Federal Reserve Bank of Dallas

Endogenous Option Pricing

Andrea Gamba and Alessio Saretto

Working Paper 2202

March 2022

Research Department

https://doi.org/10.24149/wp2202

Endogenous Option Pricing*

Andrea Gamba† and Alessio Saretto‡

March 8, 2022

Abstract

We show that a structural model of firm decisions can produce very flexible implied volatility surfaces: upward and downward sloping, u-shaped. A calibrated version of the model is able to match many unconditional financial characteristics of the average optionable stock, and can help explain how, contrary to simple economic intuition, more valuable growth and contraction options are associated with a more negatively sloped implied volatility curve (i.e., a more negatively skewed implied distribution).

JEL Classifications: G12, G32

Keywords: option pricing, risk-neutral skewness, growth options, leverage, investments

* The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Dallas or the Federal Reserve System. We thank seminar participants at the Virtual Derivatives Workshop, Virtual Finance Seminar series at the University of Lancaster, Federal Reserve Bank of Atlanta, Federal Reserve Bank of Dallas, University of Calgary, University of Wisconsin at Milwaukee, Warwick Business School, and the Washington University in St. Louis for useful comments.

† Andrea Gamba, Warwick Business School, University of Warwick, Scarman Road, Coventry, UK CV4 7AL. Email: andrea.gamba@wbs.ac.uk.

‡ Alessio Saretto, Federal Reserve Bank of Dallas, 2200 N. Pearl Street, Dallas, Texas 75201. Email: alessio.saretto@dal.frb.org.

.

1. Introduction

Stocks' implied volatility surfaces exhibit a large degree of cross-sectional and time-series variability, reflecting heterogeneous market's expectations of forward looking (risk-neutral) distributions. At any point in time the implied volatility surface of a stock for a given maturity can assume many shapes. At three months maturity, 73% of stocks exhibit an implied volatility smirk (i.e., implied volatility is monotonically decreasing with strikes), while about 18% present a smile (i.e., implied volatility decreases and then increases with strikes). In the remaining cases, the surface is either concave (i.e., an inverted smile that we call frown), or increasing with strikes. Over more than 20 years of data, the percentage of stocks that exhibit a smirk varies considerably from 41.3% to 87.7%. Substantial variation also occurs through contracts' maturities: Implied volatilities are on average decreasing with options maturities; the percentage of stocks that exhibit a smirk increases to over 80% when the implied surface is extracted from contracts that are approximately one year to maturity.

Once aggregate and industry variation is taken into account, a substantial amount of the cross-sectional and time-series heterogeneity can be explained by firm level characteristics: in particular, proxies for future investment opportunities such as the market to book ratio, are statistically highly associated with the shape/slope of the implied volatility curve. High market to book ratios are associated with more negatively sloped implied volatility curves (i.e., more negative risk-neutral skewed distributions, as in Morellec and Zhdanov, 2019). While the result is empirically very robust, it contrasts the general intuition that growth option positively increase the skewness of the return distribution (see for example, Grullon, Lyandres, and Zhdanov, 2012; Trigeorgis and Lambertides, 2014; Del Viva, Kasanen, and Trigeorgis, 2017; Bali, Del Viva, Lambertides, and Trigeorgis, 2020; Panayiotis, Bali, Kagkadis, and Lambertides, 2021).

While it is possible to explain that much variation in equity option prices in the context of (reduced form) option pricing models that allow the underlying price to follow a non-gaussian distribution, as for example by introducing jumps and stochastic volatility, such an effort would require a substantial variation in parameters across firms and time (see for example, Geske, Subrahmanyam, and Zhou, 2016; Bakshi, Cao, and Zhong, 2021), leaving unexplained what originates such parameter variability. A viable alternative is to endogenously produce variation in forward looking equity distributions in a model that structurally links asset prices to firm's policies. This approach is however also challenging. For example, the models of Geske (1979) (i.e., the compound option model), and of Toft and Prucyk (1997), who price an equity option on a firm that faces taxes and bankruptcy costs as in Leland (1994), can produce implied volatility surfaces that respond to changes in capital structure, but that are also only downward sloping, and thus do not explain curves that resembles smiles or are straight out upward sloping, or distribution that do not have a negative skewness. Notably, from an empirical point of view, the link between option prices and measures of leverage has also been debated: Toft and Prucyk (1997) and Geske, Subrahmanyam, and Zhou (2016) show that accounting for the leverage effect reduces pricing errors relative to the traditional model of Black and Scholes (1973); however, Figlewski and Wang (2000) show that the link between leverage and option prices might not be very robust. In contrast, the statistical link between the shape of the implied volatility curve and the market to book ratio is incredibly strong, in the cross-section and in the time-series.

We show that a model that incorporates the effect of both investment and financing decision on how securities are priced can produce option prices and implied distributions that are, in many respects, in line with what we observe in the data. We follow the lead of Hennessy and Whited (2005, 2007) and Zhang (2005), and adopt a model set up that is very popular in dynamic corporate and investment based asset pricing studies. Our work is thus related to recent contributions such as Morellec and Zhdanov (2019), for example, who analyze the impact of product market competition on option prices. As many papers that model production economies have pointed out (e.g., Kuehn and Schmid, 2014, for an application to corporate bond prices), investment policies play an important role in shaping the distribution of asset prices. Because of that, it is natural to ask whether such models are consistent enough to also produce sensible option prices. We show here how they do so.

Optimal investment and financing policies respond to exogenous productivity shocks to determine the forward looking distribution of stock prices. The firm can both expand or reduce its capital stock at any point in time, subject to asymmetric capital adjustment costs that make the investment process lumpy (see for example, Kogan, 2004b,a; Zhang, 2005; Cooper, 2006, for early important contributions). Thus, the firm holds an infinite stream of real "straddles", which we generally refer to as the real option (as for example in the real option models of Gu, Hackbarth, and Johnson, 2017; Aretz and Pope, 2018). At any point in time, the value of the real option is higher when the firm is either very likely to invest in the future, because of positive productivity prospects, or when it is very likely to shed some capital, because of negative future economic conditions. To keep the parallel with traditional option pricing jargon, the (present) value of the real option is large when either the call or the put real options are in the money. However, those are also the cases where the real option contains less optionality, and the real option behaves as asset in place (i.e., in the traditional option pricing jargon either the call delta is close to one or the put delta is close to minus one). The real option has instead more "optionality" when the firm is equally likely to invest or disinvest (i.e., both options are at the money) in the future. But relatively speaking that's when the real "straddle" is worth the least.

In early work, van Zwet (1964) shows that the convexity of a function of a random variable is related to the skewness of the distribution of the function of the random variable, and that more convexity implies larger skewness. van Zwet (1964) and more recently Xu (2007) are often cited in the growth option literature as a motivation for the idea that firms with more growth options (a convex function of the underlying) have more positively skewed distributions (see for example, Del Viva, Kasanen, and Trigeorgis, 2017; Bali, Del Viva, Lambertides, and Trigeorgis, 2020; Panayiotis, Bali, Kagkadis, and Lambertides, 2021). However, because firms hold both an option to invest and one to disinvest, the impact of real option on the skewness of the equity distribution, and hence on the pricing of financial options on the equity, is greater in the situations in which the real option has higher curvature/convexity. This happens when future investment policy is more uncertain, when the firm is equally likely to invest or disinvest, a situation which leads the real option to be relatively less valuable. Hence, the relation between the value of real option and the skewness of the equity return distribution is *negative*, as opposed to positive as previously postulated.

Our model, therefore, combines the leverage effect, which induces negative skewness, to the real option effect. Since the strength of these effects is state dependent, as it depends on the equilibrium choices of production capital and level of indebtedness, the resulting equilibrium option prices are also state-contingent and can generate any form of implied volatility curves (upward sloping, downward sloping, u-shaped, or even inverted u-shape) and of implied distributions (fat, long and short tails). Notably, there are many features that are missing from the model and that might be important: our economy is characterized by a homogeneous technology (i.e., all firms have the same production function), and thus does not account for the fact that differences in forward looking distributions might arise from adoption of new technologies as in Garleanu, Panaceas, and Yu (2012). Also, all firms in the economy have access to a short-term zero-coupon bond, but recent literature, as for example Chaderina, Weiss, and Zechner (2021) and Friewald, Nagler, and Wagner (2021), suggests that heterogeneity in the maturity of debt contracts might have sizable asset pricing implications.

We calibrate the model to match firm characteristics, implied volatilities, and moments of the return distribution of the average option-able stock. Our simulated firms make investment and capital structure choices in line with those observed in the data. The risk-neutral distribution instead displays negative skewness and a larger excess-kurtosis. This translates into implied volatility surfaces that are remarkably close to the data, in the shape and the frequency with which they are observed. The average implied volatility surface is downward sloping along the moneyness and along maturity levels. However the frequency of times that the curve assumes another shape is in line with what we see in the data, and in more than 20% of the cases it is either u-shaped (i.e., the traditional smile) or upward sloping.

The ability of the model to endogenously produce different return distributions is not a simple by product of the utility preferences that are underlying the stochastic discount factor that prices securities. As many other investment based asset pricing studies, we exogenously specify counter-cyclical risk premia. We employ the stochastic discount factor of Jones and Tuzel (2013), which naturally create an aversion to (left) skewness. Although it greatly aides the calibration of the model, the particular choice of the discount factor is not essential: we can also obtain our basic qualitative results with a constant (across states and time) discount rate.

We use the simulated economy to validate relationships between properties of option prices (i.e., risk-neutral skewness and implied volatilities curves) and firm characteristics that we observe in the data, focusing explicitly on the role of financial leverage and growth options. Similarly to the data, we show that a substantial amount of cross-sectional and time-series variation in the characteristics of option prices can be explained by the choices the firm makes along the path of productivity shocks that it encounters. In particular, we show that, after controlling for size, level of volatility, and leverage, the risk-neutral skewness and the steepness of the slope of the implied volatility curve (in the simulated economy) are more negative for firms with higher values of the real option and for higher market-to-book ratios, which is often used as a proxy for the former.

2. Related literature

This paper is primarily related to the strand of literature that aims at explaining equity option prices in the cross-section of stocks. Starting from the seminal work of Merton (1974), there have been a few attempts at incorporating option pricing into a structural model of the firm. Geske (1979) offers a first attempt by producing a double compound option that allows one to price a call option on the equity of a levered firm. Toft and Prucyk (1997) extends this approach to the Leland (1994) economy, thus allowing for taxes and bankruptcy costs to determine the optimal leverage policy of the firm. Geske, Subrahmanyam, and Zhou (2016) show that accounting for the leverage effect greatly reduces option pricing errors relative to the Black and Scholes (1973) model. Bai, Goldstein, and Yang (2019) show that the leverage effect is essential to explain the spread between index and individual banks equity options. Morellec and Zhdanov (2019) show risk-neutral skewness is related to the competitive landscape that surround a firm. Following Hennessy and Whited (2005, 2007), we introduce a fully dynamic model where shareholders endogenously choose production capacity, financial leverage, and default. We show that these ingredients are essential to reproduce the heterogeneity in option prices present in the data.

The leverage effect introduced by Merton (1974) has been considered in a number of applications that link volatility to stock prices/returns. Engle and Siriwardane (2017) propose a structural GARCH model that embeds the leverage effect into equity volatility forecasting models. Because we introduce a model where firms are exposed both to systematic and idiosyncratic risk, our work is also related to studies such as Duan and Wei (2009), who show the differential impact of the two sources of risk. Similar to Duan and Wei (2009), our model also implies that large variation in the prices of individual equity options is produced by realizations of aggregate risk. The leverage effect is a fundamental mechanism of modern models of credit risk. Thus, because we share many model features and because we rely on some of the same intuition our paper is also related to the rather large literature that studies corporate credit risk: from Leland (1994) to more recent contributions such as Gomes and Schmid (2021).

Our paper is also related to the very large literature that studies the impact of growth option on asset prices. From the many contributions to the understanding of the role of dynamic investment policies, among the most directly related to our paper are the works of Berk, Green, and Naik (1999) who link the predictability of stock returns to firm characteristics in a model with dynamic investments. Kogan (2004b), Kogan (2004a), Zhang (2005), Cooper (2006), Ai and Kiku (2013), Kogan and Papanikolaou (2013), Kogan and Papanikolaou (2014), and Gu, Hackbarth, and Johnson (2017) discuss the role of complete and partial investment irreversibility in shaping the risk-return profile. Aretz and Pope (2018) consider a model with both investment and disinvestment options use it to propose a rational explanation of empirical regularities in the cross-section of stock returns. Trigeorgis and Lambertides (2014) suggest an alternative measure of growth options to the book-to-market ratio and relate it to future stock returns. Del Viva, Kasanen, and Trigeorgis (2017) show that firms with more prevalent growth opportunities have more positive skewness in the return distribution. Cao, Simin, and Zhao (2008) show that growth options are mainly driven by idiosyncratic volatility. Lyandres and Zhdanov (2020) link miss-pricing to the presence of growth options, Bali, Del Viva, Lambertides, and Trigeorgis (2020) explains several stock return anomalies by linking them to options to alter the asset composition that are proxied by idiosyncratic skewness. Aguerrevere (2009) and Morellec and Zhdanov (2019) study how product market competition affect the optimal exercise of growth options.

Finally, our paper is related to the growing literature that studies the informational content of risk-neutral moments of the return distribution that can be obtained from option prices. Bakshi, Kapadia, and Madan (2003) introduce a feasible way to compute risk-neutral moments of the return distribution from option prices using a model-free approach. Bakshi, Kapadia, and Madan (2003) shows that while there is a one to one mapping between the risk-neutral distribution and the implied volatility surface, there is no unique mapping between each of the moments and implied volatilities: volatility, skewness, and kurtosis all combine to determine option prices. Dennis and Mayhew (2002) and Hansis, Schlag, and Vilkov (2010) study the relationship between risk-neutral moments and firm characteristics and find relatively contrasting results about the impact of the leverage effect.

