This is the draft version of the paper.

This version of the publication may differ from the final published version.

Contents

CrossMark

Dynamic debt runs and financial fragility: Evidence from the 2007 ABCP crisis ☆

Enrique Schroth [a], Gustavo A. Suarez [b], Lucian A. Taylor [c],[*]

[a] Cass Business School, City University, London, United Kingdom
[b] Federal Reserve Board, United States
[c] Wharton School, University of Pennsylvania, United States

ARTICLE INFO

ABSTRACT

We use the 2007 asset-backed commercial paper (ABCP) crisis as a laboratory to study the determinants of debt runs. Our model features dilution risk: maturing short-term lenders demand higher yields in compensation for being diluted by future lenders, making runs more likely. The model explains the observed tenfold increase in yield spreads leading to runs and the positive relation between yield spreads and future runs. Results from structural estimation show that runs are very sensitive to leverage, asset values, and asset liquidity, but less sensitive to the degree of maturity mismatch, the strength of guarantees, and asset volatility.

1. Introduction

Debt runs played a central role in the financial crisis of 2007–2008. Investors ran on asset-backed commercial paper (ABCP) starting in August 2007, on repo starting in September 2007, and on money market mutual funds in September 2008. Investors also ran on some large banks such as Northern Rock (September 2007) and Bear Stearns (March 2008).[1]

These events have reignited the debate about what causes runs and how we can prevent them. We contribute to this debate by measuring the sensitivity of runs to several contributing factors, including maturity mismatch, leverage, asset volatility and liquidity, and the strength of guarantees. The results help answer four questions that are vital to policy makers, regulators, bankers, and investors: How fragile are financial intermediaries? How can we design financial intermediaries ex ante to control the risk of future runs? What are the warning signs that a run is imminent? Finally, which interventions best prevent runs ex post once conditions have started deteriorating?

We address these questions by estimating a structural model of debt runs using data from the 2007 ABCP crisis. ABCP issuers, commonly referred to as conduits, are off-balance sheet investment vehicles that banks structure to invest in pools of medium- and long-term assets such as trade receivables and mortgage-backed securities (MBS).[2] A conduit finances these investments by issuing short-term ABCP to dispersed creditors and rolling it over until the conduit chooses to stop investing. The bank sponsoring the conduit provides some form of guarantee in the event that the conduit can no longer roll over its debt.

The amount of ABCP outstanding in the U.S. contracted by $370 billion (roughly one-third) between August and December of 2007. Several authors have interpreted this event as a run on debt.[3] In a debt run, creditors refuse to roll over their debt if they fear that other creditors will not roll over, in some cases even if the borrower is solvent. In the case of ABCP, roughly 40% of conduits had stopped rolling over maturing debt by the end of 2007.

ABCP provides a useful laboratory to study financial fragility for four reasons. First, since ABCP conduits perform maturity transformation, they are representative of many other financial intermediaries. Second, the simple balance sheet and operating structure of ABCP conduits lend themselves to modeling. Third, we have detailed data on the yield, maturity, size, and issuer's identity for all U.S. ABCP transactions in 2007. Because yields adjust at each

maturity date, their time series measures the conduit's health continuously and can potentially be an important lead indicator of runs. Finally, as Krishnamurthy, Nagel, and Orlov (2014) argue, the ABCP crisis was important in itself:

> The contraction in both repo and ABCP are consistent with the views of many commentators that a contraction in the short-term debt of shadow banks played an important role in the collapse of the shadow banking sector. However, it is important to note that the ABCP plays a more important role than repo in this regard.

In fact, runs on ABCP could have had a broad effect on financial intermediation through two channels. First, runs impaired ABCP conduits' ability to fund assets such as trade receivables or student loan receivables. Second, the runs on ABCP conduits forced their sponsoring banks to take troubled assets like mortgage securities back onto their own books, which impaired lending to nonfinancial firms and ultimately harmed economic activity (Irani, 2011).

Our model of ABCP conduits is based on He and Xiong (2012a). A conduit finances a long-term asset using short-term, dispersed debt with overlapping maturities. Creditors track the asset's value and optimally run as soon as the conduit's leverage crosses above an endogenous threshold. A creditor's decision to run depends on changing expectations that other creditors will run. We extend He and Xiong's (2012a) model so that debt yields are not fixed but instead vary endogenously over time, so as to make lenders indifferent between rolling over or not. This extension is necessary: we show empirically that yields on ABCP forecast runs, and yields increase exponentially leading up to runs. To have any chance of fitting these data, the model must make predictions about the time series of yields.

The model's parameters include the debt's maturity; the perceived strength of the sponsor's guarantee; and the asset's volatility, maturity, and liquidation discount in default. We observe some of these parameters directly in the data, and we estimate others using the simulated method of moments (SMM).

We find three main results. First, we show that runs are very sensitive to leverage and asset liquidity, but are less sensitive to the degree of maturity mismatch, asset volatility, and perceived guarantee strength. We measure these sensitivities by comparing simulated run probabilities between our estimated model and a counterfactual model with altered parameter values. We measure these sensitivities in both the early and late stages of a simulated crisis. In the late stages, increasing the asset's liquidation recovery rate by 1% (from 92.0% to 92.9%), while holding all else equal, lowers the probability of a run within three months from 70% to 39%. Decreasing the conduit's leverage by 1% (from 91.4% to 90.4%) has an almost identical impact. In contrast, reducing the run probability by the same amount would require either reducing asset volatility by 40%, increasing average debt maturity by 190%, reducing average asset maturity by 98%, or increasing the guarantee's expected life span by 413%.

---

[1] Brunnermeier (2009) and Krishnamurthy (2009) summarize the events of 2007–2008. We discuss the literature on ABCP below. Gorton and Metrick (2012) and Krishnamurthy, Nagel, and Orlov (2014) empirically investigate the run on repo. Martin, Skeie, and Von Thadden (2012) provide a model of repo runs. Kacperczyk and Schnabl (2013) examine the run on money market funds.

[2] One prevalent view is that ABCP conduits were essentially a way for sponsoring banks to take on systemic risk beyond regulations, without transferring the risk to ABCP investors. See Acharya and Richardson (2009), Acharya and Schnabl (2010), Acharya, Schnabl, and Suarez (2013), Brunnermeier (2009), and Shin (2009).

[3] See, for instance, Covitz, Liang, and Suarez (2013), Acharya, Schnabl, and Suarez (2013), and Gorton and Metrick (2012).

These results shed light on how regulators and bankers can manage the risk of runs, both when forming new conduits and during a crisis. For example, crisis management policies with modest effects on asset liquidity (e.g., purchasing distressed assets) or conduit leverage (e.g., injecting equity) can have substantial effects on the likelihood of runs. High ABCP yields, which result from deteriorating fundamentals, are a warning sign that a run is imminent. The model provides a quantitative mapping between these warning signs and the likelihood of runs. Of course, we do not address the feasibility or the cost of policy interventions, nor do we analyze how changing one fundamental (e.g., liquidity) could affect another (e.g., debt maturity).

The second main result is that the model can fit several features of the 2007 ABCP crisis. The model explains 73% of the sharp decline in total ABCP outstanding in the second half of 2007. For conduits offering weak guarantees to investors ('Structured investment vehicle' (SIV) and 'Extendible notes'), the model comes quite close to fitting the magnitude and timing of the dramatic run-up in yields before runs, the overall level of ABCP yield volatility and its relation to yield levels, and the relatively high likelihood that conduits recover from a run. In both simulated and actual data, the current yield level helps forecast whether a run will occur. The model's main shortcoming is that, for conduits offering strong guarantees ('Full credit' or 'Full liquidity'), it overpredicts runs when yields are high.

Our third result is theoretical. We show that introducing time-varying yields into the model typically makes runs more likely, relative to He and Xiong's (2012a) model with constant, exogenous yields. Using He and Xiong's (2012a) calibrated parameter values, we find that runs are 1.3–51 times more likely in our model than theirs. The reason, as He and Xiong (2012a) conjecture, is that the conduit must offer high yields to induce rollover when conditions deteriorate. These high yields dilute all outstanding debt that matures later. Creditors preemptively demand higher yields in compensation for the risk of future dilution. These higher yields in turn make leverage build up faster, which makes runs more likely. This new risk, which we call 'dilution risk,' can be an important driver of yields and runs.

Several papers measure the determinants of runs using a reduced-form approach. Covitz, Liang, and Suarez (2013) show that runs on ABCP conduits are negatively related to the strength of their guarantees. Calomiris and Mason (1997, 2003) show that bank runs during the Great Depression are correlated with measures of bank solvency and shocks to the aggregate, regional, and local economies. Using data on an Indian bank, Iyer and Puri (2012) show that runs are positively related to weaker deposit insurance, a shorter or shallower relationship with the bank, and runs by one's peers. Chen, Goldstein, and Jiang (2010) provide evidence of strategic complementarities in mutual funds.

We depart from the existing empirical literature by taking a structural estimation approach. The structural approach complements the reduced-form approach by overcoming certain data limitations and by imposing different identifying assumptions. The reduced-form approach requires data on the determinants of runs, many of which are difficult to obtain in the ABCP setting. For example,

data on conduit leverage and asset holdings are not publicly available.[4] We overcome this limitation by structurally estimating several run determinants. The reduced-form approach also requires a data set with sufficient variation in the determinants of runs. Finding variation is potentially a challenge in the ABCP setting because conduits resemble each other on many dimensions. The structural approach requires no heterogeneity in these determinants, as we use counterfactual analysis to measure the sensitivity of runs to their various determinants. Both approaches impose strong identifying assumptions. The reduced-form approach assumes we have exogenous variation in the determinants of runs, which is difficult to find. The structural approach assumes that the model is true. Therefore, our exercise is not meant to identify coordination failures over any alternative mechanism for the sharp decline in ABCP. However, the structural approach allows us to show that a model of coordination failures can quantitatively and jointly fit several features of the data. This paper therefore takes a step toward providing a quantitative model of financial fragility, which is crucial for guiding the management and regulation of financial intermediaries.

The paper is structured as follows. Section 2 describes the model's assumptions and solution. Section 3 discusses its predictions regarding dilution risk and the likelihood of runs. Section 4 describes the data, identification, and SMM estimation. Section 5 assesses model fit and describes our parameter estimates. Section 6 uses the estimated model to explore the determinants of runs, Section 7 discusses policy implications, and Section 8 concludes.

2. The model

We extend the model of He and Xiong (2012a) by allowing yields on short-term debt to adjust over time in response to changes in fundamentals. All assumptions below are shared with He and Xiong (2012a) unless otherwise noted. Appendix A describes the solution.

The model includes several features of ABCP conduits. The conduit finances a long-term asset using short-term, dispersed debt with overlapping maturities. The conduit must roll over this debt several times before the conduit ends, so the conduit faces rollover risk. The conduit's sponsor provides imperfect credit support if the conduit cannot roll over its paper.

2.1. Assumptions

2.1.1. Asset

At time zero, an ABCP conduit purchases a long-horizon asset. This asset represents the portfolio of assets a conduit typically buys. For the overall ABCP industry in 2007, the largest assets classes were trade receivables (14%), credit cards (12%), auto loans (11%), 'securities' (11%), commercial

---

[4] Even if asset holdings were known, measuring asset liquidity is difficult, especially for ABCP asset classes like trade receivables. Price data on trade receivables and other important asset classes are not available. Further, while we have data on credit guarantees' types, we cannot measure their perceived strength.

loans (10%), and mortgage-related assets (9%).[5] The conduit reinvests any interim cash flows from the asset. For example, the conduit could buy new trade receivables using the payouts from maturing receivables. The conduit therefore makes no net interim payouts to investors.[6] The asset produces a single net payout when the conduit matures, meaning the conduit winds down operations. The conduit matures randomly and independently at a time τ_ϕ that arrives according to a Poisson process with intensity ϕ, so the conduit's expected time until maturity is always $1/\phi$. At maturity, the asset produces a payout y_{τ_ϕ}, where y follows a geometric Brownian motion with drift μ and volatility σ:

$$\frac{dy_t}{y_t} = \mu\, dt + \sigma\, dZ_t. \tag{1}$$

Agents observe y_t at all times. All agents in the economy are risk neutral and have discount rate ρ, so the asset's value at time t is

$$F(y_t) \equiv E_t\left[e^{-(\tau_\phi - t)\rho} y_{\tau_\phi}\right] = \frac{\phi}{\rho + \phi - \mu} y_t. \tag{2}$$

2.1.2. Debt financing

The conduit finances the asset by initially borrowing \$1 from a continuum of short-term creditors. Consistent with industry practice, the conduit also issues equity to its sponsor. The conduit's debt is zero-coupon and has endogenous face value R_t per dollar loaned at date t. In contrast, debt contracts in He and Xiong (2012a) have face value normalized to one and offer exogenous interest rate r. Each debt contract in our model matures randomly and independently with probability $\delta\, dt$ in the interval $[t, t+dt]$, implying that a debt contract's average remaining maturity always equals $1/\delta$. This modeling device, which follows Calvo (1983), Blanchard (1985), and Leland (1998), reflects that ABCP conduits deliberately spread their debt maturities over time to reduce funding liquidity risk. These assumptions capture an important feature of the ABCP market, which is that before a given lender's debt matures, other lenders' debt will mature and potentially fail to roll over. Our assumptions imply that the conduit rolls over a fraction $\delta\, dt$ of its debt every instant, and the total face value of debt, D_t, fluctuates over time according to

$$dD_t = \delta D_t (R_t - 1)\, dt. \tag{3}$$

2.1.3. Runs, liquidation, and the sponsor's guarantee

As payment for a maturing loan, lenders accept a new loan with a potentially different face value. If lenders choose not to roll over, we say that they run. We assume lenders roll over if they are indifferent between rolling

over and running. If lenders run and the conduit cannot raise funds to pay off maturing lenders, then the conduit defaults. In default, the conduit sells the asset at a fraction α of its fair market price, which yields $L(y_t) \equiv \alpha F(y_t)$.

Parameter α measures the asset's liquidity in the run state.[7] Equivalently, α is the asset's recovery rate in default. Consistent with industry practice, the conduit distributes bankruptcy proceeds $L(y_t)$ to outstanding creditors pro rata, i.e., in proportion to their face value.

An ABCP conduit's sponsor provides a guarantee, which is typically a line of credit the conduit can use if it is unable to issue new paper. Section 4.1 describes the four types of guarantee in use in 2007. We follow He and Xiong (2012a) by modeling guarantees as an imperfect credit line from the sponsor. If the conduit experiences a run, it pays off maturing paper by borrowing from the sponsor at the prevailing rollover yield and maturity. The credit line therefore allows the conduit to potentially survive a run long enough for the conduit to recover and begin issuing paper again. We assume the credit line fails independently, causing default, each instant with probability $\theta\delta\, dt$. Once a run starts, the credit line is expected to last for $1/(\theta\delta)$ years, so conduits with higher values of θ have weaker guarantees.

2.2. ABCP pricing

We typically work with yield spreads, denoted r_t (in units of fraction per year), which we can compute from face values R_t (in units of dollars) using[8]

$$r_t = (R_t - 1) \times (\phi + \delta) - \rho. \tag{4}$$

Unlike in He and Xiong (2012a), debt is priced in a competitive market so that creditors exactly break even. This pricing assumption makes the debt we study here identical to the one in Leland and Toft (1996), Leland (1998), and the short-term class in He and Xiong (2012b). Specifically, the conduit sets its rollover yield spread r_t so that if creditors loan the conduit \$1 at time t, they receive a debt contract worth \$1. Intuitively, if times are bad, the conduit must issue paper at a high yield spread r_t to make creditors break even. If times are good, the new paper is almost risk free, and the new rollover yield will be close to the risk-free rate.

We assume the conduit cannot or will not issue debt with a yield spread above an exogenous cap, $\bar{r}$. This is an important assumption, which, as we show later, implies that creditors run exactly when the spread hits its cap. There are a several rationales for assuming yield spreads cannot go to infinity as conditions worsen.

---

[5] Data are from 'The ABC's of ABCP,' an unpublished document from Societe Generale. Reported portfolio weights are measured on August 31, 2007. The ambiguous 'securities' category could include mortgage-related securities.

[6] In He and Xiong (2012a), the asset pays a fixed dividend, which the conduit uses to pay a coupon on the bond. Our model assumes zero-coupon debt, since commercial paper is a discount security that pays face value at maturity with no interim coupons.

[7] Note that we do not assume that asset liquidity is constant. Instead, we assume that the creditors' expected asset liquidation value *in the run state*, i.e., if the whole asset is sold to pay off running creditors, is constant.

[8] The yield spread r_t is the interest rate (in excess of the risk-free rate, ρ) that delivers the same value as a zero-coupon bond with face value R_t, under the assumption that both bonds are paid back in full at $\tau = \min(\tau_\delta, \tau_\phi)$. Eq. (4) follows from the condition

$$E_t\left\{\int_t^\tau e^{-\rho(s-t)}(\rho + r_t)\, ds + e^{-\rho(\tau - t)}\right\} = E_t\{e^{-\rho(\tau - t)} R_t\}.$$

One rationale relates to institutional constraints. The main investors in ABCP are money market funds, which are required to invest mainly in assets with very high ratings (e.g., A-1 by Standard and Poor's (S&P) or P-1 by Moody's).[9] As an ABCP conduit's health declines and its rollover yields rise, eventually the conduit will lose its A-1/P-1 rating and its creditors will be unable to roll over its paper. Effectively, the conduit will be unable to roll over paper once yield spreads reach a certain level, i.e., a cap.

Credit rationing, as in Stiglitz and Weiss (1981), provides another rationale for capping yields. Once yield spreads exceed some level, only the conduits with very risky assets will be willing to roll over their paper. Because of this adverse selection problem, creditors will refuse to accept yield spreads above this cap.

