

Understanding Cash Flow Risk

Sebastian Gryglewicz

Erasmus University Rotterdam

Loriano Mancini

USI Lugano and Swiss Finance Institute

Erwan Morellec

EPF Lausanne, Swiss Finance Institute, and CEPR

Enrique Schroth

EDHEC Business School and CEPR

Philip Valta

University of Bern

Theory has recently shown that corporate policies should depend on firms' exposure to short- and long-lived cash flow shocks and the correlation between these shocks. We provide granular estimates of these parameters for Compustat firms using a new filter that uses only cash flow data and the theoretical restrictions of a canonical cash flow model. As predicted by theory, we find that the estimated parameters are strongly related to corporate liquidity and financing choices, that firms with a higher estimated correlation between shocks implement riskier policies, and that the sign of this correlation determines the cash flow sensitivity of cash. (*JEL* G31, G32, G35)

Received February 1, 2021; editorial decision August 18, 2021 by Editor Ralph Koijen. Authors have furnished an Internet Appendix, which is available on the Oxford University Press Web site next to the link to the final published paper online.

Starting with Gorbenko and Strebulaev (2010) and DeMarzo et al. (2012), a growing theoretical literature in corporate finance shows the importance of

We thank the editor (Ralph Koijen), two anonymous referees, Bruno Biais, Sudipto Dasgupta, Laurent Frésard, Patrick Gagliardini, Nicolae Gârleanu, Bart Lambrecht, Boris Nikolov, Gordon Phillips, Jean-Charles Rochet, Ailsa Röell, Clemens Sialm, Pascal St-Amour, Peter Swan, Elu von Thadden, Alexander Wagner, Josef Zechner, and seminar participants at EDHEC, USI Lugano, University of Oxford, University of New South Wales, University of Illinois at Urbana–Champaign, DePaul University, Warwick Business School, the CEPR Second Annual Spring Symposium in Financial Economics, European Finance Association Meetings, the European Summer Symposium in Financial Markets, and the Swiss Finance Institute Research Days for comments. Loriano Mancini and Erwan Morellec acknowledge financial support from the Swiss Finance Institute. Supplementary data can be found on *The Review of Financial Studies* web site. Send correspondence to Erwan Morellec, erwan.morellec@epfl.ch.

The Review of Financial Studies 35 (2022) 3922–3972

© The Author(s) 2021. Published by Oxford University Press on behalf of The Society for Financial Studies.

All rights reserved. For permissions, please e-mail: journals.permissions@oup.com.

doi: 10.1093/rfs/hhab127

Advance Access publication November 23, 2021

explicitly modeling firms' exposure to permanent (long-lived) and transitory (short-lived) cash flow shocks when analyzing corporate policies. Indeed, while transitory shocks affect immediate cash flows, they are uninformative about future expected profitability. By contrast, permanent shocks affect not only a firm's immediate productivity and cash flows but also its future productivity and cash flows.¹ While the decomposition of shocks between transitory and permanent components has been used productively in many areas of economics, it has been largely neglected in empirical corporate finance.² This is surprising given that theory shows corporate decisions should depend not only on the level of risk but also on its composition, as captured by firms' exposure to long- and short-lived shocks and the correlation between these shocks.

The objective of this paper is to start filling the existing void. We do so in four successive steps. First, we provide evidence that a majority of firms' operating cash flows are subject to permanent (long-lived, nonstationary) shocks, thereby providing support for cash flow models used in recent dynamic corporate finance models. Second, since permanent and transitory shocks are not separately observable, we develop a novel filter to decompose the cash flow shocks of publicly traded U.S. firms into permanent and transitory components and estimate their primitive parameters. Third, we provide granular estimates of cash flow risk parameters—that is, firms' exposure to short- and long-lived cash flow shocks and the correlation between these shocks—for a large fraction of the Compustat universe since the 1970s. Fourth, we show that the estimated parameters are strongly related to corporate liquidity and financing choices, that firms with a higher estimated correlation between shocks implement riskier policies, and that the sign of this correlation determines the cash flow sensitivity of cash, as predicted by theory.

We begin our empirical analysis by testing whether firm cash flow shocks include a permanent nonstationary component, as assumed in most recent models of investment, financing, liquidity, or compensation policies. To this end, we use two standard unit root tests, the Augmented Dickey–Fuller (ADF) and the Kwiatkowski et al. (1992) (KPSS) tests, which we implement on the individual cash flows of over 10,000 Compustat firms between 1971 and 2018.

¹ Many cash flow shocks are transitory and do not affect long-term prospects. Examples include temporary changes in demand, delays in customer payments, machine breakdowns, or supply chain disruptions. But long-term cash flows also change over time due to various firm, industry, or macroeconomic shocks that are of permanent nature. Examples include changes in technology or in consumer preferences.

² A number of asset pricing papers (see, e.g., Cochrane, 1994, Cohen, Gompers, and Vuolteenaho, 2002, Bansal, Dittmar, and Kiku, 2008, Garleanu, Panageas, and Yu, 2012) use such a decomposition to analyze stock returns and risk premia on stocks. This decomposition is also used in market microstructure to analyze price efficiency (see, e.g., Glosten and Harris, 1988, Brennan and Subrahmanyam, 1996, Boehmer and Wu, 2013). The literature on income processes also often seeks to decompose shocks into permanent and transitory components; see, e.g., Blundell, Pistaferri, and Preston (2008), Meghir and Pistaferri (2004), or Gottschalk and Moffitt (2009). The decomposition of income shocks between permanent and transitory components has found interesting applications in the life-cycle portfolio choice literature; see, e.g., Cocco, Gomes, and Maenhout (2005). In the time series literature, the permanent-transitory model is known as the unobserved component decomposition, in which the permanent part is the trend and the transitory component is named the cyclical innovation; see Hamilton (1994, Chapter 17).

Both tests reject stationarity and strongly suggest that a majority of firms' operating cash flows are, indeed, subject to permanent shocks.

Assessing the importance of permanent and transitory shocks for corporate policies faces two challenges. First, one needs to be able to identify permanent shocks separately from transitory shocks. Second, one needs to obtain reliable and granular estimates of the cash flow risk parameters to relate these to firm policies. To address these challenges, we estimate a canonical cash flow model that nests as special cases most of the cash flow models used in the recent dynamic investment, financing, liquidity, or compensation (contracting) research. In this model, firm cash flows are subject to profitability shocks that are permanent in nature. In addition, for a given level of profitability, cash flows are subject to short-term shocks, which may be purely transitory or correlated with permanent shocks.³ Cash flow risk is therefore captured by the firm's exposure to permanent shocks, its exposure to short-term shocks, and the correlation between these shocks. Importantly, positive correlation reduces risk as the firm is more likely to generate positive cash flows after positive productivity shocks have increased firm value. To identify potentially correlated permanent and short-term shocks, we develop a novel Kalman filter that is derived from this theoretical structure. We then use our filtering technique and maximum likelihood to estimate from panel data the parameters driving the cash flow shocks—that is, the volatilities of permanent and short-term shocks as well as their correlation—that best explain observed firm cash flows.

Our estimation of the parameters characterizing cash flow risk yields three striking results. First, permanent and short-term shocks are negatively correlated for most firms, suggesting large hedging needs. Second, firms that are naturally hedged due to a positive shock correlation tend both to choose riskier policies (e.g., take on more debt) and to have lower overall risk than other firms (i.e., have a lower asset or equity return volatility and a larger distance to default). By contrast, firms with a higher estimated volatility of permanent or transitory shocks adopt safer policies and yet are riskier. Third, our parameter estimates exhibit remarkable cross-sectional variation. For example, we estimate an interquartile range between 32% and 131% for the permanent shock volatility and between 12% and 60% for the temporary shock volatility. These estimates of volatilities are much more heterogeneous across the Compustat panel than the proxy used in prior research for the precautionary motive to hold cash, namely the industry cash flow volatility (see, e.g., Opler et al., 1999; Bates, Kahle, and Stulz, 2009; Graham and Leary, 2018). We also show that the joint distribution of these three estimates implies asset return

³ In pioneering work, Froot, Scharfstein, and Stein (1993) show that a firm's hedging policy varies dramatically depending on whether price shocks, which are typically short-lived, are positively or negatively correlated with longer-lasting investment opportunities. Correlation between short- and long-term shocks in the literature is indeed generally operationalized by the correlation between industry-specific investment and cash flow (see, e.g., Duchin, 2010; Acharya, Almeida, and Campello, 2007). Our cash flow model provides an alternative way to operationalize this correlation that requires weaker identifying assumptions.

volatilities that match the actual asset return volatilities of the firms in our sample.

Because the estimated parameters characterize cash flow risk and hedging needs, it is natural to explore how they relate to liquidity management policies. Our focus on liquidity management is further motivated by the central role of cash reserves and access to liquidity in ensuring firm resilience to cash flow shocks, as illustrated by the COVID-19 crisis. Due to the precautionary role of cash, models of liquidity management predict that target levels of cash reserves should increase with cash flow risk and hedging needs (see, e.g., Bolton, Chen, and Wang, 2011 or Décamps et al., 2017). In these models, firms build up cash reserves toward their target level either by retaining earnings or by raising outside equity and keeping part of the proceeds in cash reserves. In our empirical analysis, we thus focus on these two separate mechanisms to manage cash reserves. As will become clear, our tests do not require estimating the target level of cash reserves. As a result, they are unaffected by the coefficient biases that stem from trying to measure unobservable variables. Yet, because the predictions that we take to the data take the exact form of the tests that we execute, the connection between theory and tests is tight.

One way for firms to increase their cash reserves and their resilience to shocks is to retain earnings. There is considerable debate in the literature about the sign of the cash flow sensitivity of cash savings. Almeida, Campello, and Weisbach (2004) argue and provide evidence that the sensitivity is positive or zero, depending on firms' financing constraints and hedging needs. By contrast, Riddick and Whited (2009) show that this sensitivity should be negative when productivity shocks are persistent and find empirically that it is on average negative. Décamps et al. (2017) sharpen the prediction and demonstrate that the cash flow sensitivity of cash should be positive when the correlation between permanent and short-term cash flow shocks is positive, and negative otherwise. Indeed, when this correlation is positive, positive (short-term) cash flow shocks are more likely to occur simultaneously with positive (long-term) productivity shocks that increase both firm value and the marginal value of cash and, as a result, firms' incentives to save. We run various tests to verify this prediction and find that the cash flow sensitivity of cash does switch sign depending on the correlation between permanent and short-term shocks. Strikingly, the estimated sensitivities are all highly statistically significant, and all exhibit the predicted sign.

Another way for firms to replenish cash reserves is to raise new equity, as empirically shown by Kim and Weisbach (2008) or McLean (2011) and theoretically argued in Décamps et al. (2011) or Bolton, Chen, and Wang (2011). Theory predicts that firms with high permanent or transitory shock volatility should have larger (precautionary) cash reserves and issue larger amounts of new equity. In addition, they should raise equity more frequently. By contrast, firms with a high correlation between permanent and short-term cash flow shocks are more naturally hedged and should therefore hold smaller

cash reserves and issue smaller amounts of equity at a lower frequency. Our empirical analysis provides strong support for these predictions. Notably, we find a positive and significant relation between the size or frequency of equity issues and the volatility of permanent cash flow shocks. We further find that firms issue less equity when the correlation between short-term and permanent shocks is high. These results are again highly statistically significant and apply to both the size and frequency of equity issues. Our empirical tests additionally show that permanent cash flow volatility is more important than transitory cash flow volatility in explaining the cross-sectional variation in equity issuance activity, demonstrating the importance of incorporating permanent shocks in dynamic corporate finance models.

As shown by the recent COVID-19 crisis, firms may also consider issuing debt to build up their cash reserves. Our estimates can help answer whether debt-financed cash savings policies, such as the response to the COVID shock, are common or an exception to the norm. For that purpose, we explore the joint determination of equity and debt issuance as a function of our estimated cash flow risk parameters. We find that all the results for equity issues are weakened if we define the dependent variable as “external finance”—that is, the sum of equity and debt issues. Treating each policy separately, and estimating both policies simultaneously, we confirm earlier findings that the frequency and amount of equity issued decrease with the correlation between shocks and increase with their volatilities, but fail to reject that long-term debt frequency and volume of issuance are not related to the same risk parameters. Our results therefore suggest that firms resorted to long-term debt as a last resort in extremely unusual circumstances during the COVID-19 crisis, instead of using debt as a normal policy lever to manage liquidity.

Our paper relates to the growing theoretical literature on the effects of permanent and transitory shocks on corporate policies. Gorbenco and Strebulaev (2010) develop a dynamic capital structure model in which cash flows are subject not only to permanent shocks, as in Leland (1998), but also to transitory Poisson shocks. DeMarzo et al. (2012), Hoffmann and Pfeil (2010), Gryglewicz, Mayer, and Morellec (2020), and Hackbarth, Rivera, and Wong (2021) examine the effects of permanent and transitory cash flow shocks on optimal compensation and investment in dynamic moral hazard models. Décamps et al. (2017) and Bolton, Wang, and Yang (2019) examine the effects of permanent and transitory shocks on cash holdings, credit lines usage, equity issues, and risk management in models with financing frictions. We contribute to this literature by providing efficient estimates of the deep parameters of a cash flow model nesting all of these, at a granular level, and for a large fraction of the Compustat universe.

As relevant as it is to analyze the effects of transitory and permanent shocks on corporate policies, there are surprisingly only a few attempts in the empirical corporate finance literature addressing this problem. In an early study, Guay and Harford (2000) show that firms choose dividend increases to distribute relatively

permanent cash flow shocks and repurchases to distribute more transient shocks. Chang et al. (2014) decompose corporate cash flows into a transitory and a permanent component and argue that this decomposition helps understand how firms allocate cash flows and whether financial constraints matter in this allocation decision. Lee and Rui (2007) show that such a decomposition also allows determining whether share repurchases are used to pay out cash flows that are potentially transitory, thus preserving financial flexibility relative to dividends. Guiso, Pistaferri, and Schivardi (2005) examine the allocation of risk between firms and their workers and show that firms absorb transitory shocks fully but insure workers against permanent shocks only partially. Lastly, Byun, Polkovnichenko, and Rebello (2019b,a) examine the separate effects of persistent and transitory shocks on leverage and investment decisions.

Our paper advances this literature in two ways. First and more importantly, our paper is unique in that it develops a filter that is specially designed to estimate a general cash flow process used in corporate finance, while addressing the practical issues with corporate cash flow data. Instead, most existing studies use either the Hodrick–Prescott filter or the Beveridge–Nelson decomposition to separate a time series into a trend (permanent) component and a cyclical (transitory) component. The use of these filters is problematic for our purpose because they cannot handle missing values, which are pervasive in large cash flow panels, and they cannot estimate the correlations between long- and short-term shocks, which, as we show, vary widely across firms and influence corporate policies significantly. Our novel Kalman filter is free of these limitations and performs much better empirically than these standard filters when applied to corporate cash flows.⁴ Second, our paper differs from prior studies because of its focus on liquidity management policies and on identifying the differential impact of the exposure to permanent and short-term shocks and the correlation between shocks on cash savings and financing decisions.

Lastly, the type of analysis that relates cash holdings to hedging needs has precedents in the literature, such as Acharya, Almeida, and Campello (2007) and Duchin (2010). Our analysis is unique because it uses estimated deep parameters of a canonical cash flow process instead of relying on proxies for cash flow volatility and hedging needs to explain corporate cash policy. Another important difference is that these studies relate cash balances to a number of explanatory variables including hedging needs, implicitly assuming that firms are at their target level of cash holdings. Recent theory has shown, however, that due to adjustment costs, cash reserves are almost never at their target level (see Décamps et al., 2011 or Bolton, Chen, and Wang, 2011). Our paper differs from

⁴ Another problematic aspect of the HP filter (used, for instance, in Byun, Polkovnichenko, and Rebello, 2019b,a to separate the cash flow time series into a trend component and a cyclical component) is that it introduces biases and spurious effects; see, e.g., Hamilton (2018) for details. By contrast, Kalman filtering does not introduce any bias when the model is correctly specified; see, e.g., Chapter 13.4 in Hamilton (1994).

these early studies because of its focus on specific financing times when the predictions of dynamic liquidity management models are more likely to hold. Recent papers by Danis, Rettl, and Whited (2014) or Eckbo and Kisser (2021) follow a similar approach when testing dynamic capital structure theories.

1. Cash Flows Data

We begin our analysis by exploring the time series properties of corporate cash flow data. Our goal here is to understand the nature of firm-level shocks to operating cash flows and to determine which class of models best describes cash flow dynamics.

1.1 Sample

We collect accounting data for publicly traded U.S. firms from Compustat between 1971 and 2018 and stock price data from CRSP. We exclude financial services firms (SIC codes 6000 to 6999), utilities (SIC codes 4900 to 4999), and firms whose annual asset growth exceeds 500% in any given year, corresponding to the 95th percentile of the distribution of annual growth in Compustat.⁵ We convert all data into 2000 constant dollars using the GDP deflator and winsorize the firm-level variables at the 1st and 99th percentiles.

Unlike panel data studies or most structural estimations in the corporate finance literature, our goal is to estimate the deep parameters of the cash flow process with a high level of granularity. To guarantee precision in the estimation, we require firms to have sufficiently long cash flow series. Specifically, we impose that a usable firm has at least 10, not necessarily consecutive, observations.

Our final sample includes 208,605 firm-years for 10,136 firms, covering about 43% of the firms (and 92% of the market capitalization) in the Compustat universe since the 1970s. Our coverage is remarkably high, considering that over 27% of Compustat firms have at most four cash flow observations. To the best of our knowledge, Duchin (2010) is the only other study using individual cash flow moments with almost as many firms. We can achieve such a high coverage due to an advantageous feature of our Kalman filtering technique: it does not require consecutive observations. Chang et al. (2014) and Byun, Polkovnichenko, and Rebello (2019b,a) are the other three studies that decompose firm cash flows into shocks of different duration. These studies use standard filters requiring arithmetic interpolation, longer time series, or no gaps.

⁵ Typical screening criteria in the cash savings literature (e.g., Almeida, Campello, and Weisbach, 2004, Acharya, Almeida, and Campello, 2007) remove firm-year observations in which asset growth exceeds 100% to remove observations likely to reflect “large jumps in business fundamentals” that are indicative of mergers, reorganizations, or other major corporate events. Our filtering method can potentially undo the effects of removal of interim firm-year observations because missing values are imputed as part of the estimation algorithm. Hence, we proceed to remove the firm altogether. The screen results in the exclusion of 1,688 firms, which are on average almost six times smaller than those included.

Hence, they cover at most 5,803 firms (Byun, Polkovnichenko, and Rebello, 2019b). An important implication of this high coverage of the Compustat universe is that we can study firms with very volatile cash flows. Such firms are often excluded from empirical studies, and yet their behavior is informative. As we shall see, they respond to their very volatile cash flows with conservative policies, leading to much lower asset volatilities.

1.2 The operating cash flows variable

The stochastic properties of the cash flow derived from operations are key determinants of firms' policies, such as earnings retention or external financing, that aim to increase cash reserves and, therefore, the resilience to shocks. Hence, we pursue the notion of operating cash flows, defined as EBITDA minus the change in working capital.⁶ We subtract the change in working capital because this account captures the allocations of cash that are needed to sustain the firm's operations.⁷ The subtraction of the change in working capital (which does not include changes in cash or cash equivalents) is the only difference between our measure and the notion of operating income used in, for example, Hennessy and Whited (2007). We show in Section 5.2 that the estimates of the model parameters do not change significantly if we do not subtract the change in working capital.

Our sample includes 10,136 firms that vary along many dimensions. To make firms of different sizes comparable, we divide each year's operating cash flow by the firm's initial value of total assets. Importantly, this normalization does not affect the time series properties of operating cash flows because the initial value of assets is constant over time. Table 1 contains the definitions and descriptive statistics of our operating cash flow variable as well as other firm-specific characteristics that we use in the empirical analysis.