A very large literature relates risk-neutral moments to realized and expected stock returns, including, but limited to, Bali and Murray (2013), Conrad, Dittmar, and Chisels (2013), Amaya, Christoffersen, Jacobs, and Vasquez (2015), Kadan and Tang (2019), Martin and Wagner (2019), Schneider, Wagner, and Zechner (2020), and Christoffersen, Fournier, Jacobs, and Karoui (2021).

In summary, this paper follows an established literature that aims at measuring and understanding the impact of corporate policies on asset prices (among many others, see for example, Kuehn and Schmid, 2014, who use a similar model to analyze the pricing of corporate debt.) Similar to Toft and Prucyk (1997), Geske, Subrahmanyam, and Zhou (2016), and Morellec and Zhdanov (2019) we offer an alternative approach to option pricing studies that rely on exogenous specifications of stochastic properties of equity prices. While we do not believe that our approach could be as successful in delivering small pricing errors for each security as this last class of models, our calibration is remarkably close in pricing options on the average firm, and in producing, with a single set of parameters, a widespread cross-section that is entirely produced by optimal investment and capital structure decisions. That is also our main point of departure from studies such as Toft and Prucyk (1997), Geske, Subrahmanyam, and Zhou (2016), and Morellec and Zhdanov (2019): our model is calibrated to the data to obtain option prices and firm characteristics that are quantitatively close to the observed economy.

3. Data

We construct our sample of option-able stocks by combining CRSP and COMPUSTAT with OptionMetrics. To increase the frequency of observations we obtain quarterly balance sheet observations and match them to stock returns data using common filters (i.e., share code 10 and 11, no ADRs, total assets in excess of 10 million USD, etc.). We construct stock returns and accounting ratios (leverage, profitability, market-to-book) using standard definitions as in Fama and French (1992).

We then match the resulting sample with OpionMetrics. Because option data is recorded daily but firms accounting information comes at quarterly frequency, we lined up the data by averaging options data over the last three trading days of the last month of the earning reporting quarter. For each firm and quarterly reporting date we extract option prices and volatility surfaces that, at that point in time, have maturities closest to 90 days (one quarter), 180 days, and 360 days. Most of the empirical regularities do not change if we considered 30 days contract. Calibrating the model at monthly frequency is however incredibly 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 90 days options as the baseline. Having multiple maturities allows us to construct a term structure of option prices and volatilities.

We retain two sets of data which are used for different purposes. Measures that refer to the implied volatility surface (i.e., the implied volatility slopes and implied volatility shapes) are computed using OptionMetrics Volatility Surface files. Instead of averaging the volatility surface of call contracts with that of put contracts, for each firm date and contract maturity, we merge the implied volatilities obtains from calls with strike prices higher than the underlying, and from puts with strike prices below the underlying. In other words, only out-of-the money call and put contracts are used. We now have only one surface. For simplicity we refer to $IV(k)$ as the implied volatility corresponding to moneyness k: so that $IV(1)$ is the ATM implied volatility, $IV(0.8)$ is the OTM implied volatility, and $IV(1.2)$ is the ITM implied volatility. For each firm-quarter and option maturity we retain implied volatilities for the range of moneyness (i.e., ratio of strike to underlying) that goes from 0.8 to 1.2. We then construct three slope measures: the difference between $IV(1.2)$ and $IV(0.8)$ (total slope), the difference between $IV(0.8)$ and $IV(1)$ (left slope), and the difference between $IV(1.2)$ and $IV(1)$ (right slope). For each firm-quarter and option maturity we also classify the shape of the implied volatility curve into four types: the curve is downward sloping as $IV(0.8) > IV(1) > IV(1.2)$ (left smirk), the curve is upward sloping as $IV(0.8) < IV(1) < IV(1.2)$ (right smirk), the curve is u-shaped as $IV(0.8) > IV(1) < IV(1.2)$ (smile), and the curve is inverted u-shaped as $IV(0.8) < IV(1) > IV(1.2)$ (frown).

We construct model-free implied skewness and kurtosis directly from OTM call and OTM put option prices as in Bakshi, Kapadia, and Madan (2003). In particular we follow the procedure described in Hansis, Schlag, and Vilkov (2010), whose code is available on Grigory Vilkov's page. We impose several filters to limit the impact of liquidity. In particular we eliminate prices that violate arbitrage bounds, that have zero open interest, zero-bid quotes, and have both bid and ask quotes unchanged for two consecutive days. For each stock we select options that at a particular point in time have approximate maturity 90, 180, or 360 days, and have moneyness (strike divided by stock price) 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, 8 option contracts enter the calculation of risk-neutral moments.

Because we require firm-quarter observations to have valid measures of implied volatility and implied moments, the liquidity filters applied towards constructing implied moments are also implicitly applied to the implied volatility measures. The final sample is composed of 3,536 stocks and includes quarterly observations between the years 1996 and 2019.

4. Empirical evidence about the cross-section of equity option prices

The option pricing literature has mainly focused on two different ways to organize option prices for different maturities and moneyness: implied volatility surfaces and implied risk-neutral moments. Mots of these efforts have been concentrated on index options, which offer a great way to understand aggregate risk premia and investor attitudes towards risk.

We organize the data along the same lines but we focus on individual equity options. We present in this section some empirical regularities that we deem important in thinking about what an option pricing model should address.

4.1. Implied volatility surfaces

The average implied volatility surface is downward sloping along both moneyness and maturity (see Figure 2), although less pronouncedly than the index option surface. At the

Figure 1: Implied volatility surface

The figure plots the average implied volatility surface extracted from the data. The sample contains all industrial firms with options trading on their equity between 1996 and 2019. Data is sampled at quarterly frequency. A total of 3,536 firms are included.

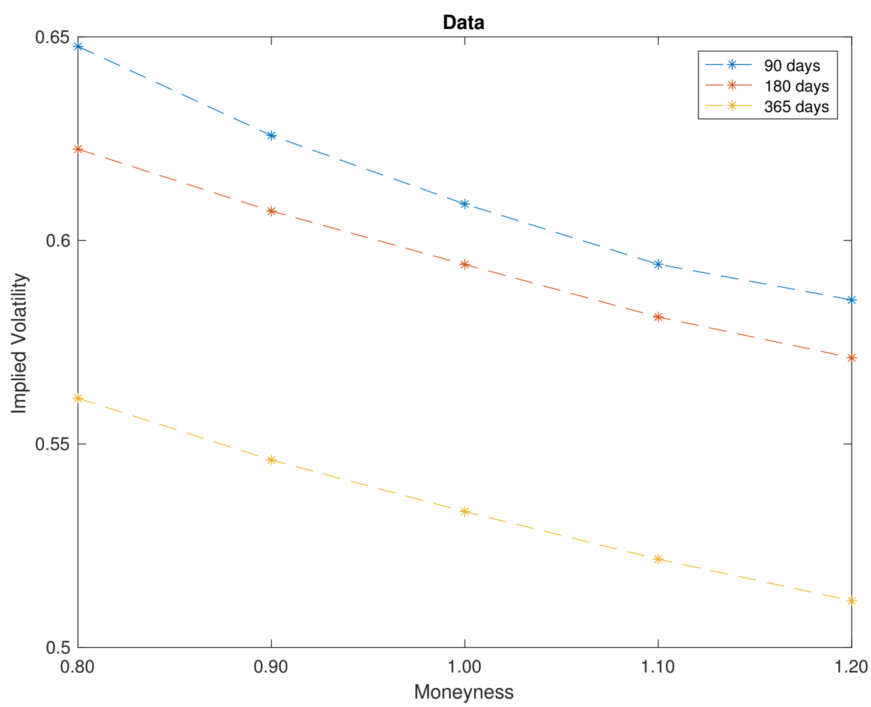

shortest maturity (that we consider) of 90 days, the average difference between $IV(0.8)$ (i.e., moneyness of 0.8) and $IV(1)$ is 3.5%, while the average difference between $IV(1)$ and $IV(1.2)$ is about 1.2%. Along maturities, the average difference between 90 and 360 days varies between 5.3% for the moneyness of 0.8, to 4% for moneyness of 1.

Figure 2: Implied volatility surface – time series

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

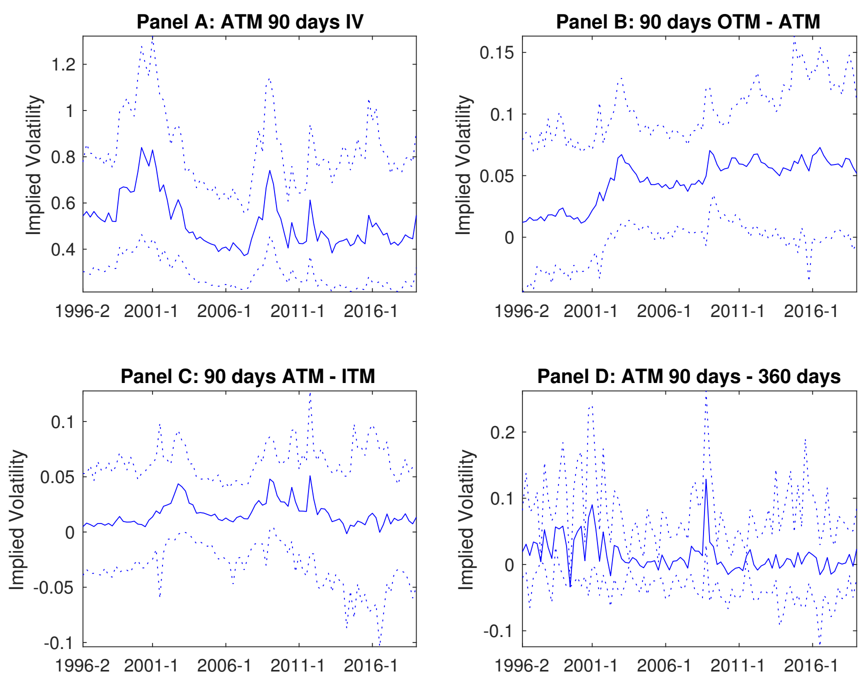

There is a considerable amount of cross-sectional and time-series variability in implied volatility surfaces. For example, the cross-sectional average ATM implied volatility varies between 90% at the height of the internet bubble crash to 35% in the middle of 2005 (see Panel A of Figure 2). At the same time, there is a fair amount of cross-sectional dispersion: for example at the hight of the financial crisis, the 95^{th} percentile of implied volatility is higher than 120%, while the 5^{th} is as low as 40%. Similarly, the left tail of the implied volatility curve can be as high as 15% and as low as -5% (see Panel B of Figure 2). The right tail varies even more from 10% to -10% (see Panel C). Similar variation can be seen even across maturities (see Panel D), where the slope of the volatility surface hovers around

4% but can be even negative (an upward sloping volatility term structure) for some stocks at particular points in time.

Figure 3: Surface types – time series

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

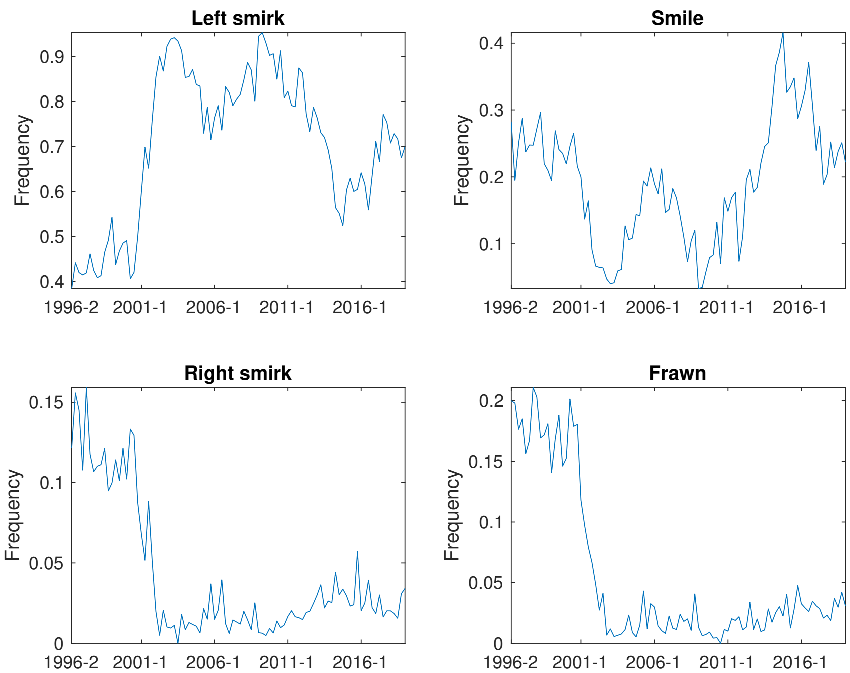

Variation in implied volatilities through time and across stocks produces also a very rich cross-section of different "shapes". We categorize the shape of the implied volatility curve into four types: left smirk (i.e., implied volatility decreasing with moneyness), smile, right smirk (i.e., implied volatility rising with moneyness), and frown (i.e., inverted smile). We plot the cross-sectional frequency of each surface type for 90 days options in Figure 3.

The most predominant surface type is a left smirk, which is observed on average 73% of the times, with a large time-series variation between 40% and 95% (see Panel A). The second most frequent surface is a smile (Panel B), which is observed on average in 18% of the cases. Right smirks and frowns are less frequent on average; they however manifest in a significant number of stocks during the years of the internet bubble (Panels C and D).

