provider would need the borrower to value a physical term loan at twice the price of the synthetic she already has, and although I cannot rule out a premium of that size, nothing in the substitution evidence supports one.

Table 5 puts the two side by side, and Figure 3 plots the comparison. The first margin, price minus expected carry, compares two prices. The second, price minus full cost, subtracts the origination-day fee, the replication-risk charge, and the sourcing leg, so it compares a price with what supplying the contract requires. The two margins behave oppositely as maturity lengthens. Price minus expected carry is small at every maturity and crosses zero, from -3.8 basis points at three months to $+3.4$ basis points at one year, so the implied rate matches the expected cost of rolling the borrow. Price minus cost widens from -49 to -100 basis points, because across these maturities full cost rises by 61.1 basis points while the implied fee rises by 10.2 basis points. Resampling months in blocks of six gives a 95% interval of $[-84.6, -62.3]$ for the six-month price-minus-cost margin and $[-16.1, +6.5]$ for price minus expected carry. The implied rate thus covers the expected carry but little of the replication risk around it.¹⁸

The 95th-percentile convention matters for magnitudes, so Table 6 re-prices the risk leg at every quantile. As the quoting quantile rises from the bottom of the wedge distribution, the margin starts positive and crosses zero at the *break-even quantile* q^* , beyond which the market's price no longer covers what the provider must hold. The break-even quantile is 24.1 at three months, 26.9 at six months, and 40.5 at twelve months. The option-implied rate therefore buys the expected carry and roughly the lowest quarter of the replication-risk distribution at the short maturities, and the lowest two-fifths at one year. This statement requires no assumption about λ , and I use it to summarize the comparison.¹⁹ The sign of the shortfall does not require the convention, since the margin at the median quantile is still -6.8 , -7.9 , and -2.7 basis points across the three maturities. These median margins are smaller than the ± 20 basis point reach of the level anchor, however, so only the margins at high quantiles are large relative to the calibration. The convention does govern the maturity profile: at the median, the gap is roughly flat across maturities, while at the 95th percentile it doubles.

The paper's bounds point in different directions. The implied rate is a lower bound on what a borrower would pay for a physical loan, the failure probability inside the sourcing leg is an upper

¹⁸Both intervals are computed on the matched stock-month sample. The price-minus-cost interval holds the charge at its full-sample value and is therefore conditional on it; Section 4.2.2 reports the charge's own sampling variation.

¹⁹Solved contract by contract, the break-even quantile is right-skewed, with a mean of 47.9 and a median of 28.6 at six months. When contracts are sorted on the fee, it falls from 53.3 in the cheapest decile to 39.1 in the most expensive, so the market covers a smaller share of the risk where the risk is widest. When they are sorted on unpredictability, however, it is flat, running from 47.3 to 49.4, so on this measure the fee partition carries the cross-section and the unpredictability partition does not.

bound against five-day threshold episodes but not against all material fails, the buy-in cost inside it is a lower bound, and the anchor has a reach of about ± 20 basis points. The implied-rate bound has no upper value to vary, so Table IA.16 varies the quoting quantile, the anchor, and the sourcing leg jointly. At six months, the margin ranges from -93.9 to $+13.6$ basis points across the corners of this box, so its sign is not robust to every combination of my own assumptions. The quantile and the anchor drive this range, while the two sourcing bounds are an order of magnitude too small to matter. With the anchor at its upper end and no sourcing leg, the median-quantile margin is positive at every maturity. The break-even quantile removes the dependence on the quoting convention but not on the anchor.

The comparison also sheds light on the difference between my estimates and those of Muravyev et al. (2022). When the implied fee is extracted by European parity without the early-exercise correction, my price-minus-spot margin is negative, whereas they report a positive premium. My screens do not cause the difference, since a construction that mimics theirs on my quotes also yields a negative margin (Table IA.17). Two candidates can account for it, the rate benchmark and early exercise. On the corrected series, price minus the spot fee is -6.0 basis points at six months and $+0.5$ basis points at twelve months, both well within the ± 20 basis point reach of the level anchor, consistent with their evidence that the options market prices the expected cost of the borrow.

Two further points qualify the comparison. First, the gap widens about thirteenfold from the cheapest fee decile to the most expensive (Panel B of Table 5), but the full cost contains a risk charge that itself rises steeply across those deciles. In a decomposition that assigns each stock the charge of its fee decile and then holds the charge fixed at its panel-wide value, roughly 80% of the cross-sectional widening disappears, so the cross-sectional result is mainly a statement about supply costs. Second, the risk charge in the full cost is a single-name figure, averaging 64 basis points at six months on the option-matched sample. Scaling it by the book-to-single ratios of Table 4 reduces the six-month cost of supply from 143 basis points to about 90 basis points for a fifty-name book and about 85 basis points in the limit, so the margin narrows from -72 to about -19 and -14 basis points. With the book charges drawn directly instead of extrapolated (Table IA.9, Panel D), the two margins are about -20 and -17 basis points. The sign does not change, but the barrier facing a large intermediary is a fraction of the one facing a desk writing a single loan.

The comparison can also be stated in units of contracts. Ex post, writing every matched six-month contract at the option-implied rate would have covered realized replication cost on about 48% of contracts, and an equal-weighted book of every matched name earns a mean margin of $+5.5$ basis points against realized carry and breaks even in 48.5% of months, with a bootstrap interval on that share of $[38.2\%, 58.3\%]$ (Table IA.9). Against realized carry, the book is therefore indistinguishable from break-even, as it should be if the options market prices the expected carry correctly. A provider who must honor the contract, however, cannot quote at the median outcome.

5.3 How the options market responds to borrowing pressure

A natural conjecture is that the options market fails to see borrowing pressure. Table 7 does not support it. Panel A regresses the margin on each measure of lending-market pressure, controlling for the fee level, so the regression tests whether pressure moves the implied rate beyond the fee that the pressure has already produced. Every measure enters positively. Utilization raises the margin by 24 basis points per standard deviation ($t = 11.3$), the cost-of-borrow score by 136 basis points ($t = 11.7$), and lender concentration by 12 basis points ($t = 3.5$). The number of brokers actively borrowing the stock also enters positively ($t = 4.6$). Conditional on the fee, the options market therefore marks the borrow up whenever the lending market tightens on any of these four measures.

The options market responds to the fee less than one for one. Panel B estimates the sensitivity of the implied fee to the actual fee across fee deciles. The slope is 0.80 on the broader sample of quotes whose spread is below half a point (0.75 on the stock-month sample of Table 5), and tighter quote screens raise it to 0.87. Part of the shortfall from one is mechanical, because the implied fee prices the average fee expected over the option’s life, while the spot fee is a point in time and fees mean-revert. Benchmarking against expected carry raises the slope from 0.75 to 0.91 across decile means and from 0.69 to 0.71 across individual contracts (Panel C). The choice of estimator therefore moves the slope by about 0.2, more than the 0.1 by which the decile-mean estimate falls short of one.²⁰ Panel D shows the same relation at its limit. In stock-months when the sourcing flag fires, the mean margin is -186 basis points against -2 basis points elsewhere, but conditional on the fee level the flag enters positively ($t = 5.39$), so these months reflect the same pass-through applied to fees that have gone to extremes.

It is harder to explain why pass-through still falls short of one. One possibility is that closing the gap requires owning the shares, since the trade is to lend the stock at the high fee while shorting synthetically at the low implied rate. If so, pass-through should be higher where lendable supply is plentiful. I find no evidence of this pattern. Across terciles of lendable supply, the slope is 0.713, 0.704, and 0.703, and when lendable supply and option-market depth enter together, depth is significant ($t = 3.35$) while supply is not ($t = -1.26$; Panel E). The concentration of institutional ownership, which would matter if the option market maker’s own borrow were the constraint, does not predict pass-through either.²¹ The remaining friction therefore appears to lie on the options side, and option-market depth predicts it.

5.3.1 What the residual gap predicts

The residual gap between the implied rate and expected carry is small, so it could be measurement error as easily as economics. Table IA.19 tests whether it predicts anything, using an ex-ante margin

²⁰Attenuation from measuring the fee as an average across outstanding loans is not the source of the remaining shortfall: instrumenting the fee raises the observation-level slope from 0.65 to between 0.71 and 0.76 (Table IA.18). Re-estimated in rolling three-year windows, the slope has ranged from 0.69 to 0.83 since 2012.

²¹Across terciles of the Herfindahl index of institutional ownership, the pass-through slope is 0.724, 0.721, and 0.669, and the observation-level interaction with the fee is 0.001 ($t = 0.01$).

whose expected-carry leg uses only originations whose paths had already completed. The results point to three pieces of information in the gap. First, a one-standard-deviation higher margin is followed by a commitment wedge 16.1 basis points larger over the next six months ($t = 5.87$), so when the options market prices the borrow above my expected-carry benchmark, carrying it does cost more, consistent with Weitzner (2023). Second, a higher margin is followed by a sourcing-failure share 0.26 percentage points lower ($t = -2.46$), against an unconditional mean below one percent, so the gap carries information about availability even though parity does not price availability directly. Third, lendable inventory grows 1.4 percent faster over the following six months per standard deviation of margin ($t = 6.13$), although this result depends on the controls. These results do not identify why the residual exists, but they indicate that it is not pure measurement error.²²

5.4 How far can short sellers route around the lending market?

A short seller facing an expensive borrow can take the position synthetically instead. If substitution were frictionless, arbitrage between the two routes would force pass-through to one, so the slopes above bound how far substitution goes in prices, and this section asks how far it goes in quantities. I measure option-market short exposure as put open interest in shares scaled by shares outstanding, which puts it in the same units as short interest. On this measure, options carry about a fifth of combined short exposure.

The main finding is a gradient in option-market depth (Table IA.20). When stocks are sorted on option open interest relative to shares outstanding, the pass-through slope is 0.53 in the shallowest tercile, 0.63 in the middle tercile, and 0.82 in the deepest ($t = 3.11$ for the interaction). Depth and the fee are not independent, and the direction of the dependence cuts against the paper: within the top fee decile, the mean fee is 770 basis points in the deepest tercile against 195 basis points in the shallowest. Forming the terciles within fee decile $\times$ month brings the tercile mean fees to 67, 71, and 93 basis points and leaves the gradient intact, at 0.58, 0.69, and 0.88 ($t = 7.26$). For a short seller in a liquid-option name, the synthetic is therefore close to a full substitute for the borrow, and the difference between the synthetic and a borrow matters mainly in shallow-option names. The synthetic fixes a borrowing cost for the life of the options at a rate that tracks the lending market closely, but it does not deliver shares or commit their availability for the term, which the physical term loan priced here does.

Quantities point the same way. A stock whose synthetic route is cheap relative to its borrow subsequently sees its option share rise. Around months in which a stock’s borrow fee more than doubles and rises by more than 50 basis points, short interest reverses from rising to falling while option short exposure accelerates. This pattern is consistent with substitution and not with common bearish demand, although the option share is already rising before the fee spike. Substitution also

²²The option-implied term structure is less informative about the fee’s path. The forward rate implied by the six- and twelve-month maturities forecasts the realized fee over months seven to twelve with a slope of 0.46 ($t = 16.0$), but its mean absolute error is 157 basis points, against 75 basis points for simply quoting today’s fee, which mirrors the forecasting result of Section 3 on the other side of the market.

has limits (Table IA.21). During the September 2008 short-sale ban, option market makers were exempt while shorting the banned names was prohibited, yet option short exposure in financial stocks, a proxy for the ban list, did not rise significantly relative to other stocks. This result is consistent with Battalio and Schultz (2011) and Grundy et al. (2012), who attribute it to hedging becoming expensive at the moment the synthetic route became the only one. During the March 2020 volatility spike, with borrow fees little changed, the option share of short exposure fell from 21.2% to 18.1% before recovering in April. These results bound the escape route but do not identify a substitution elasticity, which would require a shifter of relative cost that leaves the demand to be short unchanged.

6 Why the Market Is Thin, and What Would Change It

This section asks what the comparison between the cost of supply and the option-implied price implies for why the product is scarce, tests adverse selection as an alternative explanation, examines who would buy protection and at what price, and asks which features of the contract the cost is attached to.

6.1 Supply cost above price

A borrower values term protection in proportion to how uncertain her borrowing cost is, and the charge in Table 2 is monotone in exactly that characteristic, with the sourcing leg loading on it as well. Supply cost and willingness to pay therefore rise together in one variable. Froot (2001) documents the same pattern in catastrophe reinsurance, where the market adjusts through the quantity of protection supplied instead of through its price. Section 5.2 adds a price to this argument: the closest traded substitute is priced below the cost of supply at every maturity, by a wide margin under a prudent quoting convention, so a provider would be manufacturing at one price and selling at a lower one. This explanation does not rely on inattention or institutional inertia, and Section 4.5 shows that it is not an artifact of the interest-rate environment. An owner who supplies term from her own holdings faces the same comparison, since she bears the same wedge as an opportunity cost.

The explanation is weaker at book scale. Diversification narrows the six-month shortfall from 72 basis points to about 14 to 17 basis points, which is of the same order as the level calibration, and against realized carry an equal-weighted book is indistinguishable from break-even. For a diversified provider, the evidence therefore shows only that the trade does not cover a prudently quoted risk charge, not that it loses money on the carry or fails by a wide margin. The cost obstacle is large for a desk writing one name at a time, and it is smaller but still present for a diversified provider.

