

Production-Based Option Pricing^{*}

Andrea Gamba

Alessio Saretto

Warwick Business School

Federal Reserve Bank of Dallas[‡]

University of Warwick[†]

Abstract

We show that a structural model incorporating firms’ real options and financial leverage decisions can generate forward-looking equity distributions of various shapes, resulting in highly flexible implied volatility surfaces. In contrast, exogenously specified models with stochastic volatility and jumps can only generate smiles and downward-sloping curves. We explain the observed negative relationship between downward-sloping implied volatility surfaces and the value of real options: real options contribute positively to equity skewness, and more so the more convex their payoff, which is the case when the real option is “at-the-money” and has the lowest value.

JEL Classifications: G12, G32

Keywords: option pricing, future realized skewness, implied skewness, real options, leverage, investments

^{*}The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Dallas, the Federal Reserve System, or their staffs. We thank Adi Sunderam, Alexey Zhdanov, and seminar participants at the Frontiers in Finance, Northern Finance Association 2023, Real Options Conference 2024, the Virtual Derivatives Workshop, the Virtual Finance Seminar series at the University of Lancaster, Federal Reserve Bank of Atlanta, Federal Reserve Bank of Dallas, Texas Tech University, University of Calgary, University of Wisconsin-Milwaukee, University of Warwick, University of Birmingham, and Washington University in St. Louis for useful comments. We used Claude Code to assist with proofreading, editing, and improving clarity and flow. We reviewed all suggestions to ensure the accuracy of the arguments and to maintain their original voice and intent. All errors are our own.

[†]Scarman Road, Coventry, UK CV4 7AL. *E-mail: andrea.gamba@wbs.ac.uk*

[‡]2200 N. Pearl Street, Dallas, Texas 75201. *E-mail: alessio.saretto@dal.frb.org*

I. Introduction

Stocks' implied volatility surfaces exhibit significant variability across both cross-sectional and time-series dimensions, reflecting the heterogeneity of market expectations regarding forward-looking (risk-neutral) distributions. For instance, at three-month maturities, an average of 69% of stocks display a smirk pattern, characterized by monotonically decreasing implied volatility as strike prices increase. In contrast, 22% of stocks exhibit a smile, where implied volatility first decreases and then increases. In rarer cases, the surface is either concave, forming an inverted smile or “frown,” or it increases steadily with strike prices. Over more than two decades of data, the prevalence of smirk patterns in implied volatility surfaces has varied significantly, ranging from 40% to 94%. Additionally, implied volatility surfaces show substantial variation across contract maturities. On average, implied volatilities decline as maturities extend, with the proportion of stocks exhibiting smirks exceeding 80% for contracts approaching one year to maturity. These patterns highlight the complexity and richness of equity option price dynamics.

We show that, at low frequency, variability in implied (risk-neutral) moments derived from option prices can be explained by firm characteristics tied to investment and financing policies. We consider a class of models that endogenously generate differences in forward-looking equity price distributions by structurally linking asset prices to firm policies. For example, Morellec and Zhdanov (2019) provide a theoretical analysis of the interaction between product market competition and option prices.

Our model builds on seminal studies such as Hennessy and Whited (2005, 2007) and Zhang (2005), which provide widely used frameworks in dynamic corporate finance and investment-based asset pricing. Gamba and Saretto (2020) employ a similar setup to explain the

cross-section of CDS prices. We extend their approach by introducing a method to price equity options that mature over multiple periods and by studying higher moments of equity return distributions. By linking a structural model of the firm to option pricing, our work connects to earlier research, including Geske (1979), which provides a closed-form solution for a compound option, and Toft and Prucyk (1997), which incorporates taxes and bankruptcy costs following Leland (1994). Importantly, we integrate real investment decisions into the firm's policy set, recognizing their critical role in shaping the distribution of all corporate security prices (Kuehn and Schmid (2014), for example, study the properties of corporate debt prices through the lens of a production economy).

The combined effects of dynamic leverage and real options on equity valuations generate implied volatility surfaces and forward-looking distributions that closely align with empirical observations. As equity prices decline, leverage amplifies asset volatility, leading to increased equity volatility. At the same time, the probability of default rises, concentrating more mass on the left tail of the equity distribution. This dynamic, similar to the mechanism described by Toft and Prucyk (1997), results in more negative skewness and contributes to the prevalence of downward-sloping implied volatility surfaces, commonly known as left smirks. Conversely, real options (i.e., the flexibility to adjust productive capital in response to economic conditions) have the opposite effect. When exercised under favorable conditions for equity holders, real options shift the equity distribution's mass to the right, increasing skewness. Importantly, the magnitude of this effect is inversely related to the value of the real option. In a standard production economy, the payoff of a real option resembles that of a straddle, combining the flexibility to expand or contract operations (see, for example, Gu, Hackbarth, and Johnson, 2017; Aretz and Pope, 2018). The convexity of the real straddle is highest when the option is at-the-money and relatively

inexpensive, but convexity declines when either expansion or contraction becomes more likely. Thus, firms with more valuable real options exhibit lower skewness.

An important assumption of our model is the universality of real options. As in Hackbarth and Johnson (2015), all firms have pathways to adjust their capital stock in response to future states of nature. This contrasts with analyses that impose a dichotomy between firms with and without real options (e.g., Trigeorgis and Lambertides, 2014, and Del Viva, Kasanen, and Trigeorgis, 2017). By endogenizing the moneyness of real options as a function of firms' current conditions, our model captures the dynamic interplay between investment opportunities and market outcomes.

The key feature of the model is that, since the equilibrium choices of production capital and the level of indebtedness depend on both idiosyncratic and systematic shocks, the resulting equilibrium option prices are also state-contingent and endogenously determined. Consequently, the model can generate any form of implied volatility curve (upward sloping, downward sloping, U-shaped, or even inverted U-shaped) and implied distributions (fat tails, long tails, or short tails). With sufficient real and financing frictions, the model can replicate the observed cross-section of implied volatilities. For example, firms with low leverage and at-the-money real options (i.e., those that are as likely to invest as they are to disinvest in the next period) tend to have upward-sloping implied volatility curves.

The ability of the model to generate a variety of forward-looking distributions is crucial because it enables us to capture the significant cross-sectional variability observed in the data using a single set of calibrated parameters. For example, at a three-month maturity, simulated firms exhibit an implied volatility smirk in 70% of cases, compared with 69% in the empirical data. Additionally, 25% of the firm-period observations in the simulated economy display a smile

pattern, compared with 22% in the data; in the remaining cases, the surface is concave (a “frown”). In the model, the percentage of stocks exhibiting a smirk varies from 37% to 94% across the simulated time periods, matching the variability observed in the data (40% to 94%). Variation also occurs across different contract maturities, with implied volatilities on average declining with longer maturities at similar rates in both the model and the data.

This flexibility becomes clear when we compare our model with a reduced-form benchmark that combines the stochastic volatility of Heston (1993) with the price jumps of Merton (1976) and is calibrated to the same moments (see Section VI). Both models fit the level of the average implied volatility curve about equally well, but the benchmark understates the slope (3.3 percentage points versus 5.3 in the data and 6.5 in our model) and generates only left smirks and smiles, because, in the benchmark, shape and slope are mechanically linked. In our model, they depend on the relative strength of leverage and real options, which varies across firms.

The simulated economy allows us to investigate the equilibrium relationship between option price features (specifically implied skewness and the shape of the implied volatility curve) and firm characteristics, with a clear focus on the roles of leverage and real options. Our results, consistent with observed data, indicate that cross-sectional variation in option prices can be partly attributed to the strategic decisions firms make in response to productivity shocks. In tests of relationships that are not targeted in the calibration, we show that the model recovers the relationship between implied skewness and firm policies. Notably, after controlling for size, implied volatility, and leverage, firms with higher real option values (i.e., higher market-to-book ratios) tend to display more negative implied skewness and steeper implied volatility curves. This evidence suggests that investment and leverage decisions, which are driven by idiosyncratic shocks and the path dependence inherent in capital stock and debt adjustment frictions, play an

important role in shaping the probability distribution of future equity prices as reflected in option prices.

As every model is imperfect, so is ours. First, the technology is homogeneous: all firms share the same production function, so differences in forward-looking distributions do not arise from technology adoption (Garleanu, Panageas, and Yu, 2012) or from multiple lines of productive capital (Tuzel, 2010; Jones and Tuzel, 2013; Miao, 2019). Real option value that stems from many small options exercisable simultaneously may have different implications for skewness than the single margin in our setup. Second, firms have access only to single-period debt; the maturity structure of debt may have sizable asset pricing implications (Chaderina, Weiss, and Zechner, 2022; Friewald, Nagler, and Wagner, 2022), but incorporating multi-period debt into this class of models remains intractable. Third, the stochastic discount factor is specified exogenously, so the model cannot generate higher-order risk premia or speak to whether a skewness risk premium exists in the cross-section (see Kozhan, Neuberger, and Schneider, 2013, for the index market). Because the empirical relation between real option value and skewness is negative in both the physical and risk-neutral distributions, our results should remain valid provided the assumptions needed to generate a skewness risk premium do not alter how firms accrue real options.

II. Data

We construct our sample of U.S. optionable stocks by merging data from CRSP, COMPUSTAT quarterly files, and OptionMetrics for the period 1996–2019. We use common filters (i.e., share codes 10 and 11, total assets in excess of 10 million USD, no financial firms).

For each firm and quarterly filing date, we extract option prices that, at that point in time, have maturities closest to 91 days (one quarter), 182 days (two quarters), and 365 days (one year).¹ Having multiple maturities allows us to construct a term structure of option prices and volatilities.

We construct model-free implied skewness and kurtosis directly from OTM call and OTM put option prices as in Bakshi, Kapadia, and Madan (2003). When brevity is required, we refer to the moments as RN-Skew and RN-Kurt, where RN stands for risk-neutral. We follow the procedure described in Hansis, Schlag, and Vilkov (2010) and implement the code available on Grigory Vilkov’s webpage. We eliminate prices that violate arbitrage bounds, have zero open interest, or have zero-bid quotes. For each stock, quarterly filing date, and maturity, we keep options with moneyness between 0.7 and 1.3. We interpolate their implied volatilities in order to obtain a dense grid of prices relative to moneyness. We then compute implied moments. On average, eight option contracts enter the calculation of implied (risk-neutral) moments.

Using the interpolated data, we extract stock-level measures of implied volatilities $IV(m)$ at moneyness levels $m = \text{strike}/\text{stock price}$. Because we pool calls and puts, we label implied volatilities by moneyness: we refer to $m < 1$ as out-of-the-money (OTM), $m = 1$ as at-the-money (ATM), and $m > 1$ as in-the-money (ITM), which is the convention for puts. We focus our attention on $IV(0.8)$, $IV(1)$, and $IV(1.2)$. We construct the slope of the implied volatility curve as the log difference between $IV(1.2)$ and $IV(0.8)$. For each firm-quarter and option maturity, we

¹Most of the empirical regularities do not change if we consider 30-day contracts. Calibrating the model at monthly frequency is, however, very challenging and produces unreliable simulated economies (i.e., very small changes in some parameters produce substantial changes in the properties of the simulated data). Hence, we prefer to calibrate the model at quarterly frequency and consider 91-day options as the baseline.

also classify the shape of the implied volatility curve into four types: the curve is downward sloping if $IV(0.8) > IV(1) > IV(1.2)$ (left smirk), the curve is upward sloping if $IV(0.8) < IV(1) < IV(1.2)$ (right smirk), the curve is U-shaped if $IV(0.8) > IV(1) < IV(1.2)$ (smile), and the curve is inverted U-shaped if $IV(0.8) < IV(1) > IV(1.2)$ (frown).

For each stock and quarterly filing date t , future realized skewness is measured from daily log returns from $t + 1$ to $t + T$, where $T \in \{91, 182, 365\}$. We record the observation only when a minimum of $\{30, 60, 90\}$ daily returns are available, respectively. Future realized idiosyncratic skewness is computed by using the residuals from a market model in the skewness formula, where the market is the value-weighted aggregate portfolio as in Harvey and Siddique (2000).

The final sample is composed of 3,536 stocks and includes quarterly observations between the years 1996 and 2019.

A. Implied volatility surfaces


The average implied volatility surface is downward sloping along both moneyness and maturity (see left panel of Figure 6). This is similar, although to a lower degree, to the implied volatility surface of index options along the moneyness dimension, but opposite along the maturity dimension (e.g., Dew-Becker, Giglio, Le, and Rodriguez, 2017, show an average upward-sloping term structure of volatilities). There is a considerable amount of cross-sectional and time-series variability in implied volatility surfaces (see Appendix Figure A1).

Variation in implied volatilities through time and across stocks produces a rich cross-section of different “shapes”. We plot the cross-sectional frequency of each surface type for the 91-day maturity in Figure 1.

FIGURE 1

Surface types – time series

The figure plots the time series of frequencies of different implied volatility surface types extracted from 91-day maturity options. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.



The predominant type of surface is a left smirk, which is observed on average 69% of the time, with a large time-series variation between 40% and 94% (see Panel A). The second most frequent type is a smile (Panel B), which is observed on average in 22% of the cases. Right smirks and frowns are less frequent on average; however, they appear in a significant number of stocks during the years of the internet bubble (Panels C and D).


An alternative way to analyze the information contained in option prices is to compare option implied moments to the respective realized moments: we plot time series of cross-sectional averages, as well as the 5th and 95th percentiles, of implied volatility, skewness, and kurtosis extracted from 91-day options, and compare them to realized volatility, skewness, and kurtosis in

Figure 2. On average, implied and realized volatility tend to move in similar ways. Implied skewness becomes more negative through the sample, while realized skewness hovers around zero. Realized kurtosis is on average higher than implied kurtosis, but the latter increases the most toward the end of the sample.

FIGURE 2

Implied and realized moments

The figure plots the time series of cross-sectional averages, as well as the 5th and 95th percentiles, of implied volatility, skewness, and kurtosis extracted from 91-day maturity options (top row), and realized volatility, skewness, and kurtosis (bottom row). The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.



B. Option prices and firm characteristics

In Table I we present results of regressions where the dependent variable is either the implied skewness (columns 1 to 3) or the implied volatility (IV) slope (columns 4 to 6). In Panel A, all observations at time t are demeaned at the firm level using only information up to time t . Panel B reports results of Fama-MacBeth regressions. We present results based on traditional firm fixed effects and the Hjalmarsen (2010) model in Appendix Table A1. We also present results that include a larger set of controls and that also consider 182- and 365-day maturity options in Appendix Tables A2 and A3.