4.2. Risk-neutral moments

We plot time-series of the cross-sectional averages, as well as the 5^{th} and 95^{th} percentiles, of risk neutral skewness and kurtosis in Figure 4. At each maturity, the average risk neutral skewness is negative, with a sharp decrease around the internet bubble. Risk-neutral kurtosis is in excess of three and also increasing through the period. Large cross-sectional variation is observed and appears to be increasing through time. This is not completely surprising as the number of optionable stocks has increased substantially from a few hundred to more than one thousand (after all filters are applied) during our sample period.

Despite the fact that the time-series of cross-sectional averages are highly correlated across maturities some patterns emerge. Cross-sectional average and standard deviation of both risk neutral skewness and kurtosis tend to decrease in absolute value with maturity, which is consistent with implied volatility curves being flatter for longer maturities contracts.

Figure 4: Risk neutral moments – time series

The figure plots the time series of cross sectional averages, as well fifth and ninety-fifth percentiles, of risk-neutral skewness and kurtosis extracted from the data. The sample contains all industrial firms with options trading on their equity between 1996 and 2019. Data is sampled at quarterly frequency. A total of 3,536 firms are included.

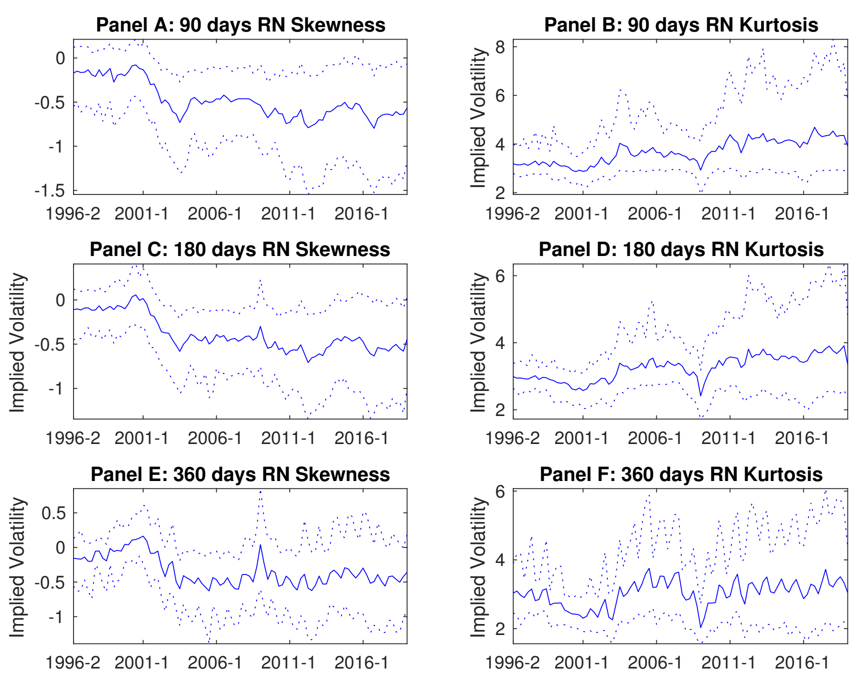

4.3. Observable characteristics

Next we show how variability in implied volatilities and distribution relates to observable firm characteristics. Table 1 reports results of panel regressions with different fixed effects specifications of several features of options prices against observable firm characteristics. Left hand side variables include the risk-neutral skewness, the slope (i.e., log difference of IV(1.2) and IV(0.8)), and the left slope (i.e., the log difference of IV(0.8)-IV(1)) extracted from 90 days to maturity options. We report results of similar regressions with a wider set of regressors, a different measure of growth option proposed by Trigeorgis and Lambertides (2014), and for 180 and 360 days maturities in the Appendix.

Table 1: Option prices and observable characteristics

The table shows regression results of several features of option prices against size, book leverage, market to book ratio, equity return, and ATM IV. Left hand side variables include the risk-neutral skewness, the slope (i.e., log difference of IV(1.2) and IV(0.8)), and the left slope (i.e., the log difference of IV(0.8)-IV(1)) of 90 day maturity options. We report parameter estimates and standard errors clustered at the firm level. Regressions include combinations of time, industry, or firm fixed effects. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data is sampled at quarterly frequency. A total of 3,536 firms are included.

	RN-Skew		Slope		Left slope	
IV	0.14	0.13	5.40	4.67	-5.45	-5.24
	(32.40)	(25.39)	(36.60)	(28.13)	(-49.55)	(-38.79)
Size	-0.13	-0.11	-3.04	-1.81	1.95	0.64
	(-20.86)	(-6.97)	(-16.95)	(-5.09)	(13.33)	(2.23)
Leverage	-0.01	-0.02	-0.17	-0.10	0.14	0.18
	(-2.97)	(-3.02)	(-1.74)	(-0.80)	(1.76)	(1.85)
Profitability	0.00	0.01	0.31	0.56	-0.40	-0.55
	(0.81)	(3.15)	(3.30)	(6.02)	(-5.64)	(-7.99)
M2B	-0.07	-0.08	-1.51	-1.45	0.95	0.86
	(-20.78)	(-15.02)	(-14.21)	(-11.30)	(11.30)	(8.11)
Time FE	X	X	X	X	X	X
Industry FE	X		X		X	
Firm FE		X		X		X
Adjusted-R^2	0.42	0.50	0.45	0.55	0.50	0.59
FE R^2	0.28	0.48	0.29	0.52	0.30	0.54

The table largely confirm the results reported in Table 3 of Morellec and Zhdanov (2019). After controlling for different fixed effects, the slope of the implied volatility curve and the risk-neutral skewness of the equity distribution are positively related to the level of volatility and profitability, and negatively related to size (i.e., natural logarithm of total assets), book leverage, and the market-to-book ratio. We are quick to note that most of the variation in the data is absorbed by fixed effects. Thus one could conjecture that it is the type of firm, rather than variation in characteristics over time that drive changes in option prices (see for example Figlewski and Wang, 2000). However, we note that covariates remain largely economically (the point estimates do not change much) and statistically significant even in the firm fixed effects regressions, thus suggesting that a relevant amount of variation in the data is attributable to time-series variation in the regressors: smaller firms with high volatility present more positively skewed risk-neutral distribution (less negatively slopped IV curves). Firms with high leverage and high market-to-book ratios present more negatively slopped implied volatility curves (more negative skewness).

Some of these relationships are not easily explainable in the context of a model such as those of Toft and Prucyk (1997) and Geske, Subrahmanyam, and Zhou (2016), which can describe very well the impact of leverage and the level of volatility, but not how size and market-to-book values affect option prices.

Overall, implied volatility surfaces or at risk-neutral moments present a pretty consistent picture of the cross-section of option prices: There is large time-series and cross-sectional variation in standardized option prices. While it is entirely possible that such variation can be explained by exogenously specifying the equity and volatility process as in Bakshi, Cao, and Zhong (2021), we propose a structural approach based on the idea that optimal firm decisions shape the physical and risk-neutral distributions of equity returns. The results reported in Table 1 are consistent with this approach.

5. Basic intuition

We convey the main economic intuition by means of a stripped-down version of the quantitative model that we present in Section 6. We start from the neoclassical model of Zhang (2005) and make some simplifications.

As in Zhang (2005), the model has infinite horizon and is in discrete time. The firm's productivity shock, z, follows a log-AR(1) process with parameters ρ and σ. The cash flow from operations is the result of applying the current productivity shock to a production function where capital exhibits decreasing returns to scale and accrues some operating costs:

$$\pi(z, k) = e^z k^\alpha - fk,$$

where $0 < \alpha < 1$ and $f \geq 0$. We model the operating cost as proportional to capital to capture the effects of mechanisms that are present in the main model such as financial leverage, and depreciation, which, contrary to a fix operating leverage, have an impact on the investment/disinvestment choice (i.e., a fix cost would not affect the investment decision).

At time $t = 1$, the equity holder invests $i = k' - k$ (i.e., there is no depreciation), and investment/disinvestment entails a quadratic asymmetric capital adjustment cost

$$h(i,k) = \frac{1}{2}\left(\theta_1 \mathbb{1}_{\{i>0\}} + \theta_2 \mathbb{1}_{\{i<0\}}\right)\left(\frac{i}{k}\right)^2 k.$$

The firm makes no other investment decisions after that but lives in perpetuity. We refer to this one time ability to change the stock of productive capital as the real option (i.e., to distinguish it from the financial option on the equity), with the general understanding it comprises the ability to both increase or decrease capital stock depending on economic conditions (i.e., the realization of the productivity shock). Investment is financed either with internal or external equity. There are no transaction costs when raising external equity financing. We assume investors do not require any risk premia and therefore securities are priced by a constant discount factor $0 < \beta < 1$.

Because we are interested in pricing financial options written on the firm's equity, we assume the following timeline: at $t = 0$, European call (and put) options are written, with strike price X, and maturity $t = 1$. The options are written after z_0 has been observed and k has been decided. At $t = 1$, the realization z_1 comes from a log-normal distribution with mean ρz_0 and standard deviation σ, and that determines the choice of k'. Hence, the value of equity at $t = 1$ is

$$S(z_1) = \max_{k'} \left\{\pi(z_1, k) - i - h(i,k) + V(z_1, k')\right\}$$

where $V(z_1, k')$ is the continuation value under the assumption that investment in all future periods is 0 (i.e., capital remains at level k'). [1] Therefore, one can think about the firm has having some capital in place plus one real (investment/disinvestment) option.

[1] Given the distributional assumptions on the productivity shock,

$$V(z_1, k') = \mathbb{E}_1\left[\sum_{j=1}^{\infty} \beta^j \pi(z_{1+j}, k')\right] = \Phi(z_1)(k')^\alpha - \frac{\beta}{1-\beta} f k'$$

where

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

Denoting $k^*(z_1)$ the solution of the equity program contingent on z_1, and the optimal investment $i^*(z_1) = k^*(z_1) - k$, the equity value is

$$S(z_1) = \pi(z_1, k) - i^*(z_1) - h(i^*(z_1), k) + V(z_1, k^*(z_1))$$

Hence, the price at $t = 0$ of a European call option with strike X maturing at $t = 1$ is

$$C_0(X) = \beta \int \max\{S(z_1) - X, 0\}\, \varphi(z_1) dz_1,$$

where $\varphi(\cdot)$ is the standard normal density.

Aside from the adjustment to the capital stock at time 1, some other elements of the model have implications for the pricing of options. In Figure 5 we present some comparative statics of the impact of relevant model parameters on the distribution of equity returns and the implied volatility surface, once we remove the investment decision. In this version of the model the firm starts with an initial capital that cannot be changed.

In the top row, we vary the coefficient of autocorrelation of the productivity shock. The case with $\rho = 1$ provides a good comparison case, as it essentially reduces to a Black and Scholes economy: the equity return is equal to 3% (i.e., $r - 0.5\sigma^2$, $r = 5\%$), equity volatility is equal to 20% (i.e., the volatility of the productivity shock), equity skewness is equal to 0, and the implied volatility curve is flat and leveled at 20%. As the autocorrelation decreases the continuation value of the firm becomes less valuable and less volatile (as a function of z_1). Because the drop in the second moment is much larger than the drop in the third moment, a decrease in autocorrelation leads to an increase in skewness, thus shifting mass to the right. As a consequence the implied volatility curve lowers and slightly steepens (positively) with moneyness.

can be easily calculated. In the equity program, the optimal k' is found solving the first order condition

$$\alpha(k')^{\alpha-1}\Phi(z_1) - \frac{\beta}{1-\beta} f = 1 + \left(\theta_1 \mathbb{1}_{\{k'>k\}} + \theta_2 \mathbb{1}_{\{k'<k\}}\right)\left(\frac{k'}{k} - 1\right)$$

which can be found numerically using Newton's method.

Figure 5: Model with no investments

In the left column, the figure plots the log equity returns between 0 and 1 for different realizations of the log productivity shock at time 1. For each case, we include the expected value, volatility, and skewness in the figure's legend. The right column displays the implied volatility of the call option prices at time 1 for moneyness levels (i.e., ratio of strike to underlying price) between 0.8 and 1.2 . For each case, we include the slope of the implied volatility curve (i.e., IV calculated for a moneyness of 1.2 minus the IV for a moneyness of 0.8), the risk-neutral skewness, and the market to book ratio. We vary the autocorrelation of the productivity shock, $\rho \in \{1, 0.95, 0.9\}$, the curvature of the production function, $\alpha \in \{0.4, 0.5, 0.6\}$, and the leverage parameter, $f = \{0, 0.15, 0.25\}$. We fix the rest of the parameters as follows: $\sigma = 20\%$, $k_0 = 3.2$, $\theta_1 = 0$, $\theta_2 = 0$, $\beta = e^{-0.05}$.