At book scale, only the sourcing leg does not shrink with the book, and it also rests on the least precise evidence: a flag that fires on 2.5% of stock-months and recovers about a quarter of material fails. Three facts limit how much this weakens the argument. The leg matters at book scale

because it fails to diversify, and this property does not depend on the flag’s accuracy as a classifier. The flag’s validity as a conditional signal, which is all the argument uses, is confirmed against SEC filings that play no part in its construction. And rebuilding the leg from realized fail episodes gives similar or larger figures, 20.1 and 26.2 basis points in the top decile under two constructions, compared with 17.6. The size of the book-scale obstacle is therefore uncertain, although its source is clear. Section 6.5 shows that cash settlement removes this leg.

6.2 Is it adverse selection instead?

The standard account of an insurance market that fails to open is informational. If a borrower who asks to fix a fee knows something about the name that the provider does not, the pool of requested contracts is adversely selected, and a provider quoting the unconditional cost loses on every one of them. Unlike the cost explanation, this rival survives diversification, because selection shifts the mean rather than the variance, so writing more contracts multiplies the loss.

The explanation is testable here because I observe the wedge on every stock-month instead of only on contracts someone requested. Sorting stock-months on the growth in utilization over the life of the term approximates what a borrower who knew her own intentions could condition on, while sorting on trailing growth captures what a provider can see at origination. Table IA.22 reports both sorts at six months. Among stock-months with a forward demand measure, the expected six-month wedge averages -3.6 basis points, against -0.8 basis points in the top quintile of forward demand growth, a shift of about three basis points. The 95th percentile in that quintile is 53.7 basis points, against 54.4 basis points across all stock-months with a wedge. The three-basis-point shift is an upper bound on the unpriceable part of selection, since the forward sort uses information unavailable at origination and assumes a borrower who times her request perfectly, and it is small relative to a price–cost gap of 72 basis points on a single position.

The larger part of selection is observable to a provider. Conditioning on trailing demand growth puts the mean six-month wedge in the highest quintile 4.8 basis points above its mean across stock-months with a trailing demand measure, and within the least predictable decile it raises the six-month charge from 200 basis points in the lowest trailing-demand quintile to 378 basis points in the highest. A lender quoting off trailing utilization is compensated for this, and a lender who ignored it would quote far too low, so observable demand pressure reinforces the cost explanation by raising the cost of supply where the product is already expensive. Forward utilization growth is only a proxy for private information, and a borrower whose information is unrelated to the path of demand would not appear in this sort. The data do not refute adverse selection. The version I can measure, however, is more than an order of magnitude smaller than the single-position gap of 72 basis points and between a sixth and a fifth of the book-scale gap of 14 to 17 basis points. It therefore cannot replace the cost explanation for a single position and could account for only part of the gap at book scale.

6.3 How much would a short seller pay?

I next ask whether some buyer should have been willing to pay the cost anyway, so that the missing market reflects gains from trade left on the table. For a diversified fund, certainty about borrowing costs is worth almost nothing. A mean–variance investor with relative risk aversion γ holding a short position of n times capital values the elimination of six-month fee risk at roughly $\gamma n \text{SD}[W^{(126)}]^2 / (4 \times 10^4)$ basis points per year on the position.²³ At the panel-wide six-month dispersion of 245 basis points, an investor with $\gamma = 3$ and a ten-percent position would pay about half a basis point, and even with $\gamma = 10$, a fifty-percent position, and twice the dispersion, she would pay about 30 basis points. Nothing in that range approaches the 424 basis points the top decile costs to supply, because fee risk is idiosyncratic and commands essentially no premium from a diversified buyer. Willingness to pay reaches the charge only for a buyer who prices tail outcomes one for one, such as a leveraged short seller facing forced cover at the 95th percentile of borrowing costs, who values certainty at approximately the quantile itself.

The shape of that demand is familiar from the literature on risk management under collateral constraints. Rampini and Viswanathan (2010) and Rampini et al. (2014) show that hedging and financing compete for the same scarce collateral, so the agents with the strongest motive to hedge are the ones least able to pay for it. The leveraged short seller values a fee guarantee at several hundred basis points because she is near her constraint. For the same reason, she is the buyer least able to post the premium up front.

The gains from trade between a diversified provider, whose top-decile charge at book scale is between a sixth and a tenth of the single-name figure, and a constrained short seller, who would pay several hundred basis points, nevertheless appear substantial. On the evidence of Section 4.4, a key obstacle between them is the sourcing leg, which, unlike fee risk, does not diversify: writing term protection at scale in the least predictable names requires guaranteed access to inventory, and that access is least reliable in the states in which the guarantee is called. That part of the obstacle comes from the obligation to *deliver shares* and is not intrinsic to protection against a fee. Section 6.5 examines contracts that remove it.

6.4 What the charge means for the short seller

Because term protection is not available at scale, a short seller bears the charge as a risk. In the least predictable decile, a six-month position opened at a borrowing cost of 545 basis points may cost 969 basis points in the tail, while in the most predictable decile protection would cost 15 basis points against a fee of 32 basis points. Certainty about borrowing costs is therefore cheap in basis points for most of the cross-section, at 15 to 34 basis points in the six most predictable deciles, and roughly as expensive as the carry for the names where short interest concentrates. She can

²³The realized borrowing cost over a six-month term, per unit of notional, is $\bar{F} \cdot 126/252$, so its dispersion as a fraction of capital is $n \text{SD}[W^{(126)}] \cdot 0.5$, and the mean–variance certainty equivalent of removing it is $\gamma/2$ times its variance. Expressed per unit of notional and annualized, that is $\gamma n \text{SD}[W]^2 \cdot 0.25 \times 10^{-4}$ basis points for SD in basis points.

turn to options: in the option-matched sample, the implied six-month rate for the top fee decile, which is a different set of stocks from the unpredictability decile above, is roughly 409 basis points against a spot fee of 528 basis points. The synthetic route is not a way around the cost, however, since pass-through reaches 0.88 where option markets are deepest. This puts a magnitude on the second short-sale constraint described in the introduction: in the three least predictable deciles, the unhedgeable component of a six-month borrowing cost runs from about three-quarters of the fee to slightly more than all of it (Table 2). I next test whether that constraint affects quantities and prices, and the answers differ in strength.

Loan tenure responds to unpredictability. A one-standard-deviation increase in unpredictability shortens the average tenure of outstanding loans by between 3.8 and 5.5 days, against a sample mean of 74 days, depending on the control set. The estimate survives a control for the fee level, a comparison of stocks with similar fees in the same month (4.4 days, $t = 9.2$), and stock fixed effects (3.8 days, $t = 7.7$). Markit’s tenure is a snapshot of outstanding loans, however, and it reflects recalls by lenders as well as decisions by borrowers. To separate these, I use the number of open loans reported alongside tenure. In a steady state with arrival rate a and daily termination hazard h , a book of N open loans with average tenure T satisfies $N = a/h$ and $T = 1/h$, so terminations run at N/T per day and arrivals at $\Delta N + N/T$. In regressions on the trailing dispersion of the commitment wedge, with stock and month effects and the full control set, tenure falls by 4.4% per standard deviation (4.0% after controlling for the size of the book, which addresses the concern that the snapshot over-represents long-lived loans), terminations rise ($t = 3.66$), arrivals do not change ($t = 0.30$), and the book shrinks ($t = -2.08$). Loans are therefore closed sooner, and they are not replaced by new loans. If lenders were recalling, the effect should concentrate where lendable supply is contracting or concentrated, but the interactions with both are precisely zero. The evidence therefore argues against lender recalls, but it does not positively identify borrowers as the ones closing loans. When the dispersion of the commitment wedge is interacted with the realized fee change, the interaction is significant ($t = 3.86$) but carries the wrong sign for a borrower-driven reading, and it is measured at the same time as the outcome.

The speed of position-building does not respond in the predicted direction. If a borrowing cost that cannot be locked in deterred short sellers from establishing positions, short interest should accumulate more slowly in high-unpredictability names. Instead, the coefficients are small but precisely estimated, at about 4 to 12 basis points of shares outstanding per standard deviation against a mean short interest ratio of about 3.5%, and they carry the opposite sign at every horizon (Table IA.23, Panel C).

Prices provide only weak evidence (Table IA.23). Unconditionally, forward returns load negatively on unpredictability, but the loading does not survive a control for the fee level, which is already known to predict returns (Jones and Lamont, 2002). Conditional sorts show a stronger pattern. When stocks are sorted within fee quintile, sixteen of twenty return spreads are indistinguishable from zero, and the four that are not are the four horizons of the most expensive quintile, where the spread between the most and least unpredictable stocks runs from -0.7 percentage points

at one month to -4.8 percentage points at twelve months. The framework points to this quintile, although I did not specify the interaction in advance, and a block bootstrap of the largest $|t|$ across the twenty cells gives $p = 0.025$. Two limitations stop me from treating this as an asset-pricing result. First, a placebo sort on total fee volatility produces a similar spread out to six months, so the returns do not separate unforecastability from fee movement. Second, the six-month spread is -1.3 percentage points ($t = -0.73$) under value weighting and changes sign once microcaps are excluded (Table IA.24). The return evidence is therefore confined to small stocks, where a constraint on arbitrage capital would bind first.²⁴

6.5 Four changes to the contract

The paper has priced one contract: a single-name, physically delivered, fixed-rate term loan, which my interviewee described as the product clients ask for. Each leg of equation (1) is tied to a different feature of that contract, so each suggests a change, and because every leg has been estimated, each change can be priced. Table 8 reports the six-month charge under each design, on the same 95% break-even convention used throughout. None of these contracts is observed, so the figures are costs of manufacture over 2007–2024, not prices at which the products would clear. The exercise establishes whether the cost obstacle is a property of one particular contract or of the underlying risk.

Cash settlement removes the sourcing leg. A term loan obliges the provider to deliver shares every day of the term, whereas a fixed-for-floating swap on the fee obliges it only to pay a rate difference. The borrower keeps its open borrow, receives the amount by which the fee it pays exceeds the fixed rate, and pays any shortfall, so nothing is delivered and nothing has to be sourced. On a single contract, this is the smallest of the four changes, worth 17.6 basis points in the top decile against a risk charge of 424 basis points. Its importance is at scale, where the sourcing leg rises to between about 30% and 42% of the risk charge, depending on whether the large-book charge is measured directly or extrapolated (Section 4.4). Cash settlement also extends the other changes to names whose shares a provider could never reliably obtain, since a provider that need not source shares can write protection on them.

Pooling exploits the idiosyncrasy of the risk. Because the commitment wedge is almost entirely idiosyncratic, protection sold on a basket costs far less than the same protection sold name by name. When the charge is computed on the equal-weighted average wedge across the names in a

²⁴Tests with outcomes other than returns are mixed. The dispersion of the commitment wedge does not predict Hou and Moskowitz (2005) price delay once standard controls are included, although fee unpredictability itself does in a horse race between the two ($t = 2.91$), and post-announcement drift shows no asymmetry with respect to the dispersion of the commitment wedge. One further return relation is statistically significant, but its sign runs against the overpricing account. In value-weighted panel regressions with stock and month effects that exclude the smallest NYSE size quintile, forward returns load positively on the trailing dispersion of the commitment wedge (1.5 percentage points over six months, $t = 2.45$). The relation does not appear in a value-weighted quintile sort on the same variable ($+0.6$ percentage points, $t = 0.35$), and I do not interpret it further.

decile each month, which does not assume independence, the top-decile six-month charge falls from 424 basis points to 61 basis points, and the ninth-decile charge falls from 150 to 51 basis points. Because fees mean-revert from high levels, the provider’s expected payout on the top-decile basket is -43 basis points, so the charge is a pure risk loading. The figure describes a book of roughly a hundred names rebalanced monthly, and a provider concentrated in fewer names would face more. It is also a 95th percentile of about two hundred monthly averages, but a block bootstrap gives a 5th-to-95th percentile interval of $[58, 64]$ basis points. This change has the largest single effect. It relies on the idiosyncrasy of the wedge, which also weakens the cost explanation at book scale.

Capping sells the tail. Section 6.3 locates the demand in the tail, among buyers who face forced cover, whereas a fixed rate sells such a buyer the whole distribution. A cap at K pays $\max(W^{(H)} - K, 0)$ and leaves the first K basis points of any increase with the borrower. The strike is pinned down by the buyer’s constraint. A short seller holding a position of n times capital who will tolerate a carry loss of L before covering is forced out when the fee rises by

$$K = \frac{10^4 L}{n \cdot H/252}, \quad (6)$$

so at six months, a half-capital position with a tolerance of fifty basis points implies a strike near 200 basis points, and a one-percent tolerance implies a strike near 400 basis points. On a single top-decile name, a cap struck 200 basis points above the origination fee costs 224 basis points, against 424 basis points for a fixed rate, and the protection it gives up lies exactly in the region the marginal buyer does not value.²⁵

In the top decile, pooling reverses the ranking. A fixed-rate book pays the average wedge, so a name whose fee falls offsets one whose fee rises, whereas a book of caps pays the average of $\max(W - K, 0)$, in which no such offsetting is possible. The top-decile pooled charge is 61 basis points for the fixed rate, against 128 basis points for a cap struck at 200 and 92 basis points for a cap struck at 500, and the bootstrap intervals do not overlap. Below the top decile, the pooled cap is cheaper than the pooled fixed rate, and in the ninth decile the cap struck at 200 costs 43 basis points against 51. In the names where protection is most wanted, capping and pooling therefore substitute for each other. A provider who can diversify should write the swap there, while a desk that cannot pool, such as one writing a single name, should write the cap. This result is unfavorable for the product. Demand is concentrated in the tail, and the contract the marginal buyer wants most there is expensive to supply at scale.