The results confirm the findings reported in Table 3 of Morellec and Zhdanov (2019). After controlling for different fixed effects, the implied skewness (or the slope of implied volatility) is positively related to ATM implied volatility and negatively related to size (i.e., the natural logarithm of total assets), book leverage, and the market-to-book ratio. The sign of the relation with profitability varies across specifications.

The leverage coefficient is consistent with the leverage channel that has tied leverage to the equity distribution since the pioneering work of Merton (1974). Market-to-book ratios (a proxy for the value of real options) have also been linked to skewness and IV smirks in prior work (e.g., Morellec and Zhdanov, 2019). We explain these relationships in the context of a structural model that endogenizes investment and debt decisions.

From a theoretical perspective, some of these relationships are not easily explained in the context of traditional structural models (*à la* Merton), such as those of Toft and Prucyk (1997) and Geske, Subrahmanyam, and Zhou (2016). These models describe the impact of leverage and

TABLE I

Option prices and firm characteristics

The table presents results of regressions where the dependent variable is either the implied skewness (columns 1 to 3) or the implied volatility slope (columns 4 to 6). Implied skewness is based on 91-day maturity options. Implied volatility (IV) slope is defined as the log difference of the ITM IV and OTM IV (multiplied by 100). We control for the level of ATM implied volatility (IV), size, book leverage, profitability, and market-to-book ratio (M2B). As an alternative to M2B, we include either Long M2B, the long-run component of M2B extracted using the methodology proposed by Rhodes-Kropf, Robinson, and Viswanathan (2005), or GO, the growth option measure constructed in a similar fashion to Trigeorgis and Lambertides (2014). Panel A reports results of a modified firm fixed-effects regression that avoids including forward-looking information: all observations at time t are demeaned at the firm level using only information up to time t . Panel B reports results of a Fama-MacBeth regression. All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications. A constant is included but not reported. Each specification in Panel A contains quarter fixed effects. We tabulate parameter estimates and t -statistics in parentheses based on standard errors clustered at the quarter and firm levels in Panel A, and heteroskedasticity-robust standard errors in Panel B. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.

	Panel A: Panel Regression					
	Implied skewness			IV Slope		
	(1)	(2)	(3)	(4)	(5)	(6)
IV	7.64 (19.39)	7.41 (17.41)	7.65 (17.79)	2.84 (19.40)	2.78 (18.53)	2.85 (19.21)
Size	-4.04 (-8.03)	-3.84 (-7.88)	-3.51 (-7.24)	-0.92 (-9.86)	-0.92 (-7.83)	-0.84 (-9.98)
Leverage	-1.26 (-2.85)	-1.51 (-5.59)	-1.27 (-4.01)	-0.18 (-3.24)	-0.25 (-4.25)	-0.19 (-3.32)
Profitability	0.05 (1.65)	0.07 (4.31)	0.05 (1.93)	0.04 (4.35)	0.04 (3.89)	0.04 (3.33)
M2B	-3.17 (-7.49)			-0.58 (-4.54)		
Long M2B		-1.92 (-7.90)			-0.53 (-8.00)	
GO			-0.85 (-4.31)			-0.36 (-5.40)
Adjusted- R^2	0.16	0.15	0.15	0.09	0.09	0.09
FE R^2	0.10	0.10	0.10	0.04	0.04	0.04

	Panel B: Fama-MacBeth Regression					
	Implied skewness			IV Slope		
	(1)	(2)	(3)	(4)	(5)	(6)
IV	10.15 (9.48)	10.25 (9.05)	10.42 (9.31)	5.22 (11.16)	5.22 (11.03)	5.28 (11.38)
Size	-13.15 (-7.44)	-11.79 (-7.35)	-11.56 (-6.97)	-1.31 (-6.05)	-1.24 (-5.73)	-1.16 (-5.38)
Leverage	-1.12 (-3.81)	-2.18 (-8.11)	-1.42 (-4.53)	-0.21 (-4.07)	-0.40 (-7.07)	-0.26 (-5.04)
Profitability	-0.25 (-1.38)	-0.19 (-1.08)	-4.35 (-11.59)	0.25 (3.86)	0.27 (3.51)	-0.44 (-4.27)
M2B	-6.92 (-12.00)			-0.83 (-8.99)		
Long M2B		-4.48 (-9.79)			-0.84 (-9.97)	
GO			-5.93 (-12.12)			-1.02 (-8.94)

the level of ATM implied volatility well, but do not explain how market-to-book values are related to option prices.

Kozhan et al. (2013) document the existence of a sizable skewness risk premium in equity index options. It is not clear whether such a risk premium also exists in the cross-section, although the central column of Figure 2 does suggest that implied skewness is more negative than realized skewness for the average stock. We do not investigate why this difference exists. However, we estimate regressions similar to those reported in Table I, where the dependent variable is the future realized skewness calculated from daily stock returns. We discuss these results in detail in Appendix B.

1. Alternatives to the market-to-book ratio as a proxy for the value of real options

The negative relationship between market-to-book and implied skewness/IV slope may be driven by overvaluation. If stocks are overvalued, they may experience negative future returns,

leading to negative forward-looking skewness. To address this possibility, we extract the long-run component of the market-to-book ratio following the procedure in Rhodes-Kropf et al. (2005), applying it to the market-to-book ratio of assets rather than equity. We then replace M2B with its long-run component (Long M2B) and report the results in columns 2 and 5 of Panel A of Table I.

We find a statistically significant negative coefficient for the long-run ratio, somewhat smaller in magnitude than the coefficient for the market-to-book ratio (-1.92 vs. -3.17 , and -0.53 vs. -0.58 , respectively). This suggests that while temporary mispricing contributes to more negative skewness, it does not fully account for the negative relation between the market-to-book ratio and implied skewness: Long M2B, which may be more directly linked to the value of real options than the market-to-book ratio, retains a significant negative coefficient.

As an alternative, we also consider the GO ratio of Trigeorgis and Lambertides (2014) and present the results in columns 3 and 6 of Table I. We find a similar negative relationship with implied skewness, though with a smaller magnitude. Since all variables are standardized, the coefficients are directly comparable across specifications.

For robustness, we report results of a Fama-MacBeth regression with the same variables in Panel B. With the exception of profitability, the coefficients remain statistically significant, with the expected signs.

III. Basic intuition

We now develop the key economic intuition using a stripped-down version of the quantitative model presented in Section V. Two economic forces shape the implied volatility surface of equity options: financial leverage and real options. We isolate each force in turn.

A. Setup

We adopt a neoclassical framework with discrete time, $t = 0, 1, 2, \dots$, and an infinite horizon. We introduce three simplifications, which do not affect the general points we make here, to derive closed-form solutions. First, we assume that all securities are priced with a constant discount factor $0 < \beta < 1$, corresponding to discounting at the risk-free rate. Second, the firm's log productivity shock, z_t , follows a random walk $z_{t+1} = z_t + \sigma \varepsilon_{t+1}$, where z_0 is known and the innovations ε_t for $t \geq 1$ are i.i.d. with a standard normal distribution. Third, the firm starts with capital stock k_1 at $t = 0$, which does not depreciate, and can adjust it only once, at date $t = 1$; thereafter, capital remains fixed.

Cash flow from operations at t is $\pi(z_t, k_t) = e^{z_t} k_t^\alpha - f k_t$, where $0 < \alpha < 1$ governs decreasing returns to scale and $f \geq 0$ is a proportional operating cost. The parameter f plays the role of leverage: as in Carlson, Fisher, and Giammarino (2004), it captures the combined effects of financial leverage and depreciation that are present in the full model.²

At $t = 1$ the equity holder invests $I_1 = k_2 - k_1$. We refer to this one-time ability to adjust productive capital as the *real option*, with the understanding that it encompasses both expansion and contraction. There are no transaction costs or other financing frictions, so investment is financed indifferently with internal or external equity.

The timeline for financial options is as follows: at $t = 0$, after z_0 has been observed, European call and put options are written on the firm's equity with strike price κ and maturity $t = 1$. At $t = 1$ the new shock z_1 is realized, and the firm chooses k_2 . The equity value at $t = 1$ is

²A fixed operating cost would produce qualitatively similar effects—higher firm risk and more negative equity skewness—but would not interact with the investment decision, so we focus on proportional costs.

$S(z_1) = \max_{k_2} \{\pi(z_1, k_1) - I_1 + V(z_1, k_2)\}$, where

$V(z_1, k_2) = \mathbb{E}_1 \left[\sum_{j=1}^{\infty} \beta^j \pi(z_{1+j}, k_2) \right] = \Phi(z_1) k_2^\alpha - \frac{\beta}{1-\beta} f k_2$ is the continuation value under the assumption that capital remains at k_2 forever.³

From the first-order condition, the optimal capital choice is $k^*(z_1) = \left(\frac{\alpha(1-\beta)\Phi(z_1)}{1-\beta+\beta f} \right)^{\frac{1}{1-\alpha}}$.

The equity value is then $S(z_1) = \pi(z_1, k_1) - I_1^*(z_1) + V(z_1, k^*(z_1))$, where $I_1^*(z_1) = k^*(z_1) - k_1$ is the optimal investment. The options expire after the investment decision, and

$C(\kappa) = \beta \int \max\{S(z_1) - \kappa, 0\} \varphi(z_1) dz_1$ is the price at $t = 0$ of a European call option with strike κ , where $\varphi(\cdot)$ is the normal density of z_1 . The moneyness of the option is the ratio of strike to underlying price, $\kappa/S(z_0)$, where $S(z_0) = \beta \mathbb{E}[S(z_1)|z_0]$. We use this framework to isolate the two forces—leverage and real options—that determine the shape of implied volatility curves.

B. The leverage effect

We first shut down the investment decision (i.e., capital is fixed at k_1) and vary the operating-cost parameter f . Figure 3 displays the results. When $f = 0$ the model reduces to a Black–Scholes economy: equity volatility equals 20% (the volatility of the productivity shock), skewness is zero, and the implied volatility curve is flat at 20%. As f increases, the probability of low or negative cash flows rises, which shifts the equity return distribution to the left, raises volatility, and generates negative skewness. These changes tilt the implied volatility curve

³Given the distributional assumptions on the productivity shock,


$$\Phi(z_1) = \sum_{j=1}^{\infty} \beta^j \mathbb{E}_1 [e^{z_{1+j}}] = e^{z_1} \sum_{j=1}^{\infty} \beta^j \left(e^{\sigma^2/2} \right)^j = e^{z_1} \frac{\beta e^{\sigma^2/2}}{1 - \beta e^{\sigma^2/2}},$$

where the sum converges provided that $\beta e^{\sigma^2/2} < 1$.

FIGURE 3

Model with no investments

On the left, we plot the equity log returns between $t = 0$ and $t = 1$ against z_1 for different values of the leverage parameter, $f \in \{0, 0.25, 0.5\}$. For each case, we report in the legend the expected value, volatility, and skewness. On the right, we display the corresponding implied volatility of the call option prices at $t = 0$ against moneyness between 0.8 and 1.2. For each case, we report in the legend the slope of the implied volatility curve (i.e., IV calculated for a moneyness of 1.2 minus IV for a moneyness of 0.8) and the implied skewness (RN-Skew). The remaining parameters are $\sigma = 20\%$, $k_1 = 1$, $z_0 = 0$, $\alpha = 0.3$, and $\beta = e^{-0.1}$.



downward—producing the left smirk commonly observed in the data. This mechanism is well understood in the literature (Geske, 1979; Toft and Prucyk, 1997; Morellec and Zhdanov, 2019): with fixed production technology, operating leverage (or, more generally, financial leverage) induces a negative slope in the implied volatility surface.


C. The real option effect

We now reintroduce the investment decision and set $f = 0$ to isolate the real option channel. Figure 4 presents the results for three values of the production-function curvature,

FIGURE 4

Model with investments

The figure plots log investment at $t = 1$ and log equity return between $t = 0$ and $t = 1$ against z_1 . We plot the implied volatility curve at $t = 0$ derived from equity financial options expiring at $t = 1$ against moneyness. We set $f = 0$ to isolate the effect of the real option, and k_1 so that the average investment is zero for $\alpha = 0.3$. We present three cases corresponding to different curvatures of the production function $\alpha \in \{0.15, 0.30, 0.45\}$. This produces variation in the incentives to invest, without affecting the distributional properties of the assets in place, $\pi(z_1, k_1)$ (i.e., the dotted lines in the three panels in the middle column are identical). We report in the legend the average investment, the average equity (log) return and its skewness, the mean and skewness of the real option value, as well as the slope and implied skewness (RN-Skew) obtained from the implied volatility surface. The other parameters are $\sigma = 20\%$, $k_1 = 4.47$, $z_0 = 0$, and $\beta = e^{-0.1}$.



	Investment	Equity Return		Real Option		Equity Option	
	Mean	Mean	Skew	PVRO	Skew	IV Slope	RN-Skew
$\alpha = 0.15$	-1.08	0.09	0.22	0.26	0.03	0.92	0.20
$\alpha = 0.30$	0.00	0.08	0.23	0.03	4.46	0.96	0.21
$\alpha = 0.45$	1.15	0.06	0.22	0.85	1.95	0.87	0.18

$\alpha \in \{0.15, 0.30, 0.45\}$, holding all other parameters fixed. We set the initial capital stock

($k_1 = 4.47$) so that for $\alpha = 0.3$ the average investment at $t = 1$ is zero.

The firm's real option can be viewed as a straddle: a combination of a call option to

expand and a put option to contract. The curvature parameter α determines the moneyness of this straddle. In the base case ($\alpha = 0.30$, in red), the firm is equally likely to expand or contract—the real straddle is at-the-money. When $\alpha = 0.45$ (black), the optimal capital stock is higher, so the firm almost always expands; the real call is deep in-the-money and the real put is out-of-the-money. The reverse holds for $\alpha = 0.15$ (blue), where the firm predominantly contracts.

A key feature of this comparison is that varying α does not alter the distribution of the return on the assets in place, $\pi(z_1, k_1)$. This can be seen directly in the figure: the dotted lines in the middle column are identical across the three panels. Any difference in equity distributions is therefore entirely due to the real option.

As with any straddle, the real option is least valuable when it is at-the-money. The legend of Figure 4 confirms this: the present value of the real option (PVRO) is 0.03 in the base case ($\alpha = 0.30$), compared with 0.26 and 0.85 for $\alpha = 0.15$ and $\alpha = 0.45$, respectively.