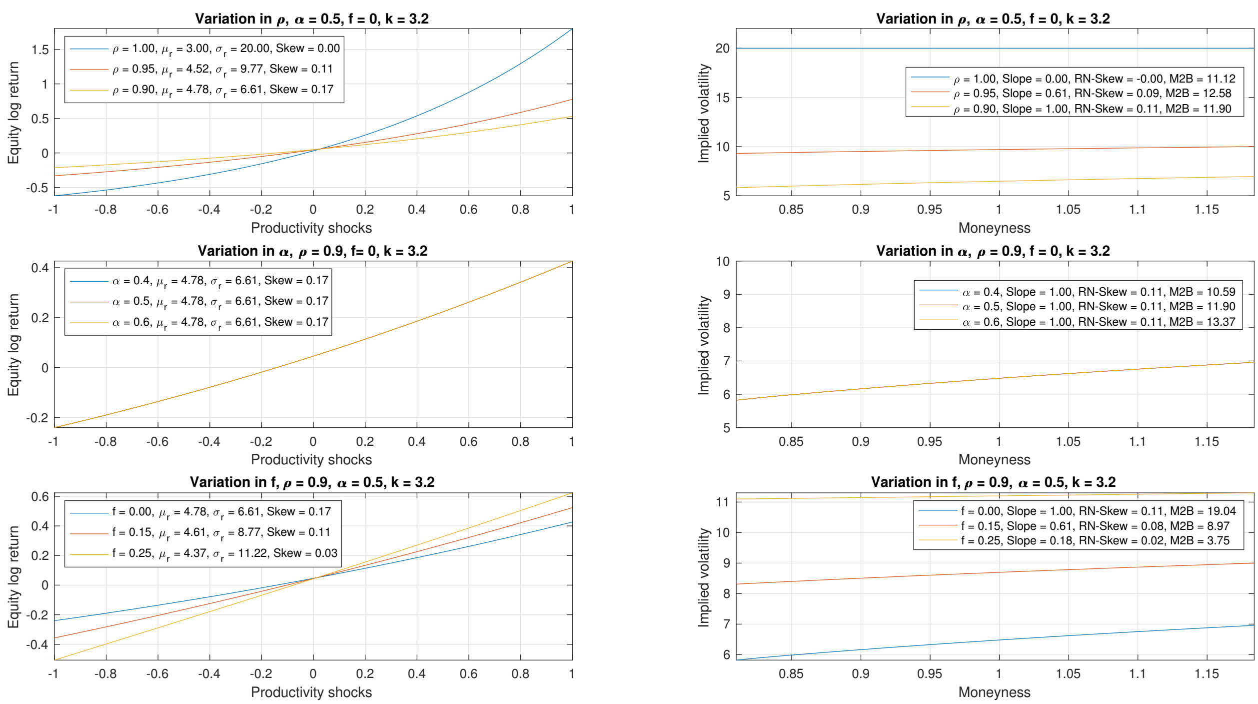

Figure 6: Model with investments

The figure plots the log investment and equity returns between 0 and 1 for different realizations of the log productivity shock at time 1, as well as the implied volatility curve at time 0 derived from options that matures at time 1. We compare the case with no investments to the solution of the model with the investment/disinvestment option with and without capital adjustment costs. We show results for a firm that will mostly invest at time 1 (i.e., $k = 2$), and for one that will mostly disinvest (i.e., $k = 6$). In the legend of the central figure we also report the present value of the real option, PVRO, which combines both the investment and the disinvestment options, and is defined as the difference between the market-to-book ratio of the solution with investment and the market-to-book ratio of the constrained model.When the firm faces adjustment costs we set $\theta_1 = 2$ and $\theta_2 = 10$. We fix the rest of the parameters as follows: $\rho = 0.9$, $x_0 = 0$, $\sigma = 20\%$, $f = 0.25$, $\beta = e^{-0.05}$.

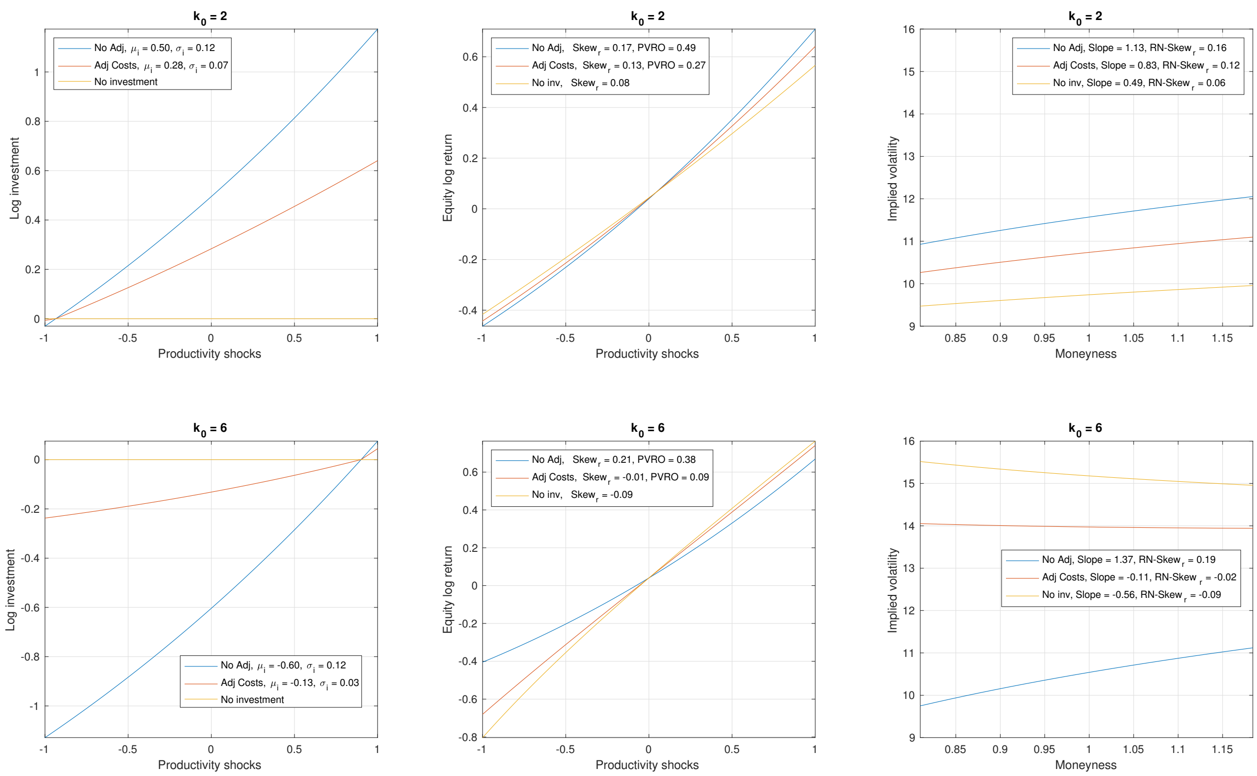

Increasing the curvature of the productivity function lowers the equity volatility, but not the skewness: the impact of α on the first and second moment of the equity distribution is proportional. Increasing leverage (i.e., f) increases equity volatility and shifts the equity distribution to the left. Similarly to Geske (1979), Toft and Prucyk (1997), and Morellec and Zhdanov (2019) this creates a downward sloping implied volatility curve. Thus, for a constant production technology and dynamic of productivity, leverage impacts the pricing of financial options.

Because adding the real option increases the convexity of the equity value, it also affects its distribution. van Zwet (1964) shows that the convexity of a function of a random variable is related to the skewness of the distribution of the function of the random variable, and that more convexity implies larger skewness. van Zwet (1964) is often cited in the growth option literature as a motivation for the idea that firms with more growth options (a convex function of the underlying) have more positively skewed distributions (see for example, Del Viva, Kasanen, and Trigeorgis, 2017; Bali, Del Viva, Lambertides, and Trigeorgis, 2020). We detail how the skewness of the equity distribution is induced by a real option that consist of a combination of an option to buy and an option to sell (i.e., a "real straddle").

In Figure 6 we show how introducing the real option affects the skewness of equity distribution and the related implied measures, in the general case where $\rho < 1$. Cognizant that investments and disinvestments do not have symmetric effects, because the capital adjustment cost function is asymmetric, we consider a firm that, given the current shock, is investing for almost all future realization of the productivity shock (i.e., $k = 2$ top row) and one that almost always disinvest (i.e., $k = 6$ bottom row). For each case, we compare the case with constant capital, to the case with the real option with and without capital adjustment costs.

We start by considering the base case scenario where the firm cannot change its capital stock at time 1 (yellow line). Going directly to the implied volatility plots, we observe that there is a size effect: The curve for the small firm is upward sloping, while the opposite is true for the large firm. This is the result of modeling operating leverage as a proportional cost, which produces the same effect of debt (without requiring further complications). The intuition is as follow: the implied volatility (at any strike) represent the volatility of the log equity price. Assume that the firm is liquidated at date 1, after the shock is realized. Then future equity value is simply the value of the cash flow realized at that point in time: $\pi_1 = e^{z_1}k^{\alpha} - fk$. This function is always convex relative to z_1. However, the convexity of $\ln(\pi_1)$ depends on how large k is (for a fixed α and f): it becomes concave for a sufficiently large k. Adding back the continuation value, so that the firm is not liquidated at time 1, makes the concavity of the value function even more parameter dependent. The convexity/concavity of the log equity price function is what determines the slope of the implied volatility curve. We can see this by comparing the yellow lines in the top and bottom right most panels of Figure 6, where as the log equity curve becomes more concave/less convex the slope of implied volatility curve becomes less positive/more negative. The leverage effect is a function of size, which is largely in line with the empirical evidence presented in Table 1, and with the existing literature.

As mentioned above, whether it implies increasing or decreasing the capital stock, the real option always adds (positive) skewness. Focus on the central panels of Figure 6. Consider first the case without adjustment costs (blue line): for the small firm (top), relative to the firm that does not invest, the ability to increase productive capital increases equity returns in good states and decreases them during bad times. However the effect is more prominent in good times, so that the net result is to shift the distribution to the right. Conversely, for the large firm (bottom), the ability to downsize increases equity returns during bad realization of the profitability shock, and decreases them in good times. The effect is much more pronounced in negative states of nature: the result is again to push more mass to the right. The resulting implied volatility curves exhibit an upward slope. Adding adjustment costs reduces future investment and disinvestments and their contribution to the equity value: since costs are asymmetric, downsizing produces a more pronounced effect. As they constrain the optimal policy, adjustment costs also dampen the skewness in the equity value function and thus produce a less skewed equity return distribution and a lower positively slopped implied volatility curve. The larger the costs, the less skewed the distribution, the lower the expected present value of the real option value. PVRO is the difference between the equity value, at time 1, when the firm can adjust the capital ratio and the corresponding value when it cannot, scaled by the current capital stock: in expected value terms, this is the difference between the respective *current* equity values scaled by the current capital stock. In other words, this is the component of the market-to-book ratio that captures only the value of the real option, controlling away for how leveraged the firm is and how large its current capital stock is.

As Table 1 shows, controlling for leverage, larger firms and firms with high market-to-book ratios tend to have more negative skewed risk-neutral distributions. We have showed above how the size effects works. We focus here on the negative relation between PVRO and RN-skewness/slope of the IV curve. To illustrate how one could obtain such relation, we compare firms of different initial productive capital, k_0, and different future prospects, which depends on the initial state x_0 (as the profitability shock is mean-reverting). Because the production function exhibits decreasing returns to scale, small firms, everything else equal, will invest more in the future when economic conditions are good, while large firms will have an incentive to reduce capital in place during bad times. Moreover, because productivity shocks are persistent, firms that face negative productivity shocks at time 0, x_0, face worse future conditions than firms that are exposed to positive shocks at time 0.

The impact of these two effects on the relation between PVRO and the shape of the future equity distribution depends on the magnitude of the adjustment costs parameters, which alter the future investment policy as shown in Figure 6. We provide here an example of how a particular choice of parameters can produce the economic forces necessary to reproduce, qualitatively at least, the relationships described in Table 1. We plot investments and implied volatility curves in Panel A and B of Figure 7 for different levels of k_0 and of x_0.

Consider first Panel A. The small firm (top row) mostly invests, and it does so even when future economic conditions are not positive. For example, consider the future scenarios that stems from the current state $x_0 = 0$: even when $x_1 < 0$, as long as it is not too negative, the optimal decision is to increase the capital stock. The large firm (bottom row) mostly disinvests, and does so even when future shocks are positive. In contrast the intermediate firm has a much more mixed optimal policy: it invests when the future shocks are positive and disinvests when they are negative (we chose k_0 in order to produce this precise split at $x_1 = 0$).

Both investment and disinvestment options are valuable. However, the extent to which they contribute to the current value of the equity depends on their “moneyness”. We can think of the small firm has currently having a deep in-the-money investment option and an out-of-the-money disinvestment option, since it will mostly invest in period 1, regardless of future economic conditions. Similarly, the large firm has an out of the money investment option and an in-the-money disinvestment option, as it mostly disinvest. In both cases, the value of the bundle is driven by the option that is probably going to be exercised. We can instead think of the intermediate firm has having both an option to invest and one to disinvest, that are almost exactly at-the-money, and thus on average produce a small expected investment that does not move the value of the firm away from the value of the constrained firm. Since, one can think about the real option as a combination of a call (invest) and a put (disinvest), the value of the combination (i.e., a straddle) depends on the absolute value of the future investment. We can see the value being very large when either the average investment or the average disinvestment are very large in the top right and bottom left plots, respectively. The current value of the real option is instead lower, when the expected investment is small (central panel).

The real option affects the equity distribution not only through how much value it produces, but also by moving value across different states of nature. As Table 2 shows, the present expected value of PVRO is not only higher when firms have more defined investment policies (current situations that lead them to either almost always invest or almost always disinvest in the future), but it is also more volatile in the same situations. For example, when the small firm currently has just observed a very good productivity shock (i.e., $k_0 = 2,\ x_0 = 0.5$), the expected value of PVRO is large at 0.76, but so is its volatility, at 0.3.

Table 2: Distributional properties of PVRO

The table reports expected value, standard deviation, and skewness of PVRO, the present value of the real option, defined as the difference between the equity value when the firm can adjust the capital ratio and the corresponding value when it cannot, scaled by the current capital stock. We vary the initial capital stock $k_0 \in \{2, 3.2, 6\}$, and the productivity shock $x_0 \in \{-0.5, 0, 0.5\}$. We fix the rest of the parameters as follows: $\rho = 0.9$, $\sigma = 20\%$, $f = 0.25$, $\theta_1 = 2$, $\theta_2 = 10$, $\beta = e^{-0.05}$.