1.3 Time series properties of operating cash flows

We run two tests to determine whether operating cash flows include a permanent (nonstationary) component, as assumed in most recent dynamic models of investment, financing, cash savings, or compensation policies (see, e.g., Leland, 1994; Carlson, Fisher, and Giammarino, 2004; Abel and Eberly, 2011; DeMarzo et al., 2012; Décamps et al., 2017):⁸ the Augmented Dickey-Fuller (ADF) test and the Kwiatkowski et al. (1992) (KPSS) test. The null

⁶ In the savings and investment literature, the maximum cash available to save is approximated by the physical flow of cash from operations minus taxes and interest expenses (Almeida, Campello, and Weisbach, 2004, Denis and Sibilkov, 2010, Hennessy and Whited, 2007, Riddick and Whited, 2009) and possibly dividends (Opler et al., 1999; Bates, Kahle, and Stulz, 2009) or extraordinary items (McLean, 2011). Chang et al. (2014) account for all flows of cash before investment, including also the net changes in working capital, deferrals, and equity transactions. These definitions are not appropriate for our study as they include a policy choice component.

⁷ The working capital account is reported differently in Compustat before and after SFAS #95 regulations. Prior to July 1988, the working capital account is reported directly (wcapc). Afterward, it needs to be constructed using the full statement of cash flows (see Chang et al., 2014).

⁸ While early dynamic contracting and liquidity models assume that cash flow shocks are purely transitory (see, e.g., DeMarzo and Sannikov, 2006, Bolton, Chen, and Wang, 2011), recent contributions have enriched these

Table 1
Definitions and descriptive statistics of variables

A. Variable definition

Variable name	Variable definition
<i>Operating cash flow</i>	EBITDA (oibdp) minus change in working capital, defined as in Table 1 of Chang et al. (2014).
<i>Cash flow-to-initial assets</i>	Ratio of <i>Operating cash flow</i> to the first observation of the book value of total assets (at)
<i>Cash savings</i>	Ratio of the change in cash holdings (che) from year $t - 1$ to t to the lagged book value of total assets (at)
<i>Gross equity issuance</i>	Ratio of the proceeds from sales or conversions of common and preferred stock (sstk) to the lagged book value of total assets (at)
<i>Equity issuance dummy</i>	Dummy variable equal to one if <i>Gross equity issuance</i> is larger than 5%, and zero otherwise
<i>Net equity issuance</i>	Proceeds from sales or conversions of common and preferred stock (sstk) minus dividends on common stock (dvc), dividends on preferred stock (dvp), and repurchased shares (prstk) divided by the lagged book value of total assets (at)
<i>Net equity issuance dummy</i>	Dummy variable equal to one if <i>Net equity issuance</i> is larger than 5%, and zero otherwise
<i>Net debt issuance</i>	Long-term debt issuance (dltris) minus long-term debt reduction (dltr) divided by the lagged book value of total assets (at)
<i>Net debt issuance dummy</i>	Dummy variable equal to one if <i>Net debt issuance</i> is larger than 5%, and zero otherwise
<i>Market-to-book ratio</i>	Book value of total assets (at) + market cap (csho*prcc_f) – book equity (ceq) divided by total assets
$\ln(\text{Total assets})$	Logarithm of the book value of total assets (at)
<i>Cash flow-to-assets</i>	Ratio of <i>Operating cash flow</i> to lagged book value of total assets (at)
<i>Industry cash flow volatility</i>	Mean of the standard deviation of firms' cash flow-to-assets ratio, for all firms in the same two-digit SIC industry and year

B. Descriptive statistics

	N	Mean	Stdev	p5	p25	p50	p75	p95
Cash flow-to-initial assets	207,554	1.507	18.403	-0.681	0.038	0.185	0.542	4.161
Cash savings	197,597	0.025	0.196	-0.156	-0.024	0.000	0.032	0.241
Change in cash-to-assets	197,549	0.000	0.093	-0.145	-0.026	0.000	0.025	0.144
Gross equity issuance	191,165	0.080	0.308	0.000	0.000	0.001	0.014	0.416
Equity issuance dummy	191,165	0.151	0.358	0.000	0.000	0.000	0.000	1.000
Net equity issuance	182,611	0.045	0.281	-0.095	-0.022	-0.001	0.003	0.337
Net equity issuance dummy	182,611	0.126	0.331	0.000	0.000	0.000	0.000	1.000
Net debt issuance	183,912	0.022	0.130	-0.104	-0.018	0.000	0.025	0.230
Net debt issuance dummy	183,912	0.190	0.392	0.000	0.000	0.000	0.000	1.000
Market-to-book ratio	181,410	2.180	2.662	0.762	1.085	1.444	2.210	5.637
$\ln(\text{Total assets})$	208,467	5.070	2.379	1.261	3.402	5.011	6.693	9.179
Cash flow-to-assets	197,168	0.074	0.263	-0.300	0.034	0.113	0.184	0.334
Industry cash flow volatility	201,480	0.177	0.188	0.060	0.091	0.134	0.204	0.382

This table presents the definitions (panel A) and the descriptive statistics (panel B) of the main variables used in the analysis. The descriptive statistics are: number of observations (N), mean, standard deviation, and the percentiles p5, p25, p50, p75, and p95. The sample covers the period 1971 to 2018.

hypothesis in the ADF test is that operating cash flows are nonstationary and have a unit root, while the alternative hypothesis is that they are stationary and follow an autoregressive process with a drift. Given the risk of failing to reject

models by adding permanent shocks. See, e.g., He (2009), Hoffmann and Pfeil (2010), DeMarzo et al. (2012), or Gryglewicz, Mayer, and Morellec (2020) for contracting papers and Décamps et al. (2017) or Bolton, Wang, and Yang (2019) for dynamic liquidity papers.

Table 2
Tests of nonstationarity of operating cash flows

ADF			KPSS			Minimum observations	Number of firms
10%	5%	1%	10%	5%	1%		
87.1	91.4	97.1	91.4	88.2	82.8	48	373
83.8	90.0	96.3	89.0	84.5	76.6	38	792
82.8	89.5	96.2	85.1	79.3	68.9	30	1,616
85.1	91.1	97.0	70.2	61.7	47.6	15	6,342

Table entries for ADF (columns 1 to 3) and for KPSS (columns 4 to 6) tests are the percentage of times that the ADF test is not rejected and the KPSS test is rejected for three different confidence levels, 10%, 5%, and 1%, respectively. The null hypothesis of the ADF test is that firm's cash flows have a unit root—that is, they are nonstationary. The null hypothesis of the KPSS test is that the firm's cash flows are stationary. There are 10,136 firms in our cash flow panel. For any given firm there is a maximum 48 yearly observations of cash flows between 1971 to 2018. The ADF and KPSS tests are run for each firm's cash flow time series when the number of available observations exceeds a given minimum (reported in "Minimum observations"). The last column reports the total number of firms tested for each required minimum number of observations. Critical values for the ADF and KPSS tests are not tabulated for firms with fewer than 15 cash flow observations.

the null hypothesis with small sample sizes, we also implement the KPSS test, in which the null hypothesis is that cash flows are stationary, while the alternative is that they follow a unit root process. In KPSS tests, the bias with small samples is toward not rejecting stationarity, that is, deeming the time series as stationary too often. We run the ADF and KPSS tests for each firm's operating cash flow time series in our panel. Table 2 summarizes the results.

The first row in panel A of Table 2 shows the results for the ADF and KPSS tests over the subsample of 373 firms that have the longest possible cash flow time series of our sample: 48 yearly observations. For 97.1% of the firms in this subsample, in which both tests have the highest power, the ADF test does not reject the null hypothesis of nonstationarity at the 1% level. For 91.4% of these firms, the KPSS tests reject the null hypothesis of stationary cash flows at the 10% level. The high rejection rate is remarkable for the KPSS test, given that rejecting the hypothesis of stationarity in relatively short samples is typically rare.

The ADF and KPSS tests results also strongly suggest nonstationarity for firms with fewer observations: For 82.8% and 85.1% of all 1,616 firms with at least 30, not necessarily consecutive, cash flow observations, the ADF tests do not reject nonstationarity while the KPSS tests reject stationarity at the 10% level, respectively. Even as the power decreases further, the KPSS tests still reject the hypothesis of stationarity at the 10% level for 70.2% of the 6,342 firms with at least 15 observations in the sample period. (The critical values for the tests cannot be computed for firms with fewer than 15 time series observations.)

In sum, there is overwhelming evidence that a majority of firms' operating cash flows are subject to permanent shocks, as in the most general, canonical cash flow models used in dynamic corporate finance. We now discuss our method to decompose cash flow shocks into permanent and transitory components and estimate their primitive parameters.

2. Estimation of the Cash Flow Model

2.1 The model

To quantify the exposure of operating cash flows to permanent and transitory shocks, we estimate (the discrete time version of) a canonical cash flow model that nests as special cases most of the cash flow models used in recent dynamic investment, financing, liquidity, or compensation models. In state space form, this cash flow model consists of the following transition and measurement equations:

$$X_t = (1 + \mu)X_{t-1} + \sigma_P X_{t-1} \varepsilon_t^P \tag{1}$$

$$Z_{i,t} = X_t + \sigma_A X_{t-1} \varepsilon_{i,t}^A, \tag{2}$$

where X_t is the unobserved asset productivity with constant growth rate μ and volatility $\sigma_P > 0$, and $Z_{i,t}$ is the operating cash flow of firm i , in year t , with short-term volatility $\sigma_A > 0$. In this model, the shock ε_t^P influences cash flows permanently by affecting the productivity of assets. The short-term shock $\varepsilon_{i,t}^A$ affects the cash flow directly and may also affect the firm's long-term prospects. Specifically, we allow short-term and permanent shocks to be correlated with correlation coefficient $\rho \in (-1, 1)$. Hence, the short-term shock can be written as

$$\varepsilon_{i,t}^A = \rho \varepsilon_t^P + \sqrt{1 - \rho^2} \varepsilon_{i,t}^T, \tag{3}$$

where $\varepsilon_{i,t}^T$ is a purely transitory shock uncorrelated with ε_t^P . Both ε_t^P and $\varepsilon_{i,t}^T$ are distributed as $\mathcal{N}(0, 1)$. When $\sigma_A = \sigma_P = 0$, cash flows follow the Gordon growth model used in textbook valuation models; see, for example, Berk and DeMarzo (2019). When $\sigma_A = 0$ and $\sigma_P > 0$, cash flows are only subject to permanent shocks. This is the discretized version of the cash flow model used in dynamic investment and capital structure models (see Abel and Eberly, 1994 or Leland, 1998). When $\mu = \sigma_P = 0$, cash flow shocks are identically and independently distributed, and follow a purely stationary process, as in early dynamic liquidity or compensation models (see Bolton, Chen, and Wang, 2011 or DeMarzo and Sannikov, 2006). When $\sigma_A > 0$ and $\sigma_P > 0$, cash flows are subject to both permanent and transitory shocks. When $\rho \neq 0$, the model captures another important dimension of risk. Indeed, ρ reflects the notion of correlation between current cash flow and investment opportunities discussed, for instance, in Froot, Scharfstein, and Stein (1993) and, hence, the firm's hedging needs.⁹

To summarize, Equations (1)–(3) capture the heterogeneity of firms' operating cash flow exposures to long-term and short-term risk via different

⁹ This general cash flow model has been proposed by Décamps et al. (2017), who show that cash policy, equity issuance and credit line usage depend on the combination of all the cash flow parameters. More recently, a similar cash flow model has been used to explain compensation policy (see, e.g., Gryglewicz, Mayer, and Morellec, 2020), debt policy (see, e.g., Bolton, Wang, and Yang, 2021), financial development (see, e.g., Rebelo, Wang, and Yang, 2020), or the horizon of corporate policies (see, e.g., Breugem, Marfe, and Zucchi, 2021).

combinations of values of σ_P , σ_A , and ρ . Estimation of these parameters with the highest possible level of granularity will enable the testing of empirical predictions of the effects that different combinations of parameter values have on corporate policies.

2.2 Estimation

The goal of our estimation is to separately identify permanent and short-term shocks and estimate the parameters (μ , σ_A , σ_P , and ρ). Because asset productivity is not observable, we represent the cash flow model in the state space form in Equations (1)–(2) and estimate the model using the most efficient method—namely, maximum likelihood with Kalman filtering. Because shocks are correlated, the standard Kalman filter is biased and inconsistent and, therefore, cannot be used. This problem is reminiscent of an endogeneity issue in regression analysis, with the major difference that the regressor (asset productivity) is unobserved and needs to be filtered out. We solve this problem by theoretically regressing $\sigma_A P_{t-1} \varepsilon_{i,t}^A$ on $\sigma_P P_{t-1} \varepsilon_t^P$ and by transforming the measurement equation (2). The new measurement error is then given by the residuals of this theoretical regression. Because of this transformation, we need to derive a novel Kalman filter that is described in Appendix A.

To provide insight into our estimation method (described in detail in Appendix A), we discuss its steps for the case with no missing observations (Appendix A.3 discusses how missing observations are handled). Using standard notation in state space models, the model in Equations (1)–(2) reads as

$$X_t = \Phi_X X_{t-1} + \omega_t \tag{4}$$

$$Z_t = H_Z X_t + u_t, \tag{5}$$

where X_t is the unobserved state process (latent asset productivity), $\Phi_X = (1 + \mu)$, ω_t is the transition shock distributed as $\mathcal{N}(0, \sigma_P^2 X_{t-1}^2)$, Z_t is the observed N -dimensional vector collecting firms' operating cash flows at time t , H_Z is an N -dimensional vector of ones, and u_t is the measurement error distributed as $\mathcal{N}(0, \sigma_A^2 X_{t-1}^2)$. In classic state space models, ω_t and u_t are uncorrelated. In our model, the correlation between permanent and short-term shocks translates into correlated ω_t and u_t . Because ρ is the contemporaneous correlation between the shock of the state variable X_t and the measurement error u_t , it cannot be handled by writing the state process in vector form, $\mathbf{X}'_t = (X_t, X_{t-1})$. Even when the state process is in vector form, the correlation between the first component of the state vector \mathbf{X}_t and the measurement error u_t is still there, making the vector-valued, standard Kalman filter not applicable. To account for this correlation, we theoretically regress u_t on ω_t and take the residual of this regression as the new measurement error. The measurement equation (5) changes as follows:

$$\begin{aligned} Z_t &= H_Z X_t + u_t + J(X_t - \Phi_X X_{t-1} - \omega_t) \\ &= H_Z^* X_t + \Phi_X^* X_{t-1} + u_t^*, \end{aligned} \tag{6}$$

where $H_Z^* = H_Z + J$, $\Phi_X^* = -J\Phi_X$ and $u_t^* = u_t - J\omega_t$. Note that Equations (5) and (6) are equivalent because $X_t - \Phi_X X_{t-1} - \omega_t = 0$. Setting $J = \mathbb{E}[u_t \omega_t | X_{t-1}] / \mathbb{E}[\omega_t^2 | X_{t-1}]$ yields that the new measurement error u_t^* is uncorrelated with X_t and X_{t-1} . Because the transformed measurement equation depends on X_{t-1} , the prediction step of Z_t and its error covariance matrix are different than in the standard Kalman filter. This difference leads to the generalized Kalman filter derived in Appendix A.2. We note that this method does not involve any approximation of the model in Equations (1)–(3). If $\rho = 0$, then $J = \Phi_X = 0$ and $H_Z^* = H_Z$, and the generalized Kalman filter reduces to the standard one.

Finally, given the model parameters ρ, σ_P, σ_A , and μ , the generalized Kalman filter recovers the unobserved state process X_t that determines the likelihood function of observed cash flows $Z_t, t = 1, \dots, T$,

$$\sum_{t=1}^T -\frac{1}{2} \left[N \log(2\pi) + \log |F_{t|t-1}| + (Z_t - \hat{Z}_{t|t-1})' F_{t|t-1}^{-1} (Z_t - \hat{Z}_{t|t-1}) \right], \quad (7)$$

where $Z_{t|t-1}$ is the one-step-ahead prediction of Z_t based on the filtered state process X_t and X_{t-1} , and $F_{t|t-1}$ is the error covariance matrix defined in Equation (A8). Model parameters are changed so as to increase the value of the log-likelihood, which requires re-running the generalized Kalman filter and re-computing the log-likelihood. The iterative procedure is repeated until convergence, which takes less than one second for a cash flow panel of $N = 10$ firms observed over $T = 48$ years. A Monte Carlo analysis, described in Appendix B, confirms that our generalized filter outperforms the standard Kalman filter, in which $\rho = 0$.¹⁰

2.3 Comparing methods to recover shocks

Our novel Kalman filter offers four main benefits over standard methods to separate a time series into a trend (persistent) component and a cyclical (transitory) component, such as the Hodrick–Prescott (HP) filter and the Beveridge–Nelson (BN) decomposition.¹¹

First, neither of these standard filters is suited to study correlations between permanent and short-term shocks. The HP filter assumes zero correlation

¹⁰ In the signal processing literature, Ma, Wang, and Chen (2010) have introduced a Kalman filter that can account for the correlation between the state variable and the measurement error. While for positive correlations their filter outperforms the standard Kalman filter, they find that for negative correlations, the two filters have a similar degree of inaccuracy. Appendix B.2 provides an in-depth discussion of the differences between the two filters and shows that our filter outperforms the filter developed by Ma, Wang, and Chen (2010).

¹¹ In the empirical corporate finance literature, the HP filter has been applied for instance by Byun, Polkovnichenko, and Rebello (2019b) and the BN decomposition by Chang et al. (2014). Given a time series $y_t, t = 1, \dots, T$, the HP and BN filters provide the additive decomposition $y_t = \tau_t + c_t$, where τ_t is identified as a trend component and c_t as a cyclical component. In the HP filter, the trend component $\tau_t, t = 1, \dots, T$, is obtained as the minimization of $\sum_{t=1}^T (y_t - \tau_t)^2 + \lambda \sum_{t=2}^{T-1} ((\tau_{t+1} - \tau_t) - (\tau_t - \tau_{t-1}))^2$, where the parameter $\lambda > 0$ controls the smoothness of the trend component. The BN filter is based on an ARMA model for $(y_t - y_{t-1})$ and identifies the trend component τ_t as a random walk with drift.

between trend and cyclical components, whereas the BN decomposition imposes perfectly correlated trend and cyclical shocks. This aspect is problematic because the shock correlation is key to capture firms' hedging needs. Second, neither is designed to recover volatilities of permanent shocks. As shown with formal tests in Section 1, permanent shocks are a major driver of firms' cash flows. Third, neither can handle missing observations. These methods were developed to decompose complete time series such as annual consumption or GDP. By contrast, Kalman filtering is by design a method of estimation and imputation of missing data (see Appendix A.3 for details). Many Compustat firms have holes in the cash flows time series. Because of these missing data, applying the HP or BN filter would force us to drop more than half of the 10,136 firms in our panel. Fourth, our method provides a decomposition while estimating the deep parameters of a canonical cash flow model. With other filtering methods, deep parameters can only be approximated with the calculation of time-varying moments over time series of filtered shocks.

To illustrate the benefits of using our method to recover cash flow shocks from uninterrupted cash flow time series (i.e., to illustrate benefits 1 and 2 above), we run a Monte Carlo simulation that consists of the following steps: (i) simulate one panel of cash flows for 10 firms over 50 years using model (1)–(3) with the parameter values in Décamps et al. (2017), that is, $\rho = -0.21$, $\sigma_P = 0.25$, $\sigma_A = 0.12$, $\mu = 0.01$; (ii) estimate the four model parameters using our method in Section 2.2, which we label as KF. For the HP and BN methods, (iii) apply the filter to each time series of cash flows to recover trend and cycle components at the firm level (using an adjusted smoothing parameter $\lambda = 6.25$ as in Byun, Polkovnichenko, and Rebello 2019b for the HP filter and an ARMA(2,2) model as in Chang et al. 2014 for the BN filter); (iv) compute the correlation between trend and cycle shocks and their standard deviations to obtain estimates of ρ , σ_P , and σ_A , and the mean of trend changes to obtain μ ; (v) average these estimates across the 10 firms; (vi) compute the estimation errors, defined as estimated values minus true parameter values, for KF, HP, and BN; (vii) repeat the procedure 1,000 times.