In these first two rationales, it is creditors who walk away from the conduit in a run. There is another rationale in which the sponsor walks away. If conditions worsen enough, rollover yields become so high that the ABCP conduit's equity is very low. The sponsor then has little incentive to keep operating the conduit. As in Leland (1998), where the default boundary is chosen to maximize the value of equity at default, the sponsors can choose the run threshold indirectly via the yield cap.[10]

A final rationale is that without a cap on yields, we find no runs. We cannot prove this result generally since we lack closed-form solutions, but for empirically relevant parameter values we show that the predicted probability of a run goes to zero as $\bar{r}$ gets large (Fig. 8). Intuitively, a higher yield cap incentivizes rollover during worse conditions. As the yield approaches infinity, the conduit effectively dilutes the earlier lenders' claims to zero, thereby transferring the entire asset to the infinitesimal maturing lender. Of course, this lender also expects to lose this claim to the next period's maturing lender, but she also expects to get the chance to roll over again with the same positive probability as all other lenders. Without a yield cap, the model approaches a Ponzi game in which creditors alternate full claims on the asset indefinitely.[11]

Since we do not know which of these rationales for a yield cap binds in the data, we treat $\bar{r}$ as a parameter to estimate. This estimate measures the minimum of the rationales' various caps. Future research could establish which cap binds empirically. The answer would help explain whether runs were the result of conduits regarding short-term funding as too expensive or of creditors considering ABCP too risky.

2.3. Discussion

Our model is one of fundamental-driven panics (e.g., Goldstein, 2014). The only variable that changes exogenously over time is the asset's fundamental value. The model therefore assumes that runs are triggered by a drop in fundamental asset value rather than by, for example, an increase in asset volatility. Asset volatility and other model parameters do affect the likelihood of runs, as we discuss later, but the model assumes these parameters remain constant over time. The model still includes an element of 'panic,' in the sense that lenders run because they fear other lenders will do the same.

The assumption that runs are triggered by a drop in asset value is strongly supported by Fig. 1, which plots price indexes in 2007 for ABCP conduits' main asset classes (solid and dashed lines), and also plots the fraction of ABCP conduits experiencing a run as defined in Section 4.1 (dashed-dot line). The figure shows that the ABX index of mortgage-related securities dropped by roughly 20 percentage points in the months *before* runs intensified.[12] Mortgage-related assets made up 9% of the ABCP industry portfolio in 2007.

We cannot rule out that some of the model's parameters changed suddenly in mid-2007. However, as we shall see below, the model fits the data remarkably well without these additional assumptions. Moreover, our parameter estimates are forward-looking: we recover the parameter values consistent with yields and run intensities well into the crisis.

Consistent with industry practice, the model assumes that the sponsor extends a line of credit to the conduit during a run. Therefore, the ABCP that matures during a run is replaced by credit-line debt. Since credit-line debt is held by the sponsor, it is not subject to coordination problems. Total debt D_t in our model is the sum of market-held ABCP and sponsor-held credit-line debt. During a run, the amount of market-held ABCP decreases while sponsor-held debt increases. On net, total debt D_t increases during a run, because sponsor debt replaces market debt issued in the past at lower yields. It is unclear whether credit-line debt is senior or junior to ABCP in practice. For tractability, we assume the two types of debt have equal seniority and maturity.[13]

Note that runs in the model are not caused by idiosyncratic liquidity shocks to creditors. If one individual creditor fails to roll over due to a liquidity shock, another creditor will take up the contract at the break-even yield, preventing a run. In

[9] At the time of the 2007 crisis, Rule 2a-7 under the Investment Company Act limited the portfolio share that registered money market mutual funds can invest in eligible securities not rated A-1/P-1 to 5% of the fund portfolio (these securities are typically rated at least A-2/P-2).

[10] Similar models to ours that incorporate an endogenous default policy in the spirit of Leland (1998) are Hugonnier, Malamud, and Morellec (2012), Décamps and Villeneuve (2007), and He and Milbradt (2012). These models, however, do not include the possibility of a binding credit supply constraint, e.g., a credit ratings requirement.

[11] A related issue is present in Hege and Mella-Barral (2005), in which creditors dilute each other repeatedly after being offered an alternating sequence of debt renegotiation options. The number of options is finite, so that the last debt renegotiation option is well defined (and given) and the problem can be solved by backward induction.

[12] The ABX index of mortgage securities is an average across tranches (from BBB- to AAA) of the ABX.HE indexes for mortgage-backed securities (MBS) originated in the first half of 2006.

[13] Making credit-line debt junior to ABCP or giving it a longer maturity than ABCP would make the credit-line debt behave somewhat like an equity stake in the conduit. Incorporating these features into the model would introduce an extra state variable to separately track the amount of ABCP and credit-line debt. Holding constant $\bar{r}$, extending the maturity of credit-line debt makes total debt increase more slowly during a run, which increases the probability that a conduit recovers from runs, hence reducing investors' willingness to run. That said, sponsors could require a higher $\bar{r}$ if their maturity is longer, which could counteract the previous effect. Making credit-line debt junior to ABCP increases ABCP investors' recovery rate in default, which would likely reduce their willingness to run. The reverse likely obtains if credit-line debt is senior to ABCP.
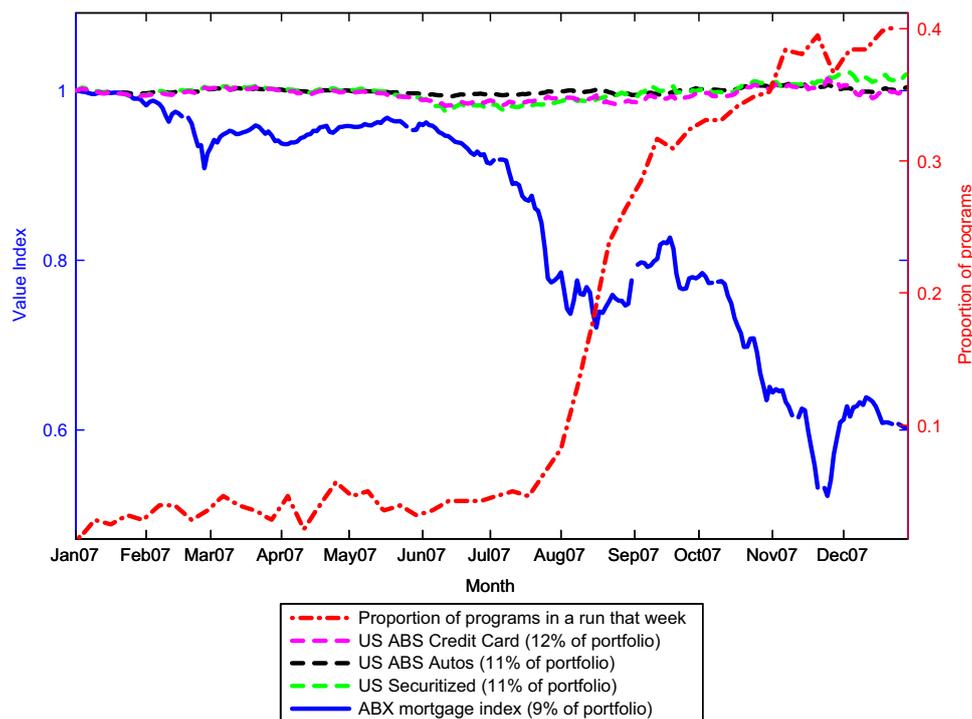
Fig. 1. This figure shows the weekly series of prices for several asset categories in the portfolio of ABCP conduits in 2007 (dashed and solid lines), as well as the proportion of ABCP programs experiencing runs in a given week (dashed-dot line). We normalize prices to $1 on January 1, 2007. The ABX index of mortgage securities is an average across tranches (from BBB- to AAA) of the ABX.HE indexes for MBS originated in the first half of 2006. Data for the remaining asset categories are from Barclays indexes. 'US Securitized' is an aggregate of U.S. asset-backed securities, commercial mortgage-backed securities, and other mortgage-backed securities; this index proxies for the ambiguous 'Securities' category, which makes up 11% of conduit assets in 2007. Portfolio weights in the legend are from 'The ABC's of ABCP,' an unpublished document from Societe Generale. Data for the proportion of runs are from the DTCC database on all issues by ABCP programs, where a run is defined as in Covitz, Liang, and Suarez (2013): an ABCP program experiences a run in a given week if either (1) more than 10% of the program's outstanding paper is scheduled to mature, yet the program does not issue new paper; or (2) the program was in a run the previous week and it does not issue new paper in the current week.

Section 4.1 we explain how our empirical definition of runs distinguishes lack of rollover for a conduit from a creditor's idiosyncratic withdrawal. Moreover, the data support our assumption that runs on ABCP in 2007 did not result from systemic liquidity shocks to creditors: runs on ABCP *preceded* runs on money market funds, ABCP conduits' main investors.

Finally, we assume a conduit never liquidates just a part of its asset to pay running creditors. It is doubtful that conduits would use partial liquidations. The Internet Appendix shows that partial asset sales, far from improving a conduit's health, actually guarantee that the run will continue.[14] The reason is that a partial asset sale automatically increases the conduit's leverage. Also, since ABCP assets like trade receivables are very illiquid, it seems plausible that conduits would wait as long as possible before liquidating them.

2.4. Model solution and examples

Appendices A and B contain details on the model's solution, including a full description of the value function, state variable dynamics, and numerical methods. This subsection describes in nontechnical terms the key features of the solution.

An infinitesimally small lender knows she will face one of three outcomes, depending on which of the following

events occurs first. The first possible outcome is that the asset matures first, delivering a total payout of $\min(D_t, y_t)$ to lenders as a group. In the second, the loan matures first, allowing the lender to choose between rolling over and running. The third, least desirable outcome is that other lenders run on the conduit, the guarantee fails, and the conduit defaults before the loan matures, which delivers $\min(D_t, \alpha F(y_t))$ to the lenders as a group. Therefore, when choosing whether to roll over, each lender must rationally anticipate other lenders' rollover choices.

As in He and Xiong (2012a), we solve for the monotone equilibrium in which lenders roll over their debt as long as the state variable does not drop below a threshold. We show that our model's only state variable is inverse leverage (x_t), which equals the ratio of the asset's fundamental value (y_t) to the conduit's total debt (D_t). Applying Ito's Lemma and Eq. (3), it is straightforward to show that inverse leverage follows:

$$\frac{dx_t}{x_t} = \mu\,dt + \sigma\,dZ_t + \delta\,dt - \delta R_t\,dt. \qquad (5)$$

In other words, the fraction change in inverse leverage equals the fraction change in the asset's value ($\mu\,dt + \sigma\,dZ_t$) plus the fraction of debt maturing ($\delta\,dt$) minus the fractional amount of new debt issued ($\delta R_t\,dt$). In equilibrium, each maturing creditor compares the conduit's current inverse leverage x_t to an endogenous, constant threshold, x^*, and the creditor runs as soon as $x_t < x^*$.

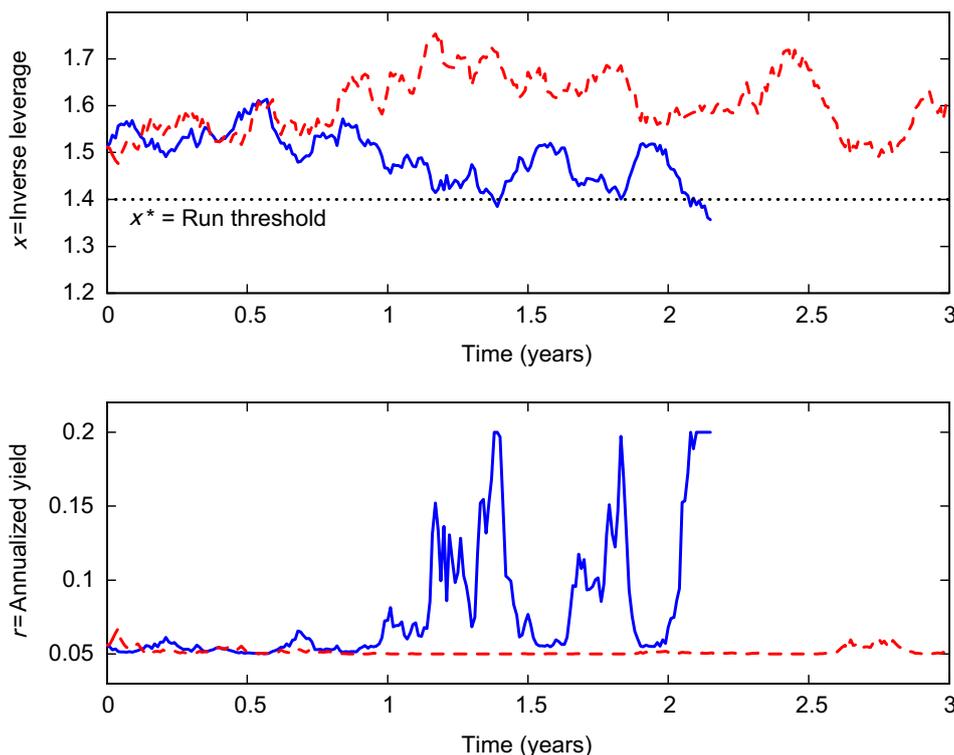[14] The Internet Appendix is available on the authors' Web sites.

Fig. 2. This figure shows the simulated paths of two conduits with the same initial leverage and parameter values. The top panel shows simulated values of x_t, inverse leverage. The dotted line denotes the run threshold. The bottom panel shows simulated paths of annual yields at rollover for the same two conduits. The risk-free rate is 5% and the cap on the rollover yield is 20%.

One important implication is that rollover yields depend on leverage (D_t/y_t) but not on y_t or D_t individually, which is intuitive. More importantly, in equilibrium, creditors will stop rolling over exactly when yield spreads reach the cap, $\bar{r}$, since they cannot be compensated for additional default risk. As a consequence, the run threshold x^* will be the assets-to-debt ratio where yields first hit their cap. Since sponsors extend the credit line at a below-market spread $\bar{r}$ during a run, the sponsors take a loss while supporting the conduit. Like He and Xiong (2012a), we find that creditors start running before a conduit becomes insolvent. The reason is that each creditor's rollover decision imposes an externality on the other creditors.

Fig. 2 illustrates how leverage and yields adjust over time. The top panel plots the time series of inverse leverage (x_t) for two simulated conduits with the same initial fundamentals but different outcomes. The flat dotted line represents x^*, the predicted run threshold. The dashed line depicts a conduit whose asset's value remains high enough so that the conduit never experiences a run, and all lenders are paid in full. The solid line represents a conduit that experiences two runs when its inverse leverage falls below x^*. During the first run, the guarantee survives long enough for the conduit to repay all running lenders and begin issuing paper again. The guarantee fails in the second run, causing the conduit to default and liquidate assets, imposing losses on some lenders.

The bottom panel of Fig. 2 shows the corresponding rollover yields for those same simulated conduits. Since the conduit represented by the dashed line remains healthy, its yield remains at or near the risk-free rate, $\rho = 5\%$. The yields of the conduit represented by the solid line spike up and become more volatile as a run becomes

imminent, eventually reaching their cap when the run begins. As soon as this conduit recovers from its first run, yields drop below the cap.

3. Flexible pricing, dilution risk, and the likelihood of runs

Allowing yields to adjust over time significantly changes the likelihood of runs, relative to the He and Xiong (2012a) model with constant yields. We compare simulated run probabilities in our model to those in He and Xiong (2012a, henceforth HX). To make the models comparable, we use the same parameter values where possible, we make the asset's initial market value the same in both models,[15] and we assume conduits in both models initially borrow $1. HX's lenders receive a face value of $1 with a fixed, exogenous yield, so in return for their initial $1 investment, lenders receive debt worth more than $1. Yields in our model are set so that lenders exactly break even, so in return for their initial $1 investment, lenders receive debt worth $1.

---

[15] The asset pays interim cash flows at rate r in He and Xiong (2012a), but our model has no interim cash flows. Setting the asset's value equal in the two models requires choosing initial fundamental y_0 by solving

$$F^{HX}(y_o^{HX}) = F(y_0)$$

$$\Rightarrow \frac{r}{\rho+\phi} + y_0^{HX}\frac{\phi}{\rho+\phi-\mu} = y_0\frac{\phi}{\rho+\phi-\mu},$$

where y_0 (y_o^{HX}) is the asset's initial fundamental value in our (HX's) model. In the first analysis we set $y_0^{HX} = 2.1$. In the second analysis we set $y_0^{HX} = 0.82$.

Table 1
The effect of flexible yields on runs.

This table compares the predictions from our model to the predictions of He and Xiong (2012a), denoted HX. Yields change over time in our model, whereas yields are constant in HX. The columns on the left use HX's calibrated values: $\rho = 1.5\%$, $\phi = 0.077$, $\alpha = 55\%$, $\sigma = 20\%$, $\mu = 1.5\%$, $y_0 = 1.4$, $\delta = 10$, and $\theta = 5$. The columns on the right use estimated parameter values for the weak-guarantee subsample in Table 4. Panel A shows the fraction of simulated conduits that experience a run in our model within one year, divided by the corresponding fraction from HX. Panel B shows the run threshold in our model (x^*) divided by the run threshold in HX. Panel C shows the conduit's initial inverse leverage in our model (x_0) divided by the conduit's initial inverse leverage in HX; these results are identical out to two digits for the three values of $\bar{r}$. The parameter $\bar{r}$ is our model's cap on yield spreads, and ρ is the risk-free rate, so $\bar{r} + \rho$ is the capped rollover yield.

Panel A: Ratio of the one-year run probability in our model to that in HX

	Using HX parameters				Using estimated parameters		
HX's fixed yield:	5%	7%	9%		$\rho + 0.1\%$	$\rho + 0.3\%$	$\rho + 0.5\%$
$\bar{r} + \rho = 15\%$	1.31	3.73	50.64		0.65	1.48	3.61
$\bar{r} + \rho = 20\%$	1.29	3.64	49.08				
$\bar{r} + \rho = 25\%$	1.28	3.56	47.65				

Panel B: Ratio of the run threshold in our model to that in HX

	Using HX parameters				Using estimated parameters		
HX's fixed yield:	5%	7%	9%		$\rho + 0.1\%$	$\rho + 0.3\%$	$\rho + 0.5\%$
$\bar{r} + \rho = 15\%$	1.43	1.77	2.29		1.22	1.24	1.26
$\bar{r} + \rho = 20\%$	1.42	1.75	2.28				
$\bar{r} + \rho = 25\%$	1.41	1.74	2.26				

Panel C: Ratio of initial inverse leverage in our model to that in HX

	Using HX parameters				Using estimated parameters		
HX's fixed yield:	5%	7%	9%		$\rho + 0.1\%$	$\rho + 0.3\%$	$\rho + 0.5\%$
	1.38	1.54	1.70		1.25	1.26	1.27

We compare the two models in Table 1. We show results using HX's calibrated parameter values (left-hand columns) and our estimated parameter values (right-hand columns). We repeat the analysis using several values of $\bar{r}$ (our model's cap on yield spreads) and HX's exogenous yield.[16]

Panel A shows that runs are 1.31–51 times more likely in our model than in HX's when we use HX's calibrated parameter values. Runs are especially more likely in our model if HX's exogenous yield is higher, because HX's investors are less willing to run on debt that offers a higher interest rate. Runs are not always more likely in our model, however. With estimated parameter values, we see that when HX's fixed yield is sufficiently low, our model produces 35% fewer runs than HX's model.