	$x_0 = -0.5$	$x_0 = 0.0$	$x_0 = 0.5$
	Expected value		
$k_0 = 2$	0.07	0.28	0.76
$k_0 = 3.2$	0.02	0.01	0.12
$k_0 = 6$	0.17	0.10	0.03
	Standard deviation		
$k_0 = 2$	0.06	0.15	0.30
$k_0 = 3.2$	0.01	0.02	0.10
$k_0 = 6$	0.04	0.03	0.02
	Skewness		
$k_0 = 2$	1.62	1.10	0.92
$k_0 = 3.2$	0.96	4.13	1.64
$k_0 = 6$	0.02	0.23	0.82

The combine effect is to produce a positive skewness of 0.92. On the opposite end, when the same small firm has just observed a negative shock, the expected value of PVRO is low and so is the volatility, but the probabilistic reallocation of value from negative to positive states of the world is larger, with a positive skewness of 1.62. Looking at the intermediate firm, we see that the highest PVRO skewness corresponds to the case $x_0 = 0$, when the investment policy is split exactly in half: half the times it invest and half the times it disinvests. Thus, for this particular choice of parameters and distribution of initial capital, there is a marked negative correlation between the expected value of PVRO and its skewness. We learn that, controlling for leverage and size, the real options add (positive) skewness to the distribution of the equity value, but they do so more prominently when there is more curvature in the real option function. To create a parallel to the Black-Scholes model, the real option can be thought as a straddle, a combination of a call and a put. The curvature of the each of the two functions (i.e., gamma) is highest when the options are exactly ATM.

Figure 7: Investments and implied volatilities

The figure plots investments at time 1 and implied volatility curves at time 0 derived from options that matures at time 1. We vary the initial capital stock $k_0 \in \{2, 3.2, 6\}$, and the productivity shock $x_0 \in \{-0.5, 0, 0.5\}$. We fix the rest of the parameters as follows: $\rho = 0.9$, $\sigma = 20\%$, $f = 0.25$, $\theta_1 = 2$, $\theta_2 = 10$, $\beta = e^{-0.05}$.

Panel A: Investments at $t = 1$

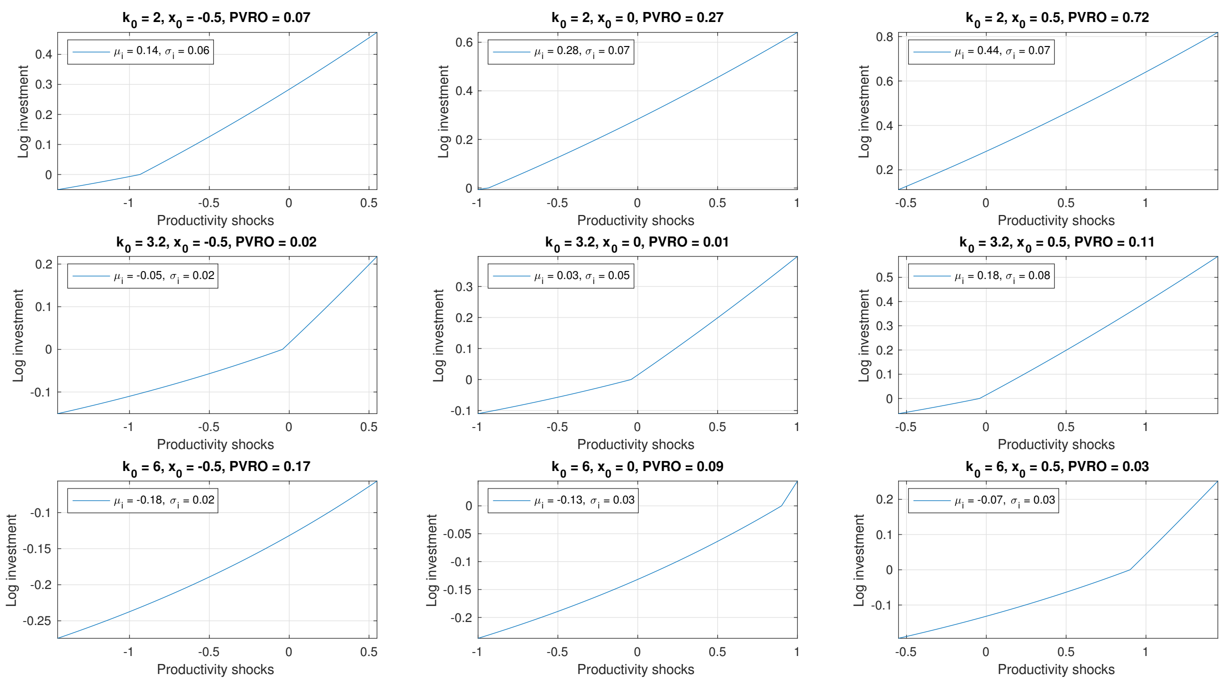

Panel B: Implied volatilities at $t = 0$

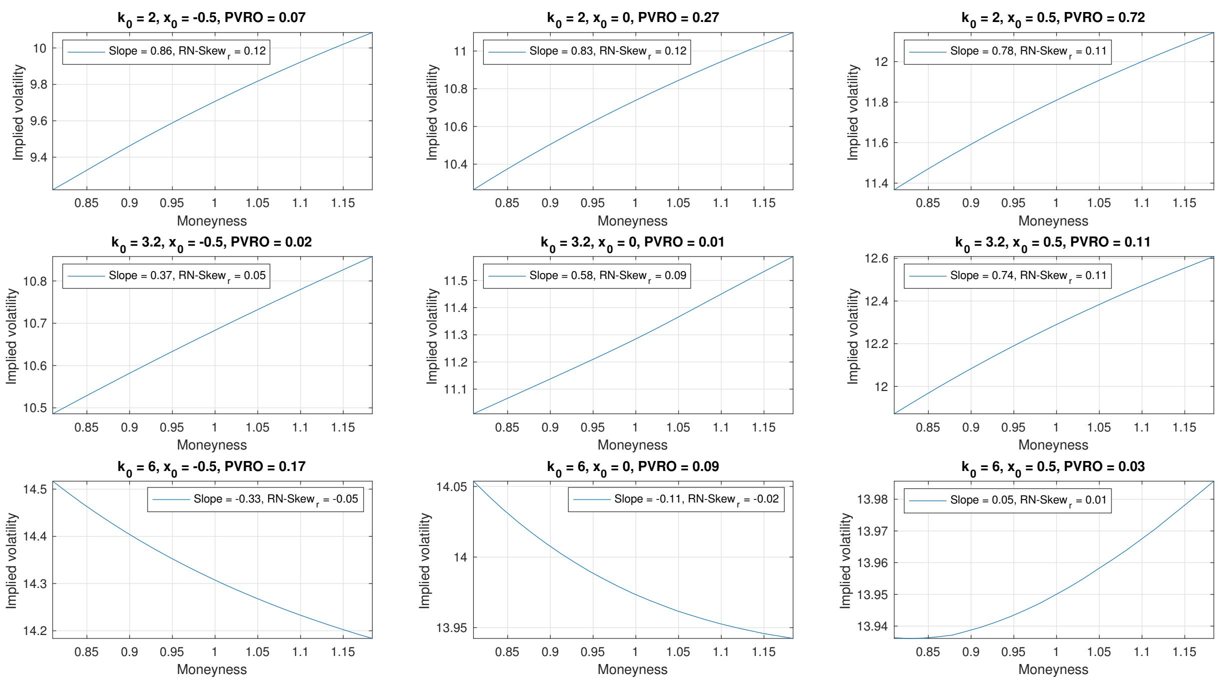

We now turn to Panel B of Figure 7. There are two economic forces affecting the shape of the implied volatility curve: leverage and the real option. Leverage induces negative skewness and a negative slope of the IV curve, the real option does the exact opposite. However, the impact of the real option on skewness is inversely related to how "in-the-money" that option is (the real option is a combination of a call and a put). Moreover, the effect of leverage is decreasing with economic conditions: when prospects are good, leverage induces less negative skewness. As the figure shows, combining those two effects together leads to a situation where, for small and large firms, there is a negative relation between the slope of the IV curve (or the risk-neutral skewness) and the PVRO. The same relation is u-shaped for the intermediate firm.

While the qualitative impact of the two effects is unchanged by parameter choices (i.e., leverage always induces negative skewness and the real option always induces positive skewness, especially when the option value is the most convex), the way they combine together to create decreasing or increasing implied volatility curves is state (in the simple model that is k_0 and z_0) and also parameter dependent. It is therefore reasonable to ask whether the same net effect can be obtained in the contest of a model that is calibrated to match other features of the data. We present an extended model, its calibration, and the relationship between properties of the risk-neutral distribution and firm characteristics within the simulated economy, in the next two sections.

6. Quantitative model

While the basic intuition developed in the previous section confirms the basic patterns in the data, whether a model of endogenous real options can give a quantitative description of the data is a question that can only be answered by calibrating a model that includes more realistic features such as; heterogeneity, endogenous default, corporate taxes, real adjustment costs, external equity financing frictions, debt adjustment costs, and considers countercyclical risk premia. The model is therefore similar, in spirit, to that of Hennessy and Whited (2007) in the description of the firm's decisions, and to those of Berk, Green, and Naik (1999), Zhang (2005) and Gomes and Schmid (2010) in the choice of a reasonably simple (exogenously specified) pricing kernel.

6.1. The economy

Information is revealed and decisions are made at a set of discrete dates $\{0, 1, \ldots, t, \ldots\}$. The time horizon is infinite. The economy is composed by a utility maximizing representative agent and a fixed number of heterogenous firms ($j = 1, \ldots, J$) that produce the same good. Firms make dynamic investment and financing decisions and are allowed to default on their obligations. Defaulted firms are restructured and then continue operations, so as to guarantee a constant number of firms in the economy. The agent consumes the dividends paid by the firms and saves by investing in the financial market. We do not close the economy and derive the equilibrium, but instead choose an exogenously specified stochastic discount factor.

There are two sources of risk that capture variation in the firm's productivity. The first, z_j, captures variations in productivity caused by firms' specific events. Idiosyncratic shocks are independent across firms, and have a common transition function $Q_z(z_j, z_j')$. z_j denotes the current (or time–t) value of the variable, and z_j' denotes the next period (or time–$(t+1)$) value.

The second source of risk, y, captures variations in productivity caused by macroeconomics events. The aggregate risk is independent of the idiosyncratic shocks and has transition function $Q_y(y, y')$. Q_z and Q_y are stationary and monotonic Markov transition functions that satisfy the Feller property. z and y have compact support. For convenience of exposition, we define the state variable $x = (y, z)$, whose transition function, $Q(x, x')$, is the product of Q_y and Q_z. As there is no risk of confusion, we drop the index j in the rest of the section.

6.2. Firm policies

We assume that firm's decisions are made to maximize shareholders' value. An intuitive description of the chronology of the firm's decision problem is presented in Figure 8. At t, the two shocks $x = (y, z)$ are realized and the firm cash flow is determined based on current capital stock, k, and total face value of debt, b. Immediately after that, the firm simultaneously chooses the new set of capital, k', and debt, b' for the period $]t, t+1]$. This decision determines P, the payout to shareholders, which can be positive (dividends and/or share repurchases) or negative (an injection of equity capital by issuing new shares).

Figure 8: Model time line

This figure offers a description of the chronology of the firm's recursive decision problem. At t, the shocks $x = (y, z)$ are realized, and the firm's cash flow is determined based on the capital stock k and the debt b, or $a = (k, b)$. Immediately after t, the firm chooses the new set of capital and debt, as the combination $a' = (k', b')$ that maximizes the value of the equity, given by the sum of the current cash flow plus the continuation value.

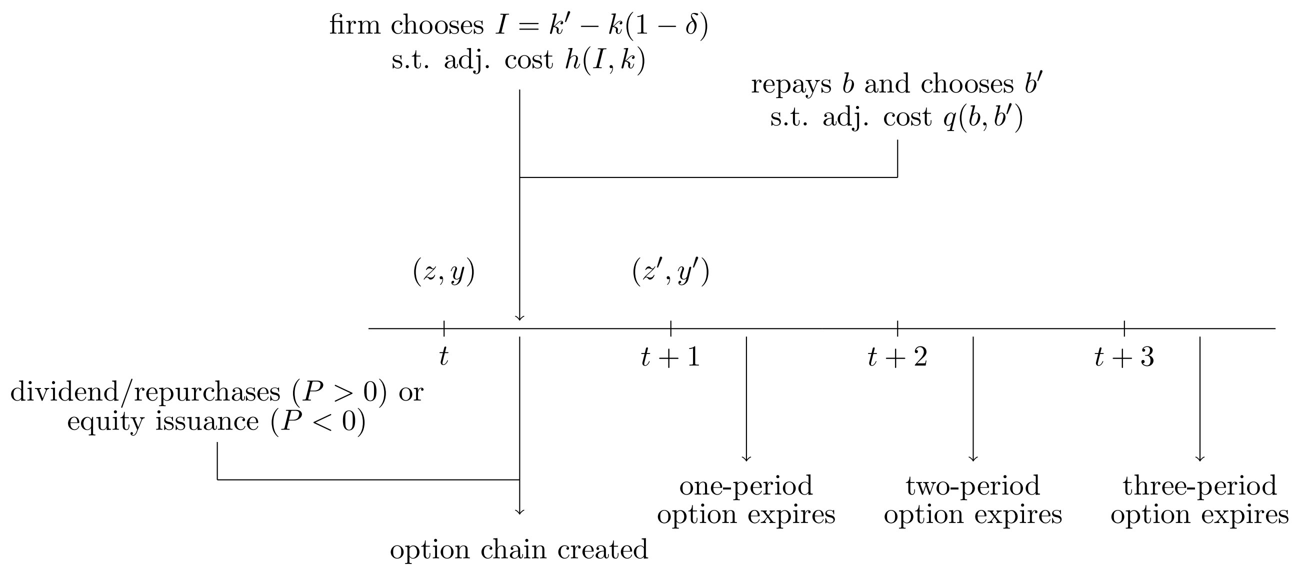

At t, the cash flow from operations (EBITDA) depends on the idiosyncratic and aggregate shocks, and on the current level of asset in place, $\pi = \pi(y, z, k) = e^{y+z}k^{\alpha} - f$, where $\alpha < 1$ models decreasing returns to scale and $f \geq 0$ is a operating cost parameter that summarizes all operating expenses excluding interest on debt.