Figure 1 summarizes the simulation results. The HP and BN methods are systematically biased and vastly inaccurate. While KF can be expected to outperform HP and BN, as KF estimates the model that is used to simulate the cash flow data, HP and BN produce largely imprecise estimates on their own. A well-known limitation of the BN decomposition is that the correlation between trend and cyclical shocks is either +1 or -1 depending on the ARMA parameters, as BN relies on a one-shock-only ARMA model. Consequently, BN is unable to estimate the correlation ρ . The HP filter imposes zero correlation and nearly always overestimates the true negative correlations, with a median bias of 0.35. While HP tends to outperform BN, both methods substantially underestimate the volatility σ_P of permanent shocks and the drift μ , and largely overestimate the volatility σ_A of short-term shocks. For example, the true value

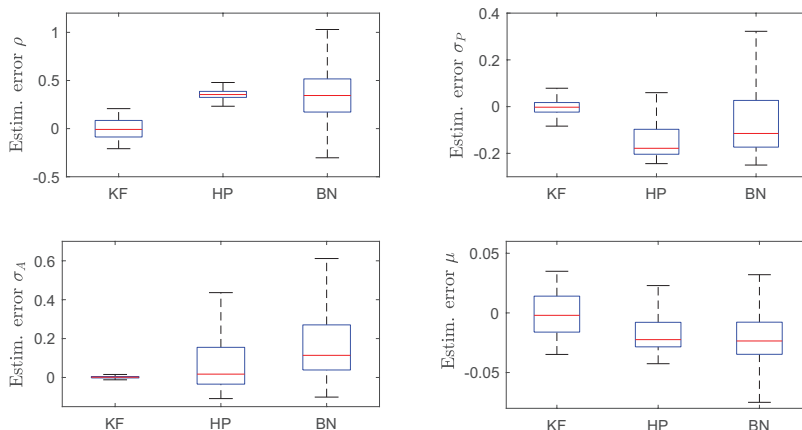


Figure 1
Estimation errors of cash flow shocks in simulated data

The model in Equations (1)–(3) with parameter values $\rho = -0.21$, $\sigma_P = 0.25$, $\sigma_A = 0.12$, $\mu = 0.01$ is used to simulate 1,000 panels of cash flows for 10 firms over 50 years. The box plots show the estimation errors of the permanent and short-term shock correlation, ρ ; the volatility of permanent and transitory shocks, σ_P and σ_A ; and the drift of the permanent process, μ . Cash flow shocks are recovered using the Kalman filter (KF), the Hodrick–Prescott filter (HP), and the Beveridge–Nelson decomposition (BN).

of σ_P is 0.25, while HP gives median estimates of 0.07 only. The KF method produces unbiased and accurate estimates of all four parameters.

Figure 2 provides an illustrative example of the performance of KF and HP. (Because HP outperforms BN, we do not report the latter in Figure 2 for readability.) For a simulated panel of cash flows, the figure shows the time series trajectory of the latent asset productivity (known in simulation), the Kalman-filtered asset productivity, and the HP-filtered trend component. The HP filter produces an estimated trajectory of the asset productivity that is too smooth, precisely because the HP filter is a cubic spline smoother. Taking the HP-filtered trend component as latent asset productivity would lead to significant underestimation of the volatility of permanent shocks; see, for instance, Hamilton (2018) for a recent discussion of the drawbacks of the HP filter. In contrast, the Kalman-filtered asset productivity closely tracks the true asset productivity. The Internet Appendix develops a simple example that shows that the HP filter introduces biases and spurious effects and that the Kalman filter does not.

In unreported Monte Carlo simulations, we draw the parameter values of $\rho, \sigma_P, \sigma_A, \mu$ from uniform distributions spanning the interquartile range of the parameter estimates reported in Table 4. In that case, simulated cash flows are more volatile than the cash flows mentioned previously. Consequently, estimation errors of HP and BN are much larger than those in Figure 1, whereas KF provides precise estimates of all model parameters.

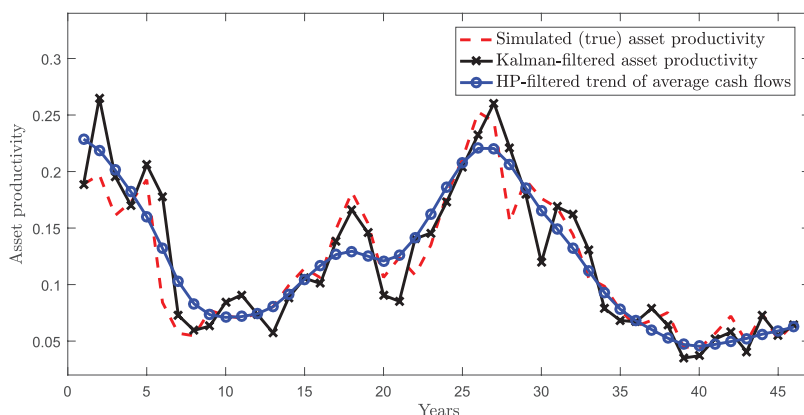


Figure 2

Generalized Kalman filter and Hodrick–Prescott filter applied to simulated cash flow data

Based on a panel of simulated cash flows from the model in Equations (1)–(3), the graph shows the time series trajectory of the true latent asset productivity, the Kalman-filtered asset productivity, and the trend component of the group average of cash flows from the HP filter. The Kalman-filtered asset productivity tracks closely the true asset productivity, while the HP filter, being a cubic spline smoother, provides a too-smooth approximation of the asset productivity.

2.4 Identifying assumptions

If there was no correlation between the shocks ($\rho=0$), the cash flow model in Equations (1)–(3) would be a classic state space model with Gaussian likelihood. The remaining three parameters would be identified from the unique global maximizer of the likelihood function (see Appendix B for details). As we show in Appendix A.4, identifying the shock correlation ρ requires imposing more structure on the model. To achieve identification, we assume in the following that transitory shocks are firm specific while permanent shocks have a “systematic” nature, in that they affect a group of firms. That is, we consider that there is a group-specific factor that moves the permanent component but not the transitory component of all firms in the group, where groups are defined in Section 2.5. Appendix A.4 indeed shows that ρ is not identified if both shocks are firm-specific.

To provide an intuition for this identifying assumption, consider a conventional model of asset productivity in which an observable process is driven by the sum of persistent shocks (modeled as an AR(1) process) and short-term shocks (modeled as a white noise process). When both shocks are unobservable, the challenge of identifying their correlation is similar to identifying the correlation ρ in Equations (1)–(3). As we formally show in Appendix A.4, the correlation between firm-specific persistent and short-term shocks is not identified because this parameter enters all autocovariances of the observed firm productivity as a multiplicative constant to both persistent and short-term volatilities, and thus the correlation cannot be disentangled from these volatilities. If instead, persistent shocks are common across firms,

as per our identifying assumption, then the time series of the cross-sectional average productivity provides additional and nonredundant moment conditions to identify the shock correlation.

The implication of this result is that the firm-by-firm estimation of the model in Equations (1)–(3) is infeasible. Individual firm estimation is also undesirable because annual cash flow time series are not very long. Our solution to achieve maximum granularity is to estimate the model parameters for the smallest possible groups of very similar firms, assuming each firm in the group is exposed to the same permanent shocks. Effectively, this assumption is significantly weaker than a common practice to assume that long-term productivity shocks are common to all firms in an industry. For example, Bates, Kahle, and Stulz (2009) use the volatility of the average cash flow over all firms in each two-digit SIC code. Similarly, Acharya, Almeida, and Campello (2007) and Duchin (2010) operationalize the firm’s hedging needs with the correlation between a firm’s current cash flow and the median or mean R&D expense over all firms with the same three-digit SIC code. As we shall see next, our estimation method achieves a very high level of precision with only 10 firms per group, rendering our commonality assumption almost innocuous.

Our identifying assumption that permanent shocks are common to a group of firms encompasses situations in which firms face common technology, regulatory, or consumer preference shocks. An alternative identifying assumption would be to consider that transitory shocks are common to a group of firms, while permanent shocks are firm-specific. This would encompass situations in which firms in the same group end up with different productivity growth paths but always face similar temporary disruptions, for example, weather shocks or common supply-chain disruptions. The econometric analysis of this alternative model—presented in the Internet Appendix—reveals two substantial problems. First, because the transitory shock is common and the asset productivity is vector valued (one asset productivity for each firm in any given group), the covariance matrix between short-term and permanent shocks is singular, which can create numerical instabilities when applying the Kalman filter. Second, missing values are pervasive in cash flow data, and it is unclear how to filter out the firm-specific asset productivity when cash flows are missing. These two problems hinder accurate estimation of this alternative model.

2.5 Grouping firms with similar cash flow dynamics

We estimate the cash flow model in (1)–(3) for each of many small groups of firms. We assume that all firms within each group g are homogeneous in that they have the same parameters μ_g , $\sigma_{P,g}$, $\sigma_{A,g}$, and ρ_g , and asset productivity, $X_{g,t}$.¹² Fitting the model to relatively small samples allows us to achieve

¹² In state space form, the cash flow model can therefore be rewritten as:

$$X_{g,t} = (1 + \mu_g) X_{g,t-1} + \sigma_{P,g} X_{g,t-1} \varepsilon_{g,t}^P$$

greater estimation accuracy because the model parameters can adjust to the data features of each specific group of firms. Moreover, we obtain a large set of estimates of the cash flow model's deep parameters, as opposed to just one or a few sets for their representative firms. Since the limiting case—which is to estimate the cash flow model firm-by-firm—is not feasible, estimation by small groups maximizes the cross sectional variation in these estimates and enables a direct test of the predicted link between deep parameter heterogeneity and corporate policies.

To group firms, we adopt two sequential criteria that are motivated by the assumption that permanent shocks are common to all firms in the group. The first is the three-digit SIC industry code. We expect firms within the same three-digit SIC industry to be exposed to similar short-term volatility (e.g., industry demand uncertainty) and similar permanent shocks (e.g., technology or regulatory shocks). The second is the firm's cash flow growth rate: Within each three-digit SIC industry, we group firms based on their average annual growth rate of cash flows. In the long run, firms with similar asset productivity will have similar average cash flow growth rates. For the precision of our parameter estimates, we impose the additional requirement that each group includes at least 10 firms whenever possible (see Appendix B). Because the number of firms in any given industry is not generally a multiple of 10, the last group of firms in each three-digit SIC code will include between 10 and 19 firms. In the rare cases in which there are fewer than 10 firms in a three-digit SIC industry, we include all firms in one group; this results in 43 groups of five to nine firms.

Applying the criteria above, our sample of 10,136 firms is split into 918 three-digit SIC/cash flow growth groups. As an example, Figure 3 shows the cash flows of one group of firms in the 100 SIC code. Missing observations in firm cash flows are evident in the interrupted time series of firm cash flows. Because our groups are relatively small, our parameter estimates can potentially exhibit substantial variation even within three-digit SIC industries.

To assess the homogeneity of firms within each group, we decompose the total variation of several firm-specific outcome and policy variables into the between- and within-group components. For each characteristic, we compare the similarity within and heterogeneity between our estimation groups to those implied by other narrow industrial classifications. Table 3 shows that, relative to the four-digit SIC or the 17 Fama and French (1997) industries, our classification produces less within-group variation for the ratios of annual sales-to-assets, earnings-to-assets, and average sales growth, as well as for key

$$Z_{i,g,t} = X_{g,t} + \sigma_{A,g} X_{g,t-1} (\rho_g \varepsilon_{g,t}^P + \sqrt{1 - \rho_g^2} \varepsilon_{i,t}^T),$$

where $X_{g,t}$ is the unobserved asset productivity with constant growth rate μ_g and volatility $\sigma_{g,P} > 0$, and $Z_{i,g,t}$ is the operating cash flow of firm i of group g , in year t , with short-term volatility $\sigma_{g,A} > 0$.

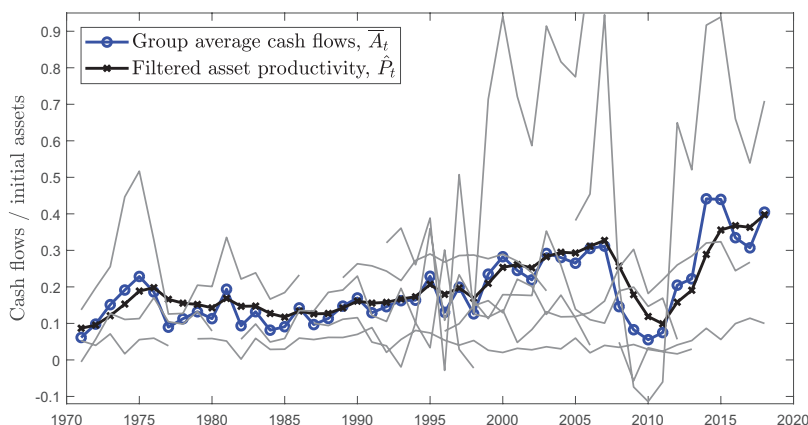


Figure 3
Cash flows for one group of firms
 This figure shows the yearly firm cash flows scaled by the initial level of assets for a select group of 10 firms in the 100 three-digit SIC code.

policy variables such as cash holdings, the rates of savings and equity issuance, the size of loans and credit lines, and the CAPEX-to-assets and debt-to-assets ratios. Our grouping also implies more within-similarity in the ratio of R&D expenditure to sales, the markups estimated following De Loecker, Eeckhout, and Unger (2020), the number of patents, and their market value, according to Kogan et al. (2017). Remarkably, grouping only by long-run similarity in the average cash flow growth rate within each three-digit SIC industry produces similarities across many other dimensions.

Table 3 shows that our grouping method also produces the most between-group variation for as many firm characteristics relative to the four-digit SIC or the 17 Fama and French (1997) industrial classifications. In a nutshell, our grouping approach produces many small and heterogeneous groups of alike firms. The implied high granularity of estimates is key for hypothesis testing.

3. Risk Parameter Estimates

Table 4 summarizes the maximum likelihood (ML) estimates of the model’s four parameters, σ_P , σ_A , ρ , and μ for all the 918 three-digit SIC/cash flow growth rate groups (panel A), their precision (panel B), and the correlation between the parameter estimates (panel C).¹³ We winsorize the estimates at the 1st and 99th percentiles when they approach their respective lower and upper

¹³ To alleviate the concern that shocks may not be normally distributed, we have also computed quasi maximum likelihood (QML) standard errors (White, 1982) as a robustness check. Based on maximum likelihood standard errors, Table 4 reports that 56%, 84%, and 76% of the estimates of the cash flow risk parameters ρ , σ_P , and σ_A , respectively, are statistically away from zero at a 5% level. Using QML standard errors yields similar results, and

Table 3
Decomposition of standard deviation by industries or estimation groups

	Standard deviation					
	Within-			Between-		
	SIC4	FF17	Groups	SIC4	FF17	Groups
1. Technology variables						
Capital-to-labor ratio	1.06	1.29	1.04	0.92	0.99	1.25
Operating leverage	0.76	0.89	0.75	0.57	0.40	0.61
2. Outcome variables						
Annual sales-to-assets	1.54	1.63	1.54	0.61	0.45	0.68
Annual earnings-to-assets	0.34	0.35	0.34	0.07	0.06	0.13
Annual sales growth	0.47	0.47	0.46	0.07	0.05	0.08
3. Policy variables						
Cash holdings	0.16	0.18	0.16	0.06	0.05	0.09
Cash savings	0.19	0.20	0.19	0.02	0.02	0.03
Net Equity issuance	0.28	0.28	0.27	0.07	0.05	0.08
Loans-to-assets	0.40	0.42	0.39	0.17	0.06	0.20
Total debt-to-assets	0.27	0.28	0.26	0.09	0.03	0.09
CAPEX-to-assets	0.07	0.07	0.07	0.03	0.03	0.04
4. Product market and innovation variables						
De Loecker et al. (2020) markups	0.38	0.44	0.39	0.13	0.12	0.25
R&D expense-to-assets	0.83	0.91	0.82	0.32	0.21	0.45
Number of patents	0.91	0.96	0.87	0.44	0.28	0.45
Kogan et al. (2017) market value of patents	1.22	1.28	1.17	0.51	0.33	0.55

This table shows the decomposition of the total standard deviation of several firm-specific outcome, financing policy, product market, and innovation variables into the between- and within-group standard deviations. Firms are grouped according to their four-digit SIC code (SIC4), their 17-industry classifications in Fama and French (1997) (FF17), or they are allocated into groups of 10 firms sorted by their average annual cash flow growth rate within each three-digit SIC code ("Groups"). The data are for all yearly observations of the 10,136 Compustat firms with at least (not necessarily consecutive) 10 years of cash flow data between 1971 and 2018. *Capital-to-labor ratio* is defined as Net PPE divided by Number of Employees; *Operating leverage* is SG&A plus Costs of Goods Sold divided by Total Assets. *Annual sales-to-assets* is Annual Sales divided by Total Assets, and *Annual earnings-to-assets* is Net Income divided by Total Assets. *Cash holdings* is Cash and Marketable Securities divided by Total Assets, and *Loans-to-assets* is the total amount principal outstanding in term loans and credit lines in Dealscan, divided by Total Assets. *Total debt-to-assets* is Short-term debt plus Long-term debt divided by Total Assets. *CAPEX-to-assets* is the Annual Capital Expense divided by Total Assets. The markups estimates come from De Loecker, Eeckhout, and Unger (2020). *R&D expense-to-assets* is the ratio of Annual R&D expense to Total Assets and the number and market value of the firm's patents from Kogan et al. (2017). All other variables are defined in Table 1.

bounds (i.e., near -1 and 1 for the shock correlation, and near zero for each volatility), and at the 10th and 90th percentiles otherwise.

3.1 Estimates of ρ

The estimates of the correlation between permanent and short-term shocks in panel A of Table 4 exhibit significant variation across groups. The 5th percentile of the estimated correlations is -0.23 , while the 95th percentile is 0.23 (with a minimum of -0.34 and a maximum of 0.58 , unreported). The median estimated correlation is -0.076 , with 80% of the estimates being

the percentages change to 60%, 74%, 73%, respectively. While some estimates of σ_P and σ_A are less significant, other estimates of ρ are more significant using QML standard errors. The inference in Table 4 is therefore robust to the misspecification of the shocks distribution.

Table 4
Summary of the parameter estimates of the cash flow model

A. Parameter estimates

	#est.	Mean	Standard deviation					p5	p25	p50	p75	p95
			Total	sd _b	sd _w							
$\hat{\rho}$	918	-0.055	0.143	0.107	0.121	-0.232	-0.122	-0.076	-0.023	0.229		
$\hat{\sigma}_P$	918	0.826	0.584	0.377	0.432	0.145	0.301	0.651	1.264	1.901		
$\hat{\sigma}_A$	918	0.755	1.411	0.905	1.225	0.029	0.060	0.093	0.357	4.486		
$\hat{\mu}$	918	0.148	0.104	0.069	0.081	0.021	0.064	0.119	0.218	0.345		

B. Values of *t*-statistics of the parameter estimates

#est.	Mean	Standard deviation	Proportion of					p5	p25	p50	p75	p95		
			<i>p</i> -values		p5	p25	p50						p75	p95
			< .05	< .01										
$\hat{\rho}$	918	-27.383	98.947	.558	.503	-94.612	-9.383	-1.058	-0.201	2.122				
$\hat{\sigma}_P$	918	37.595	59.465	.842	.757	1.635	2.612	5.425	42.477	181.635				
$\hat{\sigma}_A$	918	9.217	36.124	.758	.622	1.001	2.003	3.319	5.791	19.838				
$\hat{\mu}$	918	7.921	51.590	.464	.386	-3.415	0.808	1.368	3.937	31.186				

C. Correlations between the parameter estimates

	$\hat{\rho}$	$\hat{\sigma}_P$	$\hat{\sigma}_A$	$\hat{\mu}$
$\hat{\rho}$	1.000			
$\hat{\sigma}_P$	-0.196***	1.000		
$\hat{\sigma}_A$	-0.243***	0.445***	1.000	
$\hat{\mu}$	0.038	0.471***	0.176***	1.000

This table summarizes the maximum likelihood estimates of the cash flow model parameters, $\hat{\mu}, \hat{\sigma}_P, \hat{\sigma}_A, \hat{\rho}$, in the model

$$P_t = (1 + \mu)P_{t-1} + \sigma_P P_{t-1} \varepsilon_t^P$$

$$A_{i,t} = P_t + \sigma_A P_{t-1} \varepsilon_{i,t}^A$$

where $\varepsilon_{i,t}^A = \rho \varepsilon_t^P + \sqrt{1 - \rho^2} \varepsilon_{i,t}^T$, the correlation $\rho \in (-1, 1)$, P_t is the unobserved asset productivity, and $A_{i,t}$ are firm *i* cash flows in year *t*, for $i = 1, \dots, N$ and $N = 10$. The permanent shock ε_t^P and transitory shock $\varepsilon_{i,t}^T$ are uncorrelated and distributed as $\mathcal{N}(0, 1)$. The model parameters are estimated for each of the 918 three-digit SIC/cash flow growth groups of firms. The descriptive statistics are: number of model parameter estimates (#est.); mean; (Total) standard deviation, decomposed into between- (sd_b), and within-three-digit SIC industry (sd_w) variation; and the percentiles p5, p25, p50, p75, and p95. The sample covers the period 1971 to 2018.

negative. As shown in Section 4, this estimated correlation implies, for example, that the median firm issues approximately 5% more net equity than a firm with an estimated correlation of zero. Panel B additionally shows that close to 56% of the estimated ρ are significantly different from zero with 95% confidence.