Intuitively, flexible yields influence runs through three channels. The first two channels make runs less likely in our model compared to HX. First, conduits in our model can initially borrow at low interest rates, allowing them to start with lower leverage (Panel C in Table 1). Second, being able to raise yields in bad times helps convince lenders to roll over.

This relative advantage is especially large when HX's fixed yield is very low, which explains why our model produces fewer runs than HX only when HX's fixed yield is very low (e.g., ten basis points (b.p.) above the risk-free rate in Table 1).

The third channel makes runs more likely in our model and typically outweighs the previous two channels. Flexible yields introduce a new risk, which we call 'dilution risk,' on top of rollover risk and insolvency risk. If conditions deteriorate for a conduit, it will have to offer higher yields to induce rollover. These higher yields increase the conduit's debt by more, which dilutes earlier lenders' stakes. This effect depends strongly on the assumption that bankruptcy proceeds are distributed pro rata, consistent with ABCP industry practice. A lender deciding whether to roll over in our model anticipates the possibility of being diluted in the future if conditions worsen. The lender therefore preemptively demands a higher yield to compensate her for dilution risk. Since dilution risk increases yields for any given level of leverage, yields hit their cap at lower leverage, implying a higher run threshold for inverse leverage, x^*. Panel B in Table 1 shows that the run threshold is indeed higher in our model compared to HX, which tends to make runs more likely.

4. Estimation

This section describes the data, SMM estimator, and intuition behind the estimation method.

---

[16] In the first analysis, HX's fixed yield is centered at its calibrated value, 7%. We choose values of $\bar{r}$ much higher than HX's fixed yield, because higher values of $\bar{r}$ make runs less likely in our model, all else equal. We find that despite these high $\bar{r}$ values, runs are still more likely in our model than in HX. In the right-hand columns, $\bar{r}$ is at its estimated value, and we choose values of HX's fixed yield that are within the range of observed yields.

4.1. Data

The dataset used in this paper includes all issuance transactions in the U.S. ABCP market from the Depository Trust and Clearing Corporation (DTCC). The data contain the outstanding amount of paper for a conduit each week and the distribution of maturities and yields each day a conduit issues ABCP.

We obtain data on each conduit's guarantee type from Moody's Investors Service. ABCP conduits are structured with one of four possible types of guarantees (Acharya, Schnabl, and Suarez, 2013). In conduits structured with a full credit guarantee, the sponsor provides a line that can be drawn regardless of asset defaults. In conduits with a full liquidity guarantee, the sponsor provides a line that can be drawn as long as the assets are not in default. In SIV guarantees, only a portion of conduit liabilities are covered by the line. In conduits created to issue extendible notes, issuers have the option of extending the maturity of the paper at a prespecified penalty rate, exposing investors to asset defaults during the extension period. From the point of view of investors, full credit and full liquidity guarantees offer relatively stronger protection.

Covitz, Liang, and Suarez (2013) show that conduits with stronger guarantees experienced significantly fewer runs in 2007. For this reason, we estimate the model in two sub-samples based on guarantee strength.[17] The 'strong-guarantee' subsample contains the 191 conduits with either a full credit or full liquidity guarantee; 45% of these conduits experienced a run in 2007. The 'weak-guarantee' subsample contains the 90 conduits with either an SIV guarantee or extendible paper; 83% of these conduits experienced a run in 2007. As in many structural estimation papers (e.g., Hennessy and Whited, 2007; Strebulaev and Whited, 2012), we assume parameter values are constant within each subsample. Our parameter estimates therefore characterize an average conduit within each subsample.

We use the method of Covitz, Liang, and Suarez (2013) to identify runs in the data. More specifically, we say that conduit i is in a run in week t if either (1) more than 10% of the conduit's outstanding paper is scheduled to mature, yet the conduit does not issue new paper; or (2) the conduit was in a run in week $t-1$ and the conduit does not issue new paper in week t. We say that a conduit recovers from a run in week t if it issues paper that week but was in a run the previous week. By using the *total* amount a conduit rolls over, this definition avoids misclassifying as runs situations in which one creditor replaces another due to the first creditor's idiosyncratic liquidity needs.

We measure each conduit's rollover spread as the dollar-weighted average annualized yield for paper issued on Thursday of week t, minus the prevailing federal funds rate.[18] If the conduit did not issue paper on Thursday, we move one day ahead until finding an issuance transaction in week t.

The total amount of ABCP outstanding peaked at $1.2 trillion in late July 2007. At that time, 339 ABCP conduits operated. Yield spreads averaged five b.p. in the first half of the year. In August 2007, the amount of debt outstanding plunged by $190 billion and average spreads increased to 74 b.p.[19] Roughly 25% of ABCP conduits experienced a run in August, according to our measure. Rollover yields remained high and volatile in the second half of 2007. By the end of the year, the total amount of ABCP outstanding was 30% below its peak.

Our analysis uses all transactions from 2007. We face the trade-off that a larger sample would provide more precise estimates, but it would be harder to argue that model parameters are constant over a longer period. Year 2007 is an ideal sample because it contains many runs and also several months of pre-run data. Adding observations from 2006 would not improve precision, because yield spreads were near zero and there were no runs. Adding observations from 2008 would potentially contaminate results with effects from the Lehman Brothers failure and subsequent government interventions.

4.2. Estimator

First we explain how we measure parameters ρ, δ, ϕ, and μ directly from the data. Next we describe the SMM estimation of the four remaining parameters.

Investors' discount rate ρ is also the risk-free interest rate. We set ρ to 4.9%, the annualized yield of one-month T-bills at the beginning of 2007.

The average debt maturity in our model is $1/\delta$. We set $1/\delta$ to 0.101 years (37 days), the average maturity of ABCP as of March 2007. We use the same value of $1/\delta$ in both subsamples because there is no significant difference in maturities between them in early 2007, and, as we show later, such small differences in maturity have a very small effect on run probabilities. The assumption that δ is constant over time could be somewhat problematic given that most conduits experienced a rat-race whereby they offered shorter maturities to prevent creditors from running (Brunnermeier and Oehmke, 2013). While average rollover maturities do decrease in the six months preceding runs in our data, we find that they only drop from 38 days to 27 days. By contrast, changes in ABCP rollover yields were more dramatic, which is why we focus on time-varying yields in this paper. Extending the model to include endogenous maturity is an interesting avenue for future work.

The expected conduit life span, which corresponds to the asset's duration, is $1/\phi$. Adding the assumption that new ABCP conduits are created at a constant rate, the model predicts that the average age of conduits alive at any snapshot in time equals $1/\phi$. The average age of ABCP

---

[17] Ideally, we would estimate the model in even finer-grained subsamples, but the small number of conduits prevents us from doing so.

[18] We choose Thursday because amounts outstanding are measured at the end of Wednesday each week.

[19] Important events in early August 2007 include American Home Mortgage's declaration of bankruptcy (August 6), the halting of redemptions at three investment funds affiliated to BNP Paribas (August 9), emergency liquidity provision by the European Central Bank (ECB) (August 9) and the Federal Reserve (August 10), and the downgrade of Countrywide Financial and its drawing on bank credit lines (August 16).

conduits operating in July 2007 is 5.8 years, so we set ϕ to 1/5.8.

The parameter μ, which represents the asset's growth rate, is not identified from our data. The parameter μ is not the asset's expected return, which equals ρ (investors' discount rate). Therefore, μ would not be identified directly from average returns on ABCP assets, even if we had those data. The asset's return at τ_ϕ, the instant it matures, is positive (negative) if μ is less than (greater than) ρ.[20] These event returns could help us identify μ, but unfortunately, those data are not available either. For our main results we set $\mu = \rho$, which assumes it is neither good nor bad news for investors when the asset matures. For robustness, in the Internet Appendix we show that parameter estimates do not change significantly if we set μ to $\rho + 1\%$ per year.

We do not estimate conduits' initial leverage, $1/x_0$. Leverage at the beginning of our sample is not well identified, since yield spreads in the first half of 2007 were near zero. When yield spreads are near zero, the model's mapping from conduit leverage to yields is almost flat. Intuitively, since spreads were near zero at the beginning of 2007, we know leverage was low at that time, but we do not know how low.

Fortunately, we do not need to know $1/x_0$ to estimate the model, because the moments we use in SMM estimation are independent of its value. Some of our moments are conditional on a run starting, at which point x_t has reached x^*. Since the predicted x^* does not depend on x_0, then neither do these moments. The remaining moments (both simulated and actual) are forward-looking and only use conduit/week observations where yield spreads are at least ten b.p. per year. At spreads of ten b.p. and above, the mapping between leverage and spreads is no longer flat. Therefore, the ten b.p. threshold forces us to only use observations that exceed a certain leverage threshold. Once we condition on leverage being above this threshold, the forward-looking moments no longer depend on initial leverage, $1/x_0$, because leverage is the model's only state variable. Our simulated conduits start with $1/x_0$ low enough that their initial spreads are well below ten b.p., as they were in early 2007. We simulate a large enough sample so that we have many observations with spreads above ten b.p., which allows us to measure our simulated moments precisely.

The remaining parameters to estimate are σ (asset volatility), θ (the weakness of guarantees), α (asset liquidity), and $\bar{r}$ (the cap on yield spreads). We estimate σ as a structural parameter instead of using price data on ABCP assets, because those data are not available. Data on conduit-level asset holdings are not publicly available. Even constructing an industry-wide price index is impossible, because we lack price data for illiquid assets like trade receivables, the largest ABCP asset class.

We estimate the four remaining parameters using the simulated method of moments (SMM). This estimator chooses parameter values that minimize the distance between moments generated by the model and their sample analogs.

The following subsection defines our 13 moments and explains how they identify our parameters. Additional details are in Appendix C.

4.3. Identification and choice of moments

Since we conduct a structural estimation, identification requires choosing moments whose predicted values move in different ways with the model's parameters, and choosing enough moments so there is a unique parameter vector that makes the model fit the data as closely as possible. This subsection explains how our 13 moments vary with the four parameters. Each moment depends on all model parameters, often through the parameters' effect on leverage dynamics or the run threshold. Below, we emphasize which parameters matter most for each moment, which explains which features of the data are most important for each parameter. To illustrate, Table 2 presents the Jacobian matrix containing the derivatives of our 13 moments with respect to our four parameters.[21] For both subsamples, the Jacobians have full rank and a low condition number,[22] which confirms local identification.

4.3.1. Recoveries from runs

The first moment is the fraction of runs that are followed by a recovery, meaning that the conduit issues paper again, at least once within eight weeks of the run's start.[23] In our model, once a run starts, the probability of a recovery decreases in θ, the guarantee's weakness. Intuitively, a strong guarantee buys time for asset values to improve so the conduit can exit the run. The Jacobian in Table 2 confirms that this first moment is most sensitive to θ and fairly insensitive to other parameters.

The second moment is the average number of days until recovery for those runs that experience a recovery within eight weeks of the run's start. Conditional on a recovery within a given period, the expected time to recover is shorter for higher asset volatility σ, because higher volatility makes the conduit re-cross the run threshold sooner. Table 2 shows that this moment is indeed most sensitive to σ.

The remaining parameters have an indirect effect on our first two moments through the run threshold, x^*. In this case, however, these effects are relatively small, confirming that the recovery probability and the expected recovery time essentially identify σ and θ only.

4.3.2. Yield volatility

The moments we use to summarize yield spread volatility are the coefficients β_0 and β_1 from the following

---

[20] The asset's value immediately prior to τ_ϕ is $F(y_{\tau_\phi})$. The asset's value immediately after is y_{τ_ϕ}.

[21] We present the Jacobian evaluated at estimated parameter values for the weak-guarantee subsample. The properties of the Jacobian for the strong-guarantee subsample are very similar. In the interest of space, we report the Jacobian for the strong-guarantee subsample only in the Internet Appendix. To make the sensitivities comparable across moments, we express them as elasticities, e.g., $(\partial m_i / \partial \alpha) \times (\alpha / m_i)$ is the elasticity of the i-th moment to α.

[22] The condition number of a matrix is the ratio of its largest to smallest singular value. Large condition numbers indicate a nearly singular matrix.

[23] We find empirically that if a run is not followed by a recovery within eight weeks, then it is unlikely that the conduit will ever recover.

Table 2

Estimated Jacobian matrix.

   This table presents the estimates of the Jacobian matrix for the 13 moment conditions in our SMM estimation procedure, for the subsample of 90 ABCP conduits in 2007 with SIV or extendible credit guarantees. Moment 1 is the probability that a conduit experiences a recovery within eight weeks of a run's start. Moment 2 is the average number of days between the run's start and recovery, conditional on a recovery occurring within eight weeks of the run's start. Moments 3 and 4 are the intercept and slope from a regression of absolute changes in yield spreads on the lagged yield spread. Moments 5–7 are the intercept and slopes from a regression of yield spreads on the number of weeks relative to a run and the exponent of that same number. Moments 8–13 come from three regressions, each of the indicator $\mathbf{1}_{\{\text{run within } \tau \text{ weeks}\}}$ on the current yield spread. The three regressions use $\tau = 2$, 4, and 8 weeks. Each row of each matrix contains the estimated elasticities of the given moment with respect to the parameters across its columns. Parameters are estimated by SMM, which chooses values that minimize the distance between actual and simulated moments. Section 2 describes the model used to simulate moments. The number highlighted in each row in bold type face corresponds to the moment's highest elasticity in absolute value.

		Elasticity of moments with respect to			
		θ	σ	$\bar{r}$	α
Moments on time between run and recovery (τ):					
1.	$\Pr[\tau < 8 \text{ weeks}]$	**−0.209**	−0.106	−0.036	0.075
2.	$E[\tau \mid \tau \leq 8 \text{ weeks}]$ (in days)	−0.104	**−0.295**	−0.011	0.106
Moments from regression of $\lvert r_{it+1} - r_{it}\rvert$ on r_{it}:					
3.	Intercept	0.112	0.106	0.724	**−0.758**
4.	Slope	0.002	−0.307	0.133	**0.472**
Moments describing yield spreads leading up to runs:					
	Regression of r_{it} on $\tau[\equiv$ weeks relative to run] and $\exp(\tau)$				
5.	Intercept	0.102	−0.186	**0.990**	−0.311
6.	Slope on τ	0.133	−0.149	**0.977**	−0.870
7.	Slope on $\exp(\tau)$	0.235	−0.548	**1.093**	−0.118
	Regressions of $\mathbf{1}_{\{\text{run within } \tau \text{ weeks}\}}$ on yield spread:				
8.	Intercept ($\tau = 2$)	0.023	−0.424	0.243	**−0.446**
9.	Slope ($\tau = 2$)	−0.030	0.177	−0.055	**−0.296**
10.	Intercept ($\tau = 4$)	−0.414	0.694	−0.498	**−2.562**
11.	Slope ($\tau = 4$)	0.028	−0.023	0.057	**−0.444**
12.	Intercept ($\tau = 8$)	−0.158	0.283	**−0.571**	−0.344
13.	Slope ($\tau = 8$)	0.066	−0.067	**0.180**	−0.097

panel regression of absolute changes in yield spreads on the lagged yield spread[24]:

$$\lvert r_{it} - r_{it-1}\rvert = \beta_0 + \beta_1 r_{it-1} + \varepsilon_{it}. \qquad (6)$$

The predicted yield spread volatility is given by

$$var_t(dr_t) = \left(x_t \frac{\partial r}{\partial x}(x_t, x^*)\right)^2 \sigma^2 \, dt. \qquad (7)$$

The first term in (7) increases in the yield spread, so the model predicts that yield volatility is high when yield spreads are high. In other words, we should find $\beta_1 > 0$ in (6). The model therefore produces time-varying volatility in debt yields, even though asset volatility, σ, is constant. The term $\partial r/\partial x$ in (7) goes to zero as yield spreads approach zero. Intuitively, if a conduit's leverage is extremely low, yield spreads are near zero, and small changes in leverage still keep spreads near zero. Therefore, the model imposes $\beta_0 \approx 0$ as an overidentifying moment condition, regardless of parameter values. As a result, yield volatility

is informative about asset volatility only when spreads are high.

   Asset volatility has a direct, positive effect on yield volatility through the σ term in (7), and also a negative effect via the first term: a higher σ decreases the absolute slope $\lvert \partial r / \partial x\rvert$ for given r, provided r is high enough. For our parameter estimates, and given that we measure our moments conditional on spreads exceeding ten b.p., the second effect dominates: we find that the sensitivity of yield volatility to yield levels *decreases* with asset volatility, so σ is partially identified off its negative effect on β_1. Indeed, Table 2 confirms that β_1 depends negatively on σ.

   Note too that α has a strong positive effect on β_1. The reason is that an increase in asset liquidity decreases the run threshold, x^*, which in turn implies a higher absolute slope $\lvert \partial r / \partial x\rvert$ for any given r. Intuitively, if two conduits with assets of different liquidity have reached the same yield spreads, it must be that yields are increasing faster for the one with higher liquidity.