The capital stock of the firm might change over time. The asset depreciates both economically and for accounting purposes at a constant rate $\delta > 0$. After observing the realization of the shocks at time t, the firm chooses the new capital stock k', which will be in operation during the period $]t, t+1]$. The firm can either increase or decrease the capital stock, and the net investment equals to $I = k' - k(1-\delta)$. Similar to Abel and Eberly (1994) and many others after them, we assume that the change in capital entails an asymmetric and quadratic adjustment cost $h(I, k) = \left(\lambda_1 \mathbf{1}_{\{I>0\}} + \lambda_2 \mathbf{1}_{\{I<0\}}\right) I^2/\delta k$, where $0 < \lambda_1 < \lambda_2$ model costly reversibility, and $\mathbf{1}_{\{.\}}$ is the indicator function. The economic interpretation of λ_i, $i = 1, 2$, is straightforward: it is the per cent cost of a (dis)investment $I = \delta k$.

The debt level might also change over time. At any date, the firm can issue a one–period zero–coupon unsecured debt. As is shown in Figure 8, at time t the firm chooses the face value of the debt, b', that will be repaid at $t+1$. If the firm is solvent, the market value of the debt, $B(x, a')$, depends on the current state x and on the choices of the face value and the capital stock, $a' = (k', b')$, that are made after observing the shocks.

Changing the debt level entails a proportional adjustment cost, $\theta|b' - b|$, with $\theta \geq 0$. Since the issuance decision is contemporaneous to repayment of the nominal value of old debt b, the debt decision generates a net cash flow equal to $B(x, a') - b - \theta|b' - b|$.

We assume a linear corporate tax function with rate τ. The tax code allows deduction from the taxable income of the depreciation of assets in place, δk, and of interest expenses. Modeling deduction of the interest at maturity of the bond would entail keeping track of the value of the debt at issuance, therefore increasing the number of state variables. For the sake of numerical tractability, we assume that the expected present value of the end-of-period interest payment $b' - B(x, a')$, which we denote $H(x, a')$, can be expensed when the new debt is issued at time t. In case of linear corporate tax, and assuming knowledge of the equilibrium conditional default probability, this is equivalent to the standard case of deduction at $t+1$. The after–tax cash flow from operations plus the net proceeds from the debt decision is

$$v = v(x, a, a') = (1-\tau)\pi + \tau\delta k + \tau H(x, a') + B(x, a') - b - \theta|b' - b|. \qquad (1)$$

The cash flow to equity is therefore equal to $w = w(x, a, a') = v - I - h(I, k)$ where, on the right-hand side, the first term is the after–tax cash flow from operations and the other terms are the net proceeds from (dis)investment. If the cash flow to equity is positive, the firm pays dividends and/or repurchases shares from the current shareholders; if the cash flow to equity is negative the firm issues new shares. In the latter case, the company incurs a proportional issuance cost $\zeta \geq 0$, as only w is the actual inflow to the corporation

$$P = P(x, a, a') = w \cdot (1 + \zeta \mathbf{1}_{\{w<0\}}). \tag{2}$$

6.3. The value of corporate securities

Following Berk, Green, and Naik (1999), Zhang (2005), and Gomes and Schmid (2010), we exogenously define a pricing kernel that depends on the aggregate source of risk, y. The associated one–period stochastic discount factor $M(x, x')$ defines the risk-adjustment corresponding to a transition from the current state y to state y'. We assume that M is a continuous function of both arguments.

The firm can issue two types of securities, debt and equity, whose equilibrium prices are determined under rational expectations in a competitive market. The cum–dividend price of equity, $S(x, a)$, is the sum of current payout, P, and the present value of the expected future optimal distributions, which is equal to the next period price $S(x', a')$. Since this sum can be negative, a limited liability provision is also included (i.e., default on a value basis), in which case the firm's equity is worthless:

$$S(x, a) = \max\left\{0, \max_{a'} \left\{P(x, a, a') + \mathbb{E}_x \left[M(x, x')S(x', a')\right]\right\}\right\}. \tag{3}$$

The value function, S, is the solution of functional equation (3). We define $\omega = \omega(x, a)$ as in indicator function that captures the event of default. Note that, if $\omega = 0$, the optimal investment and financing decision is $\varphi(x, a) = a^*$, where $a^* = (k^*, b^*)$ is the optimal choice of the second argument in the max in (3). The optimal policy is therefore summarized by (ω, φ).

As for the debt contract, the end-of-period payoff to debt holders, $u(x', a')$, depends on the current policy, $a' = (k', b')$, the new realization of the shocks x', and on whether the firm is in default:

$$u(x', a') = b'(1 - \omega(x', a')) + \left[\pi' + \tau\delta k' + k'(1 - \delta)\right](1 - \eta)\omega(x', a'). \tag{4}$$

In case of default, similarly to Hennessy and Whited (2007), the bondholders receive the sum of the cash flow from operations, the depreciated book value of the asset, and the tax shield from depreciation, all net of a proportional bankruptcy cost, η. Hence, at issuance the debt value is

$$B(x, a') = \mathbb{E}_x \left[M(x, x') u(x', a') \right]. \tag{5}$$

One final item that needs to be evaluated is the expected present value of the interest payment, $H(x, a')$, which enters the determination of the after tax cash flow in (1):

$$H(x, a') = \left[b' - B(x, a') \right] \mathbb{E}_x \left[M(x, x')(1 - \omega(x', a')) \right]. \tag{6}$$

Because the interest is deductible only if the firm is not in default, the expectation term is the conditional price of a default contingent claim.

6.4. Option prices

We derive the option prices from the stock price, under the assumption that distribution to equity holders do not happen in the form of a cash dividend but are either a share repurchase or an equity issuance (when negative).[2]

Denote with $n(x, a)$ the number of outstanding shares before the current payout decision is made. The stock price of one share is

$$s(x, a) = \frac{S(x, a)}{n(x, a)}.$$

Define $S'(x, a) = S(x, a) - P(x, a, a')$ the equity value after the payout, where $a' = \varphi(x, a)$ is the optimal policy from (3).

After a payout, the firm changes the number of shares for next period to $n'(x, a)$. In particular, if $P(x, a, a') > 0$, some shares are repurchased; if $P(x, a, a') < 0$ new share are issued. The new number of shares is

$$n'(x, a) = \frac{S'(x, a)}{s(x, a)} = \frac{S'(x, a)}{S(x, a)}\, n(x, a). \tag{7}$$

While n and n' are integer numbers in real life, we assume here that $n, n' \in \mathbb{R}$.

[2] It is possible to solve the model and compute prices even when the firm pays an exogenous dividend. In that case, we are also able to price an American option.

The evolution of the number of shares is given by the application of the current optimal policy, $a' = \varphi(x, a)$, and the state transition from x to x', so that at the new state (x', a') following from (x, a),

$$n(x', a') = n'(x, \varphi(x, a)), \tag{8}$$

with $n'(x, a)$ from (7).

We assume options are on a single share of equity. For definiteness, we consider a European call option with strike κ, with payoff at maturity $\max\{s(x, a) - \kappa, 0\}$, which is based on the convention that the dividend has been paid before the option expires, and therefore the payoff is based on the ex dividend price.

Because the shares number is endogenous (i.e., it depends on the payout policy), option pricing by straightforward backward induction is numerically intractable. The drawback introduced by path dependency is due to the fact that the option price at the current state, and the stock price $s(x, a)$, is

$$c(x, a; \kappa) = \mathbb{E}_x \left[M(x, x') \max\{s(x', a') - \kappa, 0\} \right],$$

in which $a' = \varphi(x, a)$. To determine $s(x', a')$, the underlying asset of the option in state (x', a'), from $S(x', a')$ we need $n(x', a')$. However, as one can see from equation (8), $c(x', a'; \kappa)$ also depends on $n(x, a)$.

We avoid the issue of path dependency by observing that

$$\begin{aligned} c(x, a; \kappa) &= \mathbb{E}_x \left[M(x, x') \max \left\{ \frac{S(x', a')}{n(x', a')} - \kappa, 0 \right\} \right] \\ &= \frac{1}{n(x', a')} \mathbb{E}_x \left[M(x, x') \max \left\{ S(x', a') - \kappa \; n(x', a'), 0 \right\} \right]. \end{aligned}$$

From the expression above, defining the sum of prices of all options with strike k written on the firm's stock, $C(x, a; \kappa \; n(x', a')) = c(x, a; \kappa) n(x', a')$, we can write

$$C(x, a; \kappa \; n(x', a')) = \mathbb{E}_x \left[M(x, x') \max\{S(x', a') - \kappa \; n(x', a'), 0)\} \right],$$

which shows that we use backward induction to price total equity options on a predetermined set of strike prices $\mathcal{K} = \{K_1, \ldots, K_N\}$, such that for each $K \in \mathcal{K}$ we solve

$$C(x, a; K) = \mathbb{E}_x \left[M(x, x') \max\{S(x', a') - K, 0\} \right],$$

working backward from the option maturity to the current period. Given these prices, we can determine the current price of a European call option with strike price κ, by interpolating $\widehat{C}(x,a;\kappa\ n'(x,a))$ on the grid $\mathcal{K}$, and then

$$\widehat{c}(x,a;\kappa) = \frac{1}{n'(x,a)}\ \widehat{C}(x,a;\kappa\ n'(x,a)).$$

Using (7), the previous equation becomes

$$\widehat{c}(x,a;\kappa) = \frac{1}{n(x,a)}\frac{S^{ex}(x,a)}{S'(x,a)}\ \widehat{C}\left(x,a;\kappa\ n(x,a)\frac{S'(x,a)}{S^{ex}(x,a)}\right). \tag{9}$$

Given the current equity value, $S(x,a)$, our goal is to calculate the price of options on equity value at $t=0$ with maturity T and moneyness $m \in \{m_1, m_2, \ldots, m_N\}$. Where the strikes are $\mathcal{K} = \{S(x,a)m_i, i=1,\ldots,N\}$. Because the current number of shares is arbitrary, we choose $n(x,a) = S(x,a)$, which is equivalent to assuming that the current (ex dividend) stock price is \$1. Then our goal is met by solving the pricing problem

$$\widehat{c}(x,a;m) = \frac{1}{S'(x,a)}\ \widehat{C}\left(x,a;m\ S'(x,a)\right),$$

where $\widehat{c}(x,a;m)$ is the price of an European call option on a stock with current price \$1 and strike m.

6.5. Stochastic discount factor

We assume that the idiosyncratic shock z and the aggregate shock, y, follow auto-regressive processes of first order, $z' = (1-\rho_z)\overline{z} + \rho_z z + \sigma_z \varepsilon'_z$ and $y' = (1-\rho_y)\overline{y} + \rho_y y + \sigma_y \varepsilon'_y$, respectively. In the above equations, for $i = y, z$, $|\rho_i| < 1$ and ε_i are i.i.d. and obtained from a truncated standard normal distribution, so that the actual support is compact around the unconditional average. We assume that ε_z are uncorrelated across firms and time and are also uncorrelated with the aggregate shock, ε_y. The parameters ρ_z, σ_z, and $\overline{z}$ are the same for all the firms in the economy, $\overline{z}$ and $\overline{y}$ denote the long term mean of idiosyncratic risk and of macroeconomic risk, respectively, $(1-\rho_i)$ is the speed of mean reversion, and σ_i is the conditional standard deviation. With this specification, the transition function Q satisfies all the assumptions required for the existence of the value function.

Finally, we adopt the stochastic discount factor proposed by Jones and Tuzel (2013):

$$M(y,y') = \beta e^{-g(y)\varepsilon'_y - \frac{1}{2}g(y)^2\sigma_y^2},$$

with $\beta \in (0,1)$, and where the state-dependent coefficient of risk-aversion is $g(y) = \exp(\gamma_1 + \gamma_2 y)$, with $\gamma_1 > 0$ and $\gamma_2 < 1$. With this choice, the coupon is equal to the state-independent real risk-free rate, $r = 1/\beta - 1$.

Following the literature, the aggregate risk parameters are taken from Cooley and Prescott (1995) and converted to quarterly frequency. We obtain a value for the persistence of the systematic risk (ρ_x) and the aggregate volatility (σ_x) of 0.979 and 0.0072, respectively. The personal discount factor (β) is set to 0.9851, and the SDF parameters (γ_1 and γ_2) to 3.22 and -15.3, respectively. These parameters produce an annualized average real interest rate of 6.1%.

6.6. Calibration

We fix the five parameters that that describe the aggregate source of risk and the SDF, equity and debt floatation costs, and the depreciation rate (as for example, Warusawitharana and Whited, 2016). We calibrate the remaining parameters by minimizing the sum of square deviations of a set of quantities that are observable in the data and in the simulated economy.

Important objectives of the calibration exercise are that the model captures the outcomes of the decisions that firms make and that affect the relationship between the asset and the equity volatility. The model should therefore match the average (book and market) leverage ratio and the average investment as the real economy. As the relevant sources of total risk match up with the economy, firms should exhibit similar market to book ratios, and similar equity distributions in the physical measure (i.e., average, standard deviation, skewness and kurtosis of equity returns). We also calibrate the model to fit the average ATM 90 days implied volatility, as well the frequency of each implied volatility surface (i.e., left smirk, smile, right smirk, frown).