A unique feature of cash flow models that include correlated permanent and short-term shocks is that liquidity policies depend crucially on the sign of the correlation between the two shocks. Our framework enables the testing of such unique predictions because our estimates of ρ exhibit significant sign heterogeneity.

3.2 Estimates of σ_P and σ_A

The estimates of the volatility parameters in panel A also exhibit significant variation across groups. The median estimated permanent shock volatility,

$\hat{\sigma}_P$, is 65.1%. The total standard deviation of 58% is explained mostly by within rather than between three-digit SIC variation. Interestingly, since we only exclude firms with annual asset growth rates above 500%, the estimates of $\hat{\sigma}_P$ exceed 126% for 25% of the firms. That is, we can estimate and report the magnitudes of the highest permanent shocks volatilities in the Compustat universe. The median estimated short-term shocks volatility, $\hat{\sigma}_A$, is 9.3%. Estimates also vary significantly within the three-digit SIC classification and include very high values for 5% of the groups. Note that we can use our parameter estimates to recover the volatility of transitory shocks by evaluating $\hat{\sigma}_A\sqrt{1-\hat{\rho}^2}$. Because the correlation between short-term and permanent shocks is not zero for many firms, the median (mean) volatility of transitory shocks is smaller than for $\hat{\sigma}_A$: 9.2% (74.4%). The distribution is as skewed as for $\hat{\sigma}_A$ with an interquartile range between 5.9% and 34%.

We show in Section 3.6 that our estimates imply asset volatilities that are comparable to those of actual Compustat firms. As shown in Section 4, this is due to the fact that firms with higher exposure to permanent or transitory shocks engage more actively in risk management policies (broadly defined), leading to a significant smoothing of earnings and asset volatilities.

3.3 Estimates of μ

The estimates of μ exhibit an interquartile range from 6.4% to 21.8%, with an average productivity growth of 14.8%. Almost 50% of the estimates are significantly different from zero with 95% confidence. These estimates also show that our grouping procedure captures important differences in latent productivity growth rates across and within industries.

3.4 Understanding the estimates: Permanent shocks volatilities

Our cash flow model and filtering technique interpret the very high cash flow volatilities in the data as rather moderate volatilities of permanent and transitory shocks. Indeed, while the median standard deviations of annual cash flow growth in our sample is a very high 210%, we obtain median estimates of 65.1% for σ_P and 9.3% for σ_A . This inference is explained by the nonlinear interaction between permanent and transitory shocks in the cash flow model and the fact that the estimates of σ_A and σ_P are positively correlated.

To better understand how our model infers each volatility individually, consider first the case of σ_P . Let \bar{Z}_t denote the cross-sectional average cash flow within each group of $N=10$ firms, that is, $\bar{Z}_t = \sum_{i=1}^N Z_{i,t}/N$, for $t=1, \dots, T=48$ years. Let R_t denote the relative change of \bar{Z}_t . This observable rate is approximately equal to the relative change of the unobservable X_t (with equality if $N \rightarrow \infty$, and $\rho=0$ or $\sigma_A=0$), that is,

$$R_t \equiv \frac{\bar{Z}_t - \bar{Z}_{t-1}}{\bar{Z}_{t-1}} \approx \frac{X_t - X_{t-1}}{X_{t-1}}. \quad (8)$$




Figure 4
Scatter plot of permanent shock volatilities

The x-axis reports the model-based estimated volatilities of permanent shocks, $\hat{\sigma}_P$ in the model in Equations (1)–(3), for each group of firms in our sample. The y-axis reports the observed time series standard deviations of R_t in Equation (8) for $t = 1, \dots, T$, that is, the relative change of group-specific average cash flows. The sample data covers 10,136 firms, from 1971 to 2018, sorted in 918 groups.

To illustrate, Figure 3 shows \bar{Z}_t superimposed to firm cash flows for a select group of low-correlation firms ($\hat{\rho} = 0.1$): The time series trajectory of \bar{Z}_t mimics the filtered asset productivity \hat{X}_t for this group. Absent short-term shocks ($\varepsilon_{i,t}^A = 0$), the time series standard deviation of R_t would be approximately equal to σ_P , as $Z_{A_{i,t}} = X_t$ in (2). In general, the approximation in (8) implies that the estimated volatility of permanent shocks is expected to be positively related to, albeit smaller than, the time series standard deviation of R_t .

Figure 4 shows the scatter plot of the time series sample standard deviations of R_t against model-inferred volatilities of permanent shocks, $\hat{\sigma}_P$, for the 918 groups in our sample. As predicted, this figure shows a strong positive association between standard deviations of R_t and the estimates of σ_P , providing direct evidence that our estimates of σ_P are capturing the volatility of permanent shocks. Moreover, standard deviations of R_t are generally larger than the estimates of σ_P , meaning that short-term shocks are present in cash flow data.

3.5 Understanding the estimates: Cash flow shock correlation

To the best of our knowledge, there exist only two previous attempts (i.e., Duchin, 2010 and Acharya, Almeida, and Campello, 2007) in the literature to operationalize the notion of correlation between current cash flow and investment opportunities. Our estimates are the first to directly target the notion of hedging needs by way of a deep parameter of a dynamic cash flow model. Our estimates differ from those in the literature. Most importantly, Duchin (2010) and Acharya, Almeida, and Campello (2007) report mostly positive correlation between current cash flow and investment opportunities, whereas we obtain

Table 5
Estimates of the correlation between permanent and short-term cash flow shocks by industry

Fama-French industry	Number of firms	Mean	Standard deviation	p5	p25	p50	p75	p95
Food	418	-0.01	0.18	-0.19	-0.10	-0.06	0.02	0.58
Mining and Minerals	494	-0.08	0.12	-0.27	-0.15	-0.10	-0.02	0.09
Oil and Petroleum Products	723	-0.06	0.13	-0.23	-0.14	-0.08	-0.02	0.21
Textiles, Apparel and Footwear	282	-0.03	0.14	-0.20	-0.13	-0.07	0.08	0.24
Consumer Durables	373	-0.06	0.14	-0.21	-0.12	-0.09	-0.07	0.20
Chemicals	209	-0.04	0.18	-0.17	-0.11	-0.08	-0.04	0.58
Drugs, Soap, Perfumes and Tobacco	330	-0.07	0.16	-0.23	-0.18	-0.08	-0.02	0.17
Construction and Construction Materials	516	-0.04	0.15	-0.17	-0.12	-0.07	-0.02	0.29
Steel Works	192	0.05	0.26	-0.34	-0.13	-0.07	0.32	0.58
Fabricated Products	125	0.01	0.22	-0.24	-0.09	-0.07	0.12	0.58
Machinery and Business Equipment	1,431	-0.06	0.16	-0.26	-0.14	-0.08	-0.03	0.27
Automobiles	179	-0.02	0.10	-0.13	-0.09	-0.06	-0.01	0.30
Transportation	488	-0.05	0.11	-0.17	-0.10	-0.08	-0.05	0.20
Retail Stores	761	-0.04	0.11	-0.20	-0.10	-0.06	-0.01	0.20
Other	3,615	-0.06	0.13	-0.23	-0.13	-0.08	-0.04	0.16

This table summarizes the distribution of the maximum likelihood estimates of the correlation between permanent and short-term cash flow shocks, $\hat{\rho}$. This parameter is estimated, together with the other cash flow model parameters, for each of the 918 three-digit SIC/cash flow growth groups of firms, using firm-specific cash flow data from 1971 to 2018. The summaries show the number of firms, and the mean, standard deviations, and percentiles p5, p25, p50, p75, and p95 of $\hat{\rho}$ for all firms in each industry of the 17-industry classification in Fama and French (1997).

mostly negative estimates of ρ .¹⁴ An important reason for these measurement differences may be that previous proxies assume that a firm's investment opportunities are given by its industry mean or median R&D expense. Our approach does not require a proxy for the state of industry technology based on a policy variable. We are also agnostic about which firm may lead the state of technology in the industry. Under the mild assumption that technology is common to a group of only 10 firms, we find significant variation in $\hat{\rho}$ within three-digit SIC classification, with an important proportion of negative estimates—that is, high hedging needs.

An important question is which industry or firm characteristics are associated with variation in the estimates of ρ , with the objective to understand the deeper differences across firms that are captured by our estimates. Indeed, the sign and magnitude of shock correlations may be a technological characteristic intrinsic to each industry. For example, Froot, Scharfstein, and Stein (1993) conjecture that technologies requiring different degrees of operating leverage may lead to different sensitivities of investment opportunities to demand shocks and, therefore, different hedging needs.

Table 5 explores these conjectures. The table shows that the distributions of the estimates of ρ across all the 17 Fama–French industries (FF17) are

¹⁴ Duchin (2010) reports a median cash-flow-investment opportunities correlation of 0.25, i.e., a large majority of firms with a positive correlation and, hence, low hedging needs. Acharya, Almeida, and Campello (2007) do not report a summary, but one can infer that about 39% of firm-year observations have a correlation higher than 0.2 (the “Low hedging needs” subsample). Although not the focus of their paper, Chang et al. (2014) report a negative realized correlation of -0.21 between the trend and cycle components of cash flow, although theoretically their BN decomposition assumes a correlation of -1 .




Figure 5
Estimates of the correlation between permanent and short-term cash flow shocks

This figure presents the box plots of the maximum likelihood estimates of the correlation coefficient, ρ , in the cash flow model of Equations (1)–(3) for all firms within each of the 17 industries defined by Fama and French (1997). The sample data cover 10,136 firms, from 1971 to 2018.

remarkably similar: All of the medians are negative, ranging between -0.10 and -0.06 . Except for four FF17 industries, fewer than 25% of the firms in each industry have a positive ρ . However, at least 5% of firms in each FF17 industry have a positive ρ . For most industries, there exist small groups of firms with significantly high positive estimates of ρ (Figure 5). In sum, the estimates of ρ , which have been obtained independently for small groups of firms within each industry, have very similar distributions across different industries.

Next, we ask whether the estimated correlations are associated with policy choices and outcome variables within each industry. Table 6 reports the average Spearman’s rank correlation coefficients between the estimates of ρ and variables capturing risk choices and risk measures. All ranks and correlations are computed annually by industry and reported as an average over all industry-years. We compute the rank correlations for different industry definitions: three- and four-digit SIC codes, and the 17 Fama–French industries.

Table 6 highlights several interesting results. First, firms that are naturally hedged due to a positive shock correlation tend both to choose riskier policies and to have lower overall risk than other firms. Notably, the within-industry comparisons reveal that firms with the highest estimated shock correlations

Table 6
Cash flow risk parameters: Within-industry comparisons

	$\hat{\rho}$			$\hat{\sigma}_A$			$\hat{\sigma}_P$		
	SIC3	SIC4	FF17	SIC3	SIC4	FF17	SIC3	SIC4	FF17
1. Risk choices									
Capital-to-labor ratio	0.74	0.68	0.65	-0.07	-0.08	0.17	-0.13	-0.14	0.09
Operating leverage	0.73	0.66	0.63	-0.25	-0.25	-0.59	-0.17	-0.18	-0.49
Total debt-to-assets	0.76	0.70	0.67	-0.36	-0.33	-0.61	-0.47	-0.45	-0.65
Acquisitions-to-assets	0.13	0.11	-0.23	-0.08	-0.07	0.10	-0.16	-0.16	0.04
2. Risk outcomes									
Equity return volatility	-0.16	-0.12	-0.43	0.56	0.52	0.61	0.71	0.68	0.63
Asset return volatility	-0.14	-0.10	-0.45	0.62	0.58	0.66	0.77	0.76	0.70
Distance to default	0.74	0.67	0.66	-0.07	-0.15	0.00	-0.27	-0.33	-0.15
Loans maturity	-0.02	0.00	0.22	-0.30	-0.27	-0.24	-0.42	-0.40	-0.20
Loans spread	-0.02	-0.05	-0.23	0.22	0.26	-0.05	0.26	0.30	0.04

This table shows the within-industry Spearman's rank correlation between each of the firm characteristics listed and the estimates of the operating earnings' risk parameters: the volatilities of permanent shocks (σ_P), short-term shocks (σ_A), and their correlation, $\hat{\rho}$. The within-industry rank correlations are the result of sorting all firms within each industry by each parameter and each characteristic. Industries are defined using several classifications: three- and four-digit SIC codes (SIC3, SIC4), and the 17-industry classification in Fama and French (1997) (FF17). The risk parameters are estimated, together with the other cash flow model parameters, for each of the 918 three-digit SIC/cash flow growth groups of firms, using firm-specific cash flow data from 1971 to 2018. Capital-to-labor ratio is defined as Net PPE divided by the Number of Employees; Operating leverage is SG&A plus Costs of Goods Sold divided by Total Assets; Total debt-to-assets is Short-term debt plus Long-term debt divided by Total Assets; Acquisitions-to-Assets is the total value of Acquisitions in the year divided by Total Assets. Equity return volatility is the annualized standard deviation of daily stock returns; Asset return volatility and Distance to default are calculated using the Bharath and Shumway (2008) method. Loans maturity and Loans spread are the average maturities and yield spreads, respectively, of all of the sample firms' outstanding loans in Dealscan.

tend to have the largest capital-to-labor ratios, operating leverage, debt-to-asset ratios, and acquisition-to-asset ratios within their industry. Yet, despite these riskier policy choices, these firms also tend to have lower default risk (as measured using Bharath and Shumway's 2008 distance to default), lower loan spreads (despite higher loan maturities), lower equity volatility, and lower asset volatility. The results are robust to fine (SIC3, SIC4) or broad industry definitions (FF17). Second, firms with higher exposure to permanent or transitory shocks choose more conservative policies but have higher overall risk. This again applies to all the policy choices and risk measures we look at and to all industry definitions. Importantly, while Table 6 is mostly concerned with correlations, Section 4 goes one step further by showing how our deep parameter estimates can explain the dynamics of liquidity management.

3.6 Empirical versus implied asset return volatilities

An important question is whether our estimates of the characteristics of cash flow shocks are meaningful. This section employs these estimates in a setting that is independent from the estimation procedure and shows that the estimated parameters imply asset return volatilities that match the actual asset return volatilities of the firms in our sample. Importantly, empirical asset return volatilities are not used in the estimation of cash flow characteristics. Equally importantly, cash flow characteristics have been estimated using cash flow data from operating earnings, without imposing any model restriction about corporate policies.

Table 7
Model-implied asset volatilities

	N	Mean	Stdev	p5	p25	p50	p75	p95
Empirical asset return volatility	703	0.476	0.155	0.270	0.358	0.447	0.584	0.740
Model-implied asset return volatility								
Baseline parameters	703	0.559	0.317	0.166	0.307	0.495	0.716	1.340
Baseline and $r=0.05$	555	0.588	0.386	0.143	0.296	0.490	0.761	1.570
Baseline and $\Phi=0.005$	701	0.538	0.304	0.161	0.301	0.477	0.679	1.288
Baseline and $\xi^T = \xi^P = 0.6$	845	0.491	0.241	0.177	0.295	0.448	0.626	1.041

This table presents a comparison of the distributions of empirical and model-implied asset volatilities. The volatilities are estimated for each of the 918 three-digit SIC/cash flow growth groups of firms. The empirical asset return volatilities are estimated as weighted averages of equity and debt return volatilities following the approach in Bharath and Shumway (2008) using daily stock returns. Model-implied asset volatilities are calculated using the model of Décamps et al. (2017) (see Appendix C) and the estimated cash flow parameters reported in Table 4. The remaining parameters are $r=0.08$, $\lambda=0.02$, $p=1.06$, $\Phi=0.002$, $\eta^P = \eta^T = 0.4$, and $\xi^T = \xi^P = 0.4$. Model-implied asset volatilities are winsorized at p5 and p95. The descriptive statistics are: number of observations (N), mean, standard deviation, and the percentiles p5, p25, p50, p75, and p95.

To compute the model-implied asset return volatilities, we employ the model of Décamps et al. (2017) presented in Appendix C. The model uses the (continuous-time version of the) cash flow model described by Equations (1)–(3) and solves for optimal financing and liquidity policies and firm value. It thus quantitatively maps the cash flow parameters (μ , σ_A , σ_P , and ρ) to asset return volatility. We calculate asset return volatilities at the group level, consistently with the level of granularity of the estimation of cash flow parameters. The empirical asset return volatilities are estimated as weighted averages of equity and debt return volatilities following the approach in Bharath and Shumway (2008) using daily stock returns. Model-implied volatilities are averages of all firm-level volatilities within a group using the estimated cash flow parameters reported in Table 4 for each group of firms. Details are presented in Appendix C. It should be noted that the exercise we conduct here is necessarily a joint test of the estimates of the characteristics of cash flow shocks as well as of the assumptions of the corporate liquidity management model of Décamps et al. (2017).

Table 7 reports the empirical and model-implied asset return volatilities. The average and median empirical asset return volatilities are 0.476 and 0.447, respectively. In the baseline case, the model-implied volatilities are slightly higher, at 0.559 and 0.495, respectively, but very close to the empirical ones. It is remarkable that model-implied asset return volatilities appear to match actual asset return volatilities that were not used during the estimation process. Additional rows for model-implied distributions in Table 7 present calculations based on alternative model parameters. The results show that the similarity between the empirical and model-implied distributions is robust and driven by the estimates of the parameters of the cash flow dynamics rather than the assumed values for the remaining parameters.

4. Understanding Liquidity Management

This section shows how our deep parameter estimates can be used to improve our understanding of corporate policies. Because the estimated parameters

characterize cash flow risk, it is natural to explore their effects on policy choices that increase the resilience of firms to cash flow shocks, such as liquidity management. Recent dynamic liquidity management models predict that firms build up cash reserves toward their target level either by retaining earnings or by raising outside equity and keeping (part of) the proceeds in cash reserves (see, e.g. Bolton, Chen, and Wang, 2011; Bolton, Chen, and Wang, 2013; Décamps et al., 2017). In our empirical analysis, we thus focus on these two separate mechanisms to manage cash reserves. In the first part of the section, we use our estimates to address the debate on the sign of the cash flow sensitivity of cash savings. In the second and third parts of the section, we demonstrate how our estimates of cash flow risk parameters relate to firms' decisions to respectively issue equity or debt to rebuild cash buffers.

4.1 Cash savings

One way for firms to increase their cash reserves, and thus their resilience to shocks, is to retain part of their earnings. There is considerable debate in the literature as to whether the sign of the cash flow sensitivity of cash is positive or negative. In an influential paper, Almeida, Campello, and Weisbach (2004) theoretically argue and provide evidence that the sensitivity is positive, if the firm is financially constrained, or zero. Riddick and Whited (2009), by contrast, show that it can be negative if productivity shocks are sufficiently persistent. Correcting for measurement error in Tobin's q , they find a negative average sensitivity. Décamps et al. (2017) sharpen the prediction in their dynamic model, in which the cash flow sensitivity of cash is proportional to $\rho \times \frac{\sigma_P}{\sigma_A}$. That is, they demonstrate that the sign of the cash flow sensitivity of cash is equal to the sign of ρ and that its absolute value should be higher for higher ratios of σ_P/σ_A .¹⁵ Using our parameter estimates, we can directly test this prediction and address this debate.