The real option increases equity skewness because it improves equity returns in the states where capital is adjusted in the “right” direction. When the firm predominantly invests ($\alpha = 0.45$), the real option raises returns more in good states, shifting the right tail outward. When the firm predominantly disinvests ($\alpha = 0.15$), the real option raises returns more in bad states, compressing the left tail. When the straddle is at-the-money ($\alpha = 0.30$), both effects operate simultaneously: the firm expands in good states and contracts in bad states, shifting the entire distribution to the right.

The contribution of the real option to skewness depends on its convexity. van Zwet (1964) shows that a convex transformation of a random variable increases its skewness, and that more convex transformations increase it more. For option-like payoffs, convexity is concentrated around the strike, so its effect on skewness is largest when the option is at-the-money (where the

gamma of the option is largest) and declines as the option moves deeper in- or out-of-the-money. It follows that the real option adds the most skewness precisely when it is at-the-money and least valuable. The legend confirms this pattern: the skewness of the real option value is 4.46 in the base case, far exceeding the values of 0.03 and 1.95 for $\alpha = 0.15$ and $\alpha = 0.45$, respectively. The same ordering appears in equity implied skewness (0.21 versus 0.20 and 0.18), although the differences are attenuated by the assets-in-place component.

D. Summary

Two opposing forces shape implied volatility curves. Leverage induces negative skewness and a downward-sloping implied volatility surface; this effect weakens as economic conditions improve and default risk recedes. The real option induces positive skewness and an upward-sloping surface; crucially, this effect is strongest when the real option is at-the-money and therefore least valuable. The interplay of these two forces, combined with the heterogeneity in firm characteristics that determines their relative strength, generates the rich cross-sectional variation in implied volatility shapes documented in the data.

IV. Case studies


We illustrate the economic mechanism that connects real options to implied skewness with two real-world examples. Figure 5 plots smoothed time series of implied skewness for Alcoa and Oracle against their respective industry peers (based on SIC code groupings).

Around 2010, Alcoa was a large, mature manufacturing firm under pressure after the financial crisis. The aluminum industry faced weak pricing, chronic global overcapacity, and a

FIGURE 5

Case studies

The figure plots smoothed time series of implied skewness for Alcoa (left panel) and Oracle (right panel) against their respective industry peers, defined by two-digit SIC code groupings. Implied skewness is computed from 91-day maturity options following Bakshi et al. (2003). The sample period is 1996 to 2019.



cost squeeze that made traditional upstream smelting and refining unattractive. Alcoa responded with a multi-year restructuring plan that was plausible but uncertain—investors faced ambiguity over whether the plan would create value or simply manage decline in a difficult commodity cycle. The outcome became clearer over the following years and culminated in the November 2016 separation into Alcoa Corporation and Arconic.⁴ Alcoa held a restructuring option that could have led to either expansion or contraction. In other words, the real option was at-the-money. As such, the option had a large degree of convexity and contributed positively to the skewness of the equity distribution. Consistent with this interpretation, Figure 5 shows that Alcoa’s implied skewness rose relative to its peers after 2011 and declined following the 2016 break-up.

By the mid-2000s, Oracle was a large, mature incumbent with a well-established acquisition-and-integration strategy. Starting with PeopleSoft and Siebel in 2005, Oracle executed

⁴<https://news.alcoa.com/press-releases/press-release-details/2015/Alcoa-to-Separate-into-Two-Industry-Leading-Public-Companies-Completing-Successful-Multi-Year-Transformation/default.aspx>

a steady sequence of acquisitions through 2014 that broadened its applications footprint without departing from its core strategy.⁵ Oracle held a clear and very valuable in-the-money expansion option that contributed only moderately to the skewness of the equity distribution. Oracle issued some debt to finance the acquisitions, and its leverage jumped substantially from about 1.5% to 27% in the second quarter of 2006. Overall, the acquisitions and the related financing needs contributed negatively to the skewness of the equity distribution. Figure 5 shows that Oracle’s implied skewness decreased relative to its peers after 2007, bottomed out around 2012, and remained below its peers for several years.

V. Quantitative analysis

We now turn to a fully specified quantitative model that embeds the two forces identified in Section III—leverage and real options—into a realistic production economy. The model builds on Gamba and Saretto (2020). In Appendix C, we provide the full description of the economic environment, the firm’s problem, and the pricing of multi-period options within this environment. Here we summarize only the elements needed to present our contribution to option pricing within this framework.

A. Key model elements

A fixed number of independent firms operate in an infinite-horizon, discrete-time, single-good economy. Firms produce the consumption good using capital, whose productivity is stochastic and subject to both aggregate and firm-specific shocks. Capital depreciates over time.

⁵See several Oracle press releases at <https://www.oracle.com/corporate/pressrelease>.

Firms invest in new capital subject to convex adjustment costs. They cannot save, and they finance investment using operating cash flows together with a mix of equity and single-period unsecured debt. Changes in debt entail proportional adjustment costs.

Firms are taxed, and interest payments are deductible from taxable income, so debt financing is motivated by the interest tax shield. At the same time, debt exposes firms to insolvency risk. We model cash-flow insolvency, which triggers distress costs for shareholders, who may be required to raise external equity. When insolvency is severe enough that equity becomes negative, shareholders default under limited liability, and debtholders recover the firm's assets net of bankruptcy costs. A defaulted firm is restructured and continues operations, ensuring a constant number of firms in the economy. If cash flow to equity is positive, the firm repurchases shares; if it is negative, the firm issues new shares and incurs a transaction cost.⁶ Thus, the number of shares of the firm evolves endogenously as a consequence of payout and financing decisions. Firms choose investment, financing, and default policies to maximize shareholder value.

The economy features many consumers who maximize expected utility over intertemporal consumption by choosing their investments in firms' securities. Each period, they consume the returns on corporate securities and rebalance their portfolios in competitive financial markets. We assume that their decisions can be aggregated into those of a representative agent. Because this agent is well diversified, marginal utility depends only on aggregate shocks. Since the model is used to price options, we follow the traditional option pricing approach and treat options as redundant securities, so that their prices are determined based on the representative agent's marginal utility. In our economic setup, this corresponds to a partial equilibrium in which the

⁶We assume that payouts to equity holders occur exclusively through share repurchases rather than cash dividends. Option pricing becomes much more cumbersome when the firm pays endogenous dividends.

stochastic discount factor is exogenously given. We specify the stochastic discount factor to model countercyclical risk premia.

Financial markets are competitive, so corporate securities trade at fair prices: equity and debt prices reflect rational expectations about the firm's future policies. We assume that prices and policies depend only on the current state of the firm and the economy (the Markov property), and we solve for a stationary equilibrium in which equity values, debt values, and the firm's optimal policies are determined jointly.

To price options written on each firm's stock, we assume that the stock is traded and that there are no arbitrage opportunities. To match the empirical data, the model must produce option prices across different strikes and maturities. We consider European options written on a *single share* of a firm's equity, with maturities of up to four periods (quarters) and a range of strikes. We refer to this collection of option contracts as the *option chain*. For options written in a given period, strikes are set at predetermined moneyness levels, where moneyness is the ratio of the strike to the current stock price. Options are written after the firm's payout for the period, when the number of shares outstanding is known. They expire at a later date, after the new shocks are realized but before the firm makes that period's payout. Hence, the exercise payoff is based on the before-payout stock price.

In our setup, there are two main challenges in determining prices for multi-period options that are consistent with the dynamics of the firm's endogenous policy and the associated stock price, while maintaining the Markov property. The first challenge is that, in each period, different cohorts of options coexist: there are outstanding options written in previous periods, with strikes based on past stock prices, and newly written options with strikes based on the current stock price. To preserve the Markov property, one would need to track all possible strike prices for

legacy options. However, the set of strikes is infinite, as it depends on all past stock prices. The second challenge is that the stock price at each date depends on the number of shares, which evolves endogenously with the firm’s payout and financing policy. This introduces path dependence, which again conflicts with the Markov structure of the economy. To address these issues, we introduce approximations, described in Appendix C, that preserve tractability and have a small impact on option prices. Notably, the price is exact for single-period options, which we use in the model’s calibration.

B. Calibration

Our calibration should not be thought of as a structural estimation. We aim to verify three conditions: whether a standard quantitative model of the firm can reproduce implied volatility patterns observed in the data, whether the required parameters are plausible, and whether the simulated economy reproduces features of the data not used in the calibration.

We fix the five parameters governing aggregate risk and the stochastic discount factor. We take the values in Cooley and Prescott (1995), converted to quarterly frequency, and set the persistence and volatility of the aggregate shock to $\rho_y = 0.986$ and $\sigma_y = 0.006$. The risk-free discount factor is $\beta = 0.9851$, and the stochastic discount factor (SDF) parameters are $\gamma_1 = 3.22$ and $\gamma_2 = -15.3$, yielding an annualized real interest rate of 6.2%. The depreciation rate, equity flotation cost, and debt adjustment cost are taken from Warusawitharana and Whited (2016).

The remaining parameters are chosen by minimizing the sum of squared deviations between data and model moments. The simulated economy consists of 5,000 repetitions of a panel of 1,000 firms observed for 92 quarters. Numerical details are in Appendix C.D.

The calibration targets span three categories. First, the model should reproduce the outcomes of real and financial decisions: average leverage and average investment. Second, with the relevant risk sources in place, the model should generate plausible valuations and equity return distributions—that is, a realistic market-to-book ratio and realistic mean, standard deviation, skewness, and kurtosis of percentage equity returns. Third, the model should match the average 91-day implied volatility curve and the frequency of each IV surface shape (left smirk, smile, right smirk, frown).

Panel A of Table II compares the simulated economy to the data along these dimensions. Investment and financing decisions match well: for instance, the average market-to-book ratio is 1.8 in the simulation versus 2.5 in the data, and the average quarterly return is 2.5% versus 3.2%. The distribution of simulated returns has a slightly lower standard deviation but slightly higher skewness (0.77 versus 0.72) and kurtosis (5.57 versus 5.28).

The model's one-quarter IV curve closely tracks the data, with an at-the-money (ATM) IV of 0.52 compared to 0.54 empirically; the out-of-the-money (OTM) and in-the-money (ITM) levels are similarly close. The model also reproduces the heterogeneity of IV curve shapes: roughly 70% of firm-quarters display a left smirk and 25% a smile, in line with the data. Frown-shaped curves appear at the observed frequency. The one exception is the right smirk, which the calibrated parameters generate only rarely.

The calibration therefore achieves its primary goal: a dynamic structural model of the firm can replicate the key characteristics of equity option prices, including the full range of IV shapes observed in the data.

Panel B of Table II reports the calibrated parameters. The idiosyncratic productivity shock is less persistent ($\rho_z = 0.925$) and more volatile ($\sigma_z = 0.18$) than its aggregate counterpart. The

TABLE II

Model calibration

This table presents the calibration results of the firm model. In Panel A, we compare the quantities that are targeted to calibrate the model. In the left column (*Data*) we report the value of the moment conditions computed from the observed empirical sample, while in the right column (*Simulation*) we report the moment conditions computed from the simulated sample. In Panel B, we report the list of model parameters. Data are from various sources and span the period from January 1996 through December 2019.

Panel A: Calibrated quantities		
<i>Option Prices (one-quarter/91 days to maturity):</i>	Data	Simulation
Implied Volatility OTM	0.583	0.569
Implied Volatility ATM	0.541	0.519
Implied Volatility ITM	0.530	0.504
Frequency of Left Smirk	0.686	0.704
Frequency of Smile	0.220	0.250
Frequency of Right Smirk	0.041	0.000
Frequency of Frown	0.055	0.046
<i>Stock Return:</i>		
Average	0.032	0.025
Standard Deviation	0.315	0.286
Skewness	0.719	0.766
Kurtosis	5.284	5.570
<i>Firm characteristics:</i>		
Market-to-Book	2.549	1.795
Leverage	0.341	0.368
Investments	0.014	0.027

marginal corporate tax rate is $\tau = 0.12$, close to the estimates of Graham (1996a) and Graham (1996b), which average roughly 0.13 for the firms in our sample. The production-function curvature $\alpha = 0.36$ is near the value of 0.3 used in Zhang (2005) and Gomes (2001); published estimates of this parameter vary widely depending on frequency and the treatment of fixed costs. The operating cost is $f = 1.60$, or about 13% of average capital on an annualized basis; for comparison, Carlson et al. (2004) use 1.54, which corresponds to roughly 30% of median operating capacity in their model.

Panel B: Parameters			
<i>Aggregate</i>			
Systematic Productivity Autocorrelation	ρ_y		0.986
Systematic Productivity Volatility	σ_y		0.006
Risk-free Discount Factor	β		0.985
Constant Price of Risk	γ_1		3.220
State-dependent Price of Risk	γ_2		-15.300
<i>Firm-Specific</i>			
Depreciation	δ		0.030
Equity Issuance Cost	ζ		0.059
Debt Adjustment Cost	θ		0.015
Idiosyncratic Productivity Autocorrelation	ρ_z		0.925
Idiosyncratic Productivity Volatility	σ_z		0.181
Production Function	α		0.356
Operating Cost	f		1.603
Cost of Expansion	λ_1		0.500
Cost of Contraction	λ_2		0.926
Corporate Taxes	τ		0.121
Financial Distress Cost	ξ		0.111
Bankruptcy Cost	η		0.451

The bankruptcy cost $\eta = 0.45$ is within the range estimated by Glover (2016) for high-quality firms. Capital expansion costs are close to those in Zhang (2005), while contraction costs are substantially lower.⁷ The financial distress cost $\xi = 0.11$ exceeds the estimate of 0.05 obtained by Gamba and Saretto (2020) from credit default swap spreads. These relatively high bankruptcy and distress costs reflect the use of single-period debt, which is nearly risk-free and rarely triggers default. Larger cost parameters are therefore needed for these frictions to have an appreciable effect, an argument supported by Glover (2016).

⁷A direct comparison to the parameters in Zhang (2005) requires dividing by the depreciation rate. We obtain 16.7 and 30.9 for expansion and contraction costs, compared with 15 and 150 in Zhang (2005) for the same cost function.

In summary, the parameters required to match option prices are plausible and broadly consistent with estimates obtained in other settings.

C. Model performance outside the calibration

The calibration targets the one-quarter IV curve and the frequencies of IV shapes. We now assess whether the model can also account for features of the data that were not used in the calibration.