4.3.3. Yield spreads preceding runs

   Our next three moments measure average yield spreads in event time before runs. We define τ_{it} as the number of weeks relative to the run's start, and we use the subset of data from the 12 weeks preceding each run to estimate the

---

[24] Fig. 3 plots the nonparametric relation in the data and shows that the linear, parametric specification in Eq. (6) fits the data quite well. We discard observations with $r_{it} \leq 10$ basis points per year so that this moment does not depend on x_0, which we do not estimate.
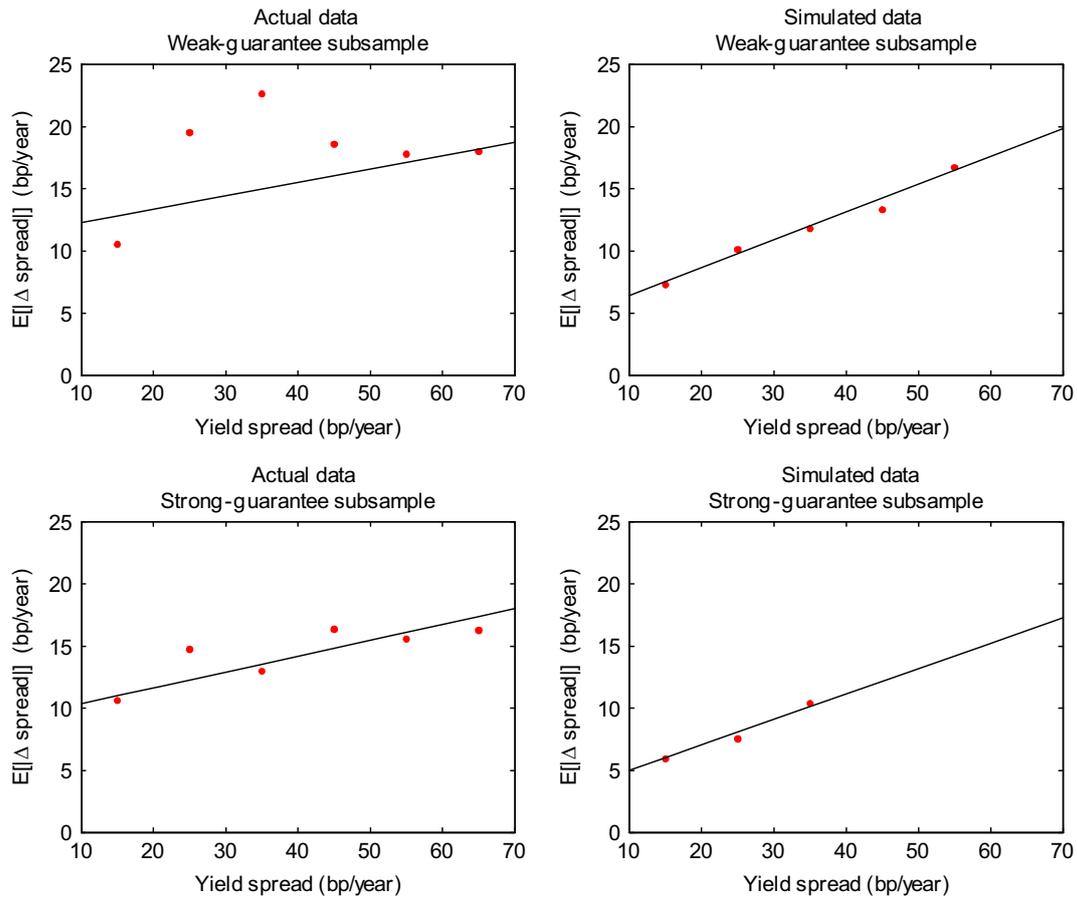
Fig. 3. This figure shows the relation between yield spread volatility and the yield spread level, both measured in basis points per year (bp/year). The vertical axis is the absolute value of one-week changes in yield spread. The horizontal axis is the lagged yield spread. The left-hand (right-hand) panels show results in actual (simulated) data. The top panels show results for the weak-guarantee subsample, which contains the 90 ABCP conduits in 2007 with extendible or SIV guarantees. The bottom panels show results for the strong-guarantee subsample, which contains the 191 ABCP conduits in 2007 with full credit or full liquidity credit guarantees. The points show local averages, and the solid line shows predicted values from a regression of absolute changes in yield spreads on the lagged yield spread. The intercept and slope from this regression provide two of the 13 moments used in the SMM estimation. The reason there are fewer points in the bottom-right panel compared to the bottom-left panel is that the estimated cap on yield spread is 36 basis points in the strong-guarantee subsample; all simulated spreads are therefore ≤ 36 basis points, whereas there are a few spreads > 36 basis points in the actual data.

regression

$$r_{it} = \gamma_0 + \gamma_1 \tau_{it} + \gamma_2 \exp(\tau_{it}) + \varepsilon_{it}. \qquad (8)$$

Fig. 4 shows that this specification fits the path of average yield spreads leading up to runs fairly well. Our next three moments are the coefficients γ_0, γ_1, and γ_2, which summarize event-time spreads.

As discussed in Section 2.4, the model predicts that a run begins the instant yield spreads hit $\bar{r}$. Since we only have weekly data, we cannot directly observe yields spreads the instant a run begins. However, yield spreads in the weeks leading up to runs are informative about the yield spread the instant a run begins. Specifically, the level (γ_0), slope (γ_1), and curvature (γ_2) of the event-time plot of average spreads before runs all increase in $\bar{r}$. The moments γ_0, γ_1, and γ_2 therefore help identify $\bar{r}$. Consistent with this reasoning, the estimated Jacobian shows that these three moments are by far most sensitive to $\bar{r}$ and therefore effectively identify this parameter.

4.3.4. Run probabilities

Our next moments are from three regressions that forecast future runs using current yield spreads. The regressions

have the form[25]

$$\mathbf{1}_{\{\text{run within } \tau \text{ weeks of } r_{it}\}} = \lambda_{0\tau} + \lambda_{1\tau} \frac{r_{it}}{maxr_i} + \varepsilon_{it}, \qquad (9)$$

where $maxr_i$ proxies for conduit i's maximum yield spread, $\bar{r}$. Since runs in the model begin as soon as r_{it} hits $\bar{r}$, the higher the fraction $r_{it}/\bar{r}$, the closer the conduit is to a run.[26] For a conduit that experiences a run, the most natural proxy for $\bar{r}$ is the conduit's maximum observed yield spread. For conduits that never experience a run or

---

[25] Following Angrist and Pischke (2009), we use ordinary least squares (OLS) rather than a probit/logit model, because OLS slopes are easier to interpret, and OLS provides the closest linear approximation of the conditional expectation function. Figs. 5 and 6 show that the linear specification in (9) fits the raw data quite well. As in regression (6), we exclude observations with $r_{it} < 10$ basis points per year; run probabilities for these observations are sensitive to the choice of initial condition x_0 in our simulations, and we want moments that do not depend on x_0.

[26] Empirically, we find some heterogeneity in $maxr_i$. Since the estimation procedure assumes $\bar{r}$ is constant across conduits, our estimated $\bar{r}$ reflects the yield-spread cap for the average conduit. We estimate in subsamples to the extent possible to accommodate parameter heterogeneity.
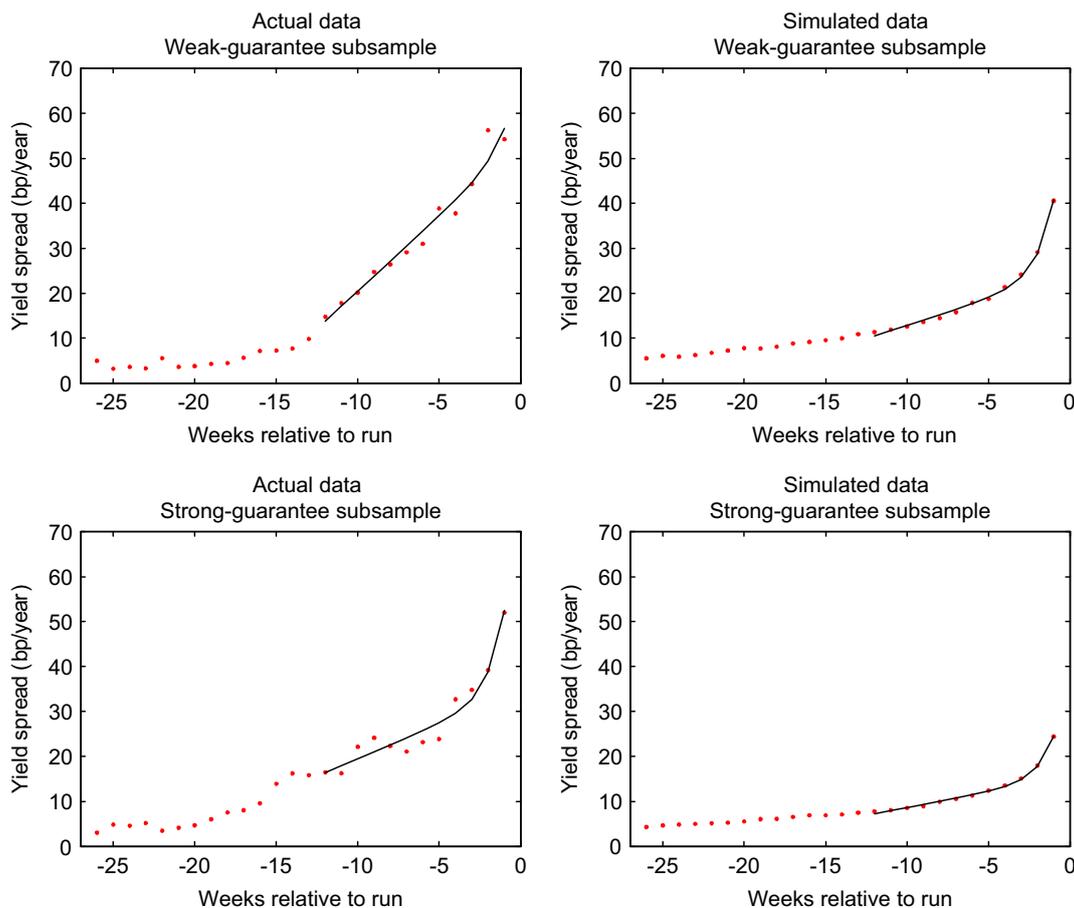
Fig. 4. This figure plots average yield spreads (basis points per year, bp/year) in event time leading up to runs in event week zero. The left-hand (right-hand) panels show results in actual (simulated) data. The top panels show results for the weak-guarantee subsample, which contains the 90 ABCP conduits in 2007 with extendible or SIV guarantees. The bottom panels show results for the strong-guarantee subsample, which contains the 191 ABCP conduits in 2007 with full credit or full liquidity credit guarantees. The points show the average yield spread in each week. The solid line shows the predicted values from the regression of yield spreads on the number of weeks relative to the run and the exponent of that same number. The intercept and two slopes from this regression provide three of the 13 moments used in SMM estimation. The solid line starts at week -12 because the regression only uses data from weeks -12 to -1.

high yield spreads, we use information from those that do experience runs: we set $maxr_i$ to the larger of the conduit's maximum observed yield spread (since the conduit's $\bar{r}$ is at least this large) and the average $maxr$ across conduits that did experience runs (a proxy for the average $\bar{r}$ in the sample). We estimate regression (9) for forecasting horizons of $\tau = 2, 4$, and 8 weeks. Our last six moments are the coefficients $\lambda_{0\tau}$ and $\lambda_{1\tau}$ from those three regressions.

The moment $\lambda_{0\tau}$ summarizes the run probability when spreads are near zero. Therefore, this moment depends negatively on the distance between x_0 and the run threshold, x^*, which is itself decreasing in α. Table 2 shows that $\lambda_{0\tau}$ is strongly decreasing in α, so these moments effectively identify α through α's effect on the run threshold. Note that a higher $\bar{r}$ also implies a lower run probability when yield spreads are near zero. However, since α has a strong effect on yield levels, which impact the drift of leverage, a lower α not only implies a lower distance to the run threshold, but also a quicker transition to it. As a consequence, the difference between a conduit that never experiences a run as opposed to a conduit that experiences a run *quickly* is more likely to be due to a difference in α rather than $\bar{r}$. Consistent with the above intuition, the

estimated Jacobian shows that $\lambda_{0\tau}$ is more sensitive to α, than $\bar{r}$ for the two- and four-week run probabilities, but not for the eight-week probabilities.

5. Estimation results

We start by assessing how well our model lines up with data from the 2007 ABCP crisis. We then present and interpret the structural parameter estimates.

5.1. Model fit

5.1.1. Moments used in SMM estimation

Table 3 compares actual and simulated values of our 13 moments. The left (right) half of the table shows moments in the weak- (strong-) guarantee subsample.

Moments 1 and 2 focus on recoveries from runs. Comparing moment one across subsamples, we see that the probability of a recovery is significantly higher ($t = 2.0$) in the strong-guarantee subsample, consistent with the model's prediction. When recoveries do occur, they arrive 17 days after the run's start, on average (moment 2). Comparing simulated and actual moments, the model
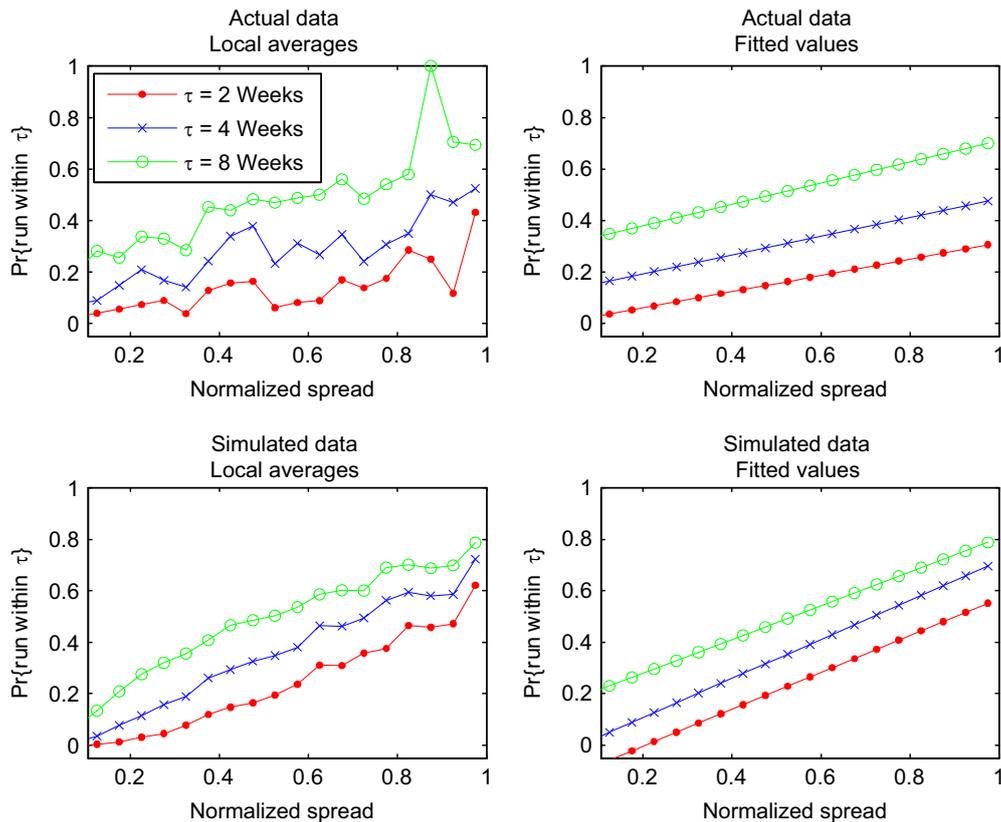
Fig. 5. This figure shows the forecasting relation between yield spreads and runs in the weak-guarantee subsample, which contains 90 ABCP conduits in 2007 with extendible or SIV credit guarantees. The top panels show the relationship using actual data, and the bottom panels show the relationship using data simulated from the model with the parameter estimates in Table 4. The vertical axis is the probability of a run within τ weeks of the current conduit/ week observation. The horizontal axis is the normalized yield, defined as the current yield spread divided by *maxr*, a proxy for the conduit's cap on yield spreads. The left-hand panels show local averages of run probabilities. The right-hand panels show the predicted values from a regression of $\mathbf{1}_{\{run\ within\ \tau\ weeks\}}$ on the normalized yield spread. We estimate this regression for forecasting horizons of $\tau=2$, 4, and 8 weeks. The intercepts and slopes from these regressions provide six of the 13 moments used in SMM estimation.

closely matches both the observed probability of a recovery and the average time until recovery.

Moments 3 and 4 measure the overall level of yield spread volatility and its sensitivity to yield spread levels. The standard errors of the actual moments indicate that moments 3 and 4 are measured fairly imprecisely, in large part because of comovement in yields across conduits. The sensitivity of yield volatility to the yield level is positive in both subsamples, consistent with the model's prediction, but the slope is statistically significant only in the strong-guarantee subsample ($t=2.13$). Although the simulated and actual moments differ in some cases by more than a factor of two, the *t*-statistics in Table 3 indicate that the differences are not statistically significant. Fig. 3 shows the nonparametric relation between yield volatility against the yield level, and it also shows the best-fit relation summarized by moments 3 and 4. We see that the model produces slightly lower yield spread volatility than we see in the data. Some of this 'extra' yield volatility in the data is likely due to measurement error.

Moments 5–7 measure yield spreads leading up to a run. The *t*-statistics show that, in each subsample, the actual and simulated moments are not significantly different from each other. The high standard errors, though, indicate that the actual moments are measured with considerable error. This

error results from the limited number of runs in our sample, and also from comovement in yields. Fig. 4 plots yield spreads in event time leading up to runs, comparing actual and simulated data, and also comparing the nonparametric pattern in the data to the moments' parametric relation. In both actual and simulated data, we see that yield spreads start below ten b.p. per year 26 weeks before runs start, then spreads increase exponentially leading up to runs. Spreads reach a higher level in actual data than in simulated data. By choosing a higher estimate of the yield cap, $\bar{r}$, we could match the actual pattern more closely, albeit at the expense of other moments. The SMM weighting matrix makes the estimator match these moments less closely since we measure them fairly imprecisely.