To analyze how the cash flow sensitivity of savings depends on hedging needs and the volatilities of cash flow shocks, we estimate the cash flow sensitivity of cash over several subsamples formed using our estimates of ρ and σ_P/σ_A . We first partition our sample based on the estimates $\hat{\rho}$. Because of estimation error in $\hat{\rho}$, we do not choose zero as the exact switching threshold of the cash flow sensitivity. Instead, we perform the tests over the subsamples of firms with $\hat{\rho}_i \leq -0.03$ (covering 73% of our sample firms) and $\hat{\rho}_i \geq 0.03$ (covering an additional 16% of our sample firms). For robustness, we also perform the tests over the subsamples $\hat{\rho}_i \leq -0.02$ and $\hat{\rho}_i \geq 0.02$, covering 92% of our sample firms.

¹⁵ The intuition for this result builds on the observation that higher productivity leads to a higher marginal value of cash and thus an increased propensity to save. In firms with positive ρ , a positive cash flow shock coincides on average with a positive productivity shock and leads to increased cash savings. Thus, positive ρ is associated with a positive cash flow sensitivity of cash. The opposite occurs for negative ρ . When σ_P is large relative to σ_A , then shocks to the propensity to save cash (proportional to productivity shocks and σ_P) are large relative to cash flow shocks (proportional to σ_A).

Sensitivity sign differences would be most clearly detected among firms with higher absolute sensitivities. Hence, we estimate the cash flow sensitivities of cash over increasingly restrictive subsamples based on the distribution of the ratio σ_P/σ_A , with values above the median, the 60th percentile, or the 70th percentile. The results are robust to measurements in the top tercile or quartile. For each of the resulting six subsamples (combining two sets of values for $\hat{\rho}$ and three sets for $\hat{\sigma}_P/\hat{\sigma}_A$), we estimate the cash savings regression in Almeida, Campello, and Weisbach (2004) and Riddick and Whited (2009):

$$\begin{aligned} \text{Cash savings}_{i,t} = & \beta_0 + \delta_i + \beta_{CF} \times \text{Cash flow-to-assets}_{i,t} + \beta_{\text{Controls}} \\ & \times \text{Controls}_{i,t-1} + u_{i,t} \end{aligned} \tag{9}$$

in which *Cash savings* is the yearly change in the stock of cash divided by total assets. Control variables include the *Market-to-book ratio* and $\ln(\text{Total assets})$.

Riddick and Whited (2009) note that this cash savings regression includes the *Market-to-book ratio*, i.e., average q , as a proxy for marginal q . As a result of measurement error in this variable, the OLS estimator of the propensity to save is biased toward zero. Consistency is achieved by additionally matching higher moments of the joint distribution of the dependent variable, *Cash savings*, and average q . Therefore, we follow Riddick and Whited (2009) and estimate the propensity to save using the fourth-order linear cumulants estimator (LC4) of Erickson, Jiang, and Whited (2014). Panel A of Table 8 presents these estimates. In follow-up research, Almeida, Campello, and Galvao (2010) discuss conditions under which instrumental variables estimators may be more robust and efficient than high-order moments-based estimators, such as LC4. We therefore also estimate Equation (9) using the dynamic panel data GMM estimator of Arellano and Bond (1991) (to address the measurement error problem in average q), implemented exactly as in Almeida, Campello, and Galvao (2010). Panel B of Table 8 presents the results.

Panel A of Table 8 shows that the estimated sensitivity, β_{CF} , is negative and statistically significant when $\hat{\rho}_i \leq -0.03$ but positive and statistically significant when $\hat{\rho}_i \geq 0.03$. The estimated sensitivity remains negative for $\hat{\rho}_i \leq -0.02$ and positive for $\hat{\rho}_i \geq 0.02$, confirming that the switching of sign in the propensity to save is not driven by observations when $\hat{\rho}$ may be close to zero. Remarkably, the estimated sensitivity, β_{CF} , exhibits the predicted sign switch for all 12 sets of test partitions. That is, we find that hedging needs (i.e., ρ) are a defining feature of the cash flow sensitivity of cash. As in panel A, the results in panel B show that the sign of the cash flow sensitivity of savings is determined by the sign of our estimate of ρ , verifying one of the unique predictions following from having correlated short-term and permanent cash flow shocks.¹⁶

¹⁶ Estimation by OLS yields the same qualitative results: the sign of the propensity to save equals the sign of $\hat{\rho}$. Confirming attenuation bias, the OLS estimates are smaller in magnitude than the LC4 or GMM-AB. We also

Table 8
Savings sensitivity of cash flow and cash flow shocks

A. LC4 estimator

Subsamples by $\hat{\rho}_i$:	Subsamples by values of $\frac{\hat{\sigma}_P}{\hat{\sigma}_A}$:			Subsamples by $\hat{\rho}_i$:	Subsamples by values of $\frac{\hat{\sigma}_P}{\hat{\sigma}_A}$:		
	Above 50%	Above 60%	Above 70%		Above 50%	Above 60%	Above 70%
$\hat{\rho}_i \leq -0.03$	-0.169*** (0.022)	-0.193*** (0.025)	-0.194*** (0.028)	$\hat{\rho}_i \leq -0.02$	-0.172*** (0.022)	-0.196*** (0.025)	-0.199*** (0.028)
	[55,734]	[44,480]	[31,723]		[57,377]	[46,123]	[33,366]
$\hat{\rho}_i \geq 0.03$	0.122*** (0.035)	0.114*** (0.036)	0.123*** (0.038)	$\hat{\rho}_i \geq 0.02$	0.105*** (0.029)	0.098*** (0.031)	0.104*** (0.033)
	[20,842]	[15,713]	[13,079]		[21,758]	[16,629]	[13,546]

B. AB-GMM estimator

$\hat{\rho}_i \leq -0.03$	-0.043*** (0.006)	-0.042*** (0.007)	-0.029*** (0.008)	$\hat{\rho}_i \leq -0.02$	-0.045*** (0.006)	-0.044*** (0.007)	-0.033*** (0.008)
	[48,228]	[38,372]	[27,279]		[49,646]	[39,790]	[28,697]
$\hat{\rho}_i \geq 0.03$	0.053*** (0.009)	0.039*** (0.010)	0.047*** (0.011)	$\hat{\rho}_i \geq 0.02$	0.051*** (0.008)	0.037*** (0.010)	0.047*** (0.011)
	[18,144]	[13,688]	[11,393]		[18,929]	[14,473]	[11,783]

This table presents estimates of the sensitivity of cash savings to cash flow, which are obtained from the slope coefficient of the regression of the yearly change in the stock of cash divided by total assets (*Cash savings*) on the firm's *Cash flow-to-assets*. Control variables include the lagged logarithm of *Total assets* and *Market-to-book ratio*. The sample period is from 1971 to 2018. The data are sorted and classified into subsamples according to the ratio $\hat{\sigma}_P/\hat{\sigma}_A$, and $\hat{\rho}$, which are the ratio of the estimated volatilities of and correlations between permanent and short-term cash flow shocks, respectively, common to all firms in the same three-digit SIC3/cash flow growth group. The coefficients in panel A are estimated using the fourth-order linear cumulants estimator (LC4) following Erickson, Jiang, and Whited (2014). In panel B, the coefficients are computed using the GMM dynamic panel data estimator proposed by Arellano and Bond (1991), following the implementation of Almeida, Campello, and Galvao (2010), using the first and second lags of the *Market-to-book ratio* as an instrument (AB-GMM). In both cases, standard errors (in parentheses) are computed using the optimal GMM weighting matrix. Please refer to Table 1 for the definition of all the variables. * $p < .1$, ** $p < .05$, *** $p < .01$

4.2 Equity issues

Another way for firms to replenish cash reserves is to raise outside funds by issuing equity, as empirically shown by Kim and Weisbach (2008) or McLean (2011). Issuance costs of securities generally deter firms from continuously raising funds to be at their target cash level. Instead, firms remain inactive (away from their target) for long spells until the benefits of raising funds to increase cash reserves outweigh the costs. Testing dynamic liquidity models thus requires distinguishing points at which firms are at (or move to) their target level of cash reserves from points at which they are not. To isolate such optimality points, we examine periods when firms simultaneously raise outside equity and increase their cash reserves. The optimality of such increases in cash reserves follows directly from dynamic inventory models because these refinancing points reflect optimal liquidity choices (see, e.g.,

explore robustness to using the OLS-IV estimator of Biorn (2000), again implemented exactly as in Almeida, Campello, and Galvao (2010). We also obtain the sign reversion result, if with less precision for the subsample with positive ρ estimates. Quantitatively, the magnitudes of the sensitivities under this estimator are also higher than OLS. Our results are also robust to using lagged instead of contemporaneous cash flow-to-assets as a regressor. All the results are included in Section 4 of the Internet Appendix.

Bolton, Chen, and Wang, 2011 or Décamps et al., 2017). Additionally, and as discussed in Danis, Retzl, and Whited (2014), “large decisions likely follow considerable deliberation, so it is hard to imagine that managers view these adjustments as suboptimal.”

Dynamic liquidity management models predict not only that firms will adjust their cash buffer infrequently but also that the frequency and size of the adjustments should be related to cash flow risk characteristics. For instance, theory predicts that firms with high permanent or transitory shock volatility should hold, on average, larger cash reserves and issue larger amounts of equity. In addition, they should raise external funds more frequently. By contrast, firms with a high correlation between permanent and short-term cash flow shocks, which are naturally hedged, should hold smaller cash reserves and issue smaller amounts of equity, and do so less frequently.¹⁷

To analyze the relation between the estimated cash flow risk parameters ($\hat{\rho}$, $\hat{\sigma}_P$, and $\hat{\sigma}_A$) and equity issuance, we use a standard liquidity regression model (as in, e.g., McLean, 2011) that we augment with our estimated parameters:

$$Y_{i,t} = \beta_0 + \delta_t + \gamma_j + \beta_\rho \times \hat{\rho}_g + \beta_{\sigma_P} \times \hat{\sigma}_{P,g} + \beta_{\sigma_A} \times \hat{\sigma}_{A,g} + \beta_{\text{Controls}} \times \text{Controls}_{i,t-1} + u_{i,t}. \quad (10)$$

We use the subscripts i for firms, g for groups of firms, and t for years. The dependent variable is either (i) gross equity issues scaled by lagged assets (*Gross equity issuance*) as in McLean (2011), (ii) gross equity issuance minus dividends and share repurchases scaled by lagged assets (*Net equity issuance*), (iii) a dummy variable equal to one if *Gross equity issuance* is larger than 5%, and zero otherwise (*Equity issuance dummy*), or (iv) a dummy variable equal to one if *Net equity issuance* is larger than 5%, and zero otherwise (*Net equity issuance dummy*). To capture the instances when firms are likely to channel the proceeds from equity issues to cash reserves, we estimate these regressions using the subsample of firm-years where firms experience a positive change in the cash-to-asset ratio (the sample median of the change of cash-to-asset ratio is zero) in that year. The results are robust to restricting the sample to firm-years in which the change in the cash-to-asset ratio is in the top tercile or in the top quartile of the change of the cash-to-asset sample distribution.

As in Bates, Kahle, and Stulz (2009), we control for the lagged *Industry cash flow volatility*, which varies yearly at the two-digit SIC industry, for firms’ growth opportunities with the residuals of the lagged market-to-book ratio

¹⁷ These firms have low hedging needs because low (high) short-term cash flows tend to occur only when productivity is declining (improving) (see, e.g., Froot, Scharfstein, and Stein, 1993; Morellec and Smith, 2007; or Décamps et al., 2017). The type of analysis that relates cash holdings to hedging needs has precedents in the literature, such as Acharya, Almeida, and Campello (2007) and Duchin (2010). Our analysis is unique because it uses estimated deep parameters of a canonical cash flow process instead of relying on proxies for cash flow volatility and hedging needs to explain corporate cash policy. Our paper also differs from prior work because of its focus on specific refinancing times when theory is more likely to hold.

(*Market-to-book ratio*), for lagged cash flow (*Cash flow-to-assets ratio*), and for lagged firm size ($\ln(\text{Total assets})$).¹⁸ The results are robust to the inclusion of other control variables such as net working capital, capital expenditures, R&D, leverage, and dividends. We include year fixed effects (δ_t) to control for time-specific shocks that affect all firms. Because our estimates of ρ , σ_P , and σ_A are constant over time (as assumed in the theoretical literature) but vary across the 918 groups used for the cash flow model estimation, we cannot include firm fixed effects in these regressions. However, the parameter estimates exhibit substantial within-industry variation. We therefore include two-digit SIC industry fixed effects (γ_j) to absorb time-invariant industry effects. Because Equation (10) uses estimates of ρ , σ_P , and σ_A , we calculate the standard errors conservatively by bootstrapping the standard errors and by clustering at the three-digit SIC level.

Table 9 presents the results. Columns 1 and 2 show the results for *Gross equity issuance*. The estimates of β_ρ are negative and significantly different from zero in both columns, consistent with the predictions of liquidity management models. The estimates are also economically significant. A one-sample-standard-deviation increase in the estimated correlation between short-term and permanent shocks ($\hat{\rho}$) is associated with a decrease in *Gross equity issuance* of 6.2% relative to the sample mean. That is, firms with high hedging needs issue significantly larger amounts of equity. In columns 3 and 4, the dependent variable is the *Gross equity issuance dummy*. The estimates of β_ρ are also negative and statistically significant, suggesting that firms with high hedging needs not only issue larger amounts of equity, but do so more frequently. We obtain very similar results using *Net equity issuance* (columns 5 and 6) and *Net equity issuance dummy* (columns 7 and 8). To our knowledge, our paper is the first that relates hedging needs to the firms' equity issuance activity.

All eight columns also show a positive and significant relation between equity issues (gross and net, size and frequency) and permanent cash flow shocks volatility. The economic magnitudes are sizable. For example, in column 1, a one-sample-standard-deviation increase in the estimated permanent volatility ($\hat{\sigma}_P$) is associated with an increase in *Gross equity issuance* of 13.7% relative to the sample mean. In six of the eight columns, the short-term cash flow shocks volatility is also positively associated with equity issues. The economic magnitudes, however, are smaller when compared with permanent cash flow shocks volatility. Our tests therefore suggest that permanent cash flow shocks are more important than transitory cash flow shocks in explaining the cross-sectional variation in equity issuance activity. Our findings for cash flow volatility are consistent with those in McLean (2011), who shows that firms

¹⁸ Controlling for the market-to-book ratio, which is typical for such regressions, is problematic in our context because, in theory, Tobin's q is a function of the cash flow shock volatilities, their correlation, and the growth rate. To address this problem, all specifications control instead for the residuals of the regression of *Market-to-Book ratio* on $\hat{\rho}$, $\hat{\sigma}_P$, $\hat{\sigma}_A$, and $\hat{\mu}$.

Table 9
Cash flow risk and equity issuance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Gross equity		Gross equity dummy		Net equity		Net equity dummy	
Hedging need ($\hat{\rho}$)	-0.036*** (0.010)	-0.029*** (0.010)	-0.036** (0.017)	-0.030** (0.015)	-0.030*** (0.010)	-0.023** (0.009)	-0.037*** (0.015)	-0.031** (0.012)
Permanent shock vol ($\hat{\sigma}_P$)	[-0.062***] (0.008)	[-0.051***] (0.005)	[-0.032***] (0.011)	[-0.027***] (0.006)	[-0.089***] (0.007)	[-0.070**] (0.004)	[-0.040**] (0.010)	[-0.034**] (0.005)
Short-term shock vol ($\hat{\sigma}_A$)	[0.137***] (0.003)	[0.100***] (0.003)	[0.202***] (0.007***)	[0.130***] (0.006***)	[0.283***] (0.004*)	[0.201***] (0.004**)	[0.228***] (0.007***)	[0.149***] (0.005***)
Industry cash flow vol	[0.057] (0.042)	[0.046] (0.011)	[0.061***] (0.065*)	[0.047***] (0.017)	[0.106*] (0.031)	[0.102**] (0.009)	[0.067***] (0.054*)	[0.054**] (0.013)
In(Total assets)	(0.029)	(0.022)	(0.035)	(0.026)	(0.027)	(0.023)	(0.030)	(0.023)
Cash flow-to-assets	-0.010*** (0.001)	-0.011*** (0.002)	-0.014*** (0.002)	-0.016*** (0.002)	-0.011*** (0.001)	-0.013*** (0.002)	-0.015*** (0.002)	-0.017*** (0.002)
Market-to-book ratio residuals	-0.388*** (0.049)	-0.375*** (0.046)	-0.250*** (0.033)	-0.236*** (0.031)	-0.376*** (0.044)	-0.362*** (0.043)	-0.251*** (0.031)	-0.235*** (0.029)
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE	N	Y	N	Y	N	Y	N	Y
Observations	76,795	76,795	76,795	76,795	73,167	73,167	73,167	73,167
Adjusted R^2	0.32	0.33	0.15	0.18	0.30	0.31	0.15	0.18

This table presents estimates from regressions of equity issuance variables on the parameters of cash flow risk. The dependent variable is *Gross equity issuance* in columns 1 and 2, *Equity issuance dummy* in columns 3 and 4, *Net equity issuance* in columns 5 and 6, and *Net equity issuance dummy* in columns 7 and 8. The sample period is from 1971 to 2018. The sample includes all firm-years in which a firm experiences a positive change in the cash-to-asset ratio. $\hat{\sigma}_P$, $\hat{\sigma}_A$, and $\hat{\rho}$ are the group-specific estimates of volatilities and the correlation between permanent and short-term cash flow shocks. All specifications include year fixed effects. Specifications 2, 4, 6, and 8 also include two-digit SIC industry fixed effects. All variables are defined in Table 1. The numbers in square brackets are economic effects, computed as the average change in the dependent variable for a one-standard-deviation change in the regressor, divided by the sample mean of the dependent variable. Standard errors (in parentheses) are bootstrapped and clustered at the three-digit SIC industry level. * $p < .1$, ** $p < .05$, *** $p < .01$.

issue equity to replenish cash reserves for precautionary reasons. However, we note that *Industry cash flow volatility* is not systematically related to equity issuance activity. That is, the variation in cash flow risk explained by our estimates of the three deep parameters subsumes and improves the explanatory power of the traditional proxy. Finally, firm size and the cash flow-to-assets ratio are negatively related to equity issuance activity, while the market-to-book ratio is positively related.

Overall, our cash flow risk parameter estimates are related to equity issuance as predicted by theory. Firms that are naturally hedged (high ρ firms) issue less equity and do so less frequently. Our findings also shed new light on the relative importance of permanent and short-term shocks for liquidity policies. Notably, permanent cash flow shocks volatility seems to have economically larger effects on equity issues than short-term shocks volatility.

4.3 Debt issues

In models of liquidity management, firms increase their cash holdings either by retaining earnings or by issuing equity (see, e.g., Décamps et al., 2011 or Bolton, Chen, and Wang, 2011). During the recent COVID crisis, a number of firms have issued long-term debt to increase their cash reserves (as documented, for instance, in Acharya and Steffen, 2020), following a massive negative cash flow shock combined with frozen equity markets. A key difference between debt and equity is that the former increases solvency risk and introduces a constant cash flow drain (due to the interest payments), which increases expected costs of refinancing (by increasing the refinancing frequency) and the optimal level of precautionary cash reserves. This suggests that debt may not be the financing vehicle of choice for managing cash reserves. Ultimately, however, whether firms use debt as a liquidity management tool is an empirical question.

The evidence so far rejects the possibility that debt may be used as a substitute for equity issuance in order to save cash outside of crisis periods. McLean (2011), for example, finds that, between 1971 and 2008, an extra dollar of net debt issuance in a fiscal year is associated with an increase of 2 cents in cash reserves at the end of the year. In comparison, an extra dollar of equity issued in a fiscal year is associated with an increase of 43 cents in cash reserves at the end of the year.¹⁹

Our estimates can help answer whether debt-financed cash savings policies, such as the response to the COVID shock, are common or an exception to the norm. For that purpose, we explore the joint determination of equity and debt

¹⁹ We have replicated McLean's result and found that this result continues to hold for the updated sample running until 2018 or for the subsample of firm-years in which the change to the cash balance is positive. We have also reestimated the McLean (2011) model but in a system of seemingly unrelated regressions, using our estimates of cash flow risk as instruments for equity and debt issuance policies. In this case, we have found an even higher savings rate following equity issues, and a negative and insignificant savings rate for debt issues. These results are presented in Section 3 of the Internet Appendix. Altogether, these results suggest that debt is not used as a liquidity management tool by most firms.