1. Term structure of implied volatilities

Figure 6 compares the average IV surface in the data (left panel) to that in the simulated economy (right panel) for maturities of one, two, and four quarters. Each curve is obtained by averaging first across time, then across firms, and finally across simulated economies.


The simulated surface closely resembles the empirical one. At every maturity the slope of the IV curve matches the data. The level of the simulated curves at one and two quarters is roughly three percentage points below the data, but the gap narrows to a few basis points at the four-quarter horizon.

Notably, the model generates a downward-sloping IV term structure without exogenously imposing a term structure of volatility. The mechanism is mean reversion: productivity shocks that drive short-horizon volatility tend to revert to long-run values, flattening the relationship between asset and equity volatility. The result is lower option prices—and therefore lower implied volatilities—at longer maturities.

FIGURE 6

Average implied volatility surface comparison

The figure plots the average implied volatility surface extracted from the data (left panel) and from the simulation (right panel). The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.



2. Implied moments across different maturities

Table III compares implied skewness and kurtosis across maturities. In the data, implied skewness is approximately flat at about -0.6 ; the model produces a similarly flat profile near -0.18 . Implied kurtosis in the data declines modestly from 4.2 at one quarter to 3.8 at four quarters. The model produces a gentler slope, at a somewhat lower level (2.82 and 2.77 at the one- and four-quarter horizons, respectively). The data exhibit more cross-sectional dispersion, as reflected by larger standard deviations and wider 5th–95th percentile ranges.

TABLE III

Implied moments

The table compares summary statistics for model-free implied skewness and kurtosis extracted from the data (left side) and from the simulated economy (right side). The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.

	Data				Model			
	91-day maturity				1 quarter			
	Average	Stand.Dev	5 th perc	95 th perc	Average	Stand.Dev	5 th perc	95 th perc
Skew	-0.62	0.49	-1.49	0.00	-0.18	0.44	-0.42	0.43
Kurt	4.23	2.23	2.71	7.65	2.82	1.39	1.95	3.16

	182-day maturity				2 quarters			
	Average	Stand.Dev	5 th perc	95 th perc	Average	Stand.Dev	5 th perc	95 th perc
Skew	-0.57	0.46	-1.37	0.03	-0.13	0.52	-0.30	0.26
Kurt	3.76	1.68	2.40	6.50	2.83	1.87	1.93	3.11

	365-day maturity				4 quarters			
	Average	Stand.Dev	5 th perc	95 th perc	Average	Stand.Dev	5 th perc	95 th perc
Skew	-0.63	0.51	-1.53	0.08	-0.18	0.14	-0.26	0.04
Kurt	3.76	1.72	2.02	7.00	2.77	0.31	2.06	3.05

3. Option prices and firm characteristics

In Section II.B we documented how option prices in the data relate to firm characteristics. Table IV reports the same regressions estimated on the simulated economy. Because the model is fully specified, we can go beyond observable proxies: in addition to the market-to-book ratio, we include the ratio of the present value of real options to the present value of assets in place (PVRO2A), computed by decomposing equity into the value from keeping capital fixed and the value from allowing future adjustments.

The qualitative patterns in the simulated economy largely mirror those in Table I. Implied skewness and the IV slope are negatively related to leverage, as in Toft and Prucyk (1997), and to

TABLE IV

Option prices and firm characteristics

The table presents regression results in the simulated economy that mirror those in the data presented in Table I. Left-hand-side variables include the implied skewness and the implied volatility slope (i.e., log difference of IV(1.2) and IV(0.8)) extracted from one-quarter-maturity options. Explanatory variables include the ATM IV, the natural logarithm of assets (Size), book leverage, profitability, the market-to-book ratio, and the ratio of the present value of the real options to the present value of the assets in place (PVRO2A). PVRO2A is constructed by decomposing the equity into the value derived from keeping current capital constant and the value derived from allowing adjustments, by either investing or disinvesting. Because simulated firms differ only in their shock histories, firm fixed effects have no natural counterpart in the simulated data. For each simulated economy, we therefore estimate a pooled regression on all firm-quarters (1,000 firms over 92 quarters) and report the average coefficient across the 5,000 economies. All right-hand-side variables are standardized. In parentheses, we report the standard deviation of the coefficients across simulated economies. The Data columns report the estimates in columns 1 and 4 of Panel A of Table I.

	Implied Skewness			IV Slope		
	Model	Data		Model	Data	
IV	27.59 (0.25)	28.62 (0.34)	7.64	0.81 (0.02)	1.55 (0.05)	2.84
Size	0.88 (0.02)	1.09 (0.03)	-4.04	0.05 (0.03)	0.07 (0.04)	-0.92
Leverage	-0.57 (0.03)	-0.69 (0.02)	-1.26	-0.05 (0.03)	-0.17 (0.02)	-0.18
Profitability	1.13 (0.61)	-4.49 (0.13)	0.05	3.03 (0.32)	-0.94 (0.03)	0.04
M2B	-6.30 (0.80)		-3.17	-4.49 (0.34)		-0.58
PVRO2A		-0.48 (0.01)			-0.04 (0.01)	

the value of real options. The magnitudes are of the same order: the coefficient on market-to-book in the skewness regression is -6.3 in the model versus -3.2 in the data, and the coefficient on leverage is -0.6 versus -1.3 . The one notable discrepancy is the sign of the size coefficient in the skewness regression, which is positive in the model but negative in the data.

The PVRO2A coefficients are negative in both the skewness and IV slope regressions, confirming the theoretical prediction of Section III: firms with more valuable real options have more negatively skewed implied distributions and steeper IV curves.

VI. Comparison to traditional option pricing models

We now compare our production-based approach to a reduced-form benchmark. The benchmark is a stochastic volatility model with jumps in the price level, combining elements of Heston (1993) and Merton (1976). We label it the Heston–Merton (HM) model.

Under the risk-neutral measure, log returns and variance evolve as

$$\begin{aligned} ds_t &= \left(r - \frac{1}{2}V_t\right) dt + \sqrt{V_t} dW_t + dL_{s,t}, \\ dV_t &= \vartheta(\bar{V} - V_t) dt + \sigma_V \sqrt{V_t} dB_t, \end{aligned}$$

where ϑ is the speed of mean reversion toward the long-run variance $\bar{V}$, and σ_V scales the volatility of variance. The Feller condition $\sigma_V^2 \leq 2\vartheta\bar{V}$ ensures that the variance process remains strictly positive. The two Brownian motions satisfy $\text{Corr}(dW_t, dB_t) = \varrho dt$, so that a negative ϱ generates the leverage effect commonly observed in equity markets. The jump component $dL_{s,t}$ is a compound Poisson process with state-dependent intensity $\iota_t = \iota_0 + \iota_1 V_t$ and normally distributed jump sizes with mean μ_s and standard deviation σ_s .

Richer reduced-form specifications exist, but the HM model provides a good balance between flexibility and parsimony: it captures stochastic volatility, leverage effects, and jumps—the three features most relevant for fitting IV surfaces—while keeping the number of parameters tractable.

We first calibrate the HM model as closely as possible to the procedure used for the production-based model and then compare the two approaches in terms of fit and economic content.

A. Calibration

A fair comparison requires calibrating the HM model with a procedure as close as possible to the one used for the production-based model. This is not entirely straightforward: several HM parameters have nearly offsetting effects on option prices, creating identification problems (Cont and Tankov, 2004). We therefore fix two parameters—the mean-reversion speed ϑ and the initial variance V_0 —and calibrate the remaining seven (ϱ , $\bar{V}$, σ_V , ι_0 , ι_1 , μ_s , σ_s) by minimizing the sum of squared deviations between model and data moments.

The target moments mirror those used for the production-based model (Table II), with one modification: the IV curve is evaluated at five moneyness points instead of three. We also target the frequency of IV surface shapes and the first four moments of realized equity returns.

The comparison requires a cross-section of firms in each model. In the production-based model, firms start from different levels of capital and debt and receive different realizations of the idiosyncratic productivity shock, while the systematic productivity shock is common to all of them. In the HM model, we construct a cross-section by drawing each firm’s initial spot variance from a distribution centered at the average implied variance in the data (V_0 in Table V). We then simulate the paths of stock prices and variances with shocks that are independent across firms, so no systematic shock links the firms over time. The result is a simulated cross-section that evolves over time, for which we compute statistics in the same way as in the data and in the production-based model.

Table V reports the calibrated parameters (Panel A) and the fit to the targeted moments (Panel B). There is no obvious benchmark against which to compare the parameters reported in Panel A.

TABLE V

Comparison with Heston–Merton model

This table compares the production-based model to a Heston–Merton (HM) stochastic volatility model with jumps. Panel A reports the HM parameters: two are fixed to resolve identification issues and seven are calibrated. Panel B compares the targeted moments across the data, the production-based model (Model), and the HM model (HM). Reported moments include the 91-day implied volatility at five moneyness levels (0.8, 0.9, 1, 1.1, 1.2), the frequency of implied volatility curve shapes (left smirk, smile, right smirk, frown), and the first four moments of quarterly realized equity returns. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.

Panel A: Parameters								
<i>Fixed</i>		<i>Calibrated</i>						
ϑ	V_0	ϱ	σ_V	$\bar{V}$	ι_0	ι_1	μ_s	σ_s
3	0.293	-0.149	0.993	0.252	0.311	1.683	-0.082	0.282
Panel B: Calibrated quantities								
						Data	Model	HM
<i>Option prices:</i>								
Implied Volatility - OTM ($m = 0.8$)						0.583	0.569	0.588
Implied Volatility - OTM ($m = 0.9$)						0.557	0.543	0.571
Implied Volatility - ATM ($m = 1$)						0.541	0.519	0.560
Implied Volatility - ITM ($m = 1.1$)						0.532	0.527	0.555
Implied Volatility - ITM ($m = 1.2$)						0.530	0.504	0.555
Frequency of Left Smirk						0.686	0.704	0.696
Frequency of Smile						0.220	0.250	0.304
Frequency of Right Smirk						0.041	0.000	0.000
Frequency of Frown						0.055	0.046	0.000
<i>Stock returns:</i>								
Average						0.032	0.025	0.023
Standard Deviation						0.315	0.286	0.228
Skewness						0.719	0.766	0.620
Kurtosis						5.284	5.570	5.226

B. Comparative analysis

Both the HM and the production-based models fit the average IV curve imperfectly. The production-based model lies below the data, whereas the HM model lies above it. For example,

at-the-money, HM produces an implied volatility of 0.560 versus 0.541 in the data, whereas the production-based model yields 0.519. A similar pattern holds across all moneyness levels.

The production-based model, however, better captures the slope of the IV curve. The spread between very-out-of-the-money and very-in-the-money implied volatilities is 5.3 percentage points in the data (0.583 minus 0.530), 6.5 in the production-based model (0.569 minus 0.504), and only 3.3 in the HM model (0.588 minus 0.555).

The most striking difference emerges in the diversity of IV surface shapes. The HM model generates only left smirks (70% of firm-quarters) and smiles (30%), producing neither right smirks nor frowns. The production-based model generates a richer mix of shapes that more closely resembles the data, where left smirks account for 69%, smiles for 22%, and frowns and right smirks together for nearly 10%.

For the first four moments of realized equity returns, the production-based model is closer to the data than the HM model. The HM model understates the mean, standard deviation, and skewness. The production-based model also understates the mean and standard deviation, but by less. Its skewness is slightly above the data (0.77 versus 0.72), whereas the HM model's skewness is below the data (0.62). Kurtosis is slightly higher in the production-based model.

Beyond the numerical fit, the two approaches differ in economic content. In the HM model, implied skewness is governed by the level of spot variance and the return–variance correlation ϱ : lower spot variance and more negative ϱ both generate more negative skewness. These relationships parallel those in the production-based model, where higher leverage induces a more negative correlation between equity returns and volatility shocks, as in the data.

The key limitation of the HM model emerges in the cross-section of IV shapes (Bakshi, Cao, and Chen, 1997). In the HM framework, a smile arises only when spot variance is

low—but low spot variance also implies more negatively skewed returns. Conversely, when spot variance is high the IV curve takes a left-smirk shape but becomes flatter. In the data, and in the production-based model, smiles are associated with more positively skewed distributions.

Similarly, there is a direct relationship between the volatility of the spot variance (σ_V) and the steepness of the IV curve. A steeper curve requires a higher σ_V . However, a steeper IV curve implies a lower percentage of smiles and a higher percentage of smirks.

The HM model therefore cannot independently vary the shape and the slope of the IV surface (Das and Sundaram, 1999). The production-based model does not face this tension: because IV shapes are driven by the interaction of leverage and real option effects, the model can generate smiles, smirks, and frowns across a wide range of economic conditions. Smiles are associated with positively skewed distributions when the firm has some leverage, but not enough to overcome the real option contribution. If the firm had very little leverage (which does not happen very frequently in the simulated economy), the IV shape could even be monotonically increasing. With the same real option but high leverage, the IV shape would be monotonically decreasing.

C. Trade-offs

The two approaches face complementary limitations. On the production side, our model—like most investment-based option pricing models—specifies the stochastic discount factor exogenously rather than deriving it from household optimization in general equilibrium (Jermann, 1998; Boldrin, Christiano, and Fisher, 2001). This preserves tractability but means the model is silent on whether risk premia are consistent with preferences and market clearing. In

addition, the predictions are sensitive to the functional form of capital adjustment costs: different specifications can alter the magnitude, and occasionally the sign, of key comparative statics (Cooper and Priestley, 2011). Finally, solving the firm's dynamic program, simulating the economy, and pricing options on the resulting equity is computationally demanding relative to the (semi-)closed-form pricing formulas available for stochastic volatility models.

On the reduced-form side, models such as the HM specification provide no economic mechanism linking option prices to observable firm characteristics (Bates, 2003)—they can fit a given surface but cannot explain why firms with different leverage or growth opportunities exhibit systematically different implied volatility shapes. As shown above, the shape and slope of the IV surface are mechanically linked, which limits the range of surface shapes the model can generate. More broadly, because reduced-form models are calibrated security by security, they offer no framework for studying cross-sectional variation in option prices.

We view the two approaches as complements: reduced-form models excel at fitting individual surfaces with minimal computational cost, while the production-based approach provides economic mechanisms and cross-sectional predictions that reduced-form models cannot deliver.