Examining moments 8–13, we see that yield spreads forecast runs in the actual data: the slope on yield spreads is significantly positive (*t*-statistics between 2.2 and 8.2) in both subsamples and at forecasting horizons of two, four, and eight weeks. The simulated slopes are also positive, so high yields also forecast runs in the model. While the model is able to fit the directional pattern in the data, the model is not able to match magnitudes exactly: the simulated and actual slopes are significantly different, and so are the intercepts. To allow for an easier comparison, Figs. 5 and 6 plot the relation between yields and run probabilities. Despite the difference in
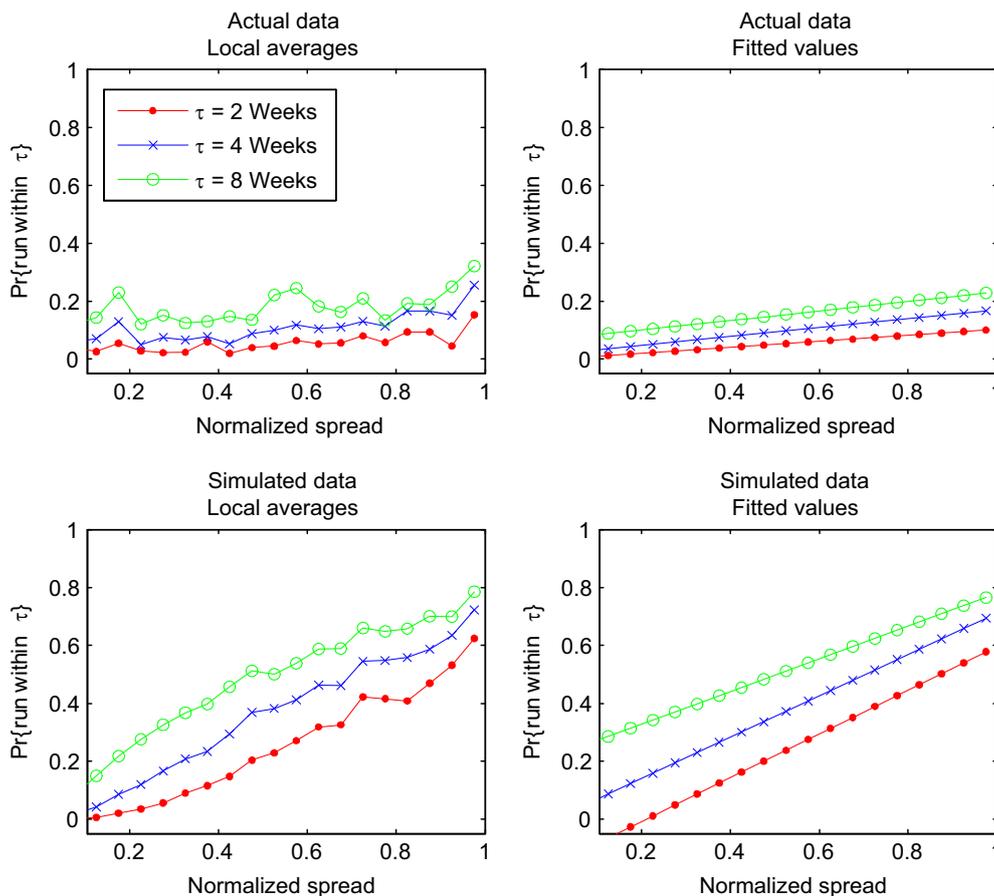
Fig. 6. This figure shows the forecasting relation between yield spreads and runs in the strong-guarantee subsample, which contains the 191 ABCP conduits in 2007 with full credit or full liquidity guarantees. Definitions are the same as in Fig. 5.

moments, the forecasting relation looks quite similar in actual and simulated data for the weak-guarantee subsample (Fig. 5).[27] The similarity is apparent both in the nonparametric plots (left panels) and parametric plots based on moments 8–13 (right panels).

These forecasting results contribute to the debate over what causes runs. The literature has been divided in two groups (see Goldstein, 2014 for a survey). The first group proposes that a run's causes are unobservable, so it is impossible to forecast or assign probabilities to runs.[28] The second group, motivated by Gorton (1988), proposes that runs are caused by deteriorating, observable fundamentals, hence we can forecast and assign probabilities to runs. Our model, along with several others, belongs to this second group.[29] We show that yields do forecast runs in our data. Further, our model of fundamental-driven runs comes close to fitting this empirical pattern, at least in one subsample.

The fit is not as close in the strong-guarantee subsample (Fig. 6). Specifically, the model struggles to explain why

conduits with strong guarantees experienced few runs even when their yield spreads reached high levels. One potential explanation is that $\bar{r}$ differs considerably across conduits in this subsample, contrary to our identifying assumption that parameters are constant within each subsample (Section 4.1). As we explain later, the probability of a run decreases in $\bar{r}$. If there are large differences in $\bar{r}$ across strong-guarantee conduits, then only those conduits with sufficiently low $\bar{r}$ experience a run. Since the moments that mainly identify $\bar{r}$ only use data from conduits that actually experience a run (Fig. 4), then the estimate of $\bar{r}$ in this subsample is representative only of such conduits, and is likely to be lower than the true $\bar{r}$ for conduits that did not experience runs. Likewise, *maxr*, our proxy for $\bar{r}$, could be too low for the high-$\bar{r}$ conduits, which would also make their normalized yields too high.[30] In other words, some high-$\bar{r}$ conduits in the strong-guarantee subsample never actually experienced spreads near their high-$\bar{r}$ values, meaning they never came close to a run. This reasoning would explain why the model over-predicts run probabilities when normalized spreads are

---

[27] The reason the moments look different while the plots look similar is that the model gets the intercept too low but the slope too high, and these two differences offset each other in the figure.

[28] This group follows Bryant (1980) and Diamond and Dybvig (1983). The unobservable is sometimes labeled a sunspot.

[29] Goldstein and Pauzner (2005), He and Xiong (2012a), and Morris and Shin (1998) also belong to this group.

[30] Recall that for conduits i that never experience runs, we set $maxr_i$ to the average of $maxr_j$ across conduits j that experience runs. With heterogeneity of $\bar{r}$ within this subsample, this average would represent the low-$\bar{r}$ conduits, so that the imputed value of *maxr* would be too low for the rest of the conduits. Normalized yields are defined in (9).

Table 3
Moments from SMM estimation.

This table shows the 13 moments used in SMM estimation. The first (last) four columns show moments for the weak- (strong-) guarantee subsample, which contains 90 ABCP conduits with extendible or SIV credit guarantees (191 conduits with full credit or full liquidity guarantees). Simulated moments are computed using the parameter estimates in Table 4. Moment 1 is the probability that a conduit recovers within eight weeks of a run's start. Moment 2 is the average number of days until the recovery, provided it occurs within eight weeks of the run's start. Moments 3 and 4 are the intercept and slope from a regression of absolute changes in yield spreads on the lagged yield spread. Moments 5–7 are the intercept and slopes from a regression of yield spreads on the number of weeks relative to a run and its exponential. Moments 8–13 come from three regressions, each of the indicator $\mathbf{1}_{\{run\ within\ \tau\ weeks\}}$ on the current yield spread, where $\tau=2$, 4, and 8 weeks. The moment condition t-statistic tests whether the actual and simulated moments are equal. The J-test is the χ^2 test for the model's overidentifying restrictions.

	Weak-guarantee subsample				Strong-guarantee subsample			
	Actual moments		Simulated moments estimate	Moments condition t-stat.	Actual moments		Simulated moments estimate	Moments condition t-stat.
	Estimate	Std. Err.			Estimate	Std. err.		
Moments on time between run and recovery (τ):								
1. $\Pr[\tau < 8\ \text{weeks}]$	0.451	0.054	0.513	−1.12	0.613	0.060	0.615	−0.04
2. $\mathrm{E}[\tau\mid\tau \le 8\ \text{weeks}]$ (in days)	17.1	1.4	15.4	1.16	17.4	0.7	17.4	0.04
Moments from regression of $\lvert r_{it+1} - r_{it}\rvert$ on r_{it}:								
3. Intercept	0.0011	0.0006	0.0004	1.21	0.0009	0.0004	0.0003	1.58
4. Slope	0.108	0.078	0.223	−1.48	0.128	0.060	0.205	−1.29
Moments describing yield spreads leading up to runs:								
Regression of r_{it} on $\tau[\equiv$ weeks relative to run] and $\exp(\tau)$								
5. Intercept	0.00538	0.00190	0.00248	1.52	0.00347	0.00130	0.00157	1.47
6. Slope on τ	0.00033	0.00012	0.00012	1.80	0.00015	0.00010	0.00007	0.85
7. Slope on $\exp(\tau)$	0.00168	0.00268	0.00467	−1.12	0.00525	0.00178	0.00260	1.49
Moments from regressions of $\mathbf{1}_{\{run\ within\ \tau\ weeks\}}$ on yield spreads								
8. Intercept ($\tau=2$)	−0.003	0.036	−0.146	3.94	−0.001	0.012	−0.158	9.64
9. Slope ($\tau=2$)	0.317	0.087	0.716	−4.54	0.104	0.040	0.755	−14.93
10. Intercept ($\tau=4$)	0.121	0.055	−0.044	2.97	0.016	0.013	0.000	0.93
11. Slope ($\tau=4$)	0.364	0.045	0.758	−8.46	0.155	0.041	0.711	−12.37
12. Intercept ($\tau=8$)	0.297	0.095	0.149	1.55	0.067	0.028	0.217	−4.89
13. Slope ($\tau=8$)	0.413	0.073	0.655	−3.24	0.165	0.034	0.561	−10.02
Test of over-identifying restrictions:								
J-test	1,111				6,771			
p-Value	0.000				0.000			

relatively high in Fig. 6. One potential solution to this problem would be to partition this subsample further using, say, proxies for conduits' $\bar{r}$. However, doing so would result in quite small subsamples.

We are not concerned about heterogeneity in the weak-guarantee subsample, because the model fits well both the yield spreads in event-time before runs (Fig. 4) and the run probabilities conditional on spreads (Fig. 5). In other words, that subsample's $\bar{r}$ estimate, which mainly relies on data from conduits that have runs, is consistent with the run probabilities for *all* conduits in the subsample.

Table 3 contains p-values for the SMM test of over-identifying restrictions, which jointly tests whether the model fits all moments. The low p-values indicate the data strongly reject the model in both subsamples. We do not interpret this result negatively, since rejection is common when trying to match many moments with few degrees of freedom. In this case, we attempt to match 13 moments with four free parameters, which is quite demanding.

To summarize, the model fits the data reasonably well in the weak-guarantee subsample. The model is able to match not just directional patterns in the data, but also magnitudes in most cases. Model fit is worse in the strong-guarantee subsample, potentially because of parameter

heterogeneity. For this reason, we only use estimates from the weak-guarantee subsample in our policy analysis.

5.1.2. Debt quantities

A key feature of the 2007 ABCP crisis is that the total quantity of ABCP contracted sharply. Covitz, Liang, and Suarez (2013) attribute 95% of the total contraction to declines in conduits that experienced runs. Next, we check whether our model can quantitatively match the contraction in ABCP observed in 2007. Since we do not use data on ABCP quantities to estimate the model, this exercise provides an out-of-sample check on the estimated model.

Fig. 7 compares the actual and predicted total ABCP outstanding during 2007. Let the dummy variable $\mathbf{1}_{\{run_{it}\}}$ denote whether conduit i is in a run in week t. We denote conduit i's ABCP held by the market in week t as D_{it}^{M} and the debt held by the sponsor as D_{it}^{S}. Total debt D in our model is the sum of ABCP and sponsor debt: $D = D^{M} + D^{S}$. We predict changes in D^{M} and D^{S} by decomposing dD_t from Eq. (3) into its two components, which yields

$$\frac{dD_{it}^{M}}{dt} = -\delta D_{it}^{M} + (1 - \mathbf{1}_{\{run_{it}\}})\delta D_{it}R_{it}, \qquad (10)$$
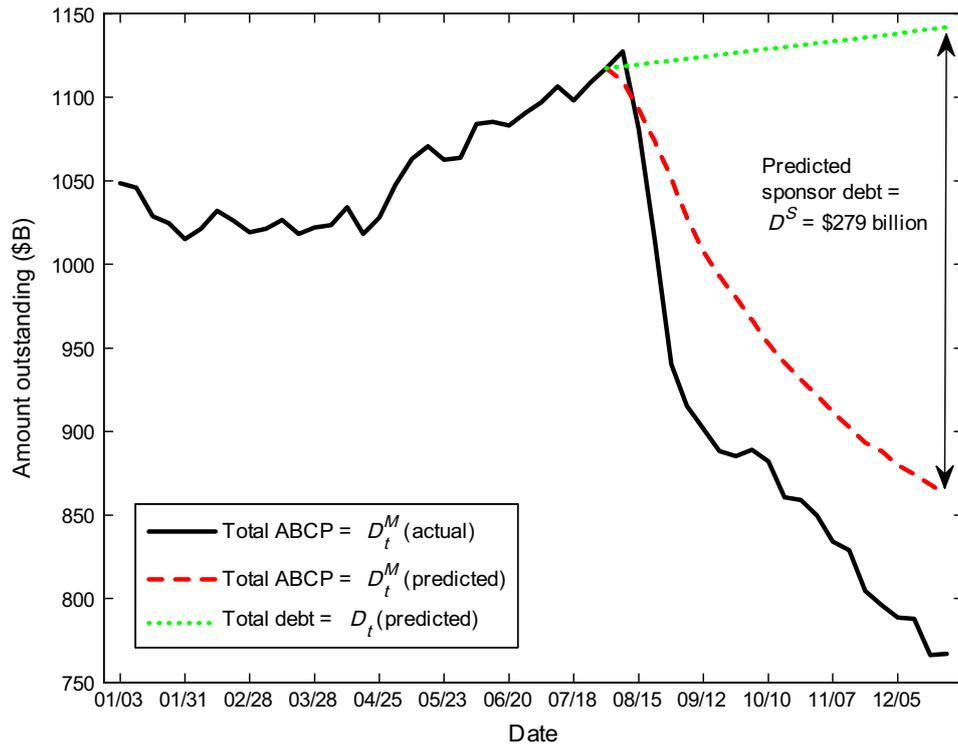
Fig. 7. This figure shows the total face value of debt outstanding during 2007. The solid line shows actual data for D_t^M, the amount of ABCP outstanding in our sample of 281 conduits. The dashed red line shows predicted values for D_t^M. The dotted green line shows the predicted value of total debt D_t, which is the sum of market ABCP debt, D^M, and debt held by the sponsor, D^S. The figure also notes that the predicted total face value of debt held by sponsors at the end of 2007 is $279 billion. We predict debt quantities after August 1, 2007 using Eqs. (10) and (11) with estimated parameters from Table 4, conduits' total face value of ABCP on August 1, and conduit/week data on rollover yields and run status throughout 2007. We assume sponsors held zero debt on August 1.

$$\frac{dD_{it}^S}{dt} = -\delta D_{it}^S + \mathbf{1}_{\{run_{it}\}}\delta D_{it}\overline{R}_i. \tag{11}$$

These equations indicate that a conduit pays off maturing debt by issuing ABCP with face value R_{it} to outside investors when not in a run, and by issuing debt with face value $\overline{R}_i$ to the sponsor during a run. Starting on August 1, 2007 when the ABCP crisis begins, we set $D_{it}^S = 0$, indicating that no backup credit lines had been used yet. We then use the equations above together with estimated parameters to forecast D_{it}^M and D_{it}^S at the conduit level. Since we cannot observe conduit-level leverage, we use data on conduits' actual rollover yield spreads r_{it} and run status $\mathbf{1}_{\{run_{it}\}}$ throughout 2007 to predict debt quantities. We aggregate across conduits to produce Fig. 7.

In both the actual data (solid black line) and predicted data (dashed red line), total ABCP declines sharply in August 2007 and continues declining through the rest of the year. The model can explain 73% of the observed decline in ABCP from August 1 to December 26, 2007. The fit is surprisingly good given that we do not use debt quantities to estimate the model. Predicted total conduit debt (dotted green line) increases steadily throughout 2007 as sponsors take increasing large stakes in their conduits to pay off maturing ABCP during runs. The model predicts that backup credit guarantees forced sponsors to bring $279 billion of conduit debt onto their balance sheets by the end of 2007. By straining the balance sheets of sponsoring banks in 2007, the ABCP crisis likely

contributed to the severe disruptions in the banking sector during the financial crisis.

5.2. Parameter estimates

Table 4 contains parameter estimates along with their standard errors.[31] We provide several consistency checks below to ensure that parameter estimates are reasonable.

The estimates of θ imply that investors expected conduits with strong (weak) guarantees to survive 262 (82) days in a run before the guarantee failed.[32] The estimate of θ is significantly higher ($t=2.04$) in the weak-guarantee subsample.[33] This result provides a useful consistency check: our parameter estimates imply weaker perceived guarantees in the subsample with weaker explicit guarantees.

The asset's estimated volatility (σ) is roughly 3.5% (4.3%) per year in the strong- (weak-) guarantee subsample. The

---

[31] Parameters' standard errors depend on the 13×13 covariance matrix for the empirical moments. When estimating this matrix, we take into account time-series autocorrelation as well cross-conduit correlation, both within moments and across moments. We also perform a two-step correction to account for measurement error in parameters ϕ and δ. Details are in Appendix C.

[32] Once a run starts, the average time until credit line failure is $1/(\theta\delta)$. As Section 4.2 explains, our estimate of $1/\delta$ is 37 days.

[33] The θ estimate in the strong-guarantee subsample could be too high, for the same reason that the subsample's estimate of $\overline{r}$ is potentially too low (Section 5.1.1). Adjusting this estimate would make the difference across subsamples even larger.

Table 4
Structural parameter estimates and predicted run thresholds.

The first two rows provide parameter estimates in the weak-guarantee subsample, which contains 90 ABCP conduits in 2007 with SIV or extendible credit guarantees. The last two rows report estimates in the strong-guarantee subsample, which contains 191 ABCP conduits in 2007 with full credit or full liquidity guarantees. Columns 2–5 report estimated structural parameters, with standard errors in parentheses. Estimation is done by SMM, which chooses parameter estimates that minimize the distance between actual and simulated moments. Section 2 describes the model used to simulate moments. Standard errors account for time-series and cross-sectional autocorrelation, both within and across the 13 moments used in estimation. Standard errors also include a two-step correction for measurement errors in parameters δ (inverse average debt maturity) and ϕ (inverse asset maturity). The last column shows the leverage threshold for runs ($1/x^*$) predicted by the model for the given parameter estimates; the model predicts that runs occur as soon as leverage exceeds this threshold.

| Subsample | Parameter estimates | | | | Predicted |
	Weakness of credit guarantee θ	Asset volatility (% per year) σ	Cap on yield spreads (b.p. per year) $\bar{r}$	Asset liquidity α	leverage threshold for runs (%) $1/x^*$
Weak guarantees	0.449 (0.133)	4.30 (0.10)	59.8 (6.7)	0.920 (0.032)	92.0
Strong guarantees	0.141 (0.045)	3.54 (0.07)	36.0 (7.0)	0.968 (0.045)	97.1

difference across subsamples is statistically significant ($t=6.3$), possibly because conduits offering weak guarantees had incentives to hold riskier assets. As a comparison to our estimates, the volatility of the ABX mortgage index in the first half of 2007 was 5.7%. The other asset categories ABCP conduits hold, such as trade receivables, are likely less volatile than the ABX was in 2007 (as suggested by Fig. 1), so an estimate for σ slightly below 5.7% seems reasonable.