Shortening the term. Tenor is also a price lever, although the pooled charge is not monotone in it. For top-decile names, the pooled fixed-rate charge is 29 basis points at one month, against

²⁵The 95th-percentile convention is informative for a cap only where the strike is breached on at least one contract in twenty. That condition fails through most of the cross-section, so the capped columns of Table 8 are reported at basket scale. The condition does hold in the top decile for strikes below its 424 basis point charge, such as the 200 basis point strike used in the text, and there the single-name figure is simply the fixed-rate charge less the strike.

61 basis points at six months and 50 basis points at one year, so a provider unwilling to write six-month protection on a hard-to-borrow name may be willing to write one-month protection and roll the offer.

What the changes are worth together. A one-month, cash-settled, fixed-rate contract written on a basket of top-decile names costs about 29 basis points, against 441 basis points for the six-month single-name delivered contract, which is a reduction from roughly four-fifths of the carry to about a twentieth of it. At the six-month tenor, the pooled cash-settled contract still costs only 61 basis points. Pooling also widens the set of names on which protection is cheap relative to the fee, bringing the eighth-decile charge to 35 basis points and the ninth-decile charge to 51 basis points, against fees of 81 and 144 basis points. Two cautions apply. The option-implied benchmark prices a fixed forward rate, so I cannot compare a redesigned product with a market price as Section 5.2 does for the standard one. And the redesigned contracts transfer less risk than the standard product. Section 6.3 suggests, however, that few buyers want full certainty at a price anyone would quote, so the marginal buyer would purchase less protection in any case. The absence of term protection is therefore a fact about the packaging of the contract as well as about the underlying risk. Pooling, capping, and shorter tenors could be offered on delivered contracts today, while cash settlement awaits a settlement rate.

6.6 Why the cheaper contract is also unwritten

A pooled, cash-settled contract costs a fraction of the product the market declines to supply, yet it is also missing. Two explanations are available, and they differ in kind.

The first answer is infrastructural, and it applies to cash settlement, which the other changes need in order to reach the hardest-to-source names. A contract that pays a rate difference needs a rate to settle against, and none exists. The fee series used here is proprietary, is an average across outstanding loans rather than a transaction price, and is not a rate any two parties could contract on. Transaction-level reporting under Rule 10c-1a would support a published daily reference rate per security from 2029, and a cash-settled fee swap would become writable once such a fixing exists. On this reading, the cheaper contract is absent because it cannot yet be settled, even though it may be economic.

The second answer is that the pooled design is already written, but sold to the other side of the market. Exclusives commit a lending portfolio forward at a guaranteed fee, which is the diversification result of Section 4.2.4 in contractual form. An exclusive, however, delivers access to inventory at a negotiated rate, not protection against that rate moving, so the technology that makes pooled fee protection cheap is in use but has not been directed at the demand identified in Section 6.3. Neither answer requires my cost estimates to be wrong, and both will be testable once loan-level terms are observable. If cash-settled or pooled fee protection is still missing after 2029, once a settlement rate can exist, the explanation will lie outside what this paper measures.

7 Conclusion

Equity lending is an open market in a product that many of its users would prefer to buy on term. This paper asks what supplying that product would cost and what the market would pay for it.

The cost has an unusual shape. Locking in the expected carry, which a term quote appears to price, costs almost nothing, since a provider quoting today’s overnight rate is approximately right on average at every horizon. Most of the cost comes from replication risk, whose dispersion grows with the horizon instead of averaging away, reaching 271 basis points at one year, and from a sourcing leg that is negligible for most stocks and severe for a few. Both legs load on how unforecastable a stock’s fee has been. The two legs behave differently for a provider. Replication risk is idiosyncratic and falls from 424 basis points on a single top-decile position to 61 basis points on a basket, while the sourcing leg, an expected cost, holds near 18 basis points at any size. Writing term protection at scale therefore requires reliable access to inventory in the names where protection is wanted, in addition to capital against fee variation, and that obligation comes from physical delivery, not from fee protection itself.

Combining the lending data with option prices lets me compare that cost with a price, which neither source permits on its own. The option-implied term rate breaks even against the cost of supply at the 24th percentile of the replication-cost distribution at three months, the 27th at six months, and the 41st at twelve months. The options market prices the carry correctly, with pass-through of about 0.91 across fee deciles once mean reversion is netted out, but it pays for little of the risk around the carry. A product whose cost of supply exceeds the price the market will pay is unlikely to be written, and I take that as the most economical account of why term equity lending is scarce. The claim is strongest for a single position and weaker for a diversified provider, for whom the gap is small and only the sourcing leg fails to diversify.

The cost I measure is the cost of one particular contract. Cash settlement and pooling take the top-decile six-month cost of the delivered contract from 441 basis points to 61 basis points, and that differently packaged product appears to be missing for two reasons: there is no rate to settle against until disclosure makes one possible, and the pooled design that does exist, the exclusive, is sold as access to inventory and not as protection against the fee. Interest rates are the main risk in the historical rebate-based product but explain little of the cross-section of risk in the fee-based one, and short sellers’ ability to use options mitigates the constraint where option markets are deep without removing it.

Several limitations bound the exercise. My fee measure is an average across outstanding loans rather than the rate on new loans, which understates day-to-day variation and makes the risk estimates conservative. The level of the option-implied rate depends on a calibration anchor and on an early-exercise correction that uses OptionMetrics’ own volatilities, and the price comparison covers only optionable stocks. The analysis also takes the borrower’s demand as given. The most important limitation is that no term quote is observed anywhere in the paper. Loan-level reporting under Rule 10c-1a will make contracted terms publicly observable from 2029. If the single-name delivered contract is still unwritten once these terms are reported, the evidence here suggests that

it will be missing because it is costliest to manufacture where protection is most demanded.

References

- Allen, F. and Gale, D. (1994). *Financial Innovation and Risk Sharing*. MIT Press, Cambridge, MA.
- Andrews, S., Lundblad, C. T., and Reed, A. V. (2026). Dancing to the same tune: Commonality in securities lending fees. Working paper. Available at SSRN: <https://ssrn.com/abstract=3690337>.
- Battalio, R. and Schultz, P. (2011). Regulatory uncertainty and market liquidity: The 2008 short sale ban’s impact on equity option markets. *Journal of Finance*, 66(6):2013–2053.
- Blocher, J., Reed, A. V., and Van Wesepe, E. D. (2013). Connecting two markets: An equilibrium framework for shorts, longs, and stock loans. *Journal of Financial Economics*, 108(2):302–322.
- Boehmer, E., Jones, C. M., and Zhang, X. (2008). Which shorts are informed? *Journal of Finance*, 63(2):491–527.
- Campbell, J. Y. and Shiller, R. J. (1991). Yield spreads and interest rate movements: A bird’s eye view. *Review of Economic Studies*, 58(3):495–514.
- Chague, F., De-Losso, R., De Genaro, A., and Giovannetti, B. (2017). Well-connected short-sellers pay lower loan fees: A market-wide analysis. *Journal of Financial Economics*, 123(3):646–670.
- Cochrane, J. H. and Piazzesi, M. (2005). Bond risk premia. *American Economic Review*, 95(1):138–160.
- Cocquemas, F. (2017). The term structure of securities lending fees. Working paper, Vanderbilt University.
- Cohen, L., Diether, K. B., and Malloy, C. J. (2007). Supply and demand shifts in the shorting market. *Journal of Finance*, 62(5):2061–2096.
- D’Avolio, G. (2002). The market for borrowing stock. *Journal of Financial Economics*, 66(2–3):271–306.
- Dew-Becker, I., Giglio, S., Le, A., and Rodriguez, M. (2017). The price of variance risk. *Journal of Financial Economics*, 123(2):225–250.
- Diether, K. B., Lee, K.-H., and Werner, I. M. (2009). Short-sale strategies and return predictability. *Review of Financial Studies*, 22(2):575–607.
- Drechsler, I., Savov, A., and Schnabl, P. (2017). The deposits channel of monetary policy. *Quarterly Journal of Economics*, 132(4):1819–1876.
- Du, W., Tepper, A., and Verdelhan, A. (2018). Deviations from covered interest rate parity. *Journal of Finance*, 73(3):915–957.
- Duffie, D. (1996). Special repo rates. *Journal of Finance*, 51(2):493–526.
- Duffie, D., Gârleanu, N., and Pedersen, L. H. (2002). Securities lending, shorting, and pricing. *Journal of Financial Economics*, 66(2–3):307–339.
- Duffie, D. and Rahi, R. (1995). Financial market innovation and security design: An introduction. *Journal of Economic Theory*, 65(1):1–42.
- Engelberg, J. E., Reed, A. V., and Ringgenberg, M. C. (2018). Short-selling risk. *Journal of Finance*, 73(2):755–786.
- Evans, R. B., Geczy, C. C., Musto, D. K., and Reed, A. V. (2009). Failure is an option: Impediments to short selling and options prices. *Review of Financial Studies*, 22(5):1955–1980.
- Fama, E. F. and Bliss, R. R. (1987). The information in long-maturity forward rates. *American Economic Review*, 77(4):680–692.

- Fotak, V., Raman, V., and Yadav, P. K. (2014). Fails-to-deliver, short selling, and market quality. *Journal of Financial Economics*, 114(3):493–516.
- Froot, K. A. (2001). The market for catastrophe risk: A clash of supply and demand. *Journal of Financial Economics*, 60(2–3):529–571.
- Geczy, C. C., Musto, D. K., and Reed, A. V. (2002). Stocks are special too: An analysis of the equity lending market. *Journal of Financial Economics*, 66(2–3):241–269.
- Giglio, S., Kelly, B. T., and Kozak, S. (2024). Equity term structures without dividend strips data. *Journal of Finance*, 79(6):4143–4196.
- Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66(2–3):361–407.
- Grundy, B. D., Lim, B., and Verwijmeren, P. (2012). Do option markets undo restrictions on short sales? Evidence from the 2008 short-sale ban. *Journal of Financial Economics*, 106(2):331–348.
- He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126(1):1–35.
- Hou, K. and Moskowitz, T. J. (2005). Market frictions, price delay, and the cross-section of expected returns. *Review of Financial Studies*, 18(3):981–1020.
- Jones, C. M. and Lamont, O. A. (2002). Short-sale constraints and stock returns. *Journal of Financial Economics*, 66(2–3):207–239.
- Jordan, B. D. and Jordan, S. D. (1997). Special repo rates: An empirical analysis. *Journal of Finance*, 52(5):2051–2072.
- Kaplan, S. N., Moskowitz, T. J., and Sensoy, B. A. (2013). The effects of stock lending on security prices: An experiment. *Journal of Finance*, 68(5):1891–1936.
- Kolasinski, A. C., Reed, A. V., and Ringgenberg, M. C. (2013). A multiple lender approach to understanding supply and search in the equity lending market. *Journal of Finance*, 68(2):559–595.
- Krishnamurthy, A. (2002). The bond/old-bond spread. *Journal of Financial Economics*, 66(2–3):463–506.
- Lustig, H., Stathopoulos, A., and Verdelhan, A. (2019). The term structure of currency carry trade risk premia. *American Economic Review*, 109(12):4142–4177.
- Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *Journal of Finance*, 32(4):1151–1168.
- Muravyev, D., Pearson, N. D., and Pollet, J. M. (2022). Is there a risk premium in the stock lending market? Evidence from equity options. *Journal of Finance*, 77(3):1787–1828.
- Nagel, S. (2016). The liquidity premium of near-money assets. *Quarterly Journal of Economics*, 131(4):1927–1971.
- Ofek, E., Richardson, M., and Whitelaw, R. F. (2004). Limited arbitrage and short sales restrictions: Evidence from the options markets. *Journal of Financial Economics*, 74(2):305–342.
- Prado, M. P., Saffi, P. A. C., and Sturgess, J. (2016). Ownership structure, limits to arbitrage, and stock returns: Evidence from equity lending markets. *Review of Financial Studies*, 29(12):3211–3244.
- Rampini, A. A., Sufi, A., and Viswanathan, S. (2014). Dynamic risk management. *Journal of Financial Economics*, 111(2):271–296.
- Rampini, A. A. and Viswanathan, S. (2010). Collateral, risk management, and the distribution of debt capacity. *Journal of Finance*, 65(6):2293–2322.

- Saffi, P. A. C. and Sigurdsson, K. (2011). Price efficiency and short selling. *Review of Financial Studies*, 24(3):821–852.
- Shiller, R. J. (1993). *Macro Markets: Creating Institutions for Managing Society's Largest Economic Risks*. Oxford University Press, Oxford.
- Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *Journal of Finance*, 52(1):35–55.
- Stambaugh, R. F., Yu, J., and Yuan, Y. (2012). The short of it: Investor sentiment and anomalies. *Journal of Financial Economics*, 104(2):288–302.
- van Binsbergen, J. H., Brandt, M. W., and Koijen, R. S. J. (2012). On the timing and pricing of dividends. *American Economic Review*, 102(4):1596–1618.
- van Binsbergen, J. H., Diamond, W. F., and Grotteria, M. (2022). Risk-free interest rates. *Journal of Financial Economics*, 143(1):1–29.
- Vayanos, D. and Weill, P.-O. (2008). A search-based theory of the on-the-run phenomenon. *Journal of Finance*, 63(3):1361–1398.
- Weitzner, G. (2023). Term structure of short selling costs. *Review of Finance*, 27(6):2125–2161.