VII. Related literature

This paper belongs to the literature that seeks to explain the cross-section of equity option prices. Building on Merton's (1974) analysis of leverage, earlier contributions incorporate option pricing into structural models of the firm, including compound-option approaches for levered equity (Geske (1979)), extensions with taxes and bankruptcy costs that endogenize leverage

(Toft and Prucyk (1997); building on Leland (1994)), and evidence that the leverage effect sharply reduces option pricing errors relative to Black and Scholes (1973) (Geske et al. (2016)).

Bai, Goldstein, and Yang (2019) show that leverage is crucial for the spread between index and single-stock bank options, while Morellec and Zhdanov (2019) link implied skewness to product-market competition.

Following Hennessy and Whited (2005, 2007), we propose a fully dynamic structural model in which shareholders endogenously choose production capacity and risky debt. We demonstrate that these two ingredients are essential for reproducing the large observed heterogeneity in equity option prices.

Our work also relates to applications of the leverage effect in volatility modeling (Engle and Siriwardane (2017)) and to studies emphasizing the different roles of systematic and idiosyncratic risk (Duan and Wei (2009)). In our model, cross-sectional variation in option prices arises from firm-specific shocks interacting with endogenous leverage and investment, while aggregate risk drives their variation over time. In this respect, the model shares intuition with the extensive structural credit-risk literature (Leland (1994); Gomes and Schmid (2021); Gamba and Saretto (2020)).

The paper further connects to the real-options literature examining how investment irreversibility and growth options affect risk and return (Berk, Green, and Naik (1999); Kogan (2001, 2004); Zhang (2005); Cooper (2006); Ai and Kiku (2013); Kogan and Papanikolaou (2013, 2014); Gu et al. (2017); Aretz and Pope (2018); Trigeorgis and Lambertides (2014); Del Viva et al. (2017); Cao, Simin, and Zhao (2008); Lyandres and Zhdanov (2020); Bali, Del Viva, Lambertides, and Trigeorgis (2020); Aguerrevere (2009); Morellec and Zhdanov

(2019)). In contrast to these studies, we emphasize the interaction of real options and financial leverage in shaping higher moments of equity return distributions.

Finally, we contribute to the literature that extracts risk-neutral moments from option prices using model-free methods (Bakshi et al. (2003)) and relates these moments to firm characteristics (Dennis and Mayhew (2002); Hansis et al. (2010)). We provide a structural link between risk-neutral moments (volatility, skewness, and kurtosis) and endogenous corporate policies.

Relative to reduced-form option pricing models and earlier structural attempts (Toft and Prucyk (1997); Geske et al. (2016); Morellec and Zhdanov (2019)), our calibrated dynamic model delivers option prices close to those observed for the average firm and, crucially, generates much of the cross-sectional dispersion solely through endogenous investment/disinvestment and capital-structure decisions.

VIII. Conclusions

Traditional option pricing models often require very strong assumptions about investor preferences and the dynamics of equity prices. We show that equity options can be priced in a production economy without assuming a process for equity prices or volatility. In our setup, the relationship between risk and value arises endogenously through a dynamic sequence of optimal investment and financing decisions made by the manager/shareholders to maximize the value of equity. Starting from the functional forms of corporate trade-offs, we derive option prices that match many properties of those observed in the cross-section of U.S. equities.

Our approach should not be seen as an attempt to offer a better option pricing model, but

rather as an attempt to provide a theoretical relationship between firm fundamentals and the price of derivatives written on the firm's equity. We believe this link is important because it connects the primitives of widely used finance models (i.e., those that price financial derivatives) to a large body of well-understood economic mechanisms that describe how firms' decisions concerning their leverage and investment affect the probability distribution of equity returns.

Ultimately, we hope to explain why option prices carry information about corporate policies and future stock prices. In standard models, options are redundant securities and should carry no such information unless markets are segmented. In our model, options remain redundant, but their prices reveal the shape of the distribution of future equity values, and that shape is determined by the firm's investment and financing decisions.

References

- Abel, A. B., and J. C. Eberly. “A unified model of investment under uncertainty.” *American Economic Review*, 85 (1994), 1369–1384.
- Aguerrevere, F. “Real options, product market competition, and asset returns.” *Journal of Finance*, 44 (2009), 957–983.
- Ai, H., and D. Kiku. “Growth to value: Option exercise and the cross section of equity returns.” *Journal of Financial Economics*, 107 (2013), 325–349.
- Aretz, K., and P. Pope. “Real options models of the firm, capacity overhang, and the cross section of stock returns.” *Journal of Finance*, 73 (2018), 1363–1415.
- Bai, J.; R. Goldstein; and F. Yang. “The leverage effect and the basket-index spread.” *Journal of Financial Economics*, 131 (2019), 180–205.
- Bakshi, G.; C. Cao; and Z. Chen. “Empirical Performance of Alternative Option Pricing Models.” *Journal of Finance*, 52 (1997), 2003–2049.
- Bakshi, G.; N. Kapadia; and D. Madan. “Stock return characteristics, skew laws, and the differential pricing of individual equity options.” *Review of Financial Studies*, 16 (2003), 637–659.
- Bali, T. G.; L. Del Viva; N. Lambertides; and L. Trigeorgis. “Growth options and related stock market anomalies: profitability, distress, lotteryiness, and volatility.” *Journal of Financial Quantitative Analysis*, 55 (2020), 2150–2180.

- Bates, D. S. “Empirical Option Pricing: A Retrospection.” *Journal of Econometrics*, 116 (2003), 387–404.
- Berk, J. B.; R. C. Green; and V. Naik. “Optimal investment, growth options, and security returns.” *Journal of Finance*, 54 (1999), 1553–1607.
- Black, F., and M. Scholes. “The pricing of options and corporate liabilities.” *Journal of Political Economy*, 81 (1973), 637–659.
- Boldrin, M.; L. J. Christiano; and J. D. M. Fisher. “Habit Persistence, Asset Returns, and the Business Cycle.” *American Economic Review*, 91 (2001), 149–166.
- Boyer, B.; T. Mitton; and K. Vorkink. “Expected Idiosyncratic Skewness.” *Review of Financial Studies*, 23 (2010), 169–202.
- Cao, C.; T. Simin; and J. Zhao. “Can growth options explain the trend in idiosyncratic risk?” *Review of Financial Studies*, 21 (2008), 2599–2633.
- Carlson, M.; A. Fisher; and R. Giammarino. “Corporate investment and asset price dynamics: Implications for the cross-section of returns.” *Journal of Finance*, 59 (2004), 2577–2603.
- Chaderina, M.; P. Weiss; and J. Zechner. “The maturity premium.” *Journal of Financial Economics*, 144 (2022), 670–694.
- Cont, R., and P. Tankov. “Non-Parametric Calibration of Jump-Diffusion Option Pricing Models.” *Journal of Computational Finance*, 7 (2004), 1–49.
- Cooley, T. F., and E. C. Prescott. In *Frontiers of Business Cycle Research*. “Economic growth and business cycles.”, T. F. Cooley, ed. Princeton University Press, Princeton - NJ (1995), 1–38.

- Cooper, I. “Asset pricing implications of non-convex adjustment costs and irreversibility of investment.” *The Journal of Finance*, 61 (2006), 139–170.
- Cooper, I., and R. Priestley. “Real Investment and Risk Dynamics.” *Journal of Financial Economics*, 101 (2011), 182–205.
- Das, S. R., and R. K. Sundaram. “Of Smiles and Smirks: A Term Structure Perspective.” *Journal of Financial and Quantitative Analysis*, 34 (1999), 211–239.
- Del Viva, L.; E. Kananen; and L. Trigeorgis. “Real options, idiosyncratic skewness, and diversification.” *Journal of Financial and Quantitative Analysis*, 52 (2017), 215–241.
- DeMarzo, P., and Z. He. “Leverage Dynamics without Commitment.” *Journal of Finance*, 76 (2021), 1195–1250.
- Dennis, P., and S. Mayhew. “Risk-neutral skewness: Evidence from stock options.” *Journal of Financial and Quantitative Analysis*, 37 (2002), 471–493.
- Dew-Becker, I.; S. Giglio; A. Le; and M. Rodriguez. “The Price of Variance Risk.” *Journal of Financial Economics*, 123 (2017), 225–250. ISSN 0304-405X.
- Duan, J.-C., and J. Wei. “Systematic risk and the price structure of individual equity options.” *Review of Financial Studies*, 22 (2009), 1981–2006.
- Engle, R. F., and E. N. Siriwardane. “Structural GARCH: The volatility-leverage connection.” *The Review of Financial Studies*, 31 (2017), 449–492.
- Fama, E. F., and K. R. French. “The cross-section of expected stock returns.” *Journal of Finance*, 47 (1992), 427–465.

- Friewald, N.; F. Nagler; and C. Wagner. “Debt Refinancing and Equity Returns.” *Journal of Finance*, 77 (2022), 2287–2329.
- Gamba, A., and A. Saretto. “Growth Options and Credit Risk.” *Management Science*, 66 (2020), 4269–4291.
- Garleanu, N.; S. Panageas; and J. Yu. “Technological growth and asset pricing.” *Journal of Finance*, 67 (2012), 1265–1292.
- Geske, R. “The valuation of compound options.” *Journal of Financial Economics*, 7 (1979), 63–81.
- Geske, R.; A. Subrahmanyam; and Y. Zhou. “Capital structure effects on the prices of equity call options.” *Journal of Financial Economics*, 121 (2016), 231–253.
- Glover, B. “The expected cost of default.” *Journal of Financial Economics*, 119 (2016), 284–299.
- Gomes, J. F. “Financing investment.” *American Economic Review*, 91 (2001), 1263–1285.
- Gomes, J. F., and L. Schmid. “Equilibrium asset pricing with leverage and default.” *Journal of Finance*, 76 (2021), 977–1018.
- Graham, J. R. “Debt and the marginal tax rate.” *Journal of Financial Economics*, 41 (1996a), 41–73.
- Graham, J. R. “Proxies for the corporate marginal tax rate.” *Journal of Financial Economics*, 42 (1996b), 187–221.
- Gu, L.; D. Hackbarth; and T. Johnson. “Inflexibility and stock returns.” *Review of Financial Studies*, 31 (2017), 278–321.

- Hackbarth, D., and T. Johnson. “Real options and risk dynamics.” *Review of Economic Studies*, 82 (2015), 1449–1482.
- Hansis, A.; C. Schlag; and G. Vilkov. “The dynamics of risk-neutral implied moments: evidence from individual options.” Frankfurt School of Finance (2010).
- Harvey, C., and A. Siddique. “Conditional Skewness in Asset Pricing Tests.” *Journal of Finance*, 55 (2000), 1263–1295.
- Hennessy, C. A., and T. Whited. “Debt dynamics.” *Journal of Finance*, 60 (2005), 1129–1165.
- Hennessy, C. A., and T. M. Whited. “How costly is external financing? Evidence from a structural estimation.” *Journal of Finance*, 62 (2007), 1705–1745.
- Heston, S. L. “A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options.” *Review of Financial Studies*, 6 (1993), 327–343.
- Hjalmarsson, E. “Predicting Global Stock Returns.” *Journal of Financial and Quantitative Analysis*, 45 (2010), 49–80.
- Jermann, U. J. “Asset Pricing in Production Economies.” *Journal of Monetary Economics*, 41 (1998), 257–275.
- Jones, C. S., and S. Tuzel. “Inventory investment and the cost of capital.” *Journal of Financial Economics*, 107 (2013), 557–579.
- Kogan, L. “An equilibrium model of irreversible investment.” *Journal of Financial Economics*, 62 (2001), 201–245.

- Kogan, L. “Asset prices and real investment.” *Journal of Financial Economics*, 73 (2004), 411–431.
- Kogan, L., and D. Papanikolaou. “Firm characteristics and stock returns: the role of investment-specific shocks.” *Review of Financial Studies*, 26 (2013), 2718–2759.
- Kogan, L., and D. Papanikolaou. “Growth opportunities, technology shocks, and asset prices.” *Journal of Finance*, 69 (2014), 675–718.
- Kozhan, R.; A. Neuberger; and P. Schneider. “Skew Risk Premia in the Equity Index Market.” *Review of Financial Studies*, 26 (2013), 2174–2203.
- Kuehn, L.-A., and L. Schmid. “Investment-based corporate bond pricing.” *Journal of Finance*, 69 (2014), 2741–2776.
- Leland, H. E. “Corporate debt value, bond covenants, and optimal capital structure.” *Journal of Finance*, 49 (1994), 1213–1252.
- Lyandres, E., Evgeny Matveyev, and A. Zhdanov. “Does the market correctly value investment options?” *Review of Finance*, 62 (2020), 1159–1201.
- Merton, R. “On the pricing of corporate debt: the risk structure of interest rates.” *Journal of Finance*, 4 (1974), 449–470.
- Merton, R. C. “Option Pricing When Underlying Stock Returns Are Discontinuous.” *Journal of Financial Economics*, 3 (1976), 125–144.
- Miao, B. Z. “Labor-Technology Substitution: Implications for Asset Pricing.” *Journal of Finance*, 74 (2019), 3347–3399.

- Mitton, T., and K. Vorkink. “Equilibrium Under-diversification and the Preference for Skewness.” *Review of Financial Studies*, 20 (2007), 1255–1288.
- Morellec, E., and A. Zhdanov. “Product market competition and option prices.” *Review of Financial Studies*, 32 (2019), 4343–4386.
- Neuberger, A. “Realized skewness.” *Review of Financial Studies*, 25 (2012), 3423–3455.
- Rhodes-Kropf, M.; D. T. Robinson; and S. Viswanathan. “Valuation waves and merger activity: The empirical evidence.” *Journal of Financial Economics*, 77 (2005), 561–603.
- Rouwenhorst, K. G. In *Frontiers of Business Cycle Research*. “Asset pricing implications of equilibrium business cycle models.”, T. F. Cooley, ed. Princeton University Press, Princeton - NJ (1995), chap. 10, 294–330.
- Toft, K. B., and B. Prucyk. “Options on leverage equity: theory and empirical tests.” *Journal of Finance*, 52 (1997), 1151–1180.
- Trigeorgis, L., and N. Lambertides. “The role of growth options in explaining stock returns.” *Journal of Financial and Quantitative Analysis*, 49 (2014), 749–771.
- Tuzel, S. “Corporate real estate holdings and the cross-section of stock returns.” *Review of Financial Studies*, 23 (2010), 2268–2302.
- van Zwet, W. “Convex transformations: A new approach to skewness and kurtosis.” *Statistica Neerlandica*, 18 (1964), 433–441.
- Warusawitharana, M., and T. Whited. “Equity market misvaluation, financing, and investment.” *Review of Financial Studies*, 29 (2016), 603–634.