The estimated cap on yield spreads, $\bar{r}$, is 36 (60) b.p. per year for the strong- (weak-) guarantee subsample. This difference in caps across subsamples is statistically significant ($t=2.7$). The reason we find a higher cap in the weak-guarantee subsample is that spreads in that subsample are higher in most of the weeks leading up to a run (Fig. 4).[34] As we explain in Section 5.1, the $\bar{r}$ estimate in the strong-guarantee sample could be too low, so it remains unclear whether $\bar{r}$ is truly lower when guarantees are stronger. Moreover, it is unclear a priori whether strong-guarantee conduits should have a lower or higher $\bar{r}$ than weak-guarantee conduits.[35]

Our estimate of α, the asset's liquidity or recovery rate in default, is 97% (92%) in the strong- (weak-) guarantee subsample. The difference across subsamples is not statistically significant. For comparison, Coval and Stafford (2007) find a 92% recovery rate for stocks during fire sales. Ellul, Jotikasthira, and Lundblad (2011) report a 93% recovery rate for corporate bonds during fire sales. For corporate defaults, Hennessy and Whited (2007) measure recovery rates of around 90%, and Andrade and Kaplan (1998) estimate recovery rates of roughly 80–90%. Our estimates are on the high end of the literature's range, which makes sense given that conduits mainly owned financial rather than real assets.

The last column in Table 4 shows the model's predicted run threshold implied by our parameter estimates. The model predicts that lenders run as soon as leverage exceeds 92% (97%) for the SIV/extendible (full credit/liquidity) subsample. As a consistency check, we compare these predicted leverage thresholds to actual data on conduit leverage, which were not used in estimation. Detailed data on conduit leverage at the time of the crisis were not available even to regulators, but from Moody's reports we can make rough estimates of leverage for 11 SIVs. These conduits' leverage ranged from 92% to 94% preceding runs in 2007,[36] which is reassuringly close to the predicted 92% threshold.

---

[34] As we explain in Section 4.3, the model uses spreads from the 12 weeks before each run to infer the spread the instant a run starts, which equals $\bar{r}$. Fig. 4 shows that the average spread in week –1 is roughly the same across subsamples. For strong-guarantee conduits, however, the spread is much lower in weeks –8 to –2, and the spread increases more convexly leading up to runs. The estimation procedure takes into account that each week's average spread is estimated with error, and that the difference in convexity across subsamples is fairly weak statistically. The lower spreads in weeks –8 to –2 dominate the higher convexity and high spread of week –1 to produce a lower estimate of $\bar{r}$ in the strong-guarantee subsample.

[35] On one hand, a weak-guarantee conduit faces a higher probability of liquidation given a run, so its expected liquidation costs given a run are higher than those of a strong-guarantee conduit. Outside the model, these relatively higher ex post costs could lead the conduit to choose a relatively higher $\bar{r}$ in order to reduce its run probability. On the other hand, a weak-guarantee sponsor expects to provide a credit line for less time once a run begins, because the credit line is expected to fail sooner. Since the sponsor lends at a below-market yield spread $\bar{r}$ during a run, it loses money while providing the credit line. Since the weak-guarantee sponsor expects to lose money for a shorter period than the strong-

(footnote continued)

guarantee sponsor during a run, it could choose a lower $\bar{r}$ and hence a higher run probability.

[36] Moody's Investors Service (2008) reports that SIVs' net asset value of capital (NAV), defined as the difference between the market value of portfolio assets minus the notional amount of senior liabilities expressed as a percentage of paid-in capital, averaged 102.5% in January 2007. NAV maps into $(F(y_t)-D_t)/K_t$ in our model, where K_t is paid-in capital. Based on individual program reports published by Moody's available for 11 SIVs, we estimate that the ratio of capital notes to senior liabilities, which maps into K_t/D_t, ranged from 6% to 9% across SIVs in December 2006. Since $F(y_t)=y_t$ when $\mu=\rho$, we obtain the formula

$$\frac{y_t}{D_t} = NAV_t \times \left(\frac{\text{capital notes}}{\text{senior liabilities}}\right)_t + 1.$$

This formula, combined with the data from Moody's reports, delivers the 92–94% range for leverage (D_t/y_t).
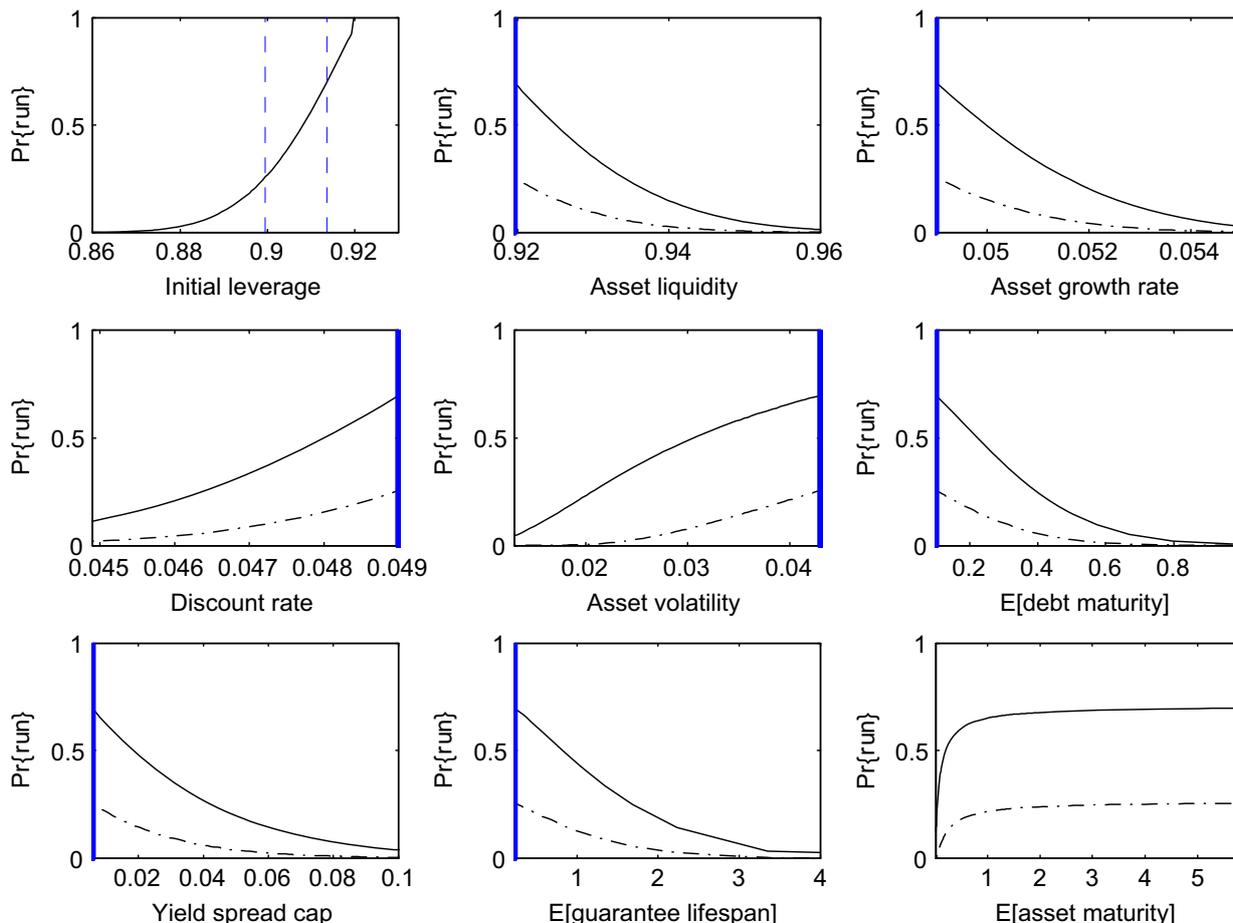
Fig. 8. This figure shows the sensitivity of run probabilities to the model's parameters: initial conduit leverage ($1/x_0$), asset liquidity (α), asset growth rate (μ), discount rate (ρ), asset volatility (σ), expected debt maturity ($1/\delta$), yield spread cap ($\bar{r}$), expected guarantee lifespan ($1/(\delta\theta)$), and expected asset maturity ($1/\phi$). Expected guarantee lifespan and debt maturity are in years. Asset growth rate and volatility, the discount rate, and the cap on yield spread are in units of fraction per year. Each panel plots the simulated three-month run probability as a function of parameter values. In each panel we set parameter values to their estimates for the weak-guarantee subsample (Table 4), then we vary the parameter indicated on the panel's horizontal axis. Initial leverage (top left panel) serves as each simulation's initial condition. The two vertical dashed blue lines in the top-left panel indicate the values we use as initial conditions in the remaining panels. The bold blue lines on the vertical axes indicate the parameters' estimated value. The solid (dashed-dot) line shows run probabilities when initial leverage is at the high (low) level indicated in the top-left panel.

6. What drives runs?

All the parameters estimated above potentially affect the likelihood of runs. In this section we use the estimated model to measure the sensitivity of run probabilities to each of these contributing factors. We also measure runs' sensitivity to changes in asset values. In the next section we discuss implications for regulators, banks, and investors.

6.1. Sensitivity of runs to model parameters

We measure the sensitivity of runs to the model's eight parameters by computing run probabilities for a range of counterfactual parameter values. Fig. 8 plots simulated three-month run probabilities.[37] In each panel we vary one parameter at a time. The bold blue line on the vertical axis marks the parameter's estimated value.

The top-left panel shows that run probabilities are extremely sensitive to conduits' leverage ($1/x_0$).[38] For instance, reducing leverage by 1% from 91.4% to 90.4% reduces the run probability by 32 percentage points from 70%, i.e., roughly a 45% decrease. Not surprisingly, the model also predicts that yield spreads are highly sensitive to leverage. For example, changing leverage from 90.4% to 92.0% is enough to increase yield spreads from ten b.p. to their estimated capped value, 59.8 b.p.

Each of the remaining panels in Fig. 8 shows results for a high-leverage (solid line) and low-leverage (dash-dot line) scenario.[39] The dashed lines in the top-left panel mark these two scenarios. The high-leverage scenario

---

[37] We tabulate these results and present one-year run probabilities in the Internet Appendix.

[38] Recall that initial leverage is not a parameter we estimate, but instead serves as the initial condition for these counterfactual simulations.

[39] High (low) initial leverage is 91.4% (90.0%), which produces a 70% (26%) three-month run probability. We choose these values because they imply that yield spreads are at 50% (10%) of their capped value, $\bar{r}$, for the high- (low-) leverage case.

Table 5
The change in asset value required to trigger a run.

  This table shows simulated changes in asset value preceding runs. We simulate conduits for one year using parameter estimates for the weak-guarantee subsample (Table 4). For these parameter values, runs occur as soon as leverage exceeds 0.92. Panel A contains the simulations' assumed initial leverage, which equals the asset's fundamental value y_0 divided by the initial face value of debt D_0. Panel B shows statistics on the percentage change in asset value y_t between the simulation's start and the run's start, for those conduits that experience a run.

Panel A: Initial leverage (D_0/y_0)						
0.85	0.86	0.87	0.88	0.89	0.9	0.91
Panel B: Percent change in asset value preceding a run						
Mean						
−4.39	−3.53	−2.69	−1.93	−1.26	−0.68	−0.26
25th percentile						
−4.97	−4.22	−3.53	−2.85	−2.24	−1.62	−0.95
Median						
−4.28	−3.47	−2.70	−2.04	−1.50	−1.05	−0.67
75th percentile						
−3.70	−2.76	−1.85	−1.04	−0.42	0.00	0.02

represents the late, severe stage of a funding crisis, while the low-leverage scenario represents an earlier, less severe stage. Since results are qualitatively similar between these two scenarios, we focus on just the high-leverage scenario for the discussion below.

The top-middle panel shows that runs are extremely sensitive to small changes in asset liquidity (α) or, equivalently, recovery rates in default. For example, increasing α by 1% from its estimated value of 0.92 reduces the run probability by 31 percentage points from 70%, roughly a 45% decrease. Intuitively, a creditor runs because she fears there will not be enough of the asset to pay her if the conduit defaults. A higher α increases creditors' recovery rate in default, reducing the incentive to run. The impact of liquidity is even more striking given that a 1% change in α amounts to an 11.5% ($1\% \times 0.92/(1-0.92)$) change in the illiquidity discount, which is well within the estimated discount variation in the literature. For example, Coval and Stafford (2007) report a standard deviation of 9.72% in the fire-sale discount of stocks. The estimates in Ellul, Jotikasthira, and Lundblad (2011) imply a standard deviation of almost 25% for fire-sale discounts on corporate bonds.

The asset's growth rate μ and the risk-free rate ρ have a smaller but still significant effect on run probabilities. Increasing μ by 1% or reducing ρ by 1% from their base-case values (both 4.9%) reduces the run probability by ten percentage points, roughly a 14.5% decrease. Increasing μ or decreasing ρ increases the asset's value, which effectively reduces the conduit's leverage. Reducing the risk-free rate also makes the conduit's leverage increase more slowly.

The remaining panels show that every parameter affects the likelihood of a run and, furthermore, is capable of bringing the run probability close to zero. However, we need larger changes in the remaining parameters to achieve large reductions in run probabilities. Reducing the probability by 31 percentage points, as achieved by increasing α by 1%, requires either reducing asset volatility σ from 4.3% to 2.6% per year (a 40% decrease); increasing average debt maturity $1/\delta$ from 37 to 107 days (a 190% increase); increasing the cap on yield spreads $\bar{r}$ from 0.6% to 2.7% per year (a 358% increase); increasing $1/(\theta \delta)$ (the

average time the guarantee survives during a run) from 82 to 422 days (a 413% increase); or reducing the asset's expected maturity $1/\phi$ from 5.8 years to 39 days (a 98% decrease). As expected, changes that reduce the mismatch between asset and debt maturity reduce the likelihood of runs. Increasing the yield cap makes runs less likely by allowing more room for yields to adjust upwards as conditions worsen; this effect outweighs the increased dilution risk that comes with a higher yield cap.

6.2. Sensitivity of runs to asset values

One way to measure the fragility of ABCP conduits in 2007 is to ask, how much did asset values need to drop to trigger a run? The answer depends on how levered conduits were before the crisis began. Since these data are not available, we answer the question for several starting leverage values. We simulate many conduits for one year with a given initial leverage value, using estimated parameter values from the weak-guarantee subsample (Table 4).

Table 5 reports statistics on the percentage change in asset values preceding runs, for several initial leverage values. The table shows that for a conduit with 85% leverage, asset values need to drop 4.39% to trigger a run, on average. For comparison, our estimated asset volatility is 4.30% per year. As expected, conduits with higher initial leverage are more fragile, in the sense that a smaller drop in asset value will trigger a run. For example, if leverage starts at 90%, asset values need drop only 0.68% to trigger a run, on average.

7. Policy discussion

There are several reasons why regulators may want to prevent runs. Runs on financial institutions could disrupt the flow of credit to nonfinancial firms that rely on intermediated finance to fund investment and operations and, thus, ultimately harm economic activity. The view that bank runs hamper economic activity is supported by evidence from banking crises in the United States (Friedman and Schwartz, 1963; Bernanke, 1983; Calomiris and Mason, 2003; Ramirez and Shively, 2012) and cross-country studies (Kaminsky and
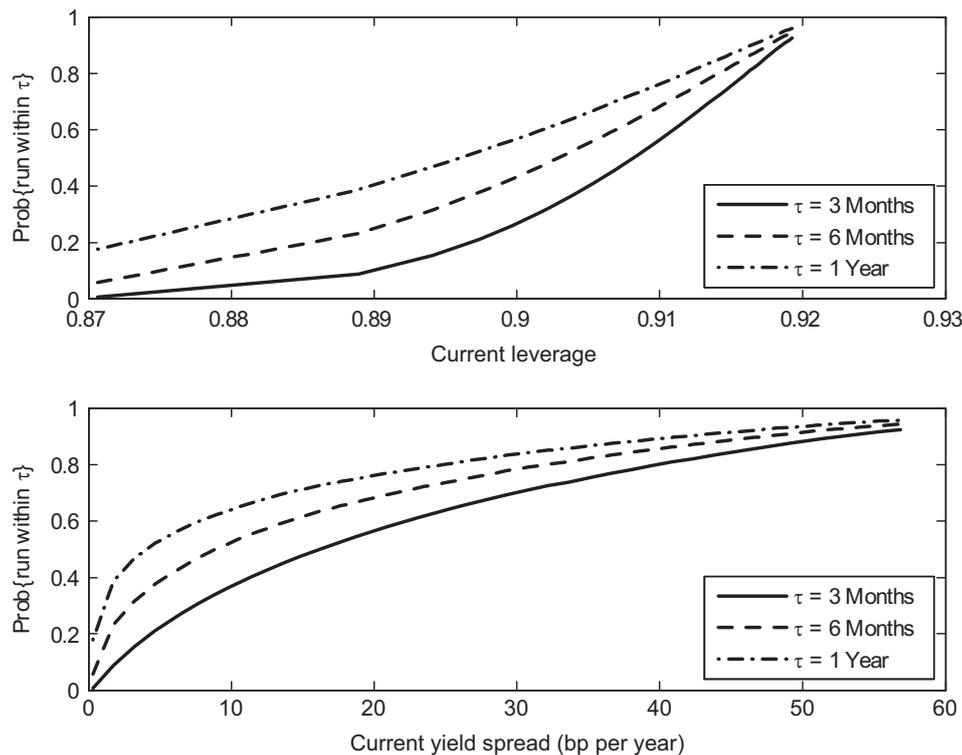
Fig. 9. The top panel plots the relation between the firm's current leverage and the probability of a run within the next three, six, and 12 months. The bottom panel shows the relation between the firm's current yield spread (basis points per year, bp/year) and the probability of a future run. Results are from model simulations using the parameter estimates for the weak-guarantee subsample, which contains the 90 ABCP conduits with SIV or extendible guarantees (Table 4).