Figure 1: **The term structure of expected lending fees.** This figure plots the mean commitment wedge $E[W^{(H)}]$ against the horizon, with one line per decile of the origination-day fee. The commitment wedge is the amount by which the realized cost of carrying an H -day term exceeds the rate observable when the provider commits, and it is expressed here as a percentage of the origination-day fee. The horizontal axis is the term on a log scale. The figure shows the curve a provider would face if she quoted today's rate and then borrowed day to day to deliver it. The pooled mean wedge, which the figure does not plot, is nearly flat, so a provider quoting the overnight rate is approximately right on average at every horizon. Conditionally, the curve reflects mean reversion: the cheapest decile's expected carry over a year is roughly double its spot fee, while the most expensive decile's is about 21% below its spot fee. Expected carry therefore contributes almost nothing to the cost of supplying term protection. Section 4.2.1 examines the risk around this average.




Figure 2: **Replication risk grows with the term.** This figure plots the dispersion of the commitment wedge against the horizon. The solid line is the standard deviation of the commitment wedge $W^{(H)} = C^{(H)} - F_t$, which is the error a provider bears if she quotes the origination-day overnight rate for an H -day term. The dashed line is the dispersion that an average of independent transitory deviations of daily fees from the origination-day rate would have, $SD[W^{(2d)}]/\sqrt{H/2}$. Both axes are logarithmic, and both series are in basis points. Under independence, the two lines would coincide. Because fee shocks are persistent, realized dispersion rises with the horizon while the benchmark falls, and the divergence is already visible at three days.




Figure 3: **The price and the cost of term protection.** This figure compares the option-implied price of term protection with its cost of supply at each maturity. The two lower series are the option-implied term borrowing rate and the spot fee at origination, which lie close together and cross between six and twelve months, and the top series is the full cost of supply, which equals the spot fee plus the replication-risk charge plus the sourcing leg, as in Table 5. The shaded region, labeled cost minus price, is the shortfall of price below cost, which Table 5 reports as price-minus-cost margins of -49 basis points at three months, -72 basis points at six months, and -100 basis points at twelve months. The implied rate rises across maturities from 70 to 80 basis points and roughly tracks expected carry, while the risk charge grows several times faster, so the shortfall widens with the term. Table 6 shows that this widening is a property of the 95th-percentile quoting convention; at the median, the shortfall is roughly flat across maturities. Section 5.2 discusses the cautions that apply to the level of the comparison.

Table 1: Summary Statistics

This table reports summary statistics for the sample. The data are Markit Securities Finance daily equity lending data merged to CRSP, July 2007 through December 2024, restricted to U.S. common stocks on the NYSE, AMEX, and NASDAQ with prices above \$5. Fees are winsorized at the 1st and 99th percentiles within month. Panel A reports the borrowing-cost measures. Panel B reports the realized cost of manufacturing a term loan, $C^{(H)}$, defined as the day-weighted average fee actually paid over an H -trading-day term and computed only where the full daily path is present. Panel C reports lending-market characteristics at the origination date. Observations are stock-months on a month-end origination grid.

	Mean	SD	p10	p50	p90	N
<i>Panel A: borrowing costs (annualized basis points)</i>						
Overnight fee at origination	132.7	564.8	20.7	30.0	178.5	454,131
Fee unpredictability	35.6	86.1	3.5	10.2	76.9	409,537
Total fee volatility (12m)	51.6	184.3	1.7	7.3	93.4	474,286
Mean absolute daily fee change (63d)	8.8	25.2	0.5	2.3	19.8	432,073
<i>Panel B: realized replication cost, $C^{(H)}$ (annualized basis points)</i>						
2 trading days	131.8	562.4	21.2	30.0	175.7	425,179
3 trading days	131.7	562.3	21.4	30.0	175.0	421,606
5 trading days	131.4	559.9	21.6	30.3	174.5	417,533
10 trading days	130.7	556.0	21.8	30.5	172.2	409,503
15 trading days	130.1	550.0	22.0	30.7	171.7	402,206
21 trading days (1 month)	129.8	543.3	22.2	30.9	172.8	394,340
42 trading days	127.1	520.5	22.6	31.1	169.1	368,316
63 trading days	125.6	506.5	22.9	31.4	168.0	361,046
126 trading days (6 months)	119.2	457.5	23.5	31.5	163.1	335,991
252 trading days (12 months)	105.1	371.9	24.3	31.3	146.3	297,480
<i>Panel C: lending-market characteristics</i>						
Utilization	0.16	0.19	0.01	0.08	0.43	542,364
Cost-of-borrow score (1–10)	1.38	1.28	1.00	1.00	2.00	452,606
Loan tenure (days)	74.44	65.35	19.95	58.41	142.91	453,304
Active brokers	16.76	7.01	7.45	16.85	25.67	454,131
Lender concentration	0.20	0.11	0.13	0.17	0.27	454,131
Share of days with no fee change	0.30	0.29	0.00	0.22	0.76	432,073
Log market capitalization	7.08	1.88	4.70	7.00	9.63	558,205

Table 2: The Break-Even Charge and Sourcing Failure by Fee Unpredictability

This table reports the break-even charge and the sourcing leg by decile of fee unpredictability. The charge is the empirical 95th percentile of $W^{(H)}$ within each cell, in basis points, that is, the markup over the origination-day overnight rate at which a provider breaks even on 95% of contracts. Deciles are formed within month on fee unpredictability, defined as the trailing dispersion of out-of-sample errors from an expanding-window fee forecasting model. The column “ $\times$ fee” expresses the charge as a fraction of the mean overnight fee in the cell. $P(> 100)$ is the share of contracts on which the realized cost exceeds the overnight rate by more than 100 basis points. The final two columns report the sourcing leg for six-month terms: the mean share of days inside the term on which shares could not be sourced, and the probability that a term contains at least one such day. A stock-day is flagged as a sourcing failure when utilization exceeds 95% or the Markit cost-of-borrow score is in its top two categories.

Decile	One month (21d)				Six months (126d)					
	Spot	Charge	$\times$ fee	$P(> 100)$	Spot	Charge	$\times$ fee	$P(> 100)$	Fail%	$P(\geq 1)$
1 (low)	32.9	8.8	0.27	0.08%	32.2	15.0	0.47	0.11%	0.02%	0.06%
2	33.3	10.3	0.31	0.10%	32.8	18.4	0.56	0.08%	0.00%	0.08%
3	35.0	11.9	0.34	0.15%	34.5	20.5	0.60	0.18%	0.00%	0.04%
4	37.0	13.8	0.37	0.22%	36.5	22.9	0.63	0.29%	0.01%	0.09%
5	39.0	16.5	0.42	0.26%	38.5	28.1	0.73	0.57%	0.01%	0.08%
6	44.7	19.8	0.44	0.35%	43.9	34.1	0.78	1.00%	0.02%	0.19%
7	56.0	26.9	0.48	0.79%	55.0	50.0	0.91	2.01%	0.08%	0.55%
8	81.5	42.7	0.52	1.67%	80.6	84.8	1.05	4.20%	0.24%	0.97%
9	145.2	74.8	0.52	3.63%	143.6	149.8	1.04	7.32%	0.69%	2.73%
10 (high)	558.8	189.7	0.34	8.44%	545.4	423.6	0.78	13.45%	5.14%	13.17%

Table 3: What Predicts Replication Risk

This table reports regressions of the absolute replication error $|W^{(H)}|$ on lagged fee unpredictability at each horizon, with $|W^{(H)}|$ winsorized at the 1st and 99th percentiles within month. The full specification controls for the fee level, total fee volatility, size, return volatility, utilization, turnover, the Markit cost-of-borrow score, average loan tenure, active brokers, lender concentration, the Engelberg et al. (2018) short-selling risk measure, and idiosyncratic volatility. All right-hand-side variables are standardized within month and lagged. Because term windows observed on a monthly grid overlap, standard errors are clustered two-way by stock and month. Two-way clustering does not cover correlation between different stocks at different months inside overlapping windows. That correlation is governed by the common component of the wedge, whose R^2 of roughly 0.006 (Section 4.2.4) makes further adjustment immaterial. The middle panel replaces month fixed effects with fee-decile $\times$ month fixed effects while keeping the full control set, so that the coefficient is identified only from stocks carrying similar fees in the same month. This specification addresses whether unpredictability is merely a proxy for being hard to borrow. The month-fixed-effects specification is the economically relevant one for a provider quoting on a stock’s characteristics. The stock-fixed-effects specification asks a different question, about within-stock timing instead of cross-sectional pricing, and is discussed in the text.

	2d	3d	5d	10d	15d	21d	42d	63d	126d	252d
<i>Month fixed effects, full controls</i>										
Unpred	1.90	1.91	2.10	2.48	2.68	3.00	3.54	3.78	4.81	6.64
<i>t</i>	7.27	7.05	6.77	6.80	6.41	6.17	5.10	4.29	3.40	3.31
<i>Fee-decile $\times$ month fixed effects, full controls</i>										
Unpred	0.94	0.87	0.94	1.13	1.25	1.48	1.73	1.76	2.41	3.48
<i>t</i>	4.29	3.90	3.67	3.67	3.58	3.68	3.03	2.39	1.93	1.97
<i>Stock and month fixed effects, full controls</i>										
Unpred	0.39	0.20	0.22	0.14	0.00	-0.24	-0.93	-1.87	-3.66	-5.52
<i>t</i>	2.11	1.01	0.97	0.53	-0.02	-0.72	-1.96	-2.91	-3.25	-3.42

Table 4: The Two Legs at the Book Level (six-month term)

This table reports the two legs of the six-month cost of supply at three book sizes, by decile of fee unpredictability. The risk leg is the charge a provider must add to break even on 95% of contracts. Because replication errors are almost entirely idiosyncratic (Section 4.2.4), the variance of an N -name average falls toward the decile's own common component, which is estimated within decile and corrected for sampling noise; the fat-tail correction from the single-name distribution is carried across book sizes. The $N = 50$ and large-book entries for the middle deciles are imprecise and non-monotone because their common component is small relative to its sampling error, so the argument rests on the top decile and its contrast with the bottom. The sourcing leg is the probability that a six-month term contains at least one day on which shares could not be sourced, multiplied by the buy-in cost of Table IA.11. The risk leg falls to between a sixth and a tenth of its single-name value across a book, whereas the sourcing leg, an expected cost, does not fall with the number of names.

Decile	Risk leg (95th pct of W)			Sourcing leg	
	Single	$N = 50$	Large book	$P(\text{fail})$	$P \times \text{buy-in}$
1 (low)	15.0	2.4	1.2	0.1%	0.1
2	18.4	5.3	4.7	0.1%	0.1
3	20.5	4.1	2.9	0.0%	0.1
4	22.9	5.7	4.7	0.1%	0.1
5	28.1	5.8	4.3	0.1%	0.1
6	34.1	6.6	4.6	0.2%	0.3
7	50.0	8.6	4.9	0.6%	0.7
8	84.8	13.7	6.6	1.0%	1.3
9	149.8	25.5	14.3	2.7%	3.6
10 (high)	423.6	72.7	41.6	13.2%	17.6

Table 5: Price Against Cost

This table compares the option-implied price of term protection with its cost of supply. The implied fee is recovered from put-call parity on matched option pairs, corrected for American early exercise (Section 5.1.1), and level-calibrated so that, across the contracts in the cheapest fee decile, it equals their realized term cost. Section 5.1 describes the screens. Expected carry is the origination-day fee plus the mean commitment wedge in the stock's fee-decile and maturity cell. Full cost is the origination-day fee plus the replication-risk charge plus the sourcing leg of equation (1). The charge is added to the origination-day fee directly, since it is quoted as a markup over that fee in the stock's unpredictability decile and maturity bucket, which avoids overstating the gap for expensive stocks whose fees revert. The origination-day fee is observed at origination, while the charge is a full-sample quantile for the stock's cell (a real-time version is reported in Table IA.9, Panel C), and the sourcing leg differs from both: it is computed as the buy-in cost times an indicator for whether *that* term actually contained a day on which shares could not be sourced, so at the level of a single row it uses information the provider would not have had at origination. Averaged over rows, the indicator has expectation $\Pr(\text{fail})$, which is the quantity equation (1) calls for, so the column means reported here are the intended object, and nothing in the paper reads the full-cost column row by row. "Price – E[carry]" compares two prices, and "price – cost" compares a price with what supplying the contract requires. Figures are in annualized basis points. Panel B reports six-month contracts, and in its cheapest decile the implied rate sits 3.6 basis points above expected carry. The two margin columns move in opposite directions as maturity lengthens because the replication-risk charge grows with the horizon faster than the implied rate does.