Zhang, L. “The value premium.” *Journal of Finance*, 60 (2005), 67–103.

Internet Appendix

A. Equity option prices

FIGURE A1

Implied volatility surface – time series

The figure plots the time series of cross-sectional averages, as well as the 5th and 95th percentiles, of implied volatility surfaces extracted from the data. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.




FIGURE A2

Implied moments – time series

The figure plots the time series of cross-sectional averages, as well as the 5th and 95th percentiles, of implied skewness and kurtosis extracted from 182-day and 365-day options. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.




TABLE A1

Option prices and firm characteristics

The table reports alternative econometric specifications to those reported in Table I. Panel A tabulates results from a traditional firm (and time) fixed-effects regression. Panel B reports regression results based on a Hjalmarsson (2010) forward demeaning regression. Variable specification and data are the same as those described in the caption of Table I. All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications.

	Panel A: Time and firm fixed effects					
	Implied skewness			IV Slope		
	(1)	(2)	(3)	(4)	(5)	(6)
IV	12.63 (14.92)	12.34 (14.31)	12.96 (14.97)	4.07 (12.74)	3.97 (12.26)	4.23 (13.28)
Size	-10.21 (-6.54)	-8.62 (-5.52)	-8.05 (-5.15)	-2.86 (-7.82)	-2.66 (-6.97)	-2.42 (-6.47)
Profitability	1.05 (3.95)	1.33 (4.54)	1.05 (3.73)	0.29 (3.77)	0.36 (4.20)	0.27 (3.34)
Leverage	-1.88 (-3.58)	-2.35 (-4.39)	-1.72 (-3.14)	-0.36 (-2.77)	-0.54 (-4.15)	-0.33 (-2.47)
M2B	-6.14 (-11.29)			-1.28 (-11.20)		
Long M2B			-2.24 (-6.40)			-0.67 (-7.48)
GO		-3.00 (-8.52)			-0.91 (-8.82)	
Adjusted- R^2	0.51	0.50	0.50	0.55	0.55	0.55
FE R^2	0.48	0.48	0.48	0.52	0.52	0.52

Panel B: Hjalmarsson (2010) forward demeaning						
	Implied skewness			IV Slope		
	(1)	(2)	(3)	(4)	(5)	(6)
IV	8.53 (8.56)	8.47 (8.87)	8.58 (8.92)	2.91 (10.47)	2.86 (10.46)	2.91 (10.58)
Size	-9.51 (-15.77)	-9.10 (-14.64)	-8.84 (-14.71)	-2.16 (-11.86)	-2.21 (-12.00)	-2.09 (-11.65)
Profitability	0.45 (1.57)	0.62 (2.40)	-0.51 (-1.49)	0.27 (3.42)	0.42 (5.70)	0.11 (1.23)
Leverage	-2.02 (-4.03)	-2.43 (-5.00)	-2.00 (-3.99)	-0.27 (-2.48)	-0.43 (-3.99)	-0.27 (-2.46)
M2B	-2.60 (-5.50)			-0.30 (-2.83)		
Long M2B		-1.84 (-5.97)			-0.69 (-8.59)	
GO			-1.12 (-2.43)			-0.22 (-2.02)
Adjusted- R^2	0.13	0.13	0.13	0.14	0.14	0.14

TABLE A2

Option prices and real options: 182- and 365-day maturities

The table shows results of Fama-MacBeth regressions of implied skewness and implied volatility slope on measures of real options. Implied skewness and implied volatility slope, which is computed as the log difference of IV(1.2) and IV(0.8), are based on 182-day and 365-day maturity options. We control for size, book leverage, profitability, the level of ATM IV, and the market-to-book ratio. All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications. We report parameter estimates (RN-Skew coefficients are multiplied by 100) and *t*-statistics, corrected for heteroskedasticity, in parentheses. Regressions include industry fixed effects. A constant is included but not reported. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.

	182 day maturity		365 day maturity	
	RN-Skew	Slope	RN-Skew	Slope
IV	12.46 (12.88)	4.02 (10.61)	0.22 (10.37)	2.98 (12.23)
Size	-13.02 (-7.21)	-2.72 (-10.96)	-0.12 (-6.77)	-0.99 (-6.22)
Leverage	-0.47 (-2.30)	-0.37 (-4.80)	-0.02 (-4.11)	-0.33 (-4.87)
Profitability	-0.29 (-1.65)	0.23 (2.76)	0.00 (0.48)	0.27 (3.57)
M2B	-5.30 (-9.42)	-1.08 (-10.92)	-0.04 (-7.01)	-0.42 (-6.55)

TABLE A3

Option prices and real options: 91-, 182-, and 365-day maturities

The table shows results of Fama-MacBeth regressions of implied skewness and implied volatility slope on measures of real options. Implied skewness and implied volatility slope are based on 91-, 182-, and 365-day maturity options. Besides the controls described in Table I, we include turnover, cumulative returns from $t - 12$ to $t - 1$ (momentum), and a Nasdaq indicator variable as in Boyer, Mitton, and Vorkink (2010), as well as a measure of R&D intensity and investment as in Del Viva et al. (2017). All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications. We report parameter estimates and t -statistics in parentheses based on heteroskedasticity-robust standard errors. Regressions include industry fixed effects. A constant is included but not reported. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.

	91 days		182 days		365 days	
	RN-Skew	Slope	RN-Skew	Slope	RN-Skew	Slope
IV	10.49 (12.01)	5.61 (12.89)	0.12 (13.63)	3.62 (10.74)	0.20 (9.69)	2.98 (13.76)
Size	-14.60 (-6.86)	-1.23 (-5.39)	-0.15 (-7.48)	-3.06 (-11.16)	-0.13 (-7.11)	-1.04 (-5.86)
Leverage	-0.47 (-1.23)	-0.12 (-2.43)	0.00 (0.49)	-0.20 (-2.31)	-0.02 (-4.27)	-0.32 (-3.89)
Profitability	-0.06 (-0.26)	0.20 (2.75)	-0.00 (-1.70)	0.01 (0.18)	0.00 (0.62)	0.16 (1.49)
M2B	-6.15 (-11.00)	-0.69 (-7.03)	-0.05 (-9.36)	-1.08 (-9.74)	-0.03 (-6.00)	-0.36 (-4.34)
Turnover	-0.58 (-1.65)	-0.17 (-1.34)	-0.00 (-0.28)	0.29 (2.39)	0.02 (3.32)	0.09 (0.71)
Momentum	-1.87 (-6.39)	-0.41 (-4.42)	-0.00 (-0.99)	0.10 (1.36)	-0.01 (-3.65)	-0.03 (-0.53)
Nasdaq	0.08 (0.21)	-0.01 (-0.15)	0.00 (0.05)	0.27 (2.27)	0.00 (0.36)	0.08 (0.73)
R&D	-0.28 (-1.27)	-0.20 (-2.85)	0.00 (2.15)	-0.17 (-2.57)	0.01 (1.29)	-0.12 (-1.50)
Investment	0.20 (0.94)	-0.04 (-0.72)	-0.00 (-0.91)	-0.10 (-2.06)	0.00 (0.40)	-0.03 (-0.50)

B. Future realized skewness

A. Real options and future realized skewness

While our theoretical analysis is mostly centered on equity option pricing, the relation between the future distribution of equity prices and real options can also be studied under the physical measure. Thus, in this section we show how variability in future realized skewness relates to measures of real options. Table B1 reports results of Fama-MacBeth regressions of future realized skewness over horizons of 91, 182, and 365 days. We present realized skewness at increasing future time horizons to mitigate the concern that measuring realized skewness from daily log returns is quite challenging (see, for example, Neuberger, 2012; Kozhan et al., 2013). All specifications include industry fixed effects. We present results for realized idiosyncratic skewness (see Mitton and Vorkink, 2007, for the different implications of total and idiosyncratic measures) from the linear and quadratic market models (as in Harvey and Siddique, 2000) in Table B2.

The regression coefficients reported in Table B1 are in line with those reported in Table I along some dimensions: they show a positive relation to the level of implied volatility (i.e., total risk) at the short horizon, a negative relation to real-option values, and an uncertain relation to profitability. Unlike implied skewness, future realized skewness is positively and statistically significantly related to leverage and shows a flat relationship with size. Including other controls that have been used to study similar measures of future realized skewness (as in Boyer et al., 2010, and Del Viva et al., 2017) leaves the coefficients on the real option measures negative and statistically significant. Replacing the dependent variable with future realized idiosyncratic skewness (see Table B2) leaves the results unchanged.

The relation between real option value and future realized skewness is negative and statistically significant at every horizon (i.e., for each future period over which realized skewness is measured). Its magnitude, however, moves in opposite directions for the two measures: the GO coefficient declines with the horizon, whereas the M2B coefficient increases.

B. Revisiting previous findings

Del Viva et al. (2017) report a *positive* correlation between future idiosyncratic skewness and the GO index. We repeat their exercise in the same sample period (i.e., between 1983 and 2011), but using our quarterly sampled data. Results are presented in Table B3.

Column 1 of Table B3, which uses the same controls as column 1 of Table 2 of Del Viva et al. (2017), shows a positive relation between future idiosyncratic skewness and GO. The coefficient is smaller than the one reported by Del Viva et al. (2017) and less statistically significant. The difference in the economic magnitude could be due to a number of choices in constructing the sample and the fact that our data are sampled at quarterly frequency. Adding controls turns the GO coefficient first insignificant and then negative and significant. Adding an interaction term between GO and an indicator variable that is set to one when equity options are available for a firm-quarter observation suggests that the relation between GO and skewness might be more negative for optionable stocks. Since the sample includes a period before 1996, when OptionMetrics first became available, this indicates that the more negative correlation is due either to stocks that are optionable after 1996 or to a difference in the entire sample before and after 1996.

Since the coefficient reported in Table B3 is about half as large as that reported in the

equivalent regression specification in Appendix Table B2, which only includes optionable stocks after 1996, we conjecture that the difference is mostly due to the sample period.

In summary, the evidence presented in Table B3 suggests that the positive relation between GO and future idiosyncratic skewness reported by Del Viva et al. (2017) is due to a multitude of factors, and cannot be found for optionable stocks in the most recent period (after 1996).

TABLE B1

Future realized skewness and real options

The table shows Fama-MacBeth regression results of future realized skewness on measures of real option value. Independent variables are measured at day t , while future realized skewness is extracted from log returns of days $t + 1$ to $t + T$, where $T = 91, 182$, or 365 . Real option measures include the market-to-book ratio and the GO ratio of Trigeorgis and Lambertides (2014). We control for size, book leverage, profitability, and the level of ATM IV as in Table I, as well as turnover, cumulative returns from $t - 12$ to $t - 1$ (momentum), a Nasdaq indicator variable as in Boyer et al. (2010), and a measure of R&D intensity and investment as in Del Viva et al. (2017). A constant is included but not reported. All regressions include industry fixed effects. All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications. We report parameter estimates and t -statistics in parentheses based on heteroskedasticity-robust standard errors. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. A total of 3,536 firms are included.

	91 days				182 days				365 days			
IV	-0.15 (-0.13)	0.11 (0.10)	2.17 (1.77)	2.68 (2.28)	-3.57 (-2.23)	-3.52 (-2.20)	-1.74 (-1.01)	-1.41 (-0.83)	-6.26 (-2.84)	-6.44 (-2.95)	-4.11 (-1.87)	-3.85 (-1.78)
Size	-1.95 (-2.11)	-1.56 (-1.64)	-0.85 (-0.86)	-0.58 (-0.57)	-2.25 (-1.72)	-1.44 (-1.08)	-1.38 (-0.91)	-0.76 (-0.50)	-2.36 (-1.39)	-0.78 (-0.47)	-1.27 (-0.69)	-0.02 (-0.01)
Leverage	1.27 (2.54)	0.76 (1.63)	1.08 (2.13)	0.46 (0.94)	2.02 (2.96)	1.68 (2.54)	1.82 (2.71)	1.45 (2.23)	2.38 (2.57)	2.17 (2.29)	1.96 (2.20)	1.82 (1.93)
Profitability	2.14 (2.50)	-2.90 (-3.13)	3.49 (3.70)	-1.92 (-1.91)	0.04 (0.04)	-3.81 (-3.12)	0.19 (0.15)	-3.71 (-2.75)	0.11 (0.08)	-4.22 (-2.33)	-0.29 (-0.21)	-3.77 (-1.88)
M2B	-4.29 (-8.12)		-3.24 (-5.13)		-4.91 (-5.50)		-3.83 (-4.51)		-7.56 (-5.05)		-6.11 (-5.04)	
GO		-8.17 (-7.04)		-9.16 (-7.46)		-5.99 (-5.76)		-6.31 (-5.62)		-5.86 (-3.65)		-4.53 (-2.83)
Turnover			-3.44 (-6.66)	-3.65 (-7.59)			-3.37 (-4.10)	-3.54 (-4.46)			-3.28 (-3.26)	-3.40 (-3.36)
Momentum			-2.56 (-3.93)	-2.89 (-4.37)			-1.99 (-2.37)	-2.60 (-2.96)			-2.97 (-3.02)	-4.19 (-3.68)
Nasdaq			0.89 (1.22)	1.15 (1.62)			0.84 (0.86)	1.01 (1.03)			1.13 (0.87)	1.16 (0.87)
R&D			0.66 (1.06)	0.92 (1.31)			-1.13 (-1.04)	-1.26 (-1.08)			-2.10 (-1.24)	-2.80 (-1.61)
Investment			-1.23 (-2.14)	-1.18 (-2.00)			-0.44 (-0.64)	-0.45 (-0.65)			-0.59 (-1.03)	-0.78 (-1.35)

TABLE B2

Future realized idiosyncratic skewness and real options

The table shows Fama-MacBeth regression results of future realized idiosyncratic skewness on measures of real option value. Independent variables are measured at day t , while future idiosyncratic skewness is extracted from the residuals of regressions of log returns on linear and quadratic functions of the corresponding log market return (as in Harvey and Siddique, 2000). All control variables are as in Table B1. A constant is included but not reported. All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications. We report parameter estimates and t -statistics in parentheses based on heteroskedasticity-robust standard errors. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data are sampled at quarterly frequency. A total of 3,536 firms are included.