Reinhart, 1999; Dell'Ariccia, Detragiache, and Rajan, 2008). Also, a run on one part of the financial system could trigger runs on other parts, amplifying the run's costs. For example, a run on ABCP could trigger a run on money market funds (the main investors in ABCP) or a run on the large banks sponsoring ABCP conduits.

Before discussing how regulators might prevent runs, we describe the warning signs that regulators, banks, and investors can use to gauge the probability of a future run. According to the model, the warning signs are high conduit leverage and high rollover yields. Furthermore, the model provides a quantitative mapping between these warning signs and the probability of runs at various horizons (Fig. 9). While conduit leverage (top panel) is a useful warning sign for a conduit's sponsor, it is less useful to regulators or investors, who currently cannot observe leverage. The bottom panel, which maps yield spreads into run probabilities, is useful to all parties.

In the previous section we measure how run probabilities change when we perturb each model parameter away from its estimated value. As in Rochet and Vives (2004), we interpret these perturbations as interventions by regulators or conduit sponsors. The estimated sensitivities are important for regulators interested in controlling the risk of runs on ABCP conduits and similar intermediaries. These sensitivities can also help sponsoring banks control the risk of runs when managing existing conduits or designing new ones.

Our analysis shows that runs are very sensitive to small changes in leverage. This result implies that conduits' sponsors can significantly reduce the probability of future

runs by including more equity in new conduits' capital structure. Regulators can achieve the same effect by placing restrictions on new conduits' leverage. Our result also suggests that once a crisis is underway, modest equity injections by either conduit sponsors or governments can make runs significantly less likely. For example, the average ABCP conduit with SIV or extendible guarantees, which held $3.9 billion in assets in January 2007, would have cut its three-month run probability by more than half with a $39 million equity injection. Such a recapitalization, which would cost just over $4 billion for all weak-guarantee ABCP conduits, would imply an expected dead-weight cost savings of $97 million per program and almost $11 billion overall.[40] Of course, a full welfare analysis would have to also take into account the positive externalities from preventing runs (e.g., allowing the sponsoring banks to keep lending to the nonfinancial sector), negative externalities (e.g., moral hazard and tax distortions), and ex post wealth transfers to the party buying discounted assets in a liquidation.

The previous analysis also shows that runs are very sensitive to α, the asset's liquidity, i.e., expected recovery rate in default. For example, increasing recovery rates by 1%, which translates into subsidizing asset values in default by $39 million for the average conduit in 2007, would have reduced the three-month run probability by

---

[40] The estimated deadweight cost savings equals the change in the probability of a run times total ABCP NAV times the liquidation discount, $1 - \alpha$.

almost a half. It is less clear how governments or sponsoring banks can improve liquidity. One possibility is that governments make a market in distressed assets or purchase them outright.

We interpret reducing the asset's volatility (σ) or increasing its growth rate (μ) as buying higher quality assets. Sponsors can clearly influence asset quality when creating new conduits, but probably not once a crisis is underway. Regulators could influence asset quality by placing credit rating restrictions on the assets conduits buy, similar to the restrictions on money market funds. However, our results show that an effective control of the run probabilities would require large changes in μ and especially σ.

Fig. 8 suggests several remaining channels for preventing runs. Increasing ϕ corresponds to the conduit buying shorter-term assets, and decreasing δ corresponds to increasing debt maturities; both reduce the degree of maturity mismatch between assets and liabilities. Reducing θ corresponds to strengthening the conduit's guarantee, which, to be credible, could not only require strengthening its legal terms, but also improving the sponsor's own financial health. Reducing ρ corresponds to reducing the federal funds rate. Our sensitivity analysis implies that interventions targeting these channels (or that increase $\bar{r}$) will have small effects on the likelihood of runs, unless the interventions can change these parameters by a large amount. For example, to match the effectiveness of a 1% equity injection in January 2007, the sponsor of an average weak-guarantee conduit would have to commit an additional $826 million to the conduit's backup credit line.[41]

There are caveats to our policy analysis. First, our analysis does not consider how policy interventions directed at ABCP conduits could spill over to other markets. To wit, an increase in the rollover yield caps ($\bar{r}$) could delay runs on ABCP conduits at the expense of weakening their sponsor's balance sheet or making money market funds take on more risk. Second, we do not address how policy interventions may affect future crises via moral hazard. Third, our sensitivity analysis is subject to a Lucas-type critique, as it does not consider how changing one parameter could affect other parameters. For example, Cheng and Milbradt (2012) endogenize the choice of the asset's growth rate and volatility of a short-term financed firm, such as an ABCP conduit, and find that a lengthening of debt maturities could lead to more risk-shifting. A comprehensive policy analysis would need to incorporate the reaction of ABCP conduit managers and investors to any interventions. We leave this analysis to future research.

---

[41] The expected payout to creditors during a run is

$$\mathrm{E}\left[\int_0^{\tau_\theta} \delta D_* e^{-\delta t}\, dt\right] = \frac{D_*}{1+\theta},$$

where D_* is the total debt outstanding when the run starts. We estimate D_* as using the NAV per conduit in January 2007 divided by $x^* = 1/0.92$. The additional capital required to sustain a stronger credit line is therefore

$$D_* \times \left(\frac{1}{1+\theta_1} - \frac{1}{1+\theta_0}\right),$$

where $\theta_1 = 0.087$ is the counterfactual value and $\theta_0 = 0.449$ is the estimated value.

8. Conclusions

We estimate a dynamic model of debt runs using data from the 2007 crisis in asset-backed commercial paper. The model allows yields to change over time, which introduces dilution risk: the conduit must offer higher yields to induce rollover if conditions worsen, which dilutes the claims of other lenders. Introducing dilution risk into the model can increase the likelihood of a run by more than an order of magnitude. Our model of fundamental-driven runs fits several features of the data, including the sharp contraction in ABCP outstanding in late 2007, the increase in ABCP yields leading up to runs, the high probability of recovery once a run starts, the positive relation between yields and the probability of future runs, the overall level of volatility in ABCP yields, and the positive relation between yield volatility and the yield level. The model fits much better in the subsample of conduits with the weakest guarantees. We find that runs are very sensitive to conduit leverage and expected asset liquidation costs. Given that leverage plays such an important role in our model, it is surprising that regulators and investors do not have systematic access to data on conduits' leverage. We find that runs are much less sensitive to the degree of maturity mismatch, the perceived strength of guarantees, and the asset's volatility. These results can help regulators and banks control the risk of runs.

Our analysis can be extended and improved in three directions. To keep the estimation tractable, we assume that yields cannot exceed an exogenous cap. It would be interesting to explore the determinants of this cap. Second, future research could ask how coordination failures affect yields, and how a single creditor might buy and renegotiate dispersed debt during a crisis. Finally, the dynamic debt runs framework, and estimation method we propose here, can potentially be used to study runs on money market funds and sovereign debt.

Appendix A. Model solution

As in He and Xiong (2012a), we focus on symmetric monotone rational expectations equilibria, in which each creditor is best-responding to all others' decision to run if and only if the fundamental asset value drops below a common threshold, y^*. To solve for our model's threshold, we show first that the creditor's value function depends only on one state variable: the conduit's (inverse) leverage, x_t, i.e., the ratio of asset value, y_t, to total debt, D_t. We start by characterizing the dynamics of the conduit's debt, then of x_t, and then solve for the threshold x^*.

A.1. Debt dynamics

Since all debt is equally likely to roll over in the next instant, regardless of when and at what yield it was originated, then the total face value of paper outstanding at t, D_t, equals the average face value of debt rolling over at time t. Debt dynamics follow (3) since a fraction $\delta\, dt$ of debt matures at each instant, and for every dollar of face value that is rolled over, the conduit issues new debt at face value R_t.

A.2. Value function

At any date τ, there are three possible payouts for a lender who holds debt with face value R_s where $s \leq \tau$:

1. The conduit matures at time $\tau = \tau_\phi$ so that the creditor is either paid in full or gets a share of the assets proportional to his face value, i.e.,

$$\frac{R_s}{D_{\tau_\phi}} \times \min\left(D_{\tau_\phi}, y_{\tau_\phi}\right) = R_s \min\left(1, \frac{y_{\tau_\phi}}{D_{\tau_\phi}}\right). \tag{12}$$

2. The conduit defaults at time $\tau = \tau_\theta$ after other creditors run and backup credit lines fail. The creditor recovers a share of the post-liquidation net present value of the asset, i.e.,

$$\frac{R_s}{D_{\tau_\theta}} \min(D_{\tau_\theta}, ly_{\tau_\theta}) = R_s \min\left(1, l\frac{y_{\tau_\theta}}{D_{\tau_\theta}}\right), \tag{13}$$

where

$$ly_{\tau_\theta} \equiv \alpha \frac{\phi}{\rho + \phi - \mu} y_{\tau_\theta}. \tag{14}$$

3. The debt contract matures at time $\tau = \tau_\delta$, allowing the creditor to choose between rolling over or running. Because the amount of debt maturing at each instant is infinitesimally small, a creditor that chooses to run will be paid off in full. If the creditor rolls over, the old loan is retired and a new loan is issued with face value R_{τ_δ}. Let $V(y_\tau, D_\tau, R_s; y^*)$ be the value in time τ of one dollar loaned at time $s \leq \tau$. The lenders' payoff in τ_δ is therefore

$$\max_{\text{roll over or run}} \{R_s V(y_{\tau_\delta}, D_{\tau_\delta}, R_{\tau_\delta}; y^*), R_s\}$$
$$= R_s \max_{\text{roll over or run}} \{V(\cdot), 1\}. \tag{15}$$

Combining these three possible payoffs, the time t value to a creditor who last loaned one dollar at time $s \leq t$ equals

$$V(y_t, D_t, R_s; y^*) = E_t\left\{e^{-\rho(\tau - t)}R_s \min\left(1, \frac{y_\tau}{D_\tau}\right)\mathbf{1}_{\{\tau = \tau_\phi\}}\right\}$$
$$+ E_t\left\{e^{-\rho(\tau - t)}R_s \min\left(1, l\frac{y_\tau}{D_\tau}\right)\mathbf{1}_{\{\tau = \tau_\theta\}}\right\}$$
$$+ E_t\left\{e^{-\rho(\tau - t)}R_s \max_{\text{rollover or run}} \{V(y_{\tau_\delta}, D_{\tau_\delta}, R_{\tau_\delta}; y^*), 1\}\mathbf{1}_{\{\tau = \tau_\delta\}}\right\}. \tag{16}$$

For $x_t \equiv y_t / D_t$, Eq. (16) simplifies to

$$V(y_t, D_t, R_s; y^*) = R_s W(x_t; x^*), \tag{17}$$

where

$$W(x_t; x^*) = E_t\{e^{-\rho(\tau - t)} \min(1, x_\tau)\mathbf{1}_{\{\tau = \tau_\phi\}}\}$$
$$+ E_t\{e^{-\rho(\tau - t)} \min(1, lx_\tau)\mathbf{1}_{\{\tau = \tau_\theta\}}\}$$
$$+ E_t\left\{e^{-\rho(\tau - t)} \max_{\text{rollover or run}} \{R_\tau W(x_\tau; x^*), 1\}\mathbf{1}_{\{\tau = \tau_\delta\}}\right\}. \tag{18}$$

$W(x_t, x^*)$ is the value at time t of each dollar of face value. This value does not depend on when the creditor last rolled over, due to the memory-less properties of the exponential distribution. Moreover, applying Ito's Lemma and Eq. (3), it is straightforward to show that inverse leverage follows:

$$\frac{dx_t}{x_t} = \mu\, dt + \sigma\, dZ_t + \delta\, dt - \delta R_t\, dt. \tag{19}$$

In other words, the fraction change in inverse leverage equals the fraction change in the asset's value ($\mu\, dt + \sigma\, dZ_t$) plus the fraction of debt maturing ($\delta\, dt$) minus the fractional amount of new debt issued ($\delta R_t\, dt$). Since the value function (18) and the dynamics of x_t are both functions of x_t only, x_t is the only state variable of the problem.

A.3. Equilibrium debt prices and run threshold

Creditors break even if for every \$1 invested in the conduit at time t, they receive a loan worth \$1. Formally, creditors break even if

$$1 = R_t W(x_t; x^*). \tag{20}$$

Since face values cannot exceed the cap, $\bar{R}$, the rollover face value is

$$R_t = \min[\bar{R}, W(x_t; x^*)^{-1}]. \tag{21}$$

The following proposition states that runs will not occur at face values R_t below the cap $\bar{R}$ but only when R_t exactly hits the cap. That is, the equilibrium condition that defines x^* is

$$\bar{R} = W(x^*; x^*)^{-1}. \tag{22}$$

Proposition 1. *Let* $R_t \equiv \min[\bar{R}, W(x_t; x^*)^{-1}]$. *Then*

$$R_t = \begin{cases} W(x_t; x^*)^{-1} & \text{if } x_t > x^*, \\ \bar{R} = W(x_t; x^*)^{-1} & \text{if } x_t = x^*, \\ \bar{R} & \text{if } x_t < x^*. \end{cases} \tag{23}$$

Proof of Proposition 1. Note first that any creditor's continuation payoff must be equal to one. By definition, for any x_t, the payoffs are

$$\max_{\text{run or roll over}} \{1, R_t W(x_t, x^*)\}$$
$$= \max_{\text{run or roll over}} \{1, \min[\bar{R}, W(x_t, x^*)^{-1}]W(x_t, x^*)\}$$
$$= \max_{\text{run or roll over}} \{1, \min[\bar{R}W(x_t, x^*), 1]\} = 1. \tag{24}$$

First, we show $R_t = \bar{R}$ if $x_t < x^*$. If $x_t < x^*$, creditors will refuse to roll over their loan at maturity. Because running gives them a payoff of one, rolling over must give them a strictly lower payoff, i.e., $R_t W(x_t, x^*) < 1$. By definition of R_t, this inequality becomes

$$\min[\bar{R}, W(x_t, x^*)^{-1}] \times W(x_t, x^*) < 1. \tag{25}$$

Since $W(x_t, x^*)^{-1} \times W(x_t, x^*) = 1$, it must be that min $[\bar{R}, W(x_t, x^*)^{-1}] = \bar{R}$. Therefore, $R_t = \bar{R}$.

Suppose that $x_t \geq x^*$. In this case, creditors choose to roll over. If they do so, their payoff must be at least as high as running, which pays one. Because their payoffs are

bounded above by one, then rolling over must always pay one. Therefore, for $x_t \geq x^*$

$$\min[\overline{R}, W(x_t, x^*)^{-1}] \times W(x_t, x^*) = 1$$
$$\Rightarrow \min[\overline{R}W(x_t, x^*), 1] = 1. \tag{26}$$

The previous equality holds if either $\overline{R}W(x_t, x^*) > 1$ for every $x \geq x^*$ or if there exist some $x' \in [x^*, \infty)$ where $\overline{R}W(x', x^*) = 1$ and $\overline{R}W(x_t, x^*) > 1$ for all other $x_t \neq x'$. Because $W(x, x^*)$ is strictly increasing in x, then x' is unique. Moreover, because $\overline{R}W(x', x^*) = 1$ is a minimum, then $x' = x^*$, i.e., the lowest point in the support. In summary, then either

$$R_t = \begin{cases} W(x_t, x^*)^{-1} > \overline{R} & \text{for all } x_t \geq x^*, \\ \overline{R} & \text{if } x_t < x^*. \end{cases} \quad [\text{case(i)}] \tag{27}$$

or

$$R_t = \begin{cases} W(x_t, x^*)^{-1} & \text{if } x_t > x^*, \\ \overline{R} & \text{if } x_t = x^*, \quad [\text{case (ii)}]. \\ \overline{R} & \text{if } x_t < x^*. \end{cases} \tag{28}$$

Next we show that case (i) cannot be true, because it implies a contradiction. In case (i) we have

$$R^* \equiv W(x^*, x^*)^{-1} < \overline{R}, \tag{29}$$

exactly at the run boundary. Hence, we have

$$1 = R^*W(x^*, x^*) < \overline{R}W(x^*, x^*). \tag{30}$$

The equality above is from the definition of R^*, and the inequality is from $W > 0$ and $R^* < \overline{R}$. By the assumed continuity of $W(x, x^*)$ at $x = x^*$, there exists a $\xi > 0$ such that for all $x' \in (x^* - \xi, x^*)$, $\overline{R}W(x', x^*) > 1$. We therefore have a contradiction: At $x' < x^*$, the investor runs (since we assume runs happen below x^*), but at x' it is not optimal to run (since $\overline{R}W(x^*, x^*)$, the payoff from rolling over at $R_t = \overline{R}$, is strictly greater than one, the payoff from running). $\square$

A.4. Analytical solution to the ODE for $W(x, x^*)$ below the run threshold

Using Eqs. (18) and (19), we can write the general Hamiltonian–Jacobi–Bellman (HJB) equation:

$$\rho W(x_t; x^*) = [\mu - \delta(R_t - 1)]x_t W_x(\cdot) + \frac{\sigma^2}{2}x_t^2 W_{xx}(\cdot)$$
$$+ \phi[\min(1, x_t) - W(\cdot)]$$
$$+ \theta\delta \mathbf{1}_{\{x_t < x^*\}}[\min(1, lx_t) - W(\cdot)]$$
$$+ \delta\left[\max_{\text{rollover or run}}\{R_t W(x_t; x^*), 1\} - W(\cdot)\right]. \tag{31}$$

For a given threshold x^*, the HJB equation can be solved analytically for $x_t < x^* \Longleftrightarrow R_t = \overline{R} < W(x_t, x^*)^{-1}$. The general solution to this Ordinary Differential Equation (ODE) is

$$W(x, x^*) = d_1 x^\eta + d_2 x^{-\gamma} - \frac{a_5}{a_3} - \frac{a_4}{a_3 + a_1}x, \tag{32}$$

for

$$\eta \equiv \frac{1}{2a_2}\left(a_2 - a_1 + \sqrt{(a_2 - a_1)^2 - 4a_3a_2}\right) > 1,$$
$$-\gamma \equiv \frac{1}{2a_2}\left(a_2 - a_1 - \sqrt{(a_2 - a_1)^2 - 4a_3a_2}\right) < -1,$$
$$a_1 = (\mu + \delta - \delta\overline{R}),$$

$$a_2 = \frac{\sigma^2}{2} > 0,$$
$$a_3 = -(\phi + \rho + \theta\delta + \delta) < 0,$$
$$a_4 = \theta\delta l \mathbf{1}_{\{x \leq 1/l\}} + \phi \mathbf{1}_{\{x \leq 1\}} \geq 0,$$
$$a_5 = \delta + \theta\delta l \mathbf{1}_{\{x \geq 1/l\}} + \phi \mathbf{1}_{\{x \geq 1\}} > 0;$$

with limits

$$W_x(0, x^*) = \frac{\theta\delta\left(\frac{\alpha\phi}{\phi + \rho - \mu}\right) + \phi}{\delta\overline{R} + (\phi + \rho - \mu) + \theta\delta} > 0, \tag{33}$$

$$W(0, x^*) = \frac{\delta}{\phi + \rho + \theta\delta + \delta} > 0, \tag{34}$$

$$\lim_{x \to \infty} W(x, x^*) = \frac{\phi + \delta}{\phi + \delta + \rho}, \tag{35}$$

and where the values of d_1 and d_2 are obtained by imposing smooth-pasting and value-matching at $x = 1, x = 1/l$, and on x^*.