	<i>N</i>	Implied	Spot	Full cost	Price – E[carry]	Price – cost
<i>Panel A: by maturity</i>						
3 months	91,102	69.6	75.1	118.8	–3.8	–49.2
6 months	120,286	71.1	77.1	143.2	–3.3	–72.1
12 months	54,166	79.8	79.3	179.9	3.4	–100.1
<i>Panel B: by decile of the spot fee, six-month contracts</i>						
1 (cheap)	12,126	35.8	22.2	63.8	3.6	–28.0
2	12,204	37.2	26.1	63.0	5.0	–25.8
3	12,107	34.8	28.1	65.8	1.7	–31.0
4	12,326	34.9	29.9	69.1	0.5	–34.2
5	12,144	37.3	31.8	72.8	1.3	–35.5
6	12,388	38.0	34.3	77.7	0.9	–39.7
7	12,486	40.2	38.1	86.1	0.2	–45.9
8	12,321	43.5	46.0	104.6	–2.3	–61.2
9	12,362	68.0	76.1	179.7	–16.7	–111.6
10 (expensive)	9,822	408.9	527.6	773.9	–32.6	–365.0

Table 6: The Margin as a Function of the Quoting Convention

This table reports the cost of supply and the price-minus-cost margin when the replication-risk leg is set at the stated quantile of the commitment wedge instead of the 95th percentile used in Table 5. Because that percentile is a prudent quoting convention and not an estimate of a risk premium (Section 4.2.2), the full schedule is the paper’s central comparison. The bottom block states the comparison without a convention. q^* is the quantile at which the margin crosses zero, solved exactly by bisection on the same stock-month sample as the rows above; the market’s price covers the expected carry plus the risk distribution up to q^* and no further. The coverage rate is the share of individual matched contracts on which writing at the option-implied rate would have covered realized cost. The coverage rate counts contracts, whereas q^* averages basis points, so the two weight the cross-section differently and answer different questions. The coverage rate is computed on a smaller set of contracts, 90,441 at three months, 112,319 at six months, and 45,718 at twelve months. At the median, a natural central benchmark for a diversifiable risk, the margin is already negative at every maturity at the baseline anchor, although these median margins lie within the anchor’s ± 20 basis point reach. The convention governs the maturity profile: the gap widens with the term at the 95th percentile and is roughly flat at the median. Figures are in basis points on the option-matched sample.

Quoting quantile	Three months		Six months		Twelve months	
	Cost	Margin	Cost	Margin	Cost	Margin
Median	76.4	−6.8	79.1	−7.9	82.6	−2.7
75th percentile	80.7	−11.1	84.7	−13.5	89.8	−10.0
90th percentile	93.3	−23.7	103.4	−32.2	120.6	−40.8
95th percentile (the paper)	118.8	−49.2	143.2	−72.1	179.9	−100.1
<i>Embedding no convention</i>						
Break-even quantile q^*		24.1		26.9		40.5
Ex-post coverage rate (%)		46.1		47.9		51.5

Table 7: How the Options Market Responds to Borrowing Pressure

This table examines how the options market responds to borrowing pressure. The second column reports t -statistics in Panels A and D and in the interaction rows of Panel E, and observation-level slopes in Panel C. The first-column slopes in Panels B and C and in the three tercile rows of Panel E are fitted on ten fee-decile cell means and carry no sampling statement, and Table IA.18 reports winsorized observation-level estimates with inference. Panel A regresses the six-month price-minus-expected-carry margin on each measure of lending-market pressure, one at a time, with month fixed effects, the standard controls, and two-way clustering; regressors are standardized within month, so coefficients are in basis points per cross-sectional standard deviation. Panel B varies the quote-quality screen; a slope of one would mean that the two markets price the borrow identically. Panel B is estimated on all quotes passing each screen (140,242 at the baseline screen), which is broader than the stock-month sample of Table 5 used in Panel C, so its baseline slope differs from Panel C's first row. Panel C replaces the spot benchmark with expected carry, formed four ways, and its observation-level slopes use the 100,784 contracts with a nonmissing benchmark (83,485 for the ex-ante fits). Panel D reports the margin in stock-months where the sourcing flag of Table 2 fires. Panel E re-estimates pass-through within terciles of lendable supply and then includes supply and option depth together. Section 5.3 discusses the panels in order.

	Estimate	t or obs.-level slope
<i>Panel A: does the options market see lending-market pressure? (basis points per SD)</i>		
Utilization	24.02	11.27
Cost-of-borrow score	136.23	11.74
Lender concentration	12.32	3.49
Active brokers	6.26	4.59
<i>Panel B: sensitivity of the implied fee to the actual fee, by quote quality (slope)</i>		
Baseline (spread < 0.50)	0.795	
Spread < 0.15	0.828	
Spread < 0.15, open interest ≥ 100	0.832	
Spread < 0.10, open interest ≥ 500	0.869	
<i>Panel C: removing mean reversion from the benchmark (slope)</i>		
Against the spot fee	0.749	0.685
Against expected carry, decile $\times$ maturity cell means	0.906	0.714
Against expected carry, fitted on the log fee (ex ante)	0.820	0.753
adding unpredictability to the fit (ex ante)	0.819	0.758
memo: the log-fee fit in sample	0.851	0.795
<i>Panel D: the margin when shares cannot be sourced (basis points)</i>		
No sourcing failure ($N = 109,490$)	-2.4	
Sourcing failure ($N = 2,829$)	-186.4	
median	-43.2	
tight quotes only	-151.2	
Sourcing failure, controlling for the fee level	150.9	5.39
<i>Panel E: pass-through by lendable supply (slope)</i>		
Low supply (18% of market cap)	0.713	
Middle (31% of market cap)	0.704	
High supply (42% of market cap)	0.703	
Interaction, high supply $\times$ fee	-0.128	-0.96
with option depth included	-0.129	-1.26
option depth, same specification	0.341	3.35

Table 8: What Alternative Contract Designs Would Cost to Supply

This table reports the six-month charge in basis points under four contract designs, by within-month decile of fee unpredictability, on the 95% break-even convention of Table 2. *Fee* is the mean origination-day borrowing cost. *Fixed rate, single* is the charge on one name and reproduces Table 2. *Fixed rate, basket* is the 95th percentile of the equal-weighted average wedge across the decile's names within a month; it is computed empirically and does not assume independence. The cap columns report the 95th percentile of the monthly equal-weighted average of $\max(W^{(126)} - K, 0)$ across the same names, so each name's payout is truncated before averaging. A cap pays only when the realized cost of rolling exceeds the origination-day fee by more than K . Cap figures are reported at basket scale because the 95th-percentile convention is degenerate for a deep out-of-the-money cap on a single name. *Sourcing* is the leg that a physically delivered contract carries and a cash-settled one does not, and it should be added to the fixed-rate columns for a delivered loan. Section 6.5 discusses the four changes, and the horizon dimension is reported there instead of being tabulated.

Decile	Fee	Fixed rate		Cap, basket		Sourcing
		Single	Basket	200bp	500bp	
1 (low)	32	15.0	11.4	1.5	0.2	0.1
2	33	18.4	14.1	1.4	0.0	0.1
3	34	20.5	12.2	0.9	0.0	0.1
4	37	22.9	13.6	1.4	0.0	0.1
5	39	28.1	15.6	3.8	1.8	0.1
6	44	34.1	16.9	4.2	1.9	0.3
7	55	50.0	25.5	10.4	7.1	0.7
8	81	84.8	35.4	20.3	14.1	1.3
9	144	149.8	50.5	43.1	27.9	3.6
10 (high)	545	423.6	60.5	127.5	92.1	17.6

Internet Appendix

A Missing Market: The Cost of Term Lending in Equity Markets

Not for publication

IA.1 Supporting Tables

This appendix collects the supporting tables cited in the main text. Table IA.1 defines the constructed variables, and the remaining tables follow in the order in which they are first cited.

Table IA.1: Variable and Terminology Definitions

This table defines the objects constructed in the paper and lists the section that builds each one. Fees and charges are in annualized basis points throughout. The three legs of equation (1) organize the analysis. A provider selling a fixed rate for H days must pay the floating fee every day of the term, and the cost of doing so is the expected carry, plus the risk that the realized cost departs from it, plus the possibility of not being able to source the shares at all.

Term	Definition	Section
<i>Market terminology</i>		
General collateral	Easy to borrow; lends at a low, stable fee. Most of the market.	2
Special	Borrowing demand high enough to command a materially higher fee.	2
Warm	The band between general collateral and special.	2
Re-rate	The lender's right to reset the fee on an open loan, typically daily.	2
Recall	The lender's right to demand the shares back.	2
Buy-in	Forced purchase in the market when a borrower cannot return recalled shares.	2
Rebate	Under cash collateral, the rate returned on the borrower's cash. The fee is the short rate less the rebate, so a fee is a spread net of policy.	3.2
<i>The three legs of a term quote</i>		
Expected carry	The average daily fee a provider expects to pay while rolling the borrow over the term.	4.1
Replication risk	The risk that the fee actually paid departs from that expectation. Measured by the dispersion of $W^{(H)}$.	4.2.1
Sourcing leg	The buy-in cost times $\Pr(\text{fail})$, the chance that on some day the shares cannot be borrowed at any price.	4.3
<i>Constructed quantities</i>		
$F_{i,t}$ (SAF)	Loan fee: the simple average fee across all outstanding loans on stock i at t .	2.4
Utilization	Value of shares on loan over value available to lend; the gauge of tightness.	2.4
$C_{i,t}^{(H)}$	Replication cost: the day-weighted average fee actually paid over an H -day term.	3
$W_{i,t}^{(H)}$	Commitment wedge: $C_{i,t}^{(H)} - F_{i,t}$, the amount by which the realized cost of carrying the term exceeds the rate observable when the provider commits.	3
The charge	The empirical 95th percentile of $W^{(H)}$ within a cell: the markup over the origination-day rate at which a provider breaks even on 95% of contracts. A quoting convention, not a price of risk.	4.2.2
Fee unpredictability	Trailing 24-month dispersion of out-of-sample errors from a fee forecasting model estimated on an expanding window. The paper's sorting variable.	4.2.2
Sourcing flag	A stock-day on which utilization exceeds 95% or the Markit cost-of-borrow score is in its top two categories.	4.3
Implied fee	Term borrowing rate recovered from put-call parity on matched option pairs, corrected for American early exercise, and level-calibrated so that, across the contracts in the cheapest fee decile, it equals their realized term cost.	5.1
Price – cost margin	Implied fee less the full cost of supply: the origination-day fee plus the charge plus the sourcing leg. Negative means the market pays less than manufacturing costs.	5.2

Table IA.2: The Term Structure of Expected Lending Fees

This table reports the term structure of expected lending fees as the mean commitment wedge $W^{(H)} = C^{(H)} - F_t$ in basis points, where $C^{(H)}$ is the day-weighted average fee actually paid over an H -trading-day term and F_t is the overnight fee on the origination day. Deciles are formed within month on the origination-day overnight fee. The Spot column is the mean origination-day fee in each decile, and the final row pools all stocks. A positive entry means that the expected cost of carrying the term exceeds the rate observable at commitment.

Decile	Spot	2d	3d	5d	10d	21d	42d	63d	126d	252d
1 (cheap)	21.3	4.52	4.92	5.96	7.09	8.85	10.62	11.97	15.77	21.48
2	25.6	1.90	2.10	2.59	3.06	3.89	4.93	5.76	7.67	11.14
3	28.7	0.89	1.02	1.37	1.71	2.55	3.37	3.90	5.37	8.29
4	31.3	0.57	0.65	0.94	1.35	2.25	2.85	3.25	4.60	7.35
5	34.0	0.45	0.52	0.81	1.16	1.73	2.19	2.63	3.95	6.43
6	38.2	-0.17	-0.11	0.16	0.42	1.11	1.44	2.06	3.47	6.14
7	43.3	-0.30	-0.29	-0.11	0.04	0.80	1.53	2.10	3.76	6.99
8	62.2	-2.28	-2.50	-2.22	-2.36	-1.45	-0.49	0.58	4.24	12.10
9	111.9	-1.64	-1.50	-1.26	-0.84	0.43	2.54	5.13	13.94	24.28
10 (expensive)	931.2	-3.28	-4.76	-6.69	-14.45	-28.31	-49.65	-71.42	-119.85	-198.96
All	132.7	0.05	0.00	0.15	-0.26	-0.74	-1.85	-3.04	-4.84	-7.16

Table IA.3: The Sourcing Flag and the Charge, Rebuilt from the Daily File

This table rebuilds the sourcing flag and the charge directly from the daily lending file. The reconstruction mirrors the panel construction, using running sums within stock, month-end originations, and a complete 126-day path with a calendar span of at most 200 days. Because the reconstruction applies its own path filter, its baseline differs slightly from Table 2, with a top-decile failure probability of 11.8% compared with 13.2%. Panels A and B vary the severity of the sourcing flag, and the concentration across deciles is preserved at every severity. Panel C reports the 95th-percentile six-month charge built from unwinsorized fees. It preserves the shape and the roughly 28-fold spread of the winsorized profile in Table 2. Because it uses unwinsorized fees and its own path filter, its levels are not directly comparable with those in that table, and the panel shows only that winsorization does not produce the shape of the charge.

	Dec. 1	Dec. 3	Dec. 5	Dec. 8	Dec. 10
<i>Panel A: probability a six-month term contains a failure day (%), by flag severity</i>					
Loose (Util > 90 or DCBS ≥ 8)	0.1	0.1	0.2	1.3	16.7
Baseline (Util > 95 or DCBS ≥ 9)	0.1	0.0	0.1	0.8	11.8
Tight (Util > 98 or DCBS = 10)	0.1	0.0	0.0	0.4	7.4
<i>Panel B: mean share of failure days inside the term (%), by flag severity</i>					
Loose	0.02	0.01	0.04	0.34	6.90
Baseline	0.02	0.00	0.02	0.17	4.47
Tight	0.01	0.00	0.01	0.08	2.55
<i>Panel C: six-month break-even charge from unwinsorized fees (p95 of W, bp)</i>					
Unwinsorized	14.8	20.5	26.6	72.7	411.2

Table IA.4: The Forecasting Model Loses to the Naive Quote

This table compares a fee forecasting model with the naive quote. The model is an expanding-window panel regression, with stock fixed effects and a minimum estimation window of 36 months, of the realized term cost $C^{(H)}$ on the origination-day overnight fee, lagged fee volatility, utilization, short interest, size, return volatility, turnover, and a lagged market-wide fee. It is re-estimated each month using only data whose outcome was fully observed at the time of the forecast. The naive quote is the origination-day overnight rate. The out-of-sample R^2 is measured against the naive quote, so a negative value means that the model does worse than quoting today's rate for the whole term. For this reason, the paper builds the risk charge on $W^{(H)}$, the error on the naive quote, and not on a model residual.