	Idiosyncratic skewness from linear market model											
	91 days				182 days				365 days			
IV	-0.34 (-0.27)	-0.09 (-0.07)	2.62 (1.90)	3.19 (2.41)	-4.29 (-2.37)	-4.26 (-2.36)	-1.89 (-0.97)	-1.50 (-0.77)	-7.45 (-2.95)	-7.75 (-3.12)	-4.84 (-1.90)	-4.53 (-1.80)
Size	-1.82 (-1.99)	-1.37 (-1.47)	-0.81 (-0.81)	-0.50 (-0.48)	-2.16 (-1.62)	-1.21 (-0.89)	-1.51 (-0.97)	-0.77 (-0.49)	-3.22 (-1.67)	-1.21 (-0.64)	-2.35 (-1.11)	-0.76 (-0.38)
Leverage	1.53 (2.95)	0.86 (1.75)	1.38 (2.76)	0.57 (1.17)	2.54 (3.62)	2.09 (3.08)	2.42 (3.56)	1.94 (2.95)	3.34 (3.31)	3.05 (3.01)	3.02 (3.19)	2.81 (2.86)
Profitability	2.84 (2.90)	-3.56 (-3.73)	4.77 (4.27)	-2.04 (-1.85)	-0.32 (-0.25)	-5.15 (-3.92)	0.06 (0.04)	-4.77 (-3.15)	-1.05 (-0.63)	-6.38 (-3.15)	-1.43 (-0.87)	-5.76 (-2.52)
M2B	-5.41 (-8.18)		-4.23 (-5.39)		-6.05 (-5.97)		-4.80 (-4.93)		-9.59 (-5.33)		-7.81 (-5.17)	
GO		-10.33 (-7.22)		-11.51 (-7.56)		-7.47 (-5.77)		-7.83 (-5.74)		-7.16 (-3.65)		-5.62 (-2.92)
Turnover			-4.35 (-6.64)	-4.64 (-7.33)			-4.36 (-4.46)	-4.61 (-4.81)			-3.88 (-3.04)	-4.09 (-3.20)
Momentum			-3.15 (-4.62)	-3.60 (-5.17)			-2.39 (-2.43)	-3.19 (-3.12)			-3.50 (-3.04)	-5.14 (-3.94)
Nasdaq			0.38 (0.48)	0.68 (0.88)			0.23 (0.21)	0.41 (0.38)			0.85 (0.55)	0.89 (0.56)
R&D			1.24 (1.93)	1.59 (2.24)			-1.03 (-0.87)	-1.09 (-0.86)			-2.49 (-1.34)	-3.37 (-1.73)
Investment			-1.40 (-2.07)	-1.32 (-1.93)			-0.40 (-0.50)	-0.39 (-0.49)			-0.59 (-0.97)	-0.79 (-1.31)

	Idiosyncratic skewness from quadratic market model											
	91 days				182 days				365 days			
IV	-0.30 (-0.24)	-0.06 (-0.05)	2.50 (1.85)	3.05 (2.33)	-4.15 (-2.30)	-4.12 (-2.30)	-1.84 (-0.94)	-1.46 (-0.74)	-7.42 (-2.95)	-7.74 (-3.13)	-4.85 (-1.90)	-4.55 (-1.80)
Size	-1.83 (-2.06)	-1.39 (-1.52)	-0.86 (-0.88)	-0.55 (-0.55)	-2.27 (-1.72)	-1.31 (-0.96)	-1.66 (-1.08)	-0.90 (-0.58)	-3.45 (-1.81)	-1.41 (-0.76)	-2.68 (-1.29)	-1.06 (-0.54)
Leverage	1.47 (2.93)	0.83 (1.74)	1.31 (2.72)	0.54 (1.13)	2.64 (3.79)	2.19 (3.27)	2.49 (3.68)	2.02 (3.09)	3.52 (3.50)	3.23 (3.22)	3.20 (3.37)	2.99 (3.07)
Profitability	2.64 (2.73)	-3.58 (-3.77)	4.49 (3.99)	-2.13 (-1.93)	-0.48 (-0.37)	-5.29 (-4.00)	-0.14 (-0.10)	-4.92 (-3.22)	-1.19 (-0.71)	-6.57 (-3.22)	-1.60 (-0.97)	-5.95 (-2.58)
M2B	-5.34 (-8.45)		-4.17 (-5.49)		-6.10 (-6.02)		-4.87 (-4.98)		-9.75 (-5.43)		-7.97 (-5.26)	
GO		-10.00 (-7.06)		-11.15 (-7.31)		-7.42 (-5.57)		-7.72 (-5.53)		-7.21 (-3.57)		-5.60 (-2.85)
Turnover			-4.30 (-6.66)	-4.58 (-7.32)			-4.30 (-4.48)	-4.56 (-4.83)			-3.75 (-2.96)	-3.97 (-3.14)
Momentum			-3.03 (-4.32)	-3.47 (-4.86)			-2.32 (-2.29)	-3.14 (-2.99)			-3.44 (-2.93)	-5.12 (-3.92)
Nasdaq			0.37 (0.48)	0.65 (0.86)			0.15 (0.14)	0.32 (0.30)			0.70 (0.46)	0.73 (0.47)
R&D			1.18 (1.88)	1.52 (2.18)			-1.09 (-0.93)	-1.17 (-0.93)			-2.56 (-1.38)	-3.46 (-1.78)
Investment			-1.20 (-1.79)	-1.12 (-1.68)			-0.29 (-0.36)	-0.28 (-0.35)			-0.50 (-0.82)	-0.69 (-1.16)

TABLE B3

Future skewness and real options in sample period of Del Viva et al. (2017)

The table shows Fama-MacBeth regression results of future realized idiosyncratic skewness on the GO ratio of Trigeorgis and Lambertides (2014). Independent variables are measured at day t , while future idiosyncratic skewness is extracted from the residuals of regressions of log returns on the corresponding log market return of days $t + 1$ to $t + 365$. We control for size, book leverage, and profitability as in Table I, as well as a lag of the dependent variable, past realized volatility from the previous 90 days, quarterly average turnover, cumulative returns from $t - 12$ to $t - 1$ (momentum), a Nasdaq indicator variable as in Boyer et al. (2010), and a measure of R&D intensity and investment as in Del Viva et al. (2017). In the last specification, we also include an indicator variable that is set to one when equity options are available for a firm-quarter observation. A constant is included but not reported. All right-hand-side variables are winsorized at the 1st and 99th percentiles and subsequently standardized to aid comparability across specifications. We report parameter estimates and t -statistics in parentheses based on heteroskedasticity-robust standard errors. The sample contains all non-financial firms between 1983 and 2011. Data are sampled at quarterly frequency. A total of 5,616 firms are included.

	(1)	(2)	(3)	(4)	(5)
Past Skew		1.66 (2.80)	1.30 (2.11)	0.57 (0.87)	0.90 (1.08)
Vol		7.12 (4.12)	8.78 (5.28)	3.77 (2.58)	3.01 (1.61)
Size				-11.29 (-5.85)	-10.78 (-4.56)
Leverage				1.71 (1.63)	0.33 (0.29)
Profitability				-2.35 (-0.92)	-0.23 (-0.07)
GO	2.46 (1.67)	-0.83 (-0.81)	-1.35 (-1.48)	-3.37 (-2.91)	-1.12 (-0.53)
Optionable					5.88 (0.84)
Optionable $\times$ GO					-9.91 (-1.45)
Turnover			-11.21 (-6.74)	-8.47 (-5.31)	-6.87 (-3.81)
Momentum			-7.90 (-9.68)	-8.18 (-9.77)	-8.11 (-7.32)
Nasdaq			-0.58 (-0.28)	-4.33 (-2.57)	-2.66 (-1.47)
R&D	0.32 (0.24)	1.63 (1.34)	1.95 (1.73)	1.07 (0.98)	0.85 (0.64)
Investment	-2.77 (-4.16)	-2.80 (-3.65)	-1.89 (-2.85)	-1.84 (-3.06)	-1.29 (-2.10)

C. Full model description

This appendix provides the complete specification of the quantitative model summarized in Section V, including the economic environment, the firm's problem, and the pricing of multi-period options.

A. Economic setup

In the economy, each firm is exposed to an aggregate shock y and an idiosyncratic shock z . These follow the processes $z_{t+1} = \rho_z z_t + \sigma_z \varepsilon_{t+1}^z$ and $y_{t+1} = \rho_y y_t + \sigma_y \varepsilon_{t+1}^y$, respectively, where $|\rho_i| < 1$ and $\sigma_i > 0$, for $i = y, z$. The random variables ε_t^i have compact support and are i.i.d. for $i = y, z$. All firms share the same parameters ρ_z and σ_z . For convenience, let $x_t = (y_t, z_t)$ denote the vector of exogenous state variables. Its transition function is $Q(x_t, x_{t+1})$, which, given the independence of y and z , factorizes as the product of Q_y and Q_z . The transition function Q is assumed stationary, monotone, and Markov, and it satisfies the Feller property. These properties are, together with other properties discussed below, sufficient to ensure the existence of a solution to the firm valuation problem.

Following Jones and Tuzel (2013), we specify the one-period stochastic discount factor for a transition from y_t to y_{t+1} as

$$\mathcal{M}_{t+1} = \beta \exp \left(-g(y_t) \sigma_y \varepsilon_{t+1}^y - \frac{1}{2} g(y_t)^2 \sigma_y^2 \right),$$

where $0 < \beta < 1$. The state-dependent market price of risk is $g(y_t) = \exp(\gamma_1 + \gamma_2 y_t)$, with $\gamma_1 > 0$ and $\gamma_2 < 0$. The implied risk-free rate is state-independent and equal to $r = 1/\beta - 1$.⁸

B. The model of the firm

The timeline is as follows: x_t is realized at t and, given the legacy capital stock k_t , the cash flow from operations is $\pi_t = \pi(y_t, z_t, k_t) = e^{y_t + z_t} k_t^\alpha - f$. The condition $\alpha < 1$ implies decreasing returns to scale, and $f \geq 0$ summarizes operating expenses.⁹ The firm can issue one-period zero-coupon unsecured bonds.

For convenience, we collect the endogenous state variables in $a_t = (k_t, b_t)$, where b_t is the legacy total debt obligation. Based on a_t and the realization of π_t , at t the firm may be solvent or insolvent. If it is solvent, the firm chooses the new capital stock k_{t+1} and debt b_{t+1} . The firm's investment is $I_t = k_{t+1} - (1 - \delta)k_t$, which can be positive or negative, where $\delta > 0$ is a constant depreciation rate. As in Abel and Eberly (1994), changes in capital are subject to an asymmetric cost $h(I_t, k_t) = (\lambda_1 \mathbf{1}_{\{I_t > 0\}} + \lambda_2 \mathbf{1}_{\{I_t < 0\}}) \frac{I_t^2}{\delta k_t}$, where $0 < \lambda_1 < \lambda_2$ capture costly partial reversibility of investment.

The market value of new debt, $B(x_t, a_{t+1})$, depends on the current state x_t and on the firm's choices $a_{t+1} = (k_{t+1}, b_{t+1})$. Changes in debt are subject to a proportional adjustment cost

⁸This holds exactly when the shocks are normally distributed. With the discretized shocks used in the numerical solution, the risk-free rate differs from $1/\beta - 1$ by a negligible amount.

⁹Unlike in the simplified model in Section III, operating leverage is independent of capital stock and financial leverage is endogenously determined.

$q(b_t, b_{t+1}) = \theta|b_{t+1} - b_t|$, with $\theta \geq 0$.¹⁰ The choice of new debt generates a net cash flow to equity equal to $B(x_t, a_{t+1}) - b_t - q(b_t, b_{t+1})$.

The corporate tax function is linear with rate τ . We assume that $H(x_t, a_{t+1})$, the expected present value of the end-of-period interest payment $b_{t+1} - B(x_t, a_{t+1})$, can be expensed for tax purposes at time t , when the new debt is chosen. This is equivalent to expensing interest at maturity, at $t + 1$, if investors know the equilibrium conditional default probability.

At t , the after-tax cash flow from operations and financing decisions is

$$v_t = v(x_t, a_t, a_{t+1}) = (1 - \tau)\pi_t + \tau\delta k_t + \tau H(x_t, a_{t+1}) + B(x_t, a_{t+1}) - b_t - q(b_t, b_{t+1}).$$

If $v_t < 0$, the firm is in financial distress, although not yet in default. In this case, shareholders, conditional on survival, may inject new equity subject to a proportional transaction cost $\xi \geq 0$.

Overall, equity holders receive $w_t = w(x_t, a_t, a_{t+1}) = v_t (1 + \xi \mathbf{1}_{\{v_t < 0\}}) - I_t - h(I_t, k_t)$, that is, the after-tax cash flow from operations and financing decisions minus investment. A positive w_t leads to share repurchases, while a negative w_t requires equity issuance with proportional cost $\zeta \geq 0$. The net payout is $P(x_t, a_t, a_{t+1}) = w_t (1 + \zeta \mathbf{1}_{\{w_t < 0\}})$.

The prices of debt and equity are determined in competitive markets under rational

¹⁰A debt adjustment cost is required to generate persistence in debt when using single-period debt. The use of single-period debt avoids the commitment problems associated with long-term debt, as analyzed by DeMarzo and He (2021). While these problems would affect security prices and therefore option values, they would also make the model intractable.

expectations. The cum-payout equity price at t satisfies the Bellman equation

$$(1) \quad S(x_t, a_t) = \max \left\{ 0, \max_{a_{t+1}} [P(x_t, a_t, a_{t+1}) + \mathbb{E}_t(\mathcal{M}_{t+1} S(x_{t+1}, a_{t+1}))] \right\},$$

where $\mathbb{E}_t[\cdot] = \mathbb{E}[\cdot \mid x_t]$. In (1), the outer maximization captures endogenous default, and we

denote by $\omega_t = \omega(x_t, a_t)$ the default indicator. If the firm is solvent, the optimal policy is

$a_{t+1}^* = \varphi(x_t, a_t) = \varphi_t$, where $a_{t+1}^* = (k_{t+1}^*, b_{t+1}^*)$. Define the optimal net payout

$P_t = P(x_t, a_t, \varphi_t)$. If $P_t > 0$, this corresponds to share repurchases; if $P_t < 0$, it corresponds to equity issuance. The optimal policy is summarized by (ω_t, φ_t) .