Table 10
Cash flow risk, debt, and equity issuance

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Net external financing		Net external financing dummy		Net equity	Net debt	Net equity dummy	Net debt dummy
Hedging need ($\hat{\rho}$)	-0.023** (0.011)	-0.023** (0.011)	-0.023 (0.016)	-0.023 (0.016)	-0.023*** (0.007)	-0.001 (0.003)	-0.031*** (0.009)	-0.004 (0.010)
Permanent shock vol ($\hat{\sigma}_P$)	[-0.050**] 0.023*** (0.004)	[-0.050**] 0.023*** (0.004)	[-0.012] 0.036*** (0.006)	[-0.012] 0.036*** (0.006)	[-0.067***] 0.017*** (0.003)	[-0.003] 0.004*** (0.001)	[-0.033***] 0.034*** (0.003)	[-0.003] 0.004 (0.003)
Short-term shock vol ($\hat{\sigma}_A$)	[0.205***] 0.003* (0.002)	[0.205***] 0.003* (0.002)	[0.074***] 0.005** (0.002)	[0.074***] 0.005** (0.002)	[0.199***] 0.004*** (0.001)	[0.161***] 0.001 (0.000)	[0.145***] 0.006*** (0.001)	[0.013] 0.001 (0.001)
Industry cash flow vol	[0.070*] 0.008 (0.002)	[0.070*] 0.008 (0.002)	[0.022**] 0.013 (0.002)	[0.022**] 0.013 (0.002)	[0.104***] 0.007 (0.001)	[0.012] -0.003 (0.000)	[0.055***] 0.010 (0.001)	[0.005] -0.027*** (0.005)
ln(Total assets)	-0.011*** (0.002)	-0.011*** (0.002)	-0.007*** (0.021)	-0.007*** (0.021)	-0.013*** (0.006)	0.002*** (0.003)	-0.016*** (0.008)	0.010 (0.009)
Cash flow-to-assets	-0.375*** (0.045)	-0.375*** (0.045)	-0.192*** (0.029)	-0.192*** (0.029)	-0.360*** (0.004)	-0.016*** (0.002)	-0.235*** (0.006)	-0.0040 (0.007)
Market-to-book ratio residuals	0.036*** (0.002)	0.036*** (0.002)	0.024*** (0.002)	0.024*** (0.002)	0.031*** (0.000)	0.005*** (0.000)	0.023*** (0.001)	0.006*** (0.001)
Year FE	Y	Y	Y	Y	Y	Y	Y	Y
Industry FE	N	Y	N	Y	Y	Y	Y	Y
Observations	69,085	69,085	69,085	69,085	69,085	69,085	69,085	69,085
Adjusted R^2	0.30	0.30	0.08	0.08	0.31	0.02	0.18	0.01

This table presents estimates from regressions of equity or debt issuance variables on the parameters of cash flow risk. The dependent variable is *Net external financing*, which equals *Net equity issuance* plus *Net debt issuance*, in columns 1 and 2; *Net external financing dummy* in columns 3 and 4, which equals one if *Net external financing* is positive and zero otherwise; *Net equity issuance* in column 5; *Net debt issuance* in column 6; *Net equity issuance dummy* in column 7; and *Net debt issuance dummy* in column 8. The models in columns 5 and 6 or 7 and 8 are estimated as a system of seemingly unrelated regressions. The sample period is from 1971 to 2018. The sample includes all firm-years in which a firm experiences a positive change in the cash-to-asset ratio. $\hat{\sigma}_P$, $\hat{\sigma}_A$, and $\hat{\rho}$ are the group-specific estimates of volatilities and the correlation between permanent and short-term cash flow shocks. All variables are defined in Table 1. The numbers in square brackets are economic effects, computed as the average change in the dependent variable for a one-standard-deviation change in the regressor, divided by the sample mean of the dependent variable. Standard errors (in parentheses) are bootstrapped and clustered at the three-digit SIC industry level. * $p < .1$, ** $p < .05$, *** $p < .01$.

issuance as a function of our estimated cash flow risk parameters. Table 10 shows the results of estimating Equation (10), replacing the dependent variable with different definitions of external financing. We find that all the results for equity issuance, shown in Table 9, are weakened if we define the dependent variable as “External finance”—that is, the sum of equity and debt issuance (columns 1–4). Treating each policy separately, and estimating both policies simultaneously, we confirm earlier results that the frequency and amount of equity issued decrease with the correlation parameter, ρ , and increase with the volatility of permanent shocks, σ_p , (columns 5 and 7), but fail to reject that long-term debt frequency and volume of issuance are not related to the same risk parameters (columns 6 and 8). To conclude, using our granular estimates of the three cash flow risk parameters, we find no evidence that firms treat equity and debt issuance as substitutes when choosing liquidity policies *ex ante* as a function of cash flow risk. Our results suggest that long-term debt may well be used to add liquidity, but only as an *ex post* response to severe cash flow shocks when issuing equity is not possible.

5. Robustness

5.1 Selection issues

Our filter has the ability to deal with missing cash flow observations, and, as a result, our sample has the maximum possible coverage of Compustat firms with at least 10 cash flow data points. Additionally, it would be possible to classify up to 7,239 more firms with eight or nine available observations into one of our estimation groups. Imputing the group parameter estimates to each of these firms virtually implies the same distribution of parameter estimates in a larger sample covering 73% of the firms in the Compustat universe.

Our range of estimates is not representative of the very short-lived firms in Compustat. We use simulation to try to understand the range of parameter values that these firms may have. Using the Décamps et al. (2017) model (summarized in Appendix C), we have confirmed that firms with relatively high ρ outlast otherwise similar firms with the lowest values of ρ , despite the fact that high ρ firms may simultaneously experience negative productivity and short-term shocks. Again, high ρ firms have a natural hedge advantage in that low (high) short-term cash flows tend to occur only when productivity is declining (improving) and the firm has lower (higher) needs for cash. Hence, in all likelihood our sample is excluding firms with the most negative values of ρ in their respective industries.

5.2 On the cash flow measure

The definitions of “cash flow” in the literature differ depending on whether the application intends to capture cash flow from operations, which is typically taken as given, or free cash flow available for savings and distributions to claimholders, which may include components related to debt or payout policies.

Table 11
Estimates of the cash flow model parameters under a different operating cash flow definition

A. Parameter estimates

	#est.	Mean	Standard deviation					p5	p25	p50	p75	p95
			Total	sd _b	sd _w							
$\hat{\rho}$	918	-0.050	0.121	0.090	0.105	-0.208	-0.103	-0.065	-0.028	0.180		
$\hat{\sigma}_P$	918	0.609	0.461	0.287	0.351	0.116	0.217	0.438	0.940	1.491		
$\hat{\sigma}_A$	918	0.412	0.711	0.438	0.614	0.033	0.064	0.097	0.208	2.264		
$\hat{\mu}$	918	0.117	0.081	0.052	0.063	0.012	0.056	0.097	0.172	0.271		

B. Values of *t*-statistics of the parameter estimates

	#est.	Mean	Standard deviation	Proportion of <i>p</i> -values		p5	p25	p50	p75	p95
				< .05	< .01					
$\hat{\rho}$	918	-19.648	72.458	.411	.380	-96.799	-5.568	-0.718	-0.176	3.189
$\hat{\sigma}_P$	918	85.926	292.648	.826	.723	1.169	2.436	4.579	29.213	383.462
$\hat{\sigma}_A$	918	6.001	12.608	.757	.646	0.994	2.003	3.316	5.157	15.295
$\hat{\mu}$	918	5.990	21.995	.420	.334	-3.552	0.845	1.474	3.108	28.355

C. Correlations between the parameter estimates

	$\hat{\rho}$	$\hat{\sigma}_P$	$\hat{\sigma}_A$	$\hat{\mu}$
$\hat{\rho}$	1.000			
$\hat{\sigma}_P$	-0.140***	1.000		
$\hat{\sigma}_A$	-0.172***	0.471***	1.000	
$\hat{\mu}$	0.035	0.406***	0.093***	1.000

This table summarizes the maximum likelihood estimates of the cash flow model parameters, $\hat{\mu}$, $\hat{\sigma}_P$, $\hat{\sigma}_A$, and $\hat{\rho}$ when operating cash flow is defined as EBITDA. The model parameters are estimated for each of the 918 three-digit SIC/cash flow growth groups of firms. The descriptive statistics are: number of model parameter estimates (#est.); mean; (Total) standard deviation, decomposed into between- (sd_b), and within-three-digit SIC industry (sd_w) variation; and the percentiles p5, p25, p50, p75, and p95. The sample covers the period 1971 to 2018.

In line with structural corporate finance, we focus on the cash flow from operations and our only departure is to subtract changes in working capital, which is a necessary cash expense to sustain operations.

Table 11 summarizes the parameter estimates if we define operating cash flows as EBITDA. Without the adjustment for changes in working capital, the average estimates of the short-term volatility, σ_A , and long-term volatility, σ_P , become smaller. Hence, cross-sectional differences in the working capital account, contain useful information about short-term (and long-term) variability. However, the estimates in Table 11 replicate the cash savings policy implications of Table 8, suggesting that both operating cash flow definitions agree on the sorting by all three dimensions of cash flow risk.

We note finally that our estimates of cash flow parameters can be used for many other applications on hypothesis testing or numerical analysis of dynamic structural corporate finance models. For each application, the researcher may change the definition of the operating cash flow variable in consistency with the theory. Naturally, we can apply our filter to any definition as long as we interpret the parameters correctly in each case.

5.3 Alternative firm grouping

We perform two additional exercises to assess the appropriateness of our method of grouping firms with similar cash flow dynamics. First, we reestimate the cash flow model in Equations (1)–(3) using all firms in each four-digit SIC industry classification and then test whether these parameter estimates can better explain cash savings and equity issuance policies. The grouping at the four-digit SIC level is less granular and, as shown in Table 3, results in much fewer estimation groups and with less between- but more within-group heterogeneity. As a result, the distributions of cash flow model parameters estimated at the four-digit SIC level have similar averages but exhibit less dispersion. The estimates of ρ at the four-digit SIC level still explain well the sign reversion of the cash flow sensitivity of savings, but their coefficients in the equity issuance regressions have a weaker statistical and economic significance, if preserving the right signs. These results underscore the importance of achieving high granularity in the estimation of the operating cash flow process.

Second, we conduct a placebo test and reestimate the model parameters in Equations (1)–(3) for 10,000 groups of 10 firms selected at random and with replacement from our dataset. Because each firm can belong to several groups, we (i) match each firm to the parameter estimates of only one randomly selected group, or (ii) we randomly select 918 of the 10,000 groups such that no firm belongs to more than one group. Using both approaches, we then test whether the cross-sectional variation in these estimates explains differences in equity issuance and cash savings. In both cases, the parameter estimates are not related to savings or equity issuance policies in any clear and systematic way. We conclude that our grouping method captures well the inherent similarities of the firms' cash flow dynamics and risk.

There exist in the literature other alternatives to classify a firm based on its product space overlap with others. Such is the case of the text-based product market similarity classification (TNIC) proposed by Hoberg and Phillips (2016). Applying the TNIC classification to our analysis is challenging for at least two reasons. First, this classification is available only since 1988, implying a loss of 16 years of data or one-third of our sample period. Second, the TNIC classification is firm-centric, with each firm's TNIC set changing from one year to the next as the firm or its competitors enter and exit individual product markets from its whole product market space. Adopting our model to this data structure would require different assumptions and estimation procedure. We view this extension of our framework as an interesting avenue for future research.

6. Conclusion

We estimate a canonical cash flow model that combines productivity shocks, which have permanent effects, and short-term cash flow shocks, which may be purely transitory or correlated with permanent shocks. Efficient estimation of this model is achieved with a high level of granularity for a large fraction

of the Compustat universe since the 1970s. Estimates of the model parameters provide a rich summary of a given firm's cash flow risk, as captured by its exposure to permanent and short-term shocks and the correlation between these shocks. The estimated correlation between permanent and short-term shocks, a key parameter that summarizes a firm's hedging needs, is a powerful indicator of within-industry differences in risk taking. In addition, the sign of this correlation determines the cash flow sensitivity of cash, as predicted by theory. The estimates of short- and long-term shock volatilities and of the correlation between shocks explain corporate liquidity management policies better than usual risk proxies.

Our empirical analysis is based on a cash flow model that assumes constant risk parameters and ignores potential feedback effects of agency conflicts and financing policies on cash flow dynamics. One way to relax these assumptions would be to conduct a structural estimation of a richer model with endogenous financing and investment. This would also potentially allow estimating the magnitude and heterogeneity of financing and investment frictions. Another promising avenue for future research would be to adapt our cash flow model and estimation method to allow for a different structure of transitory and permanent shocks. One could, for instance, consider that permanent and transitory shocks follow a factor model. This would require showing that this model is well identified and that its deep parameters can be estimated accurately with a newly derived Kalman filter. We leave these promising extensions for future research.

Appendix

Appendix A. Kalman Filter and Maximum Likelihood Estimation

This appendix provides a detailed exposition of the model estimation approach used in Section 2. We first describe the state space model and then derive the Kalman filter to compute the likelihood of cash flow data.

A.1 The State Space Model The state space model in Equations (1)–(2) consists of a transition equation and a measurement equation. The transition equation describes the discrete-time dynamics of the latent state process, which is the unobserved asset productivity X_t . The measurement equation describes the relation between the state process and the observed cash flows of firms that share the same asset productivity. To facilitate the exposition, we use a standard notation in state space models, and present the model as if missing observations were absent (Appendix B.3 discusses how we handle missing observations).

Let X_t denote the asset productivity in year t . The transition equation (1) can be rewritten as

$$X_t = \Phi_X X_{t-1} + \omega_t, \tag{A1}$$

where $\Phi_X = (1 + \mu)$, $\omega_t = \sigma_P X_{t-1} \varepsilon_t^P$ and $\varepsilon_t^P \sim \mathcal{N}(0, 1)$. Thus, $\omega_t \sim \mathcal{N}(0, Q_t)$, where $Q_t = \sigma_P^2 X_{t-1}^2$, and the error term ε_t^P is the permanent shock to cash flows.

Let $Z_{i,t}$ denote the cash flows of firm i in year t , that is, we set $Z_{i,t} = A_{i,t}$, and $Z_t = (Z_{1,t}, \dots, Z_{N,t})'$ to be the $N \times 1$ vector collecting the cash flows of the N firms that share the same asset productivity, where $'$ denotes transposition. The measurement equation in (2) can be written in vector form as

$$Z_t = H_Z X_t + u_t, \tag{A2}$$

where the i -th element is $Z_{i,t} = X_t + u_{i,t}$, $u_{i,t} = \sigma_A X_{t-1} \varepsilon_{i,t}^A$, and $\varepsilon_{i,t}^A \sim \mathcal{N}(0, 1)$ is the short-term shock to cash flows. In (A2), $H_Z = \mathbf{1}$, where $\mathbf{1} = (1, \dots, 1)'$.

In classic applications of state space models, u_t is merely a measurement error of X_t , and it is assumed to be uncorrelated with X_t . In contrast, because permanent and short-term shocks in model (1)–(2) are correlated, u_t and X_t turn out to be correlated. Specifically, the correlation between $u_{i,t}$ and X_t is equal to ρ and enters the short-term shock $\varepsilon_{i,t}^A = \rho \varepsilon_t^P + \sqrt{1 - \rho^2} \varepsilon_{i,t}^T$, where $\varepsilon_{i,t}^T \sim \mathcal{N}(0, 1)$ is the transitory shock, uncorrelated with ε_t^P . Thus,

$$Cov[X_t, u_{i,t} | X_{t-1}] = \mathbb{E}[\omega_t u_{i,t} | X_{t-1}] = \mathbb{E}[\sigma_P X_{t-1} \varepsilon_t^P \sigma_A X_{t-1} \varepsilon_{i,t}^A | X_{t-1}] = \rho \sigma_P \sigma_A X_{t-1}^2.$$

Collecting the transitory shocks of the N firms in $\varepsilon_t^T = (\varepsilon_{1,t}^T, \dots, \varepsilon_{N,t}^T)'$, the error term $u_t = \sigma_A X_{t-1} (\rho \varepsilon_t^P \mathbf{1} + \sqrt{1 - \rho^2} \varepsilon_t^T) \sim \mathcal{N}(0, \Omega_t)$, where $\Omega_t = \sigma_A^2 X_{t-1}^2 (\rho^2 \mathbf{1}\mathbf{1}' + (1 - \rho^2) I_N)$, and I_N is the $N \times N$ identity matrix.

The correlation between u_t and X_t makes the standard Kalman filter biased and inconsistent. To overcome this issue, we transform the measurement equation as follows:

$$\begin{aligned} Z_t &= H_Z X_t + u_t + J(X_t - \Phi_X X_{t-1} - \omega_t) \\ &= (H_Z + J)X_t - J\Phi_X X_{t-1} + u_t - J\omega_t \\ &= H_Z^* X_t + \Phi_X^* X_{t-1} + u_t^*, \end{aligned} \tag{A3}$$

where $H_Z^* = H_Z + J$, $\Phi_X^* = -J\Phi_X$, $u_t^* = u_t - J\omega_t$, and J is a $N \times 1$ vector that will be defined shortly. In the first equation, the third term on the right-hand side is zero by definition of the transition equation (A1). This means that the transformed measurement equation (A3) is an exact alternative representation of the measurement equation (A2). Importantly, the vector J is defined such that the transformed measurement error u_t^* is uncorrelated with X_t :

$$Cov[X_t, u_t^* | X_{t-1}] = \mathbb{E}[\omega_t u_t^* | X_{t-1}] = \mathbb{E}[\omega_t u_t | X_{t-1}] - J\mathbb{E}[\omega_t^2 | X_{t-1}] = 0. \tag{A4}$$

Solving the last equation for J gives $J = \mathbb{E}[\omega_t u_t | X_{t-1}] / \mathbb{E}[\omega_t^2 | X_{t-1}]$. In the state space model (A1)–(A2), J takes a simple form, that is, $J = \rho \sigma_A / \sigma_P \mathbf{1}$.

Plugging J in u_t^* clarifies why u_t^* is uncorrelated with X_t in Equation (A3):

$$u_t^* = u_t - J\omega_t = \sigma_A X_{t-1} \varepsilon_t^A - \rho \frac{\sigma_A}{\sigma_P} \mathbf{1} \sigma_P X_{t-1} \varepsilon_t^P = \sigma_A X_{t-1} (\varepsilon_t^A - \rho \mathbf{1} \varepsilon_t^P) = \sigma_A X_{t-1} \sqrt{1 - \rho^2} \varepsilon_t^T,$$

where ε_t^T is by definition uncorrelated with X_t , and we use $\varepsilon_t^A = (\varepsilon_{1,t}^A, \dots, \varepsilon_{N,t}^A)'$. The error term u_t^* is by definition uncorrelated with X_{t-1} too.

The transformation of the measurement equation in (A3) can be applied to more general state space models to handle the correlation between state variables and measurement errors. For example, if X_t is a $k \times 1$ state variable, then $J = \mathbb{E}[u_t \omega_t' | X_{t-1}] \mathbb{E}[\omega_t \omega_t' | X_{t-1}]^{-1}$, which is an $N \times k$ matrix. Also, J could be time varying when the conditional expectations are state dependent.

In the signal processing literature, Ma, Wang, and Chen (2010) suggest transforming the transition equation to account for the correlation between measurement and transition errors in state space models. We use a different approach and transform the measurement equation which results in a stable Kalman filter for the state space model in Equations (1)–(2). See Appendix A.5 for details.

A.2 The Generalized Kalman Filter Because the transformed measurement equation (A3) features X_{t-1} in the right-hand side, it is necessary to re-derive the Kalman filter to filter out the latent state process.