Horizon	Corr(fitted, actual)	RMSE, model	RMSE, naive	OOS R^2 vs naive	N
21 days	0.957	91.2	85.3	-0.144	335,618
42 days	0.932	106.0	97.5	-0.183	311,021
63 days	0.904	120.2	108.9	-0.218	302,783
126 days	0.836	132.2	110.9	-0.422	273,270
252 days	0.729	132.2	110.6	-0.430	223,844

Table IA.5: Replication Risk Does Not Diversify Over the Term

This table reports the dispersion of the commitment wedge at each horizon and its ratio to an independence benchmark. $SD[W^{(H)}]$ is the standard deviation of the commitment wedge in basis points. The i.i.d. benchmark is $SD[W^{(2d)}]/\sqrt{H/2}$, which is the dispersion that an average of independent transitory daily deviations would have. The within-stock ratio computes each stock's own time-series dispersion, requiring at least 24 origination months, and averages across stocks. Panel B repeats the pooled ratio within three subperiods. Figure 2 plots the $SD[W]$ and benchmark rows of Panel A.

	2d	3d	5d	10d	15d	21d	42d	63d	126d	252d
<i>Panel A: pooled</i>										
N (000s)	425	422	418	410	402	394	368	361	336	297
$SD[W]$	94.6	98.0	104.8	120.6	133.1	147.1	178.5	206.1	244.8	270.9
i.i.d. benchmark	94.6	77.2	59.8	42.3	34.5	29.2	20.6	16.9	11.9	8.4
Ratio	1.00	1.27	1.75	2.85	3.85	5.04	8.65	12.23	20.54	32.15
Ratio, within stock	1.00	1.27	1.76	2.83	3.79	4.89	8.06	11.00	18.47	29.00
<i>Panel B: by subperiod, ratio to the independent-deviation benchmark</i>										
2007–2012	1.00	1.32	1.76	2.94	4.01	5.16	9.03	12.85	24.27	43.33
2013–2018	1.00	1.26	1.64	2.59	3.43	4.48	7.80	11.14	20.16	34.98
2019–2024	1.00	1.27	1.75	2.85	3.83	4.99	8.49	11.94	19.78	30.84
Full sample	1.00	1.27	1.75	2.85	3.85	5.04	8.65	12.23	20.54	32.15

Table IA.6: The Divergence Is Not a Reporting Artifact

This table tests whether the divergence between realized dispersion and the independence benchmark is an artifact of fee reporting. Each cell is the ratio of realized dispersion in the commitment wedge to the dispersion that independent transitory deviations of daily fees from the origination-day rate would produce, $SD[W^{(H)}]/(SD[W^{(b)}]\sqrt{b/H})$. Because SAF averages across outstanding loans, it is smoothed, which attenuates short-horizon variation and would mechanically inflate any ratio anchored at a short base b . Panel A re-bases the benchmark, varying b from two to twenty-one days on a fixed sample of stock-months in which every horizon is observed. Panel A is a restatement and not a test, since the entries satisfy $\text{ratio}(H;b) = \text{ratio}(H;2)/\text{ratio}(b;2)$ identically and therefore decline as b lengthens. It shows the level of the ratio when short-horizon dispersion, which smoothing would attenuate, is not used at all. Panels B and C are the tests. Panel B re-estimates the full curve within terciles of the share of the prior sixty-three trading days on which the fee did not move, with each tercile anchored to its own base dispersion. Panel C restricts the sample to stock-months whose fee moved on more than 75% of days. Under the artifact hypothesis, the divergence would be largest in the stalest tercile; instead, it is largest in the tercile whose fee moves most often.

Anchor / sample	2d	3d	5d	10d	15d	21d	42d	63d	126d	252d
<i>Panel A: moving the anchor</i>										
anchored at 2 days	1.00	1.30	1.84	3.12	4.31	5.73	10.45	14.74	26.31	45.87
anchored at 5 days			1.00	1.70	2.35	3.12	5.68	8.01	14.30	24.93
anchored at 10 days				1.00	1.38	1.84	3.35	4.72	8.42	14.68
anchored at 21 days						1.00	1.82	2.57	4.59	8.00
<i>Panel B: by how often the fee moves</i>										
fee moves most often	1.00	1.35	2.01	3.54	5.01	6.90	12.69	18.37	31.60	49.86
middle tercile	1.00	1.27	1.70	2.72	3.69	4.84	8.02	11.28	19.57	32.07
fee moves least often	1.00	1.27	1.76	2.93	3.93	5.13	9.21	13.37	22.84	38.40
<i>Panel C: actively repriced only</i>										
fee moves on >75% of days	1.00	1.30	1.83	3.08	4.25	5.74	10.20	14.71	25.28	40.00

Table IA.7: An AR(1) Reproduces the Curve

This table compares the realized ratio of wedge dispersion to the independence benchmark with the ratio implied by a stock-by-stock AR(1). The first row repeats the within-stock estimates of Table IA.5. The second row calibrates an AR(1) to each stock's own daily fee series, evaluates equation (5) at that stock's persistence, and averages the resulting curves across stocks. The stationary variance cancels from the ratio, so each stock's implied curve depends only on its estimated ρ . Persistence is estimated on 5,870 stocks with at least 250 daily observations (mean $\rho = 0.798$, median 0.902). The two rows agree at one year, and the AR(1) slightly over-predicts at intermediate horizons, so the growth of replication risk with the horizon is a restatement of fee persistence. The magnitude of this risk comes from the realized measurement. As Section 4.2.1 shows, parameters estimated on one half of the sample do not predict a stock's term structure of risk in the other half.

Benchmark	2d	3d	5d	10d	15d	21d	42d	63d	126d	252d
Realized (within stock)	1.00	1.27	1.76	2.83	3.79	4.89	8.06	11.00	18.47	29.00
AR(1), stock-by-stock	1.00	1.32	1.91	3.22	4.37	5.61	9.21	12.15	19.03	28.95

Table IA.8: The Charge Holding the Fee Level Fixed

This table reports the six-month break-even charge by decile of fee unpredictability under three sorts. The first block forms deciles within month among stock-months with a complete six-month path, whereas Table 2 forms them among all stock-months with an unpredictability measure, so the first block's endpoints (14.6 and 387.4) sit slightly inside those of Table 2 (15.0 and 423.6). The second block forms deciles within fee decile $\times$ month, so every comparison is between stocks that cost about the same to borrow at the same date. This block tests whether the ordering simply restates being hard to borrow. The charge still rises monotonically, and the ratio of the charge to the fee rises from 0.30 to 0.70, almost monotonically, whereas it is hump-shaped in the unconditional sort. The spread falls from roughly twenty-seven-fold to sixfold, however, so most of the magnitude in Table 2 reflects the fee level and not unpredictability. The third block repeats the unconditional sort using a twelve-month trailing window for the unpredictability measure instead of a twenty-four-month window (Section 4.2.2). The two windows give the same profile, so the choice of window does not drive the result. Figures are in basis points.

Decile	Unconditional sort			Within fee decile $\times$ month			12-month window	
	Spot	Charge	$\times$ fee	Spot	Charge	$\times$ fee	Spot	Charge
1 (low)	31.9	14.6	0.46	68.6	20.6	0.30	31.2	15.0
2	32.6	17.5	0.54	65.0	25.6	0.39	32.0	17.2
3	34.2	20.0	0.59	68.3	28.9	0.42	33.5	19.7
4	36.0	21.9	0.61	72.0	31.2	0.43	35.4	21.3
5	37.8	25.9	0.68	78.9	41.4	0.52	37.3	25.3
6	42.2	33.0	0.78	85.2	46.2	0.54	40.6	30.4
7	51.2	43.5	0.85	93.9	49.1	0.52	49.2	42.4
8	70.5	71.2	1.01	105.5	60.2	0.57	63.6	67.6
9	124.2	129.3	1.04	132.6	73.2	0.55	107.8	117.1
10 (high)	490.1	387.4	0.79	179.1	125.0	0.70	481.1	388.1

Table IA.9: Robustness of the Charge and the Break-Even Statistics

This table reports robustness checks for the charge and the break-even statistics. Panels A and B use a matched six-month sample of 58,324 stock-month contracts in the 160–200 day bucket. This sample is narrower than the 112,319 contracts behind the coverage rate in Table 6, so the coverage rate at a zero shift, 47.6%, differs slightly from the 47.9% reported there. Panel A adds the stated shift to the option-implied level and recomputes three statistics: the share of contracts whose implied rate covered the realized replication cost, the mean monthly margin of an equal-weighted book, and the share of months in which that book margin is positive. Panel B reports a moving-block bootstrap over calendar months. Panel C computes the 95th-percentile charge on an expanding window, re-estimated each January from 2010 using only information available at origination, on the subsample with nonmissing unpredictability. Its full-sample row matches the first block of Table IA.8 and differs slightly from Table 2, which forms deciles among all stock-months with an unpredictability measure instead of only those with a complete six-month path. Realized coverage is the share of subsequent contracts whose wedge fell below the real-time charge. Panel D draws twenty equal-weighted fifty-name books per month within the top unpredictability decile and reports the 95th percentile of the book-average wedge across book-months, together with the same quantile for the full decile-month mean. Both exceed the extrapolation in Table 4, because adverse months are shared across a book’s positions. Panel E compares stock-months whose two-day path exists but whose twelve-month path does not with stock-months whose paths are complete.

<i>Panel A: coverage and book statistics under anchor shifts (bp added to the implied level)</i>					
	−20	−10	0	+10	+20
Coverage rate (%)	38.9	43.2	47.6	52.0	56.4
Book mean margin (bp)	−14.5	−4.5	5.5	15.5	25.5
Book break-even (% of months)	31.9	40.7	48.5	54.9	62.3
<i>Panel B: moving-block bootstrap, 95% confidence intervals (block length six months, 1,000 reps)</i>					
Coverage rate	[43.5%, 52.0%]				
Book break-even share	[38.2%, 58.3%]				
<i>Panel C: real-time charge, expanding window re-estimated each January, 2010–2024 (six months, bp)</i>					
	Dec. 1	Dec. 3	Dec. 5	Dec. 8	Dec. 10
Real-time charge	10.9	14.1	16.6	40.5	341.1
Full-sample charge	14.6	20.0	25.9	71.2	387.4
Realized coverage (%)	91.0	90.8	90.1	90.7	94.3
<i>Panel D: empirical book-level 95th percentiles, top unpredictability decile (six months, bp)</i>					
Fifty-name books (drawn)	78.6 (extrapolated: 72.7)				
Decile-month mean (large book)	60.5 (extrapolated: 41.6)				
<i>Panel E: complete-path attrition between the 2d and 252d horizons</i>					
Share of 2d observations without a 252d path	30.5%				
Mean origination fee, dropped vs. kept (bp)	176.1 vs. 112.2				
Mean unpredictability decile, dropped vs. kept	5.92 vs. 5.35				
Top-decile observations missing a 252d path	37.4%				

Table IA.10: Tightening the Fee Control

This table re-estimates the six-month specification of Table 3 as the nonparametric fee control is tightened from deciles to ventiles to percentiles, each interacted with month. Everything else is held fixed: the left-hand variable ($|W^{(126)}|$, winsorized within month), the control block, and two-way clustering by stock and month. The first row therefore reproduces the corresponding entry of Table 3 exactly, and the rows differ only in the partition. The final three columns report how much fee variation each partition leaves behind, namely the mean origination-day fee in the lowest and highest unpredictability deciles formed *within* the partition cell, and their ratio. The coefficient falls roughly in step with the fee variation from which it is identified. Section 4.2.3 interprets this as evidence that at six months the design cannot separate unforecastability from the fee level. The finer partitions are estimated at the six-month horizon only.

Fee control	Unpredictability		N	Mean fee within cell (bps)		
	b	t		Low	High	Ratio
Fee decile $\times$ month (the paper)	2.41	1.93	255,459	71.4	176.2	2.47
Fee ventile $\times$ month	1.11	1.04	255,459	89.2	126.2	1.41
Fee percentile $\times$ month	0.81	0.96	255,283	88.7	98.9	1.12

Table IA.11: What a Buy-In Costs

This table estimates the cost of a buy-in. A stock-day is flagged as a sourcing failure when utilization exceeds 95% or the Markit cost-of-borrow score is in its top two categories. The table reports the quoted half-spread and the cumulative return over the following five trading days, on failure days and on days when sourcing is available. The buy-in cost is the quoted half-spread on failure days plus the excess five-day return on failure days relative to other days. The comparison characterizes the states in which a buy-in occurs; it does not estimate the causal cost of executing one. The state-conditional cost is the relevant object, because a provider forced to cover faces the cost of the environment that forced the trade and not an average execution cost. Market impact is not observable in these data and is excluded, so the total is a lower bound.