We report parameter values and quantities used for calibration in Panel A of Table 3. The firm-specific productivity shock is less persistent (0.91 versus 0.98) and more volatile (0.19 versus 0.01) than the aggregate shock. The estimated marginal corporate tax rate, τ, is 0.120, close to the estimates produced by Graham (1996a) and Graham (1996b) (i.e., average of approximately 13% for our sample). The estimate for the production function parameter α is 0.56. There are large bounds around figures reported in the literature, which are largely affected by the frequency at which models are calibrated and what type of fixed costs (proportional or not) are considered. Our value is close to the 0.3 figure used in Zhang (2005) and Gomes (2001). We estimate the operating cost to 4.32 (unit of capital), which translates to an annualized value of approximately 35% of the average capital. The calibrated value of the bankruptcy cost parameter, η, is 0.284, which is in the range of the

Table 3: Model calibration

This table presents the calibration results of the firm model. In Panel A, we report the list of model parameters. In Panel B, we compare the quantities that are weighted 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 (*Model*) we report the moment conditions computed from the simulated sample. Data is from various sources and spans the period between January 1996 throughout December 2019.

Panel A: Parameters		
Aggregate		
Systematic Productivity Autocorrelation	ρ_x	0.970
Systematic Productivity Volatility	σ_x	0.013
Discount Factor	β	0.985
Constant Price of Risk Parameter	g_0	3.220
Time-varying Price of Risk Parameter	g_1	-15.300
Firm Specific		
Depreciation	δ	0.050
Equity Issuance Cost	ζ	0.018
Debt Adjustment Cost	θ	0.009
Idiosyncratic Productivity Autocorrelation	ρ_y	0.918
Idiosyncratic Productivity Volatility	σ_y	0.197
Production Function	α	0.560
Fix Cost	f	4.327
Cost of Expansion	λ_1	0.272
Cost of Contraction	λ_2	0.820
Corporate Taxes	τ	0.120
Bankruptcy Cost	η	0.284

firm's average default costs estimated by Glover (2016). There is not a direct benchmark for the two capital adjustment costs.

Panel B: Calibrated quantities		
Option Prices (90 days to maturity):	Data	Simulation
Implied Volatility OTM	0.648	0.647
Implied Volatility ATM	0.609	0.603
Implied Volatility ITM	0.585	0.589
Percentage Left Smirk	0.703	0.723
Percentage Smile	0.193	0.239
Percentage Right Smirk	0.044	0.001
Percentage Frown	0.062	0.037
Stock Return:		
Average	0.025	0.037
Standard Deviation	0.345	0.363
Skewness	0.649	0.540
Kurtosis	4.167	4.135
Firm characteristics:		
Market-to-Book	2.570	2.022
Leverage	0.504	0.479
Investments	0.043	0.053

In Panel B of Table 3, we compare the simulated economy to the real data along the dimensions used to calibrate the mode. The investment and financing choices of the average simulated firm reflects well those of real firms (investment and leverage are really close). Valuations are also appropriately close, as well the physical distribution of percentage equity returns. Average option prices are also relatively well matched as is the frequency of implied volatility shapes: the average implied volatility curve at 90 days ATM is close the the equivalent in the data. Moreover, the model can create enough heterogeneity in the IV curve shapes that it matches very closed what observed in the data: about 70% of the time the curve is downward sloping with smiles, and about 20% of the time it is "smiling". Thus, similarly to Geske (1979) and Toft and Prucyk (1997) who both incorporate leverage, the model can generate average downward sloping curves across moneyness levels. Differently from those other models, our set up can also create other IV surfaces.

7. Comparison of simulated and observed option prices

While it is remarkable that the model can do match the 90 days IV curve and the frequencies of various IV shapes, it is also true that we used those quantities as part of the calibration exercise. In this section, we present comparisons of the simulated economy with the real one along other dimensions.

7.1. Term structure of implied volatilities

We start by comparing the average IV surfaces across all maturities considered (90, 180, and 360 days). Please remember that the model is only calibrated to fit the 90 days curve. Figure 9 juxtaposes the curves extracted from the data (left panel) to those extracted from the simulation. To obtain each curve, we first average across time, then across firms, and eventually across simulations.

Figure 9: 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 industrial firms with options trading on their equity between 1996 and 2019. Data is sampled at quarterly frequency. A total of 3,536 firms are included.

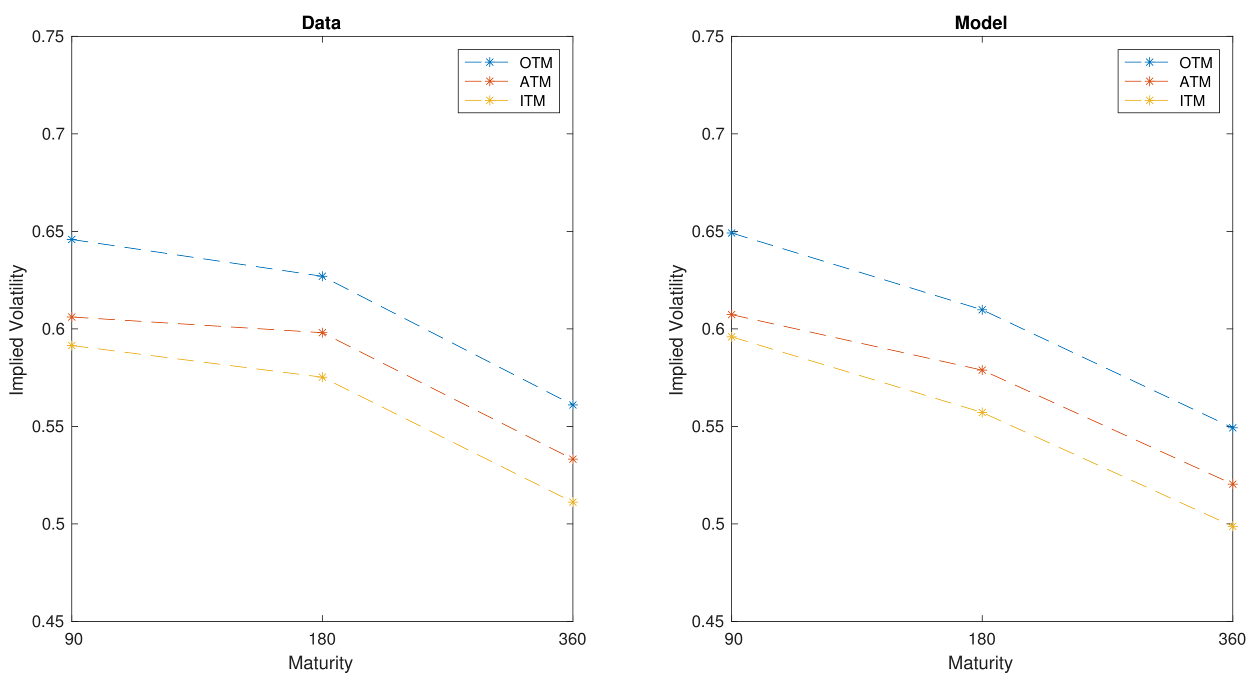

Interestingly, the model can also generate a downward sloping surface across maturities without exogenously imposing a term-structure of volatility. Productivity shocks that affect the firm's value at short horizon tend to revert towards long run values, and as that happens the relationship between asset and equity volatility flattens. The total effect is to decrease prices for options at longer maturities, and henceforth producing a decreasing volatility surface. As Figure 9 shows, the model is able to replicate this feature of the data quite well.

7.2. Risk-neutral moments across different maturities

As Panel B of Table 3 shows, the moments of the physical distribution of stock returns match quite well with the corresponding quantities in the data. Table 4 confirms that the implied higher moments of the risk-neutral distribution match as well. While skewness is relatively flat across maturities (i.e., slowly decreasing in the data and moderately increasing in the model), the model can replicate the downward sloping feature of implied kurtosis, almost perfectly. As one might expect, there is more heterogeneity in the data, as evidence but larger standard deviations. Nonetheless the ranges of the variables compare quite favorably.

Table 4: Model free risk-neutral moments

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

	Data				Model			
	90 days maturity							
	Aver.	S.Dev	5^{th} perc	95^{th} perc	Aver.	St.Dev	5^{th} perc	95^{th} perc
Skewness	-0.51	0.38	-1.15	0.01	-0.45	0.15	-0.73	-0.31
Kurtosis	3.76	1.28	2.66	5.84	3.79	0.46	2.76	4.35
	180 days maturity							
	Aver.	S.Dev	5^{th} perc	95^{th} perc	Aver.	St.Dev	5^{th} perc	95^{th} perc
Skewness	-0.42	0.35	-1.01	0.08	-0.48	0.19	-0.71	-0.27
Kurtosis	3.32	0.93	2.28	4.83	3.44	0.59	1.94	3.94
	360 days maturity							
	Aver.	S.Dev	5^{th} perc	95^{th} perc	Aver.	St.Dev	5^{th} perc	95^{th} perc
Skewness	-0.41	0.39	-1.03	0.20	-0.50	0.28	-0.73	0.03
Kurtosis	3.10	0.96	1.94	4.75	3.11	0.74	1.44	3.84

7.3. Cross-sectional regressions

As the firm parameters are determined by the calibration exercised, variability in the simulated economy in terms of implied volatility shapes is dictated by the optimal choices made by the firm relative to the realizations of the exogenous variables and the current state of capital and debt. Ultimately those choices determine the equity value relative to the capital in place and optimal amount of leverage. We estimate here the same linearized relationships that we presented in Section 4.3.

Table 5: Option prices and firm characteristics

The table presents regression results in the simulated economy, that mirrors those in the data presented in Table 1. Left hand side variables include the risk-neutral skewness, the slope (i.e., log difference of IV(1.2) and IV(0.8)), and the left slope (i.e., the log difference of IV(0.8)-IV(1)) of 90 day maturity options. Dependent 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 option to assets (PVRO2A). PVRO is constructed by decomposing the equity into the value derived from keeping current capital level constant and the value derived from making adjustments, by either investing or disinvesting. Since fixed effects are not very meaningful in simulated data, reported coefficient are obtained from a simulated Fama-MacBeth regression, where slopes obtained from cross-sectional regressions (all firms observations in one period in one simulated economy) are averaged first through time and then through simulated economies. Standard errors are obtained from considering deviations around the mean across simulated economies.

	RN-Skew		Slope		Left slope	
IV	0.41	0.41	-0.17	-0.24	-2.22	-2.09
	(115.15)	(119.50)	(-9.09)	(-13.69)	(-85.53)	(-85.82)
Size	0.01	0.00	-0.83	-0.92	-0.37	-0.32
	(34.73)	(5.35)	(-71.93)	(-73.73)	(-32.72)	(-29.45)
Book Leverage	-0.03	-0.03	-0.20	-0.47	0.66	0.99
	(-42.51)	(-116.96)	(-17.44)	(-55.80)	(63.51)	(70.30)
Profitability	0.07	0.05	2.29	1.93	-4.04	-4.15
	(16.30)	(16.01)	(20.47)	(25.27)	(-119.62)	(-158.29)
M2B	-0.07		-3.33		3.47	
	(-12.22)		(-33.40)		(77.62)	
PVRO2A		-0.05		-2.00		2.33
		(-10.91)		(-59.89)		(62.26)

The results reported in Table 5 largely mirror those reported in Table 1. In the simulated economy, the skewness of the risk-neutral distribution and the slope of the IV curve are negatively related to leverage, as in Toft and Prucyk (1997), and to the value of the real option, whether that is proxied by the market to book ratio, or measured exactly by the present value of the real option, which we obtain in the model by separating the contribution to the equity value of assets in place and of future investments and disinvestments.

8. Conclusions

Traditional option pricing models often requires very strong assumptions about investor preferences and the dynamic of equity prices. We show that equity options can be priced in a production economy where we do not make strong exogenous assumptions about equity and volatility. In our set up the relation between risk and value arises endogenously through a dynamic sequence of optimal decisions that maximize the value of the firm. We derive option prices that match many properties of those observed in the cross-section of US equities starting from a different set of assumptions that specify the functional forms of corporate trade-offs.

Our approach is not a better option pricing model, but rather an attempt to provide a link between fundamentals and derivative pricing. We think that such link is important as it relates the primitives of the most successful finance models (i.e., those that price financial derivatives) to a large body of well understood economic mechanisms that describe the decision-making process within a typical firm.

Ultimately, we hope to provide an explanation for why option prices contain forward looking information about stock prices and corporate policies, despite being classically derived in models where such links should be uninformative unless one assumes some form of market segmentation.

References

Abel, Andrew B., and Janice C. Eberly, 1994, A unified model of investment under uncertainty, *American Economic Review* 85, 1369–1384 https://doi.org/10.3386/w4296.

Aguerrevere, Felipe, 2009, Real options, product market competition, and asset returns, *Journal of Finance* 44, 957–983 https://doi.org/10.1111/j.1540-6261.2009.01454.x.

Ai, Hengjie, and Dana Kiku, 2013, Growth to value: Option exercise and the cross section of equity returns, *Journal of Financial Economics* 107, 325–349 https://doi.org/10.1016/j.jfineco.2012.08.009.

Amaya, Diego, Peter Christoffersen, Kris Jacobs, and Aurelio Vasquez, 2015, Does realized skewness predict the cross-section of equity returns?, *Journal of Financial Economics* 118, 135 – 167 https://doi.org/10.1016/j.jfineco.2015.02.009.

Aretz, Kevin, and Peter Pope, 2018, Real options models of the firm, capacity overhang, and the cross section of stock returns, *Journal of Finance* 73, 1363–1415 https://doi.org/10.1111/jofi.12617.

Bai, Jennie, Robert Goldstein, and Fan Yang, 2019, The leverage effect and the basket-index spread, *Journal of Financial Economics* 131, 180–205 https://doi.org/10.1016/j.jfineco.2018.07.015.