Let $\hat{X}_{t|t-1} = \mathbb{E}_{t-1}[X_t]$ and $\hat{Z}_{t|t-1} = \mathbb{E}_{t-1}[Z_t]$ denote the expectation of X_t and Z_t , respectively, using information up to and including time $t - 1$, and let $V_{t|t-1}$ and $F_{t|t-1}$ denote the corresponding (a priori) error variance and error covariance matrix. Furthermore, let $\hat{X}_t = \mathbb{E}_t[X_t]$ denote the expectation of X_t including information at time t , and let V_t denote the (a posteriori) error variance.

The Kalman filter consists of two steps, that is, prediction and update. In the prediction step, $\hat{X}_{t|t-1}$ and $V_{t|t-1}$ are given by

$$\hat{X}_{t|t-1} = \Phi_X \hat{X}_{t-1} \tag{A5}$$

$$V_{t|t-1} = \Phi_X V_{t-1} \Phi_X + Q_t. \tag{A6}$$

and $\hat{Z}_{t|t-1}$ and $F_{t|t-1}$ are in turn given by

$$\hat{Z}_{t|t-1} = H_Z \hat{X}_{t|t-1} + \Phi_X^* \hat{X}_{t-1} \tag{A7}$$

$$F_{t|t-1} = H_Z^* V_{t|t-1} H_Z^{*'} + \Phi_X^* V_{t-1} \Phi_X^{*'} + \Omega_t^*. \tag{A8}$$

Because the transition equation (A1) is standard, $\hat{X}_{t|t-1}$ and $V_{t|t-1}$ take the usual forms in Kalman filtering. The transformed measurement equation (A3) changes $\hat{Z}_{t|t-1}$ and $F_{t|t-1}$ relative to standard Kalman filtering, with the additional terms in Φ_X^* .

In the update step, the estimate of the state process X_t is refined based on the difference between the observed and predicted values of Z_t , with \hat{X}_t and V_t given by

$$\hat{X}_t = \hat{X}_{t|t-1} + G_t'(Z_t - \hat{Z}_{t|t-1}) \tag{A9}$$

$$V_t = V_{t|t-1} - 2V_{t|t-1}H_Z^{*'}G_t + G_t'F_{t|t-1}G_t, \tag{A10}$$

where G_t is an $N \times 1$ vector called Kalman gain, which is determined by minimizing V_t with respect to G_t . Solving the first-order condition $\partial V_t / \partial G_t = 0$ for G_t gives $G_t' = V_{t|t-1}H_Z^{*'}F_{t|t-1}^{-1}$. This choice of G_t minimizes V_t because $\partial^2 V_t / (\partial G_t \partial G_t') = 2F_{t|t-1}$ is positive definite.

Model estimation is achieved by maximizing the log-likelihood of cash flows data of N firms over T periods with respect to the model parameters μ , σ_P , σ_A , and ρ . Specifically, for fixed model parameters the generalized Kalman filter (A5)–(A10) is run to compute the log-likelihood

$$\sum_{t=1}^T -\frac{1}{2} \left[N \log(2\pi) + \log |F_{t|t-1}| + (Z_t - \hat{Z}_{t|t-1})' F_{t|t-1}^{-1} (Z_t - \hat{Z}_{t|t-1}) \right]. \tag{A11}$$

Model parameters are changed as to increase the value of the log-likelihood, which then requires to re-running the generalized Kalman filter, and re-computing the log-likelihood.²⁰ The iterative procedure is repeated until convergence of the numerical likelihood search. As mentioned in Section 2.2, on a common laptop computer, it takes less than one second to fit the model to a panel of 10 firm cash flows observed over 46 years.

A.3 Missing Observations Handled with Kalman Filtering A prominent feature of cash flow data are missing observations. In our panel, 56% of firm-year observations are missing relative to the full balanced panel. Although our Kalman filter is different from the standard one, missing values can be handled using the usual method in Kalman filtering; see Section 3 in Shumway and Stoffer (1982). For completeness we briefly recall the procedure.

Suppose that there are no missing observations in year t . Then, the measurement equation (A3) holds. That is, Z_t collects the cash flows of all the N firms in a year t . Suppose now that the cash flow data of some firms in year t are missing. The idea is to “select” the components of Z_t corresponding to firms with observed (not missing) cash flow data. This task is achieved by simply using a matrix S_t consisting of zeros and ones with dimension $M_t \times N$, where M_t is the number of firms with observed cash flow data. To illustrate, consider an extreme and unrealistic case in

²⁰ The starting value of the state process X is set equal to the average of cash flows at $t=1$, but is then optimized during the likelihood search.

which only the cash flow of the first firm in Z_t is available in year t . In that case, $S_t = (1, 0, \dots, 0)$ is a $1 \times N$ row vector, $M_t = 1$ and $S_t Z_t$ is the cash flow of that firm. If cash flows of all N firms are available in year t , then S_t is a $N \times N$ identity matrix.

The procedure to compute the log-likelihood with missing observations is as follows. First, for each year t , construct the matrix S_t based on the position of observed cash flows in Z_t . Then, pre-multiply both sides of equation (A3) by S_t and use this measurement equation to run the generalized Kalman filter. Finally, compute the log-likelihood in (A11) replacing N by M_t .

The matrix S_t is time dependent and needs to be computed for each year t . This time dependence allows the procedure to accommodate missing observations in different positions of the cash flow panel as well as entry and exit of firms in the panel.

A.4 Identification of Shock Correlation To illustrate the issue of the identification of shock correlation, we use a popular model in the corporate finance literature. The model features persistent and short-term productivity shocks at a firm level. Denote $y_{i,t}$ an observable productivity or cash flow process for a given firm i ,

$$y_{i,t} = \varepsilon_{i,t} + \sigma_v v_{i,t}, \tag{A12}$$

where $\varepsilon_{i,t}$ is an AR(1) process, namely $\varepsilon_{i,t} = \beta \varepsilon_{i,t-1} + \sigma_\eta \eta_{i,t}$ and $0 < \beta < 1$ and $\sigma_\eta > 0$. The process $\eta_{i,t} \sim i.i.d. \mathcal{N}(0, 1)$ models persistent (or often called long-term) shocks. The process $v_{i,t} \sim i.i.d. \mathcal{N}(0, 1)$ models short-term shocks and $\sigma_v > 0$. Both shocks are firm-specific. The usual assumption in the literature is that these shocks are uncorrelated. We consider instead the case in which these shocks are correlated, $corr[\eta_{i,t}, v_{i,t}] = \rho$. It is perhaps surprising that even observing an infinite time series of $y_{i,t}$, the correlation ρ (and other model parameters) cannot be identified. In what follows, we formally prove this result.²¹

A time series model is identified when the system of equations, matching population and model-based autocovariances, can be solved uniquely for the model parameters. The unknowns in this system are the model parameters. Population autocovariances are (asymptotically) known. Define the autocovariance function as $\gamma(h) = Cov[y_{i,t}, y_{i,t-h}]$, for $h = 0, 1, \dots$, then

$$\gamma(0) = \frac{\sigma_\eta^2}{1 - \beta^2} + \sigma_v^2 + 2\rho\sigma_\eta\sigma_v \tag{A13}$$

$$\gamma(h) = \beta^h \left[\frac{\sigma_\eta^2}{1 - \beta^2} + \rho\sigma_\eta\sigma_v \right]. \tag{A14}$$

β can be easily identified from the decay of the autocovariance function $\gamma(h)$, say from the equation $\gamma(h_2)/\gamma(h_1) = \beta^{h_2-h_1}$ for $h_2 > h_1 \geq 1$. It is therefore taken as known in the discussion that follows. Although there is an infinite number of equations (A14) for $h \geq 1$, effectively, (A13)–(A14) is a system of two equations in three unknowns, $\rho, \sigma_\eta, \sigma_v$, and the model in (A12) is not identified.

To see the lack of identification of the shock correlation, suppose for simplicity that $\rho > 0$. Solving (A13) for σ_v (which then admits only one real and positive solution) and plugging this solution in (A14) gives

$$\gamma(h) = \beta^h \left[\frac{\sigma_\eta^2}{1 - \beta^2} + \rho\sigma_\eta \left(\sqrt{\rho^2\sigma_\eta^2 + \left(\gamma(0) - \frac{\sigma_\eta^2}{1 - \beta^2} \right)} - \rho\sigma_\eta \right) \right],$$

which is effectively one equation in two unknowns, ρ and σ_η . Therefore, ρ and σ_η are not identified.

²¹ A proof that the shock correlation in Equations (1)–(3) is not identified when permanent shocks are firm-specific is beyond the scope of this paper given that the model features a multiplicative, nonstationary process.

Consider now the model

$$y_{i,t} = \varepsilon_t + \sigma_v v_{i,t} \tag{A15}$$

in which the persistent shock $\varepsilon_t = \beta \varepsilon_{t-1} + \sigma_\eta \eta_t$ is not firm-specific but common across firms, and is still correlated with the short-term shock, $corr[\eta_t, v_{i,t}] = \rho$. The persistent shock ε_t plays the role of a “systemic factor” for the firms’ productivity. The short-term shock can be decomposed as $v_{i,t} = \rho \varepsilon_t + \sqrt{1 - \rho^2} v_{i,t}^T$, where $v_{i,t}^T$ is the firm-specific transitory shock, uncorrelated with ε_t . The cross-sectional mean, \bar{y}_t , is such that

$$\bar{y}_t = \frac{1}{N} \sum_{i=1}^N y_{i,t} = \frac{1}{N} \sum_{i=1}^N (\varepsilon_t + \sigma_v v_{i,t}) = \frac{1}{N} \sum_{i=1}^N (\varepsilon_t + \sigma_v (\rho \varepsilon_t + \sqrt{1 - \rho^2} v_{i,t}^T))$$

and therefore when $N \rightarrow \infty$, $\bar{y}_t = \varepsilon_t (1 + \rho \sigma_v)$. The last equation indicates that if, for example, $\rho > 0$, then firms’ productivity loads more on the systemic factor ε_t relative to the case when $\rho = 0$. Importantly, the autocovariances of \bar{y}_t provide additional moment conditions to identify the model in (A15). In essence, additional information from the cross-section of firms allows to identify the model and in particular the shock correlation. Denote $\bar{\gamma}(0) = V[\bar{y}_t]$, then

$$\bar{\gamma}(0) = \frac{\sigma_\eta^2}{1 - \beta^2} (1 + \rho \sigma_v)^2. \tag{A16}$$

The moment conditions (A13), (A14), and (A16) provide a system of three equations in three unknowns, ρ , σ_η , and σ_v , to identify the model in (A15). This system can be solved as follows. Solving (A14) with respect to $\rho \sigma_v$ and plugging this solution in (A16) gives a quadratic equation in which σ_η is the only unknown. Ensuring that only one real and positive solution exists, identifies σ_η . The difference between (A13) and (A14) gives

$$\gamma(0) - \frac{\gamma(h)}{\beta^h} = \sigma_v^2 + \rho \sigma_\eta \sigma_v. \tag{A17}$$

Matching the expression of $\rho \sigma_\eta \sigma_v$ from (A17) and from (A14) gives a linear equation in which σ_v^2 is the only unknown, identifying this parameter. Having identified σ_η and σ_v , (A13) can be used to identify ρ .

In sum, a model in which persistent and short-term shocks are both firm-specific is not identified. Instead, assuming that persistent shocks are common across firms allows to identify the shock correlation, because these common shocks would behave like a systemic factor for firms’ productivity.

Appendix B. Monte Carlo Analysis of Estimation Accuracy

B.1 Comparison with the Standard Kalman Filter To check the accuracy of our estimation method, we conduct a Monte Carlo simulation. In the cash flow model given by Equations (1) and (3), we set the parameters ρ , σ_A , σ_p , and μ to their respective average estimated values as in Table 4. We then use the model to simulate 10,000 panels of cash flows. As in our empirical analysis with Compustat data, each simulated panel consists of the cash flows of 10 firms over 48 years. For each simulated panel we estimate the model in Equations (1)–(3) using maximum likelihood with our generalized Kalman filter, as described in Appendix A.2. As a benchmark method, we also estimate the cash flow model using maximum likelihood with a standard Kalman filter.

The standard filter presumes that the correlation between permanent and short-term cash flow shocks is zero, and thus delivers no estimate of ρ . As a measure of estimation accuracy, for each estimated parameter $\hat{\theta} = \{\hat{\rho}, \hat{\sigma}_A, \hat{\sigma}_p, \hat{\mu}\}$, we compute the mean square error (MSE), that is, $\sum_{j=1}^{10000} (\hat{\theta}_j - \theta_0)^2 / 10000$, where $\hat{\theta}_j$ is the estimate of the parameter θ based on the j -th simulated panel of cash flows and θ_0 is the true parameter value. To compare MSEs across parameters, we report the relative MSE, that is, the MSE divided by the absolute value of θ_0 .

Underscoring the accuracy of our estimation method, the MSEs of the maximum likelihood estimates of the parameters μ , σ_ρ and σ_A with the generalized Kalman filter are, respectively, 0.360, 0.085 and 0.106. The MSEs of the maximum likelihood estimation with the standard Kalman filter are an order of magnitude larger than those with the generalized Kalman filter. The ratios between the two MSE's are 2.2, 1.7 and 2.1, respectively. Hence, our method is uniformly more accurate than maximum likelihood with a standard Kalman filter, often by a large extent. Finally, the MSE of ρ based on maximum likelihood with the generalized Kalman filter is 0.034, which is even smaller than the MSEs of the other parameters. As mentioned above, maximum likelihood with standard Kalman filter provides no estimate of ρ . We also experimented with smaller group sizes. Estimation results were less accurate.

In sum, the Monte Carlo simulation above shows that maximum likelihood with the generalized Kalman filter delivers accurate estimates of the cash flow model in Equations (1) and (3) while outperforming the maximum likelihood estimator with a standard Kalman filter. The latter method is not suited to handle the correlation between permanent and short-term shocks Equations in (1)–(3).

B.2 Comparison with Ma, Wang, and Chen (2010) Ma, Wang, and Chen (2010) have introduced a Kalman filter that can account for the correlation between the state variable and the measurement error. In a classic engineering setting (namely to recover location and velocity of an object in a one-dimensional motion), they compare their filter to the standard Kalman filter, which presumes zero correlation between the state variable and measurement error. While for positive correlations their filter outperforms the standard Kalman filter, they find that for negative correlations the two filters have a similar degree of inaccuracy. This suggests that their filter may not be reliable over the full range of possible correlation values.

How do the two filters differ in their construction and in their performance? To compare the filters, we allow both the state variable X_t and the measurements Z_t to be vector-valued. We first present the approach developed by Ma, Wang, and Chen (2010), and then discuss our approach, highlighting the difference between the two. In a nutshell, the former approach rearranges the *transition equation* by regressing the transition shock on the measurement error. Our approach follows the “opposite route”: we rearrange the *measurement equation* by regressing the measurement error on the transition shock.

In the signal processing literature (see Ma, Wang, and Chen (2010)), the transition equation, $X_{t+1} = \Phi_X X_t + \omega_t$, is rearranged by adding the zero term from the measurement equation, $Z_t - H_Z X_t - u_t = 0$, as follows:

$$\begin{aligned} X_{t+1} &= \Phi_X X_t + \omega_t \\ &= \Phi_X X_t + \omega_t + K(Z_t - H_Z X_t - u_t) \\ &= (\Phi_X - K H_Z)X_t + K Z_t + (\omega_t - K u_t) \\ &= (\Phi_X - K H_Z)X_t + K Z_t + \omega_t^*, \end{aligned}$$

where $\omega_t^* = \omega_t - K u_t$ is the new transition shock.²² Then K is chosen such that the new transition shock, ω_t^* , and the measurement error, u_t , are uncorrelated

$$0 = Cov[\omega_t^*, u_t'] = \mathbb{E}[(\omega_t - K u_t)u_t'] = \mathbb{E}[\omega_t u_t'] - K \mathbb{E}[u_t u_t']$$

and therefore $K = E[\omega_t u_t'] (\mathbb{E}[u_t u_t'])^{-1}$. That is, ω_t is theoretically regressed on u_t .

Our approach is to rearrange the measurement equation, $Z_t = H_Z X_t + u_t$, by adding the zero term from the transition equation, $X_t - \Phi_X X_{t-1} + \omega_t = 0$, as follows:

$$Z_t = H_Z X_t + u_t$$

²² The transition shock of X_{t+1} is denoted by ω_t . This is merely a convention. The shock ω_t is moving the state variable from X_t to X_{t+1} in discrete time.

$$\begin{aligned}
 &= H_Z X_t + u_t + J(X_t - \Phi_X X_{t-1} + \omega_t) \\
 &= (H_Z + J)X_t - J\Phi_X X_{t-1} + (u_t - J\omega_t) \\
 &= H_Z^* X_t + \Phi_X^* X_{t-1} + u_t^*,
 \end{aligned}$$

where $u_t^* = u_t - J\omega_t$ is the new transition shock. J is chosen such that

$$0 = Cov[u_t^*, \omega_t'] = E[(u_t - J\omega_t)\omega_t'] = E[u_t\omega_t'] - J E[\omega_t\omega_t']$$

and therefore $J = E[u_t\omega_t'](E[\omega_t\omega_t'])^{-1}$. That is, we theoretically regress u_t on ω_t to orthogonalize the measurement error to the transition shock.

The two methods generate two different Kalman filters. For example, the one-step-ahead prediction of Z_t in Ma, Wang, and Chen (2010) is the usual $\hat{Z}_{t|t-1} = H_Z \hat{X}_{t-1}$, whereas in our approach it is $\hat{Z}_{t|t-1} = H_Z^* \hat{X}_{t-1} + \Phi_X^* \hat{X}_{t-1}$, where \hat{X}_{t-1} is the updated value of the state process given the available observations up to and including time $t - 1$. This quantity is not present in the Ma, Wang, and Chen’s (2010) filter (nor in the standard Kalman filter), and changes the inner workings of the filter (prediction and update steps).

There are two main theoretical differences between our and Ma, Wang, and Chen’s (2010) filter, and both differences induce a better performance of our filter. First, as mentioned in the previous paragraph and in contrast to Ma, Wang, and Chen (2010), the updated value \hat{X}_{t-1} enters the one-step-ahead prediction of Z_t in our filter as $\hat{Z}_{t|t-1} = H_Z^* \hat{X}_{t-1} + \Phi_X^* \hat{X}_{t-1}$. Consequently, the a posteriori covariance matrix V_{t-1} of $(X_{t-1} - \hat{X}_{t-1})$ enters the a priori covariance matrix $F_{t|t-1}$ of the prediction error $(Z_t - \hat{Z}_{t|t-1})$. Relative to the corresponding covariance matrix in Ma, Wang, and Chen (2010), the covariance matrix $F_{t|t-1}$ has an additional term, $\Phi_X^* V_{t-1} \Phi_X^{* \prime}$, that changes the Kalman gain $G_t' = V_{t|t-1} H_Z^{* \prime} F_{t|t-1}^{-1}$ that in turn changes the update of the latent state $\hat{X}_t = \hat{X}_{t|t-1} + G_t'(Z_t - \hat{Z}_{t|t-1})$. Effectively, our filter appears to make a more efficient use of the available information (observations Z_t and filter-based quantities \hat{X}_{t-1} and V_{t-1}) to extract the latent state X_t . Second, in our filter the covariance matrix of the transformed measurement error u_t^* is diagonal, $V[u_t^*] = \sigma_A^2(1 - \rho^2)I_N X_{t-1}^2$, where I_N is the N -dimensional identity matrix. In contrast, the covariance matrix of the measurement error in the Ma, Wang, and Chen’s (2010) filter is a full matrix, $V[u_t] = \sigma_A^2((1 - \rho^2)I_N + \rho^2 \mathbf{1}\mathbf{1}')X_{t-1}^2$. As a consequence, our filter induces an “efficient rotation” of the vector-valued measurement error that results in an homoscedastic error covariance matrix. Because the Kalman filter relies on linear projections (e.g., projecting the measurement error u_t on the transition shock ω_t), the homoscedastic feature of the error covariance matrix allows to achieve efficiency gains in our filter. Moreover, when the shock correlation ρ approaches +1 or -1, the variance $V[u_t^*]$ of the measurement error tends to zero in our filter, increasing the precision of the filter, whereas that is not the case in Ma, Wang, and Chen’s (2010) filter. Finally, the diagonal covariance matrix in our filter enhances also the computational efficiency of its inverse, making the covariance matrix of the prediction error “less singular” in applications.