	Sourcing available	Sourcing unavailable	Difference
Quoted half-spread (bps)	32.9	93.6	60.8
Return, next 5 days (bps)	32.2	72.2	39.9
<i>Estimated buy-in cost (lower bound)</i>			133.6

Table IA.12: The Charge and the Sourcing Rate by Policy-Rate Regime

This table reports the six-month break-even charge and the share of terms containing a sourcing failure, split by the change in the federal funds rate over the six months following origination. In *flat* months the rate moves by less than 25 basis points (123 months), and in *moving* months it moves by at least 50 basis points (63 months). The two right-hand columns split the moving months by direction. Panel A sorts by decile of fee unpredictability and shows that the cross-sectional ordering is preserved within both regimes. Panel B sorts by the absolute fee at origination. This comparison identifies the effect, since it holds fixed the quantity in which the regimes differ, and it shows that the regime effect is large below 100 basis points and absent between 100 and 1,000 basis points. Figures are in basis points except where noted, and Section 4.5 discusses the table.

Decile	Flat rates ($ \Delta FF_{6m} < 25\text{bp}$)			Moving rates ($\geq 50\text{bp}$)			Charge by direction	
	Spot	Charge	$P(\text{fail})$	Spot	Charge	$P(\text{fail})$	Rising	Falling
<i>Panel A: by decile of fee unpredictability</i>								
1 (low)	31	9.3	0.1%	34	21.4	0.1%	24.3	12.3
2	32	10.8	0.1%	34	27.4	0.1%	31.5	12.3
3	33	12.8	0.1%	37	28.6	0.0%	33.5	12.3
4	35	13.9	0.1%	39	33.6	0.1%	38.6	14.3
5	36	16.3	0.1%	42	41.0	0.1%	47.2	14.9
6	40	21.6	0.1%	48	49.1	0.3%	58.9	17.4
7	50	31.8	0.5%	62	67.6	0.7%	84.6	27.3
8	69	64.0	0.7%	98	110.3	1.4%	136.6	36.9
9	112	116.5	1.9%	185	187.2	4.0%	258.2	84.7
10 (high)	441	369.0	10.4%	702	588.8	17.9%	629.2	487.7
All	85	33.2	1.3%	109	62.2	2.0%	76.0	26.1
<i>Panel B: by the absolute fee at origination</i>								
< 25	21	18.6	0.1%	21	52.0	0.2%	69.6	21.6
25–100	38	22.3	0.2%	44	40.1	0.4%	46.0	22.8
100–300	172	258.0	3.3%	170	240.2	4.0%	289.6	164.4
300–1000	539	692.4	15.5%	529	762.5	13.8%	772.4	662.0
≥ 1000	2536	755.6	60.9%	2908	1479.9	66.5%	1441.1	1879.8

Table IA.13: Monthly Aggregates and the Rate Environment

This table reports time-series regressions of monthly aggregates on interest-rate and volatility variables, one at a time, with Newey–West standard errors at six lags and t -statistics in parentheses. The row labeled “ R^2 , all five jointly” reports the R^2 when all five enter together. Regressors are standardized, so coefficients are per one standard deviation. Row labels give the units of each variable before standardization. The dependent variables are the cross-sectional 95th percentile of $W^{(126)}$ for all stocks and for the top unpredictability decile, the share of six-month terms containing a sourcing-failure day, and the mean and median price-minus-cost margin of Section 5.2. The charge and the failure rate load strongly on the rate environment. The price–cost margin, on which the paper’s conclusion rests, loads significantly on the next-six-month change in the funds rate, and its median also loads on the level of the funds rate, but the joint R^2 is only 0.06 for the mean and 0.09 for the median. The sample contains only two near-zero-rate episodes, 2009–2015 and 2020–2021, so the level loadings are identified from few independent episodes and should be interpreted with caution.

	Six-month charge (q_{95} of W)		Sourcing failure (%)		Price – cost	
	All	Top decile	All	Top decile	Mean	Median
Fed funds level (%)	16.2	168.9	0.5	6.7	5.0	8.4
	(2.75)	(5.59)	(2.11)	(7.44)	(1.42)	(3.06)
Δ Fed funds, next 6m (bp)	13.6	69.3	0.4	2.8	–13.1	–10.5
	(2.25)	(1.78)	(3.06)	(2.77)	(–3.80)	(–3.56)
$ \Delta$ Fed funds , next 6m (bp)	12.7	80.2	0.3	2.8	–3.2	–0.6
	(2.38)	(2.23)	(1.34)	(2.72)	(–0.56)	(–0.13)
1y yield vol, trailing 60d (bp/day)	17.2	175.3	0.4	7.0	–1.0	3.3
	(2.63)	(5.72)	(1.71)	(7.66)	(–0.19)	(0.75)
VIX	–3.1	–9.0	–0.2	–0.1	–0.4	–3.5
	(–0.78)	(–0.19)	(–1.35)	(–0.08)	(–0.05)	(–0.53)
R^2 , all five jointly	0.55	0.40	0.46	0.75	0.06	0.09
Months	204	167	204	167	167	167

Table IA.14: Reweighting the Option Sample to the Full Panel

This table reweights the option-matched sample to the full lending panel. The six-month option-implied fee covers 26.5% of lending-panel stock-months. A logit of coverage on the fee level (in logs and levels), fee unpredictability, utilization, the cost-of-borrow score, size, return volatility, and loan tenure, each standardized within month, with month dummies, yields a coverage propensity. Covered observations are weighted by the inverse of that propensity, which reproduces the full panel under selection on observables. Weights are capped at their 99th percentile; the cap binds for 1.0% of covered observations and mechanically damps the reweighting, so part of the stability shown is by construction. The table reports mean six-month margins (price minus spot, price minus expected carry, and price minus full cost) in annualized basis points, unweighted and reweighted, overall and by decile of the spot fee formed on the full panel within month. It also reports the pass-through slope estimated from unweighted and inverse-propensity-weighted decile cell means. Price minus full cost here is built from expected carry in the stock's fee-decile cell plus the risk loading net of that carry and the sourcing leg, so its unweighted level differs slightly from Table 5, which adds the charge directly to the spot fee.

	Price – spot		Price – E[carry]		Price – full cost	
	Unwtd	IPW	Unwtd	IPW	Unwtd	IPW
All covered stock-months	–6.0	–7.4	–3.3	–5.9	–68.9	–63.1
Decile 1 (cheap)	12.8	11.0	3.4	1.5	–36.4	–34.1
Decile 2	8.9	9.2	2.4	2.6	–32.7	–29.0
Decile 3	3.9	2.6	–1.1	–2.5	–37.7	–35.5
Decile 4	3.0	2.7	–1.3	–1.6	–39.5	–35.6
Decile 5	3.7	2.1	0.2	–1.3	–41.5	–38.2
Decile 6	2.1	–1.1	–0.5	–3.7	–46.8	–45.2
Decile 7	–0.5	–5.6	–2.8	–7.9	–55.2	–54.6
Decile 8	0.1	–7.6	4.1	–1.8	–64.4	–63.3
Decile 9	–2.3	–14.9	9.5	–2.3	–101.6	–106.7
Decile 10 (expensive)	–131.4	–144.9	–62.8	–77.8	–327.1	–353.1
Pass-through slope: 0.750 unweighted, 0.674 reweighted						

Table IA.15: American Early Exercise in the Option-Implied Fee

This table reports the American early-exercise correction to the option-implied fee. Put–call parity is a European relation, while single-name equity options are American. For each matched pair, I price the American and European call and put on a CRR binomial tree under a continuous yield equal to the current estimate of the borrow fee, subtract the resulting early-exercise premia from the observed mid quotes, re-apply parity, and iterate toward a fixed point in the fee over four passes. Volatilities are OptionMetrics’ `civ` and `piv`. The premium columns are in percent of the spot price. “European” is the implied fee as extracted in Section 5.1, and “de-Americanized” is the corrected series. Each series is anchored separately on the within-month cheapest fee decile at its own maturity, so the reported difference is net of any common level effect. Δ margin is the change in price minus expected carry, computed on contracts with a nonmissing expected carry, so it can differ slightly from the difference of the two fee columns. The two premium columns contain the identifying pattern: the call premium rises steeply from the ninth to the tenth fee decile, because a large borrow fee is economically a large dividend yield, while the put premium, which is driven by the interest rate, does not. Figures are at the contract level, so stock-months with many matched contracts receive more weight than in the stock-month sample of Table 5. One caveat applies. OptionMetrics computes its implied volatilities under its own American model and its own borrow assumption, so the volatility input is not independent of the quantity being corrected. Scaling the volatility input by 0.75 or 1.25 moves the top-decile correction by under three basis points, against a stock-month correction of 44 basis points (63 basis points at the contract level shown here). This effect is small relative to the correction but not zero, and a fully independent volatility would require re-solving the whole surface under the borrow costs I measure.

	N	Early-exercise premium		Implied fee and margin (bps)		
		Call (% of S)	Put (% of S)	European	De-Americanized	Δ margin
<i>Panel A: Three months</i>						
All contracts	126,492	0.035	0.050	100.6	109.4	+8.6
Fee decile 9	11,093	0.025	0.050	63.4	72.4	+8.9
Fee decile 10 (expensive)	11,159	0.186	0.038	611.1	685.6	+74.5
<i>Panel B: Six months</i>						
All contracts	183,542	0.049	0.092	95.3	101.7	+6.4
Fee decile 9	16,783	0.034	0.094	69.3	73.4	+4.1
Fee decile 10 (expensive)	16,996	0.263	0.069	478.1	541.1	+63.1
<i>Panel C: Twelve months</i>						
All contracts	82,771	0.064	0.178	103.7	103.0	−1.1
Fee decile 9	7,343	0.041	0.162	81.9	79.7	−2.2
Fee decile 10 (expensive)	7,667	0.378	0.111	449.2	505.1	+56.0

Table IA.16: The Margin at Every Corner of the Paper’s Own Bounds

This table evaluates the price-minus-cost margin at three, six, and twelve months when the quoting quantile, the level anchor, and the sourcing leg are varied jointly. The quoting quantile moves between the median and the 95th percentile (Table 6). The level anchor moves across its ± 20 basis point reach (Panel A of Table IA.9). The sourcing leg moves between zero and twice the estimated buy-in cost. Zero is a lower bound by construction, and twice the buy-in cost is an upper bound that allows for the market impact the estimate excludes. Nothing is re-estimated, since every input is a number reported elsewhere in the paper. The first row is the 95th-percentile margin from the finer quantile grid behind the corners, which differs from the 95th-percentile rows of Tables 5 and 6 by at most 0.3 basis points. The corner least favorable to the paper’s conclusion combines the median quantile, the anchor at its high end, and no sourcing leg. The corner most favorable combines the 95th percentile, the anchor at its low end, and a doubled buy-in cost. The row “Signed across the box” reports whether the margin is negative at both corners. The final row repeats the break-even quantile of Table 6, which does not depend on the quoting convention but does depend on the anchor.

	Three	Six	Twelve
	months		
Published margin (95th percentile)	−49.3	−72.2	−99.8
<i>Least favorable corner</i>			
Median quantile, anchor +20, no sourcing leg	14.3	13.6	20.1
<i>Most favorable corner</i>			
95th percentile, anchor −20, buy-in doubled	−70.5	−93.9	−122.2
Signed across the box	No	No	No
Break-even quantile q^*	24.1	26.9	40.5

Table IA.17: Reconciling the Sign of the Margin with the Literature

This table compares my price-minus-spot margin with that of Muravyev et al. (2022), who report a positive option-implied borrowing premium. When the implied fee is extracted by European parity without the early-exercise correction of Section 5.1.1, my price-minus-spot margin is negative, and the table locates the source of that difference on the uncorrected series. On the corrected series, which the paper reports throughout, price minus spot is close to zero and well within the reach of the level calibration (Section 5.2). The first four rows, S0 to S3, build the implied fee in the manner of Muravyev et al. (2022), including dividend payers with a present-value dividend adjustment and maturities from twenty days, and then apply my screens cumulatively; a screen that caused the sign would flip it. None does, since the baseline is already negative and every screen shrinks the gap. The next three rows split the baseline by the two screens, reporting short-dated contracts, dividend payers, and non-payers separately, which shows the margins on the observations each screen removes or keeps. The last four rows split the no-dividend sample by rate regime. The early-exercise value of American puts is negligible when short rates are near zero, so if unmodeled put exercise drove the sign, the near-zero-rate years would show the smallest margins; instead, the two near-zero-rate periods, 2009–2015 and 2020–2021, show margins at least as large in magnitude as 2016–2019 and 2023–2024. Two sources can account for the level difference on the uncorrected series. The first is the interest-rate benchmark, which the calibration of Section 5 absorbs: I impose the OptionMetrics zero curve, Muravyev et al. (2022) estimate rates from option prices, and option-embedded rates exceed Treasury-based curves (van Binsbergen et al., 2022). The second is the American early-exercise premium, which Section 5.1.1 corrects. The +14.4 basis point shift in row S3 is the cheapest-decile calibration re-estimated on this table’s broader sample, anchored to the decile’s overnight fee; it is not comparable to the maturity-specific shifts used in the main analysis, which are anchored to realized term cost on the screened sample. Margins are in basis points, and the correlation is between the implied and actual fees.