At issuance, the debt price at t is

$$(2) \quad B(x_t, a_{t+1}) = \mathbb{E}_t [\mathcal{M}_{t+1} \{b_{t+1}(1 - \omega_{t+1}) + [\pi_{t+1} + \tau \delta k_{t+1} + (1 - \delta)k_{t+1}] (1 - \eta)\omega_{t+1}\}],$$

which reflects the firm's optimal policy. In (2), if the firm is solvent ($\omega_{t+1} = 0$), debtholders are paid the total obligation b_{t+1} ; if default occurs ($\omega_{t+1} = 1$), debtholders receive the sum of operating cash flow, the depreciated asset value, and the depreciation tax shield, net of a proportional bankruptcy cost η .

Finally, the expected present value of the interest payment is

$$(3) \quad H(x_t, a_{t+1}) = [b_{t+1} - B(x_t, a_{t+1})] \mathbb{E}_t [\mathcal{M}_{t+1} (1 - \omega_{t+1})],$$

which is the price of a default-contingent claim that pays the interest if the firm is solvent at $t + 1$.

To price single-share options, we need the law of motion for the share count. Let n_t be the number of shares outstanding at time t , before the payout decision. The pre-payout price per share

is $s(x_t, a_t) = S(x_t, a_t)/n_t$. Because we assume that only share repurchases are used to disburse cash to shareholders, the payout decision implies¹¹


$$n_{t+1} = \frac{S(x_t, a_t) - P_t}{s(x_t, a_t)} = \frac{\tilde{S}(x_t, a_t)}{s(x_t, a_t)} = \left(1 - \frac{P_t}{S(x_t, a_t)}\right) n_t,$$

where $\tilde{S}(x_t, a_t) = S(x_t, a_t) - P_t$ denotes the after-payout equity value. A timeline of the firm's decisions is shown in Figure C.1.

FIGURE C.1

Model timeline

At t , the number of shares outstanding is n_t . Given the legacy state $a_t = (k_t, b_t)$, the firm observes $x_t = (y_t, z_t)$. It chooses next-period capital $k_{t+1}^* = I_t + (1 - \delta)k_t$, incurring an adjustment cost $h(I_t, k_t)$. The firm repays b_t and chooses b_{t+1}^* , incurring an adjustment cost $q(b_t, b_{t+1}^*)$. Hence, it determines $a_{t+1}^* = \varphi(x_t, a_t)$, and the payout is $P_t = P(x_t, a_t, \varphi(x_t, a_t))$. The number of shares for next period becomes n_{t+1} .



Solving the model entails finding a solution to the Bellman equation (1), subject to the constraints in (2) and (3). This amounts to determining the prices $S(\cdot)$, $B(\cdot)$, and $H(\cdot)$, as well as the optimal policies $(\omega(\cdot), \varphi(\cdot))$, for each possible state of the economy. The existence and uniqueness of a solution are guaranteed by the properties of the transition function $Q(x_t, x_{t+1})$ and the compactness of the domain for capital stock and debt, as shown in Gamba and Saretto (2020).

¹¹While n_t and n_{t+1} are integer-valued in reality, we assume $n_t, n_{t+1} \in \mathbb{R}$ for analytical convenience.

C. Option prices

We focus on pricing an option chain with moneyness $m \in \{m_1, m_2, \dots, m_N\}$, so that the strike of each of these options is $\kappa = m s(x_t, a_t)$, where $s(x_t, a_t)$ is the per-share price at the date of creation. When a European call option¹² with strike κ and maturity $i = 1, \dots, 4$ is written at t , its price is $c(x_t, a_t; \kappa, i) = \mathbb{E}_t [\mathcal{M}_{t+i} \max \{s(x_{t+i}, a_{t+i}) - \kappa, 0\}]$, where x_{t+i} is determined by the evolution of the exogenous state until maturity, $\mathcal{M}_{t+i}$ is the stochastic discount factor used to discount payoffs at $t + i$ determined by the evolution of the aggregate shock, and a_{t+i} is defined recursively as $a_{t+j+1} = \varphi(x_{t+j}, a_{t+j})$, $j = 0, \dots, i - 1$ along the realized path of x_t . If a company defaults at a given t , all call and put options written on the company's stock are canceled. We omit default policy from the above expression of the option price for simplicity. However, our calculations correctly incorporate the default event in the option price.

The set of moneyness levels is predetermined, but different option chains are created along different paths of the stock price. Given the Markovian structure of our economy, in principle at each (x_t, a_t) we would be required to price up to four (past and current) vintages of options, with strikes κ based on infinitely many past stock prices, and four residual maturities. This is clearly not feasible, and some simplifying assumptions are in order to make the problem tractable within our stationary solution.


We start by pricing all options currently traded at t , without consideration of their vintage, with the same strike κ based on a recursive expression that is flexible enough to accommodate

¹²The arguments presented below are the same for European put options.

FIGURE C.2

Pricing of multi-period single-share options

At time t , the number of shares is n_t . Given a_t , the firm observes x_t and chooses a_{t+1} , and the resulting net payout is P_t . Hence, the new number of shares is n_{t+1} . Options on one share of the equity are created with many strikes and four expiration dates (end of period one, end of period two, and so on). We refer to this collection of options as the *option chain*. Options expire after the realization of the new state and before the payout at that date.



different possible (residual) maturities:¹³

$$(4) \quad c(x_t, a_t; \kappa) = \mathbb{E}_t [\mathcal{M}_{t+1} (\chi_{t+1} \max \{s(x_{t+1}, a_{t+1}) - \kappa, 0\} + (1 - \chi_{t+1}) c(x_{t+1}, a_{t+1}; \kappa))],$$

where $\chi_{t+1} = 1$ if the option expires at the end of the period. When one period remains, the problem in (4) is straightforward, as there is no continuation value. The payoff

$s(x_{t+1}, a_{t+1}) = S(x_{t+1}, a_{t+1})/n_{t+1}$ depends on a_{t+1} and n_{t+1} , which are known at t . Hence, the calculation of $c(x_t, a_t; \kappa)$ is exact, as it requires a simple expectation with respect to x_{t+1} .

However, if the maturity of the options is longer than one period, the payout-driven evolution of the share count makes the problem in (4) path dependent and prevents the use of standard backward induction. Path dependence arises because the continuation value

$c(x_{t+1}, a_{t+1}; \kappa)$ is based on x_{t+i} for $i \geq 2$, depending on option maturity. For instance, if maturity

¹³For this reason, we drop for the moment the dependence on i from the notation $c(x_t, a_t; \kappa, i)$.

is $t + 2$, the payoff is based on $s(x_{t+2}, a_{t+2}) = S(x_{t+2}, a_{t+2})/n_{t+2}$, where $a_{t+2} = \varphi(x_{t+1}, a_{t+1})$ and $n_{t+2} = \left(1 - \frac{P_{t+1}}{S(x_{t+1}, a_{t+1})}\right) n_{t+1}$. In this case, both n_{t+2} and a_{t+2} , and therefore $s(x_{t+2}, a_{t+2})$, depend on x_{t+1} . This means that the payoff of the options at $t + 2$ is not a function of x_{t+2} alone, but depends on the full path $\{x_{t+1}, x_{t+2}\}$, which makes backward induction impossible.¹⁴ This issue compounds over each intermediate period for maturities longer than two periods. The standard fix for path dependence—adding n_t to the state space—is infeasible because there is no reason to assume that the attainable share count is bounded. Also, this would condition the result on a given initial n_0 , which can only be set arbitrarily.

In what follows, we address the issue in three steps: a change of underlying asset, an interpolation step that recovers single-share prices, and a normalization. As a first step, we substitute $s(x_{t+1}, a_{t+1}) = S(x_{t+1}, a_{t+1})/n_{t+1}$ into (4), multiply both sides by n_{t+1} , and then define $C(x_t, a_t; \kappa n_{t+1}) \equiv c(x_t, a_t; \kappa) n_{t+1}$, obtaining

$$C(x_t, a_t; \kappa n_{t+1}) = \mathbb{E}_t [\mathcal{M}_{t+1} (\chi_{t+1} \max \{S(x_{t+1}, a_{t+1}) - \kappa n_{t+1}, 0\} + (1 - \chi_{t+1}) C(x_{t+1}, a_{t+1}; \kappa n_{t+1}))].$$

This equation reformulates (4) in terms of total equity. Since $S(x_{t+1}, a_{t+1})$ does not depend on share count, this eliminates the path dependence arising from $s(x_{t+1}, a_{t+1})$ noted above. At the same time, this formulation makes option prices fully dependent on the firm's endogenous policy through the evolution of total equity value. The strike price of options on total equity is set at κn_{t+1} , which is assumed constant until the option expires.¹⁵

¹⁴To put it differently, to find $c(x_t, a_t; \kappa)$ in (4), one should know not just (x_t, a_t) , but also x_{t-1} .

¹⁵Reformulating the original problem, which is based on the stock price, in terms of total equity is not equivalent. For equivalence, the strike price would need to evolve stochastically with the share count, which is not the case here. Instead, this is an *approximation* for options with maturity longer than one period. Given the original κ , it assumes

After this reformulation, at each (x_t, a_t) and for a given n_{t+1} , there are still infinitely many strikes κ to price, because the strikes of options written in earlier periods depend on past stock prices. Because solving the recursion separately for each strike is infeasible, we introduce an auxiliary problem that prices options on a fixed grid of strikes and recovers $C(x_t, a_t; \cdot)$ by interpolation. The auxiliary problem is defined on a predetermined and time-independent grid of strikes $\mathcal{K} = \{K_1, \dots, K_N\}$. For each $K_j \in \mathcal{K}$, we use backward induction to compute

(5)

$$C(x_t, a_t; K_j) = \mathbb{E}_t [\mathcal{M}_{t+1} (\chi_{t+1} \max \{S(x_{t+1}, a_{t+1}) - K_j, 0\} + (1 - \chi_{t+1}) C(x_{t+1}, a_{t+1}; K_j))].$$

For each K_j , $C(x_t, a_t; K_j)$ correctly incorporates the firm's optimal policy at all future dates. For any possible κ n_{t+1} , we then interpolate $C(\cdot)$ from (5), using the grid $\mathcal{K}$, obtaining the approximation $\widehat{C}(x_t, a_t; \kappa n_{t+1})$.¹⁶ By definition, $C(x_t, a_t; \kappa n_{t+1}) \equiv c(x_t, a_t; \kappa) n_{t+1}$. Therefore, the single-share option price is approximated by

$$(6) \quad \widehat{c}(x_t, a_t; \kappa) = \frac{1}{n_{t+1}} \widehat{C}(x_t, a_t; \kappa n_{t+1}).$$

Since (6) still depends on n_{t+1} , a final step is required to obtain a workable expression.

that options written on total equity have a strike price equal to κn_{t+1} until maturity, even though the share count will change endogenously to n_{t+i} , $i = 2, \dots, 4$.

¹⁶For instance, using linear interpolation,

$$\widehat{C}(x_t, a_t; \kappa n_{t+1}) = C(x_t, a_t; K_j) + \frac{\kappa n_{t+1} - K_j}{K_{j+1} - K_j} [C(x_t, a_t; K_{j+1}) - C(x_t, a_t; K_j)],$$

where j satisfies $K_j < \kappa n_{t+1} < K_{j+1}$. This step requires that the grid $\mathcal{K}$ is wide enough to encompass all relevant κ and dense enough to deliver a good approximation.

Because we specified $\kappa = m s(x_t, a_t)$, the strike on the right-hand side of (6) becomes

$\kappa n_{t+1} = m \tilde{S}(x_t, a_t)$,¹⁷ which is independent of the share count. Using $n_{t+1} = \frac{\tilde{S}(x_t, a_t)}{s(x_t, a_t)}$, we obtain

$$\frac{\widehat{c}(x_t, a_t; m s(x_t, a_t))}{s(x_t, a_t)} = \frac{1}{\tilde{S}(x_t, a_t)} \widehat{C}(x_t, a_t; m \tilde{S}(x_t, a_t)),$$

which is independent of n_t . The last equation shows that the ratio of the option price to the current share price is the natural object to compute and match to the data. A convenient normalization is to set $n_t = S(x_t, a_t)$, so that $s(x_t, a_t) = 1$ and $\kappa = m$. Hence, the price of an option with moneyness m and maturity $i = 1, \dots, 4$ is

$$\widehat{c}(x_t, a_t; m, i) = \frac{1}{\tilde{S}(x_t, a_t)} \widehat{C}(x_t, a_t; m \tilde{S}(x_t, a_t), i).$$

The three-step procedure described above leads to an option pricing expression that is solely based on the outputs of the firm valuation problem, and yields an exact price for one-period options and a close approximation for longer maturities.

D. Numerical settings

The solution to the Bellman equation (1) with the constraints in equations (2) and (3), and the related optimal policy, are obtained by discretizing the state space of $x = (y, z)$ and the control variables $a = (k, b)$. Because the stochastic process of systematic risk is highly persistent, we discretize the two exogenous processes using the numerical approach proposed by Rouwenhorst (1995). The aggregate shock y is discretized with $N_y = 9$ points and the

¹⁷Indeed, $\kappa n_{t+1} = m s(x_t, a_t) n_{t+1} = m s(x_t, a_t) \frac{\tilde{S}(x_t, a_t)}{s(x_t, a_t)} = m \tilde{S}(x_t, a_t)$, since $n_{t+1} = \frac{\tilde{S}(x_t, a_t)}{s(x_t, a_t)}$.

idiosyncratic shock z with $N_z = 37$. The discretized set of values for capital stock is

$\{k_j = k_u(1 - \delta)^j \mid j = 1, \dots, N_k\}$, and the set of discrete debt levels is

$\{b_j = j \frac{b_u}{N_b} \mid j = 1, \dots, N_b\}$ with $N_k = 33$ and $N_b = 15$.

The fixed point of (1), S , is found using a value function iteration algorithm, and the algorithm is halted when the maximum change in S between two iterations is below the tolerance 10^{-4} . We generate 5,000 simulated economies, each characterized by a different history of the aggregate state variable. Each sample is composed of a panel of 1,000 firms, each characterized by an independent history of the idiosyncratic variable. To make the simulated economy comparable to the observed data, each firm's history, after the first 20 periods are discarded, is composed of 92 periods. The desired quantities are obtained by applying the optimal policy at each step. In particular, at each step, option chains with four maturities (1, 2, 3, and 4 quarters) and three moneyness levels are evaluated using the approach described in Section C.C. If a company defaults at a given t , it is restarted at the steady state for k_{t+1} and b_{t+1} . In case of default, all call and put options written on the company's stock are canceled.