Case 1: For $0 < x^* \leq 1$

The solution is

$$W(x, x^*) = A_1 x^\eta - \frac{a_5}{a_3} - \frac{a_4}{a_3 + a_1}x \quad \text{for } x \leq x^*,$$
$$A_1 = \left[\frac{1}{\overline{R}} + \frac{a_5}{a_3}\right](x^*)^{-\eta} + \frac{a_4}{a_3 + a_1}(x^*)^{1-\eta}. \tag{36}$$

Case 2: For $1 < x^* \leq 1/l$

The solution is

$$W(x, x^*) = \begin{cases} A_2 x^\eta - \frac{a_5}{a_3} - \frac{a_4}{a_3 + a_1}x & \text{for } x \leq 1, \\ B_1 x^\eta + B_2 x^{-\gamma} - \frac{b_5}{a_3} - \frac{b_4}{a_3 + a_1}x & \text{for } 1 < x \leq x^*, \end{cases}$$

$$A_2 = \left(\frac{1}{\overline{R}} + \frac{b_5}{a_3}\right)(x^*)^{-\eta} + \frac{b_4}{a_3 + a_1}(x^*)^{1-\eta} - B_2(x^*)^{-\gamma-\eta}$$
$$- \frac{\phi}{\gamma + \eta}\left[\frac{\gamma}{a_3} - \frac{\gamma + 1}{a_3 + a_1}\right],$$
$$B_2 = \frac{\phi}{\gamma + \eta}\left[\frac{(1 - \eta)}{a_3 + a_1} + \frac{\eta}{a_3}\right],$$
$$B_1 = A_2 + \frac{\phi}{\gamma + \eta}\left[\frac{\gamma}{a_3} - \frac{\gamma + 1}{a_3 + a_1}\right]. \tag{37}$$

Case 3: Can $x^* > 1/l$?

For some parameter values, case 2 implies that creditors run on a solvent conduit. That is, if $F(y_t) > D_t$ (which implies $x > 1$), then the conduit is able to repay all the debt if the asset is not liquidated at a discount. But creditors will not run frantically, i.e., if the post-liquidation value of the asset $\alpha F(y_t)$, is larger than D_t. Intuitively, a creditor cannot best-respond to other creditors' decisions to run on a super-solvent conduit, because rolling over guarantees the creditor a full payment even if the asset is liquidated. Formally, smooth-pasting at $W(1/l, x^*)$ implies that $W_x < 0$ for all $1/l < x < x^*$, i.e., frantic runs require that bond values decrease in asset values.

Appendix B. Numerical solution of the value function and run threshold

This appendix describes the algorithm we use to solve numerically for the value function $W(x; x^*)$ and the run

threshold x^*. The solution for W satisfies the HJB equation, value-matching and smooth-pasting for W everywhere (including at $x = x^*$), the limit condition $\lim_{x \to \infty} W(x, x^*)$, and the condition $W(x^*, x^*) = 1/\overline{R}$. The algorithm follows the following steps:

1. Guess a value for x^*.
2. Compute the solution to the HJB for $x \leq x^*$, using the analytical solution in Appendix A.
3. Solve W numerically for $x > x^*$, as follows:

   (a) Using the standard method, reduce the order of the ODE by introducing a new variable Z:

   $$W_x \equiv Z, \tag{38}$$

   $$\begin{aligned} Z_x = W_{xx} &= -\frac{2[\mu + \delta]}{\sigma^2}\frac{Z}{x} + \frac{2\delta}{\sigma^2}\frac{Z}{xW} - \frac{2\phi}{\sigma^2}\frac{\min(1, x)}{x^2} \\ &+ \frac{2(\rho + \phi + \delta)}{\sigma^2}\frac{W}{x^2} - \frac{2\delta}{\sigma^2}\frac{1}{x^2}. \end{aligned} \tag{39}$$

   (b) Solve analytically for $W(x^*; x^*, A(x^*))$ and $W_x(x^*; x^*, A(x^*)) = Z(x^*; x^*, A(x^*))$, using the solutions for W in Appendix A.
   (c) Using the initial conditions in step (b), numerically integrate the system of ODEs in step (a) for $x \in [x^*, \overline{x}]$, where $\overline{x}$ is a very large value of x that approximates $x = \infty$.
4. Check whether the numerical solution for $W(\overline{x}, x_*)$ is sufficiently close to its known limit, derived in Appendix A. If so, we have found the equilibrium threshold x^*. If not, return to step 1.

Appendix C. Details on SMM estimation

This appendix summarizes the SMM estimation procedure, which closely follows DeAngelo, DeAngelo, and Whited (2011). Additional details are in Strebulaev and Whited (2012) and Erickson and Whited (2012).

The goal is to estimate parameters $b = (\alpha, \sigma, \theta, \overline{r})$ by matching a vector of simulated moments as closely as possible to the corresponding vector of data moments. Let x_i denote a data vector and $y_{ik}(b)$ denote a simulated vector from simulation k, where $i = 1, \ldots, n$ indexes conduit/week observations and $k = 1, \ldots, K$ indexes simulations. We use $K = 20$; Michaelides and Ng (2000) find that a sample at least ten times larger than the empirical sample delivers good finite-sample performance. We simulate each conduit for $T = 8.7$ years, which is 1.5 times $1/\phi$, the asset's expected lifetime. When a conduit's asset matures, the conduit drops out of the sample and is not replaced. All simulated conduits begin with initial inverse leverage $x_0 = x^* + \frac{1}{2}\sigma\sqrt{T}$. This initial value of x_0 is far enough from x^* so that initial spreads are well below the ten b.p. threshold used to compute our moments, yet x_0 is close enough to x^* so that a large enough fraction of simulated conduits eventually cross the ten b.p. threshold. Of course, some simulated conduits experience lucky outcomes so

that their spreads never exceed ten b.p. per year. These conduits' observations do not contribute to any of our simulated moments.

We denote simulated moments as $h(y_{ik}(b))$ and data moments as $h(x_i)$. All the moments we use, including means and variances, can be expressed as slopes from OLS regressions. The SMM estimate of b is

$$\widehat{b} = \arg\min_b g_n(b)' \widehat{W}_n g_n(b), \tag{40}$$

where

$$g_n(b) = n^{-1} \sum_{i=1}^{n} \left[h(x_i) - \frac{1}{K} \sum_{k=1}^{K} h(y_{ik}(b)) \right] \tag{41}$$

is the difference between data moments and simulated moments, and $\widehat{W}_n$ is a positive definite weighting matrix that converges in probability to a deterministic positive definite matrix W. The efficient weighting matrix is the inverse of the sample covariance matrix of the moments. Since our 13×13 covariance is estimated with considerable noise, we use only its diagonal blocks when computing the weighting matrix,[42] and we divide the diagonal elements by the number of elements in each group of moments to apply roughly equal weight to our four sets of moments. For instance, we divide the elements of $\widehat{W}_n$ corresponding to moments 8–13 (the forecasting regressions) by six.

The estimator's asymptotic distribution is

$$\sqrt{n}(\widehat{b} - b) \xrightarrow{d} N(0, avar(\widehat{b})), \tag{42}$$

where

$$\begin{aligned} avar(\widehat{b}) = \left(1 + \frac{1}{K}\right) &\left[\frac{\partial g_n(b)'}{\partial b} W \frac{\partial g_n(b)}{\partial b'}\right]^{-1} \\ &\times \left[\frac{\partial g_n(b)'}{\partial b} W\Omega W \frac{\partial g_n(b)}{\partial b'}\right]\left[\frac{\partial g_n(b)'}{\partial b} W \frac{\partial g_n(b)}{\partial b'}\right]^{-1}. \end{aligned} \tag{43}$$

We estimate Ω using the Generalized Method of Moments (GMM) while taking into account heteroskedasticity and serial correlation, which is equivalent to computing heteroskedasticity-robust standard errors from a system of seemingly unrelated OLS regressions. Since our moments are slopes from a system of OLS regressions, this approach is equivalent to the influence-function approach of Erickson and Whited (2000) in the special case in which observations are independently distributed.[43] Since our empirical observations are not necessarily i.i.d., we estimate Ω while allowing correlation in regression disturbances both (1) within and across regressions and (2) within and across conduits, as long as observations are near each other in calendar time. The GMM approach

---

[42] We use a 2×2 block for moments 1 and 2 (recoveries from runs), a 2×2 block for moments 3 and 4 (volatility regression), a 3×3 block for moments 5–7 (event-time regression), and three 2×2 blocks for each of the forecasting regressions.

[43] To estimate slope β in regression $y_i = x_i'\beta + u_i$, the influence function is $\psi_i = (X'X)^{-1}x_i u_i$, and the heteroskedasticity-robust GMM or OLS asymptotic covariance of $\widehat{\beta}$ is $(X'X)^{-1}E[X'uu'X](X'X)^{-1}$. This covariance equals the covariance from the influence-function approach, $E[\psi_i\psi_i']$, in the special case where $E[u_iu_{j \neq i}] = 0$.

takes into account that different regressions use different sets of conduit/week observations. We use the eigenvalue method of Rousseuw and Molenberghs (1994) to guarantee that $\widehat{\Omega}$ is positive definite. We adjust parameters' standard errors for first-stage estimation error in parameters ϕ and δ using the method of Newey and McFadden (1994).

References

Acharya, V., Richardson, M., 2009. Causes of the financial crisis. Critical Review 21, 195–210.

Acharya, V., Schnabl, P., 2010. Do global banks spread global imbalances? ABCP during the financial crisis of 2007–2009. IMF Economic Review 58, 37–73.

Acharya, V., Schnabl, P., Suarez, G., 2013. Securitization without risk transfer. Journal of Financial Economics 107, 515–536.

Andrade, G., Kaplan, S., 1998. How costly is financial (not economic) distress? Evidence from highly leveraged transactions that became distressed. Journal of Finance 53, 1443–1493.

Angrist, J., Pischke, J., 2009. Mostly Harmless Econometrics. Princeton University Press, Princeton.

Bernanke, B., 1983. Nonmonetary effects of the financial crisis in the propagation of the Great Depression. American Economic Review 73, 257–276.

Blanchard, O., 1985. Debt, deficits, and finite horizons. Journal of Political Economy 93, 223–247.

Brunnermeier, M., 2009. Deciphering the liquidity and credit crunch 2007–2008. Journal of Economic Perspectives 23, 77–100.

Brunnermeier, M., Oehmke, M., 2013. The maturity rat race. Journal of Finance 68, 483–521.

Bryant, J., 1980. A model of reserves, bank runs, and deposit insurance. Journal of Finance 4, 335–344.

Calomiris, C., Mason, J., 1997. Contagion and bank failures during the Great Depression: The June 1932 Chicago banking panic. American Economic Review 87, 863–883.

Calomiris, C., Mason, J., 2003. Fundamentals, panics, and bank distress during the Depression. American Economic Review 93, 1615–1647.

Calvo, G., 1983. Staggered prices in a utility-maximizing framework. Journal of Monetary Economics 12, 383–398.

Chen, Q., Goldstein, I., Jiang, W., 2010. Payoff complementarities and financial fragility: evidence from mutual fund flows. Journal of Financial Economics 97, 239–262.

Cheng, I., Milbradt, K., 2012. The hazards of debt: rollover freezes, incentives, and bailouts. Review of Financial Studies 25, 1070–1110.

Coval, J., Stafford, E., 2007. Asset fire sales (and purchases) in equity markets. Journal of Financial Economics 86, 479–512.

Covitz, D., Liang, N., Suarez, G., 2013. The evolution of a financial crisis: collapse of the asset-backed commercial paper market. Journal of Finance 68, 815–848.

DeAngelo, H., DeAngelo, L., Whited, T., 2011. Capital structure dynamics and transitory debt. Journal of Financial Economics 99, 235–261.

Décamps, J., Villeneuve, S., 2007. Optimal dividend policy and growth option. Finance and Stochastics 11, 3–27.

Dell'Ariccia, G., Detragiache, E., Rajan, R., 2008. The real effect of banking crises. Journal of Financial Intermediation 17, 89–112.

Diamond, D., Dybvig, P., 1983. Bank runs, deposit insurance, and liquidity. Journal of Political Economy 91, 401–409.

Ellul, A., Jotikasthira, C., Lundblad, C., 2011. Regulatory pressure and fire sales in the corporate bond markets. Journal of Financial Economics 101, 596–620.

Erickson, T., Whited, T., 2000. Measurement error and the relationship between investment and q. Journal of Political Economy 108, 1027–1057.

Erickson, T., Whited, T., 2012. Treating measurement error in Tobin's q. Review of Financial Studies 25, 1286–1329.

Friedman, M., Schwartz, A., 1963. A Monetary History of the United States, 1867–1960. Princeton University Press, Princeton.

Goldstein, I., 2014. Fundamentals or panic: lessons from the empirical literature on financial crises. Encyclopedia of Financial Globalization, forthcoming.

Goldstein, I., Pauzner, A., 2005. Demand deposit contracts and the probability of bank runs. Journal of Finance 60, 1293–1328.

Gorton, G., 1988. Banking panics and business cycles. Oxford Economic Papers 40, 751–781.

Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. Journal of Financial Economics 104, 425–451.

He, Z., Milbradt, K., 2012. Endogenous liquidity and defaultable bonds. Unpublished working paper. University of Chicago, Massachusetts Institute of Technology.

He, Z., Xiong, W., 2012a. Dynamic debt runs. Review of Financial Studies 25, 1799–1843.

He, Z., Xiong, W., 2012b. Rollover risk and credit risk. Journal of Finance 67, 391–429.

Hege, U., Mella-Barral, P., 2005. Repeated dilution of diffusively held debt. Journal of Business 78, 737–786.

Hennessy, C., Whited, T., 2007. How costly is external financing? Evidence from a structural estimation. Journal of Finance 62, 1705–1745.

Hugonnier, J., Malamud, S., Morellec, E., 2012. Capital supply uncertainty, cash holdings, and investment. Unpublished working paper. Ecole Polytechnique Fédérale de Lausanne.

Irani, R., 2011. Bank health and corporate liquidity provision. Unpublished working paper. New York University.

Iyer, R., Puri, M., 2012. Understanding bank runs: the importance of depositor–bank relationships and networks. American Economic Review 102, 1414–1445.

Kacperczyk, M., Schnabl, P., 2013. How safe are money market funds? Quarterly Journal of Economics 128, 1073–1122.

Kaminsky, G., Reinhart, C., 1999. The twin crises: the causes of banking and balance-of-payments problems. American Economic Review 89, 473–500.

Krishnamurthy, A., 2009. How debt markets have malfunctioned in the crisis. Journal of Economic Perspectives 24, 3–28.

Krishnamurthy, A., Nagel, S., Orlov, D., 2014. Sizing up repo. Journal of Finance, forthcoming.

Leland, H., 1998. Agency costs, risk management and optimal capital structure. Journal of Finance 53, 1213–1243.

Leland, H., Toft, K., 1996. Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads. Journal of Finance 51, 987–1019.

Martin, A., Skeie, D., Von Thadden, E., 2012. Repo Runs. Federal Reserve Bank of New York Staff Report No. 444.

Michaelides, A., Ng, S., 2000. Estimating the rational expectations model of speculative storage: a Monte Carlo comparison of three simulation estimators. Journal of Econometrics 96, 231–266.

Moody's Investors Service, 2008. Moody's update on structured investment vehicles.

Morris, S., Shin, H., 1998. Unique equilibrium in a model of self-fulfilling currency attacks. American Economic Review 88, 587–597.

Newey, W., McFadden, D., 1994. Large sample estimation and hypothesis testing. In: Engle, R., McFadden, D. (Eds.), Handbook of Econometrics, Elsevier, Amsterdam, pp. 2111–2245.

Ramirez, C., Shively, P., 2012. The effect of bank failures on economic activity: evidence from U.S. states in the early 20th century. Journal of Money, Credit and Banking 44, 433–455.

Rochet, J., Vives, X., 2004. Coordination failures and the lender of last resort: Was Bagehot right after all?. Journal of the European Economic Association 2, 1116–1147.

Rousseuw, P., Molenberghs, G., 1994. The shape of correlation matrices. The American Statistician 48, 276–279.

Shin, H., 2009. Financial intermediation and the post-crisis financial system. Unpublished working paper. Princeton University.

Stiglitz, J., Weiss, A., 1981. Credit rationing in markets with imperfect information. American Economic Review 71, 393–410.

Strebulaev, I., Whited, T., 2012. Dynamic models and structural estimation in corporate finance. Foundations and Trends in Finance 6, 1–163.