Sample	N	Mean	Median	Corr.
S0 MPP-like (div-adj, 20–400d)	1,162,266	−62.2	−54.5	0.59
S1 + ≥ 70 days	704,901	−55.0	−48.1	0.64
S2 + no dividend in life	358,041	−48.0	−37.5	0.68
S3 + calibration (shift +14.4 bp)	358,041	−33.7	−23.1	0.68
S0 restricted to < 70 days only	457,365	−73.3	−77.6	0.59
S0 dividend payers only	449,359	−79.1	−63.5	0.29
S0 non-payers only, all maturities	712,907	−51.5	−46.1	0.61
2009–2015 (r near zero)	116,518	−46.0	−37.4	0.61
2016–2019	85,235	−45.4	−37.4	0.71
2020–2021 (r near zero)	56,194	−87.8	−67.6	0.64
2023–2024 (r $\sim 5\%$)	58,392	−30.8	−19.8	0.77

Table IA.18: The Pass-Through Slope under Measurement Error

This table examines whether measurement error attenuates the pass-through slope. SAF is an average across outstanding loans and is therefore a noisy proxy for the marginal borrow rate, and classical measurement error in the regressor would attenuate the pass-through slope toward zero. Each row estimates the sensitivity of the six-month option-implied fee to the actual fee on the option-covered sample, winsorized at the 1st and 99th percentiles within month. The OLS row includes month fixed effects, with standard errors clustered two-way by stock and month. The IV rows are two-stage least squares with month dummies and the same clustering, and the first-stage F statistics come from the corresponding first stage. The one-month-lagged fee is correlated with the persistent component of the true fee but not with current-month recording noise. Because lending-market tightness moves the implied fee conditional on the fee (Table 7, Panel A), the lagged utilization and cost-of-borrow instruments do not satisfy a strict exclusion restriction, so the row that uses them is an upper-bound diagnostic and not a clean instrument. Because fees mean-revert, the instruments also remove transitory true variation, so the instrumented slopes are comparable to the reversion-adjusted estimate in Panel C of Table 7. The grouped rows apply the paper's cell-mean estimator with deciles formed on the current fee and on the lagged fee. The reverse regression gives the errors-in-variables upper bound $1/b$, which is loose because the implied fee is itself noisy.

	Slope	t	First-stage F	N
OLS, observation level, month FE	0.649	17.7		120,286
IV: one-month-lagged fee	0.710	34.7	546	118,710
IV: lagged utilization and cost-of-borrow score	0.757	46.3	1,117	115,876
Cell means, deciles of the current fee	0.749			10
Cell means, deciles of the lagged fee	0.780			10
Reverse-regression bound ($1/b$)	2.112			120,286

Table IA.19: What the Price–Cost Gap Forecasts

This table reports forward regressions of each outcome on the six-month price-minus-expected-carry margin, standardized within month, so coefficients are per cross-sectional standard deviation. The margin is the *ex-ante* version: its expected-carry leg uses only originations whose full path had been completed by the quoting month, so nothing in the predictor is unobservable at the time of the quote. The memo row repeats the first test with the in-sample margin, whose carry leg uses the whole sample, and shows that the ex-ante restriction changes the estimate very little. Controls are the log fee, fee unpredictability, utilization, log market capitalization, and return volatility. All specifications include month fixed effects, and standard errors are clustered two-way by stock and month. The outcomes are the realized commitment wedge over the term in basis points, the share of days inside the term on which sourcing fails in percentage points, and six-month forward growth in lendable supply and in the number of active brokers in percent. The two quantity outcomes are significant only with controls, and Section 5.3.1 interprets them accordingly.

Outcome over the next six months	No controls		With controls		N
	b	t	b	t	
Realized commitment wedge $W^{(126)}$ (bps)	29.8	7.76	16.1	5.87	97,979
Share of days sourcing fails (pp)	−0.67	−3.34	−0.26	−2.46	97,979
Growth in lendable supply (%)	0.57	1.35	1.36	6.13	87,928
Growth in active brokers (%)	0.07	0.43	0.32	2.07	87,928
<i>memo: $W^{(126)}$ on the in-sample margin</i>			16.3	5.87	97,979

Table IA.21: Substitution When the Options Route Reprices

This table examines substitution between the lending market and the options market around episodes in which one route reprices. In Panel A, an implied-fee shock is a month in which the six-month option-implied borrow rate rises by more than 50 basis points while the actual fee moves by less than 25 basis points in absolute value, and an actual-fee spike is a month in which the borrow fee more than doubles and rises by more than 50 basis points. Panel A regresses the forward three- and six-month change in the option share on both indicators, the level of the share, size, return volatility, and turnover, with month fixed effects and two-way clustered standard errors. Substitution predicts opposite signs on the two indicators, whereas common bearish sentiment predicts the same sign. Panel B compares financials (SIC 6000–6999, a proxy for the ban list) with all other stocks around the September 2008 short-sale ban; the difference-in-differences compares the ban months (September–October) with June–August within stock groups. Panel C reports monthly aggregates through the COVID volatility spike, including the borrow fee, the option-implied fee, short interest, option short exposure, and the option share.

Panel A: implied-fee shocks and actual-fee spikes

	3-month change		6-month change	
	<i>b</i>	<i>t</i>	<i>b</i>	<i>t</i>
Implied-fee shock	−0.0002	−0.12	−0.0039	−1.62
Actual-fee spike	0.0088	3.42	0.0118	3.14

Panel B: the September 2008 ban (banned financials vs. others)

DiD, option short exposure	<i>b</i> = 0.0010, <i>t</i> = 1.01, <i>N</i> = 8,042
DiD, option share	<i>b</i> = 0.0101, <i>t</i> = 1.27, <i>N</i> = 7,727
DiD, short interest	<i>b</i> = −0.0035, <i>t</i> = −1.37, <i>N</i> = 7,767

Panel C: COVID (options route reprices without a ban)

	2019m12	2020m1	2020m2	2020m3	2020m4	2020m5
Borrow fee (bp)	68.3	86.1	83.4	73.9	77.9	81.0
Option-implied fee (bp)	81.4	100.9	253.5	137.9	40.7	122.3
Short interest	0.0407	0.0416	0.0445	0.0456	0.0373	0.0380
Option short exposure	0.0113	0.0100	0.0111	0.0094	0.0111	0.0119
Option share	0.2168	0.2194	0.2122	0.1806	0.2050	0.2146

Table IA.22: Selection: The Wedge on the Loans a Borrower Would Choose

This table reports the mean and 95th percentile of the commitment wedge $W^{(H)}$, in basis points, by quintile of demand growth. Panel A sorts on the change in utilization over the one month (for one-month terms) or six months (for six-month terms) before origination, which a provider observes and can quote against. Panel B sorts on growth over the life of the term itself, which approximates what a borrower who knew her own intentions could condition on. Panel B uses information unavailable at origination by construction, so the implied selection cost is an upper bound on the adverse-selection channel and not an estimate of it. The unconditional row pools every stock-month with a wedge at that horizon, including those without a demand-growth measure, so it is not the pooled value of either panel's quintiles. At six months, the pooled mean of the quintile rows is 1.0 basis points in Panel A and -3.6 basis points in Panel B. Section 6.2 discusses the double sort of trailing demand growth against fee unpredictability.

Demand quintile	One month (21d)			Six months (126d)		
	$E[W]$	95th pct	N	$E[W]$	95th pct	N
<i>Panel A: trailing demand growth (observable at origination)</i>						
1 (falling)	-4.59	27.5	78,110	-4.66	38.7	63,000
2	-0.66	22.4	73,997	0.60	33.3	58,715
3	0.17	25.5	69,702	1.23	35.1	56,236
4	-0.22	21.9	74,877	1.63	31.8	61,059
5 (rising)	0.89	32.6	81,192	5.76	64.8	69,043
<i>Panel B: forward demand growth over the term (borrower's information)</i>						
1 (falling)	-7.18	26.3	78,954	-15.64	38.7	67,641
2	-0.89	22.5	74,664	-0.09	32.6	61,354
3	-0.26	25.4	69,799	0.76	35.3	57,353
4	-0.90	20.8	74,860	-1.00	29.0	61,717
5 (rising)	0.31	32.9	80,502	-0.77	53.7	68,135
Unconditional	-0.74	31.6	394,340	-4.84	54.4	335,991

Table IA.23: Do the Costs of Term Protection Show Up in Prices?

This table examines whether fee unpredictability shows up in returns and in the buildup of short interest. Panel A reports monthly cross-sectional (Fama–MacBeth) regressions of forward returns on lagged standardized fee unpredictability, with Newey–West standard errors on the resulting coefficient series; the control set adds the standardized spot fee to the usual cross-sectional controls. Panel B sorts stocks into fee quintiles each month and, within each fee quintile, into unpredictability quintiles, and it reports the return on the most unpredictable quintile less the return on the least unpredictable quintile. The sort is run at every horizon, and the t -statistics beneath the spreads are Newey–West statistics on the monthly spread series. Panel C regresses the forward change in the short interest ratio on the same predictor and the Panel A controls with month fixed effects, clustering by stock and by month. Panel C uses the raw short interest field, which is scaled down by a factor of one thousand (see the note to Table IA.20), so a coefficient of 10^{-6} corresponds to 10 basis points of shares outstanding. A *negative* coefficient would indicate that short interest builds more slowly where the borrowing cost cannot be fixed, which is the dynamic-constraint prediction. Table IA.24 bounds how much weight the surviving cells of Panel B can carry. Returns and spreads are in percentage points. The table reports the return and short-interest specifications discussed in Section 6.4, including those that find nothing; the remaining tests are reported in the text.

	Unpredictability alone		With fee level and controls		
Horizon	b	t	b	t	
<i>Panel A: Fama–MacBeth, forward returns on lagged unpredictability (percentage points per s.d.)</i>					
One month	−0.10	−2.61	0.02	0.26	
Three months	−0.26	−2.54	0.02	0.13	
Six months	−0.47	−2.70	−0.05	−0.17	
Twelve months	−0.59	−2.22	−0.12	−0.37	
memo: fee level alone, six months	−0.66	−3.28			
<i>Panel B: return spread, most less least unpredictable quintile, within fee quintile (percentage points)</i>					
Horizon	Q1 (cheap)	Q2	Q3	Q4	Q5 (expensive)
One month	−0.05	−0.16	−0.07	−0.08	−0.72
Three months	−0.36	−0.42	−0.04	0.09	−1.83
Six months	−0.95	−0.95	0.15	0.12	−3.86
Twelve months	−2.04	−1.25	0.46	0.33	−4.78
<i>t-statistics, Newey–West</i>					
One month	−0.33	−0.96	−0.36	−0.37	−3.80
Three months	−0.85	−0.80	−0.06	0.17	−3.73
Six months	−1.18	−0.95	0.14	0.15	−4.07
Twelve months	−1.48	−0.80	0.26	0.23	−3.12
<i>Panel C: forward change in the short interest ratio on lagged unpredictability ($b \times 10^6$)</i>					
Horizon	b	t	N		
Three months	0.39	4.11	298,400		
Six months	0.74	4.77	287,014		
Twelve months	1.24	4.79	265,789		

Table IA.24: How Much Weight the Return Result Carries

This table reports tests that bound the cells of Panel B of Table IA.23 that are distinguishable from zero, namely the spreads between the most and least unpredictable stocks within the most expensive fee quintile. The *placebo* block runs the identical sort on total fee volatility in place of unpredictability, which asks whether returns separate unforecastability from movement in the way Section 4.2.3 shows the costs do. Both sorts are formed on the same stock-months, because the two measures use trailing windows of different lengths and running each on its own maximal sample would overstate the difference between them. The *orthogonalized* block sorts on the residual from a monthly regression of unpredictability on total volatility, which captures the incremental content of unforecastability. The *multiple testing* row takes the largest $|t|$ across all twenty cells of Panel B and reports the share of moving-block bootstrap replicates, under a null imposed by demeaning each monthly spread series, whose own largest $|t|$ exceeds it. Blocks are six months long to respect the overlap in forward returns. This test is computed on the 160 months for which all twenty series are present, so the observed statistic is 3.83 instead of the 4.07 in Panel B. The *construction* block re-forms the top-quintile spread under alternative portfolio conventions: value weighting by market capitalization at formation; value weighting after also excluding stocks below the NYSE 20th percentile of size; the same with the fee quintile taken at an NYSE breakpoint; and equal weighting with the borrow fee netted out of the return. All t -statistics are Newey–West with six lags, and Section 6.4 discusses the table.

<i>Placebo: the same sort on total fee volatility, sample-matched (percentage points, top fee quintile)</i>				
Horizon	Unpredictability	t	Total volatility	t
One month	−0.53	−2.66	−0.68	−2.91
Three months	−1.27	−2.62	−1.55	−2.36
Six months	−2.77	−3.21	−2.84	−2.42
Twelve months	−3.99	−2.82	−3.59	−1.90
<i>Unpredictability orthogonalized against total volatility, same sample</i>				
One month	−0.22	−1.26		
Three months	−0.73	−1.45		
Six months	−1.47	−1.52		
Twelve months	−4.19	−2.41		
<i>Multiple testing across all twenty cells of Panel B</i>				
Largest $ t $, and its bootstrapped p	3.83	0.025	2,000 block-resamples	
<i>Portfolio construction, top fee quintile (percentage points)</i>				
Construction	6m b	t	12m b	t
Equal-weighted (Panel B)	−3.86	−4.07	−4.78	−3.12
Value-weighted	−1.28	−0.73	1.07	0.34
Value-weighted, no microcaps	0.45	0.20	3.74	0.90
Value-weighted, no microcaps, NYSE breakpoint	0.79	0.38	1.15	0.29
Equal-weighted, net of the borrow fee	−6.99	−5.79	−10.52	−5.51

References

- Muravyev, D., Pearson, N. D., and Pollet, J. M. (2022). Is there a risk premium in the stock lending market? Evidence from equity options. *Journal of Finance*, 77(3):1787–1828.
- van Binsbergen, J. H., Diamond, W. F., and Grotteria, M. (2022). Risk-free interest rates. *Journal of Financial Economics*, 143(1):1–29.