To compare the performance of the two filters, we run the following Monte Carlo simulation:

1. Fix the parameters of the cash flow model (1)–(2) in the paper.
2. Simulate one panel of cash flows of 10 firms over 50 years.
3. Apply each of the two Kalman filters above to the simulated cash flow panel to extract the trajectory of the true state process, $X_t, t = 1, \dots, 50$ (known in simulation).
4. Compute the root mean square error of the filtered trajectory, $\sqrt{\sum_{t=1}^{50} (X_t - \hat{X}_t)^2 / 50}$, for each filter, where \hat{X}_t is the Kalman-filtered value.
5. Repeat 10,000 times the steps 2 to 4.

Note that we use the true model parameters to apply both filters in step 3. Hence, this simulation exercise is purely a comparison of the two filters (designed to accommodate the shock correlation),

as there is no estimation error. We select three different sets of parameter values in step 1, and repeat 10,000 times steps 2 to 4 for each set of parameters. First, we set the model parameters in step 1 equal to the median estimates (p50 in Panel A in Table 4). Our filter performs better than the filter of Ma, Wang, and Chen (2010) 62% of the times in terms of root mean square error. Although the two filters are the same when $\rho=0$, the simulation evidence indicates that our filter is more accurate even for a moderate value of the shock correlation ($\rho=-0.076$). Second, we set the parameter values equal to the first quartile estimates (p5 in panel A in Table 4), and find that our filter outperforms 97% of the times the filter of Ma, Wang, and Chen (2010). In this case, the shock correlation is equal to $\rho=-0.232$. Given that most Compustat firms in our sample have a negative ρ (high hedging needs), we view the high performance of our filter in this setting as an important benefit. Third, we set the model parameters equal to the third quartile estimates (p95 in panel A in Table 4), where $\rho=0.229$. In this case, our filter outperforms the filter of Ma, Wang, and Chen (2010) 77% of the times. We conclude that our filter is more accurate and can be used for any value of the shock correlation.

C. Asset Growth Volatilities

This exercise illustrates a numerical application of our decomposition. The purpose is twofold. First, and as a validation of our results, we ask whether the parameter estimates in Table 4 imply asset volatilities that are comparable to those of actual Compustat firms. It is not clear this should be the case: our estimator recovers cash flow parameters from cash flow data only, and does not impose the restrictions from corporate policies that feed from those parameters to predict asset values and volatilities. Second, we also test whether our implied asset volatilities are robust, in the sense that they are not driven by the financing model assumptions but instead by inference based only on cash flow data, using the model in Equations (1)–(2). As noted by Gorbenko and Strebulaev (2010), an important limitation of standard EBIT models with only permanent shocks (e.g., Leland (1994)) is that asset growth volatilities are equal to cash flow volatilities. Hence, this exercise evaluates the extent to which the cash flow structure with permanent and transitory shocks helps reconcile the large differences between relatively high cash flow volatilities and the much more moderate asset volatilities in the data.

C.1 Cash flow model To compute the model-implied asset volatilities, we employ a version of the model of Décamps et al. (2017). In this continuous time model, operating revenue is subject to permanent and transitory shocks. Asset productivity $X=(X_t)_{t \geq 0}$ is governed by the geometric Brownian motion

$$dX_t = \mu X_t dt + \sigma_P X_t dW_t^P, \tag{C1}$$

where μ and $\sigma_P > 0$ are constant and $W^P=(W_t^P)_{t \geq 0}$ is a standard Brownian motion. Therefore, asset productivity is nonstationary and features permanent shocks. In addition to these shocks, cash flows are subject to short-term shocks. For a given firm, the cash flows dZ_t are proportional to X_t but uncertain and governed by

$$dZ_t = X_t dt + \sigma_A X_t dW_t^A, \tag{C2}$$

where $\sigma_A > 0$ is constant and $W^A=(W_t^A)_{t \geq 0}$ is a standard Brownian motion. W^A and W^P can be correlated with correlation coefficient ρ , in that

$$\mathbb{E}[dW_t^P dW_t^A] = \rho dt, \text{ with } \rho \in (-1, 1). \tag{C3}$$

The specification for cash flow dynamics in (C1) and (C2) nests those in traditional dynamic corporate finance models. If $\sigma_A=0$, we obtain the model with time-varying profitability applied extensively in dynamic capital structure models (see Goldstein, Ju, and Leland 2001; Strebulaev 2007; or Morellec, Nikolov, and Schürhoff 2012) and real-options models (see Abel and Eberly 1994; Carlson, Fisher, and Giammarino 2006; or Morellec and Schürhoff 2011). If $\mu=\sigma_P=0$, we obtain the stationary framework of dynamic agency models (see DeMarzo and Sannikov 2006 or

DeMarzo et al. 2012) and liquidity management models (see Décamps et al. 2011; Bolton, Chen, and Wang 2011; or Hugonnier, Malamud, and Morellec 2015).

The transition equation (1) is a simple Euler discretization of Equation (C1). The measurement equation (2) is related to an Euler discretization of Equation (C2) and it is obtained by setting $Z_{i,t}$ equal to the cash flow accumulated over year t , for firm i . In (2), X_{t-1} and not X_t enters the error term $P_{t-1}\sigma_A \varepsilon_{i,t}^A$ for it to have zero mean.

With this specification, the firm's cash flow over the time interval $[t, t+dt]$ is given by

$$dZ_t = X_t dt + \sigma_A X_t (\rho dW_t^P + \sqrt{1-\rho^2} dW_t^T), \tag{C4}$$

where $W^T = (W_t^T)_{t \geq 0}$ is a Brownian motion independent from W^P . This decomposition implies that short-term cash flow shocks dW_t^A consist of transitory shocks dW_t^T and permanent shocks dW_t^P .

C.2 Management's Optimization Problem Short-term shocks expose the firm to potential losses that can be covered using cash reserves or new equity financing. Specifically, management is allowed to retain earnings inside the firm and we denote by M_t the firm's cash holdings at any time $t > 0$. Cash reserves earn a rate of return $r - \lambda$ inside the firm, where $\lambda > 0$ is a cost of holding liquidity. The firm can also raise additional funds from investors. External equity financing is costly with a fixed cost ϕP_t and a proportional cost θ . The dynamics of cash reserves are then given by:

$$dM_t = (r - \lambda)M_t dt + \left(dt + \sigma_A \rho dW_t^P + \sqrt{1-\rho^2} dW_t^T \right) P_t + \frac{dE_t}{\theta} - d\Phi_t - dL_t, \tag{C5}$$

where E_t , Φ_t , and L_t are nondecreasing processes that represents the cumulative gross external financing, the cumulative fixed cost of financing, and the cumulative dividend paid to shareholders. Equation (C5) is an accounting identity that shows that cash reserves increase with the interest earned on cash holdings (first term on the right-hand side), with the firm's earnings (second term), and with net external equity (third and fourth terms), and they decrease with payouts (last term).

Management chooses the cash savings/payout and equity financing policies to maximize shareholder value. There are two state variables for the firm's optimization problem: Profitability X_t and the cash balance M_t . We can thus write this problem as

$$V(x, m) = \sup_{L, E} \mathbb{E}_{p, m} \left[\int_0^\infty e^{-rt} (dL_t - dE_t) \right], \tag{C6}$$

where x and m denote realizations of X and M at time $t=0$. Décamps et al. (2017) show that there exists a unique solution to this optimization problem and characterize firm value and optimal policies in their Proposition 1.

C.3 Implied Volatilities To compute the model-implied asset volatilities, we use the cash flow process estimated previously and solve for optimal financing and liquidity policies and firm value. We thus quantitatively link the cash flow parameters (μ , σ_A , σ_P , and ρ) to asset volatility.

We calculate asset volatilities at the group level, consistently with the level of granularity of the estimation of cash flow parameters. The empirical asset volatilities are estimated as weighted averages of equity and debt volatilities following the approach in Bharath and Shumway (2008) using daily stock returns. Model-implied volatilities are averages of all firm-level volatilities within a group using the estimated cash flow parameters reported in Table 4 for each group of firms. We drop groups of firms with insufficient stock price data or where the model cannot be solved given parameter values, winsorizing actual and predicted asset volatilities at the 5th and 95th percentiles.²³

²³ Remaining parameter values follow Décamps et al. (2017) and Bolton, Chen, and Wang (2013): the risk-free rate $r=0.08$, the carry cost of cash $\lambda=0.02$, proportional and fixed equity issuance costs $p=1.06$ and $\Phi=0.002$, market price of risk of temporary and permanent shocks $\eta^T = \eta^P = 0.4$, and the correlation of temporary and permanent shocks with market shocks $\xi^T = \xi^P = 0.4$; cash holdings are assumed at the target level.

Table 7 presents a comparison of the empirical and model-implied asset volatilities. The average and median empirical asset volatilities are 0.476 and 0.447, respectively. In the baseline case, the model-implied volatilities are slightly higher, at 0.559 and 0.495, respectively, but very close to the empirical ones. It is remarkable that model-implied asset volatilities appear to match actual asset volatilities that were not used during the estimation process. While similar at the center of the distribution, the model-implied volatilities tend to be somewhat more extreme in the tails compared to the empirical ones. This suggests that there could be some forces beyond those in the Décamps et al. (2017) model moderating the volatility of real firms' asset values. Additional rows for model-implied distributions in Table 7 present calculations based on departures from the baseline parameters. The results show that the similarity between the empirical and model-implied distributions is robust and driven by the estimates of the parameters of the cash flow dynamics rather than the assumed values for the remaining parameters.

References

- Abel, A., and J. Eberly. 1994. A unified model of investment under uncertainty. *American Economic Review* 84:1369–84.
- . 2011. How Q and cash flow affect investment without frictions: An analytic explanation. *Review Economic Studies* 78:1179–200.
- Acharya, V., H. Almeida, and M. Campello. 2007. Is cash negative debt? A hedging perspective on corporate financial policies. *Journal of Financial Intermediation* 16:515–54.
- Acharya, V., and S. Steffen. 2020. The risk of being a fallen angel and the corporate dash for cash in the midst of COVID. *Review of Corporate Finance Studies* 9:430–71.
- Almeida, H., M. Campello, and A. Galvao. 2010. Measurement errors in investment equations. *Review of Financial Studies* 23:3279–328.
- Almeida, H., M. Campello, and M. Weisbach. 2004. The cash flow sensitivity of cash. *Journal of Finance* 59:1777–804.
- Arellano, M., and S. Bond. 1991. Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations. *Review of Economic Studies* 58:277–97.
- Bansal, R., R. Dittmar, and D. Kiku. 2008. Cointegration and consumption risks in asset returns. *Review of Financial Studies* 22:1343–75.
- Bates, T., K. Kahle, and R. Stulz. 2009. Why do U.S. firms hold so much more cash than they used to? *Journal of Finance* 64:1985–2021.
- Berk, J., and P. DeMarzo. 2019. *Corporate finance*. 5th ed. Harlow: Pearson.
- Bharath, S., and T. Shumway. 2008. Forecasting default with the Merton distance-to-default model. *Review of Financial Studies* 21:1339–69.
- Biorn, E. 2000. Panel data with measurement errors: Instrumental variables and GMM procedures combining levels and differences. *Econometric Reviews* 19:391–424.
- Blundell, R., L. Pistaferri, and I. Preston. 2008. Consumption inequality and partial insurance. *American Economic Review* 98:1887–921.
- Boehmer, E., and J. Wu. 2013. Short selling and the price discovery process. *Review of Financial Studies* 26:287–322.
- Bolton, P., H. Chen, and N. Wang. 2011. A unified theory of Tobin's q, corporate investment, financing, and risk management. *Journal of Finance* 66:1545–78.
- . 2013. Market timing, investment, and risk management. *Journal of Financial Economics* 109:40–62.

- Bolton, P., N. Wang, and J. Yang. 2019. Optimal contracting, corporate finance, and valuation with inalienable human capital. *Journal of Finance* 74:1363–429.
- . 2021. Leverage dynamics and financial flexibility. Working Paper, Columbia University.
- Brennan, M., and A. Subrahmanyam. 1996. Market microstructure and asset pricing: On the compensation for illiquidity in stock returns. *Journal of Financial Economics* 41:441–64.
- Breugem, M., R. Marfe, and F. Zucchi. 2021. Corporate policies and the term structure of risks. Working Paper, Collegio Carlo Alberto.
- Byun, S., V. Polkovnichenko, and M. Rebello. 2019a. Composition of cash flow shocks, firm investment, and cash holdings. Working Paper, University of Texas at Dallas.
- . 2019b. The long and short of cash flow shocks and debt financing. Working Paper, University of Texas at Dallas.
- Carlson, M., A. Fisher, and R. Giammarino. 2004. Corporate investment and asset price dynamics: Implications for the cross-section of returns. *Journal of Finance* 59:2577–603.
- . 2006. Corporate investment and asset price dynamics: Implications for SEO event-studies and long-run performance. *Journal of Finance* 61:1009–34.
- Chang, X., S. Dasgupta, G. Wong, and J. Yao. 2014. Cash-flow sensitivities and the allocation of internal cash flow. *Review of Financial Studies* 27:3628–57.
- Cocco, J., F. Gomes, and P. Maenhout. 2005. Consumption and portfolio choice over the life cycle. *Review of Financial Studies* 18:491–533.
- Cochrane, J. 1994. Permanent and transitory components of gnp and stock prices. *Quarterly Journal of Economics* 109:241–65.
- Cohen, R., P. Gompers, and T. Vuolteenaho. 2002. Who underreacts to cash-flow news? Evidence from trading between individuals and institutions. *Journal of Financial Economics* 66:409–62.
- Danis, A., D. Retzl, and T. Whited. 2014. Refinancing, profitability, and capital structure. *Journal of Financial Economics* 114:424–43.
- De Loecker, J., J. Eeckhout, and G. Unger. 2020. The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics* 135:561–644.
- Décamps, J.-P., S. Gryglewicz, E. Morellec, and S. Villeneuve. 2017. Corporate policies with permanent and transitory shocks. *Review of Financial Studies* 30:162–210.
- Décamps, J.-P., T. Mariotti, J.-C. Rochet, and S. Villeneuve. 2011. Free cash flow, issuance costs, and stock prices. *Journal of Finance* 66:1501–44.
- DeMarzo, P., M. Fishman, Z. He, and N. Wang. 2012. Dynamic agency and the q theory of investment. *Journal of Finance* 67:2295–340.
- DeMarzo, P., and Y. Sannikov. 2006. Optimal security design and dynamic capital structure in a continuous-time agency model. *Journal of Finance* 61:2681–724.
- Denis, D., and V. Sibilkov. 2010. Financial constraints, investment, and the value of cash holdings. *Review of Financial Studies* 23:247–69.
- Duchin, R. 2010. Cash holdings and corporate diversification. *Journal of Finance* 65:955–92.
- Eckbo, E., and M. Kissler. 2021. The leverage-profitability puzzle resurrected. *Review of Finance* 25:1089–128.
- Erickson, T., C. Jiang, and T. Whited. 2014. Minimum distance estimation of the error-in-variables model using linear cumulant equations. *Journal of Econometrics* 183:211–21.
- Fama, E. F., and K. R. French. 1997. Industry costs of equity. *Journal of Financial Economics* 43:153–93.

- Froot, K., D. Scharfstein, and J. Stein. 1993. Risk management: Coordinating corporate investment and financing policies. *Journal of Finance* 48:1629–58.
- Garleanu, N., S. Panageas, and J. Yu. 2012. Technological growth and asset pricing. *Journal of Finance* 67:1265–92.
- Glosten, L., and L. Harris. 1988. Estimating the components of the bid/ask spread. *Journal of Financial Economics* 21:123–42.
- Goldstein, R., N. Ju, and H. Leland. 2001. An EBIT-based model of dynamic capital structure. *Journal of Business* 74:483–512.
- Gorbenko, A. S., and I. A. Strebulaev. 2010. Temporary versus permanent shocks: Explaining corporate financial policies. *Review of Financial Studies* 23:2591–647.
- Gottschalk, P., and R. Moffitt. 2009. The rising instability of U.S. earnings. *Journal of Economic Perspectives* 23:3–24.
- Graham, J., and M. Leary. 2018. The evolution of corporate cash. *Review of Financial Studies* 31:4288–333.
- Gryglewicz, S., S. Mayer, and E. Morellec. 2020. Agency conflicts and short- versus long-termism in corporate policies. *Journal of Financial Economics* 136:718–42.
- Guay, W., and J. Harford. 2000. The cash-flow permanence and information content of dividend increases versus repurchases. *Journal of Financial Economics* 57:385–415.
- Guiso, L., L. Pistaferri, and F. Schivardi. 2005. Insurance within the firm. *Journal of Political Economy* 113:1054–87.
- Hackbarth, D., A. Rivera, and T.-Y. Wong. 2021. Optimal short-termism. *Management Science* Forthcoming.
- Hamilton, J. 1994. *Time series analysis*. Princeton, NJ: Princeton University Press.
- . 2018. Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics* 100:831–43.
- He, Z. 2009. Optimal executive compensation when firm size follows a geometric Brownian motion. *Review of Financial Studies* 22:859–92.
- Hennessy, C., and T. Whited. 2007. How costly is external financing? Evidence from a structural estimation. *Journal of Finance* 62:1705–45.
- Hoberg, G., and G. Phillips. 2016. Text-based network industries and endogenous product differentiation. *Journal of Political Economy* 124:1423–65.
- Hoffmann, F., and S. Pfeil. 2010. Reward for luck in a dynamic agency model. *Review of Financial Studies* 23:3329–45.
- Hugonnier, J., S. Malamud, and E. Morellec. 2015. Capital supply uncertainty, cash holdings, and investment. *Review of Financial Studies* 28:391–445.
- Kim, W., and M. Weisbach. 2008. Motivations for public equity offers: An international perspective. *Journal of Financial Economics* 87:281–307.
- Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman. 2017. Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics* 132:665–712.
- Kwiatkowski, D., P. Phillips, P. Schmidt, and Y. Shin. 1992. Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics* 54:159–78.
- Lee, B.-S., and O. M. Rui. 2007. Time-series behavior of share repurchases and dividends. *Journal of Financial and Quantitative Analysis* 42:119–42.
- Leland, H. E. 1994. Corporate debt value, bond covenants, and agency costs. *Journal of Finance* 49:1213–52.
- . 1998. Agency costs, risk management, and capital structure. *Journal of Finance* 53:1213–43.

- Ma, L., H. Wang, and J. Chen. 2010. Analysis of Kalman filter with correlated noises under different dependence. *Journal of Information and Computational Science* 5:1147–54.
- McLean, D. 2011. Share issuance and cash savings. *Journal of Financial Economics* 99:693–715.
- Meghir, C., and L. Pistaferri. 2004. Income variance dynamics and heterogeneity. *Econometrica* 72:1–32.
- Morellec, E., B. Nikolov, and N. Schürhoff. 2012. Corporate governance and capital structure dynamics. *Journal of Finance* 67:803–48.
- Morellec, E., and N. Schürhoff. 2011. Corporate investment and financing under asymmetric information. *Journal of Financial Economics* 99:262–88.
- Morellec, E., and C. Smith. 2007. Agency conflicts and risk management. *Review of Finance* 11:1–23.
- Opler, T., L. Pinkowitz, R. Stulz, and R. Williamson. 1999. The determinants and implications of corporate cash holdings. *Journal of Financial Economics* 52:3–46.
- Rebelo, S., N. Wang, and J. Yang. 2020. Rare disasters, financial development, and sovereign debt. Working Paper, Northwestern University.
- Riddick, L., and T. Whited. 2009. The corporate propensity to save. *Journal of Finance* 64:1729–66.
- Shumway, R., and D. Stoffer. 1982. An approach to time series smoothing and forecasting using the EM algorithm. *Journal of Time Series Analysis* 3:253–64.
- Strebulaev, I. A. 2007. Do tests of capital structure theory mean what they say? *Journal of Finance* 62:1747–87.
- White, H. 1982. Maximum likelihood estimation of misspecified models. *Econometrica* 50:1–25.