Bakshi, Gurdip, Charles Cao, and Zhaodong (Ken) Zhong, 2021, Assessing models of individual equity option prices, *Review of Quantitative Finance and Accounting* 57 https://doi.org/10.1007/s11156-020-00951-4.

Bakshi, G., N. Kapadia, and D. Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *Review of Financial Studies* 16, 637–659 https://doi.org/10.1093/rfs/16.1.0101.

Bali, Turan, and Scott Murray, 2013, Does risk-neutral skewness predict the cross-section of equity option portfolio returns?, *Journal of Financial and Quantitative Analysis* 48, 1145–1171 https://doi.org/10.1017/S0022109013000410.

Bali, Turan G., Luca Del Viva, Neophytos Lambertides, and Lenos Trigeorgis, 2020, Growth options and related stock market anomalies: profitability, distress, lotteryness, and volatility, *Journal of Financial Quantitative Analysis* 55, 2150–2180 https://doi.org/10.1017/S0022109019000619.

Berk, Jonathan B., Richard C. Green, and Vasant Naik, 1999, Optimal investment, growth options, and security returns, *Journal of Finance* 54, 1553–1607 https://doi.org/10.1111/0022-1082.00161.

Black, Fisher, and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–659 https://doi.org/10.1086/260062.

Cao, Charles, Tim Simin, and Jing Zhao, 2008, Can growth options explain the trend in idiosyncratic risk?, *Review of Financial Studies* 21, 2599–2633 https://doi.org/10.1093/rfs/hhl039.

Chaderina, Maria, Patrick Weiss, and Josef Zechner, 2021, The maturity premium, *Journal of Financial Economics* forthcoming https://doi.org/10.1016/j.jfineco.2021.07.008.

Christoffersen, Peter, Mathieu Fournier, Kris Jacobs, and Mehdi Karoui, 2021, Option-based estimation of the price of coskewness and cokurtosis risk, *Journal of Financial and Quantitative Analysis* 56, 65–91 https://doi.org/10.1017/S002210902000023X.

Conrad, Jennifer, Robert Dittmar, and Eric Chisels, 2013, Ex ante skewness and expected stock returns, *Journal of Finance* 68, 85–124 https://doi.org/10.1111/j.1540-6261.2012.01795.x.

Cooley, T. F., and E. C. Prescott, 1995, Economic growth and business cycles, in Thomas F. Cooley, ed.: *Frontiers of Business Cycle Research* . pp. 1–38 (Princeton University Press: Princeton - NJ).

Cooper, Ilan, 2006, Asset pricing implications of non-convex adjustment costs and irreversibility of investment, *The Journal of Finance* 61, 139–170 https://doi.org/10.1111/j.1540-6261.2006.00832.x.

Del Viva, Luca, Eero Kasanen, and Lenos Trigeorgis, 2017, Real options, idiosyncratic skewness, and diversification, *Journal of Financial and Quantitative Analysis* 52, 215241 https://doi.org/10.1017/S0022109016000703.

Dennis, Patrick, and Stewart Mayhew, 2002, Risk-neutral skewness: Evidence from stock options, *Journal of Financial and Quantitative Analysis* 37, 471–493 https://doi.org/10.2307/3594989.

Duan, Jin-Chuan, and Jason Wei, 2009, Systematic risk and the price structure of individual equity options, *Review of Financial Studies* 22, 1981–2006 https://doi.org/10.1093/rfs/hhn057.

Engle, Robert F., and Emil N. Siriwardane, 2017, Structural garch: The volatility-leverage connection, *The Review of Financial Studies* 31, 449–492 https://doi.org/10.1093/rfs/hhx099.

Fama, Eugene Francis, and Kenneth R. French, 1992, The cross-section of expected stock returns, *Journal of Finance* 47, 427–465.

Figlewski, Stephen, and Xiaozu Wang, 2000, Is the "leverage effect" a leverage effect?, Discussion paper New York University.

Friewald, Nils, Florian Nagler, and Christian Wagner, 2021, Debt refinancing and equity returns, *Journal of Finance* forthcoming http://doi.org/10.2139/ssrn.2684462.

Garleanu, Nicolae, Stavros Panaceas, and Jianfeng Yu, 2012, Technological growth and asset pricing, *Journal of Finance* 67, 1265–1292 https://doi.org/10.1111/j.1540-6261.2012.01747.x.

Geske, Robert, 1979, The valuation of compound options, *Journal of Financial Economics* 7, 63–81 https://doi.org/10.1016/0304-405X(79)90022-9.

——— , Avanidhar Subrahmanyam, and Yi Zhou, 2016, Capital structure effects on the prices of equity call options, *Journal of Financial Economics* 121, 231–253 https://doi.org/10.1016/j.jfineco.2016.03.009.

Glover, Brent, 2016, The expected cost of default, *Journal of Financial Economics* 119, 284–299 https://doi.org/10.1016/j.jfineco.2015.09.007.

Gomes, Joao F., 2001, Financing investment, *American Economic Review* 91, 1263–1285 https://doi.org/10.1257/aer.91.5.1263.

——— , and Lukas Schmid, 2010, Levered returns, *Journal of Finance* 65, 467–494 https://doi.org/10.1111/j.1540-6261.2009.01541.x.

——— , 2021, Equilibrium asset pricing with leverage and default, *Journal of Finance* 76, 977–1018 https://doi.org/10.1111/jofi.12987.

Graham, John R., 1996a, Debt and the marginal tax rate, *Journal of Financial Economics* 41, 41–73 https://doi.org/10.1016/0304-405X(95)00857-B.

——— , 1996b, Proxies for the corporate marginal tax rate, *Journal of Financial Economics* 42, 187–221 https://doi.org/10.1016/0304-405X(96)00879-3.

Grullon, Gustavo, Evgeny Lyandres, and Alexey Zhdanov, 2012, Real options, volatility, and stock returns, *Journal of Finance* 62, 1499–1537 https://doi.org/10.1111/j.1540-6261.2012.01754.x.

Gu, Lifeng, Dirk Hackbarth, and Tim Johnson, 2017, Inflexibility and stock returns, *Review of Financial Studies* 31, 278–321 https://doi.org/10.1093/rfs/hhx092.

Hansis, A., C. Schlag, and G. Vilkov, 2010, The dynamics of risk-neutral implied moments: evidence from individual options, Discussion paper Frankfurt School of Finance.

Hennessy, Christopher A., and Tony Whited, 2005, Debt dynamics, *Journal of Finance* 60, 1129–1165 https://doi.org/10.1111/j.1540-6261.2005.00758.x.

Hennessy, Christopher A., and Tony M. Whited, 2007, How costly is external financing? evidence from a structural estimation, *Journal of Finance* 62, 1705–1745 https://doi.org/10.1111/j.1540-6261.2007.01255.x.

Jones, Christopher S., and Selale Tuzel, 2013, Inventory investment and the cost of capital, *Journal of Financial Economics* 107, 557–579 https://doi.org/10.1016/j.jfineco.2012.09.001.

Kadan, Ohad, and Xiaoxiao Tang, 2019, A bound on expected stock returns, *The Review of Financial Studies* 33, 1565–1617 https://doi.org/10.1093/rfs/hhz075.

Kogan, Leonid, 2004a, Asset prices and real investment, *Journal of Financial Economics* 73, 411–431 https://doi.org/10.1016/j.jfineco.2003.12.002.

——— , 2004b, An equilibrium model of irreversible investment, *Journal of Financial Economics* 73, 411–431 https://doi.org/10.1016/S0304-405X(01)00077-0.

——— , and Dimitris Papanikolaou, 2013, Firm characteristics and stock returns: the role of investment-specific shocks, *Review of Financial Studies* 26, 2718–2759 https://doi.org/10.1093/rfs/hht026.

——— , 2014, Growth opportunities, technology shocks, and asset prices, *Journal of Finance* 69, 675–718 https://doi.org/10.1111/jofi.12136.

Kuehn, Lars-Alexander, and Lukas Schmid, 2014, Investment-based corporate bond pricing, *Journal of Finance* 69, 2741–2776 https://doi.org/10.1111/jofi.12204.

Leland, Hayne E., 1994, Corporate debt value, bond covenants, and optimal capital structure, *Journal of Finance* 49, 1213–1252 https://doi.org/10.2307/2329184.

Lyandres, Evgeny Matveyev, Egor, and Alexey Zhdanov, 2020, Does the market correctly value investment options?, *Review of Finance* 62, 1159–1201 https://doi.org/10.1093/rof/rfaa002.

Martin, Ian W. R., and Christian Wagner, 2019, What is the expected return on a stock?, *The Journal of Finance* 74, 1887–1929 https://doi.org/10.1111/jofi.12778.

Merton, Robert, 1974, On the pricing of corporate debt: the risk structure of interest rates, *Journal of Finance* 4, 449–470 https://doi.org/10.2307/2978814.

Morellec, Erwan, and Alexey Zhdanov, 2019, Product market competition and option prices, *Review of Financial Studies* 32, 4343–4386 https://doi.org/10.1093/rfs/hhz027.

Panayiotis, Andreou, Turan G. Bali, Anastasios Kagkadis, and Neophytos Lambertides, 2021, Firms growth potential and option returns, Discussion paper Georgetown University.

Schneider, Paul, Christian Wagner, and Josef Zechner, 2020, Low risk anomalies?, *Journal of Finance* 75, 2673–2718 https://doi.org/10.1111/jofi.12910.

Toft, Klauss Bjerre, and Brian Prucyk, 1997, Options on leverage equity: theory and empirical tests, *Journal of Finance* 52, 1151–1180 https://doi.org/10.2307/2329519.

Trigeorgis, Lenos, and Neophytos Lambertides, 2014, The role of growth options in explaining stock returns, *Journal of Financial and Quantitative Analysis* 49, 749–771 https://doi.org/10.1017/S0022109014000118.

van Zwet, Willem, 1964, Convex transformations: A new approach to skewness and kurtosis, *Statistica Neerlandica* 18, 433–441 https://doi.org/10.1111/j.1467-9574.1964.tb00530.x.

Warusawitharana, Missaka, and Tony Whited, 2016, Equity market misvaluation, financing, and investment, *Review of Financial Studies* 29, 603–634 https://doi.org/10.1093/rfs/hhv066.

Xu, Jianguo, 2007, Price convexity and skewness, *Journal of Finance* 62, 2521–2552 https://doi.org/10.1111/j.1540-6261.2007.01283.x.

Zhang, Lu, 2005, The value premium, *Journal of Finance* 60, 67–103 https://doi.org/10.1111/j.1540-6261.2005.00725.x.

A. Appendix tables and figures

Table A1: Option prices and observable characteristics

The table shows regression results of several features of option prices against size, book leverage, market to book ratio, profitability return, and ATM IV, the growth option variable of (GO), the six months cumulative return, and operating leverage (defined as ...). The left hand side variable is the risk-neutral skewness of 90 day maturity options (Panel A), 180 day maturity options (Panel B), and 360 day maturity options (Panel c). We report parameter estimates and standard errors clustered at the firm level. Regressions include time and firm fixed effects. The sample contains all non-financial firms with options trading on their equity between 1996 and 2019. Data is sampled at quarterly frequency. A total of 3,536 firms are included.

Panel A: RN-Skew 90 day maturity

	(1)	(2)	(3)	(4)	(5)
IV	0.13	0.13	0.13	0.13	0.13
	(25.39)	(25.49)	(25.22)	(20.11)	(25.02)
Size	-0.11	-0.07	-0.11	-0.12	-0.11
	(-6.97)	(-4.77)	(-7.33)	(-6.52)	(-6.94)
Leverage	-0.02	-0.02	-0.02	-0.02	-0.02
	(-3.02)	(-3.07)	(-3.12)	(-2.74)	(-2.99)
Profitability	0.01	-0.04	-0.01	0.01	0.01
	(3.15)	(-6.95)	(-2.42)	(3.59)	(3.15)
M2B	-0.08		-0.07	-0.08	-0.08
	(-15.02)		(-13.78)	(-12.30)	(-15.16)
GO		-0.05	-0.03		
		(-11.42)	(-7.26)		
Return6				-0.04	
				(-18.80)	
Operating Leverage					0.00
					(0.28)
Adjusted-R^2	0.50	0.50	0.50	0.51	0.50
FE R^2	0.48	0.48	0.48	0.48	0.48

Panel B: RN-Skew 180 day maturity

	(1)	(2)	(3)	(4)	(5)
IV	0.16	0.16	0.16	0.16	0.16
	(32.72)	(32.80)	(32.64)	(28.22)	(31.98)
Size	-0.11	-0.08	-0.11	-0.11	-0.11
	(-7.92)	(-5.80)	(-8.19)	(-7.10)	(-7.86)
Leverage	-0.01	-0.01	-0.01	-0.01	-0.01
	(-2.52)	(-2.54)	(-2.59)	(-2.48)	(-2.45)
Profitability	0.01	-0.03	-0.01	0.01	0.01
	(2.74)	(-4.84)	(-1.21)	(2.85)	(2.83)
M2B	-0.06		-0.05	-0.06	-0.06
	(-12.49)		(-11.68)	(-10.57)	(-12.71)
GO		-0.04	-0.02		
		(-8.61)	(-5.12)		
Return6				-0.02	
				(-10.76)	
Operating Leverage					0.00
					(0.34)
Adjusted-R^2	0.61	0.61	0.61	0.62	0.61
FE R^2	0.57	0.57	0.57	0.58	0.57

Panel C: RN-Skew 360 day maturity

	(1)	(2)	(3)	(4)	(5)
IV	0.29	0.29	0.28	0.29	0.28
	(32.61)	(32.93)	(32.92)	(30.40)	(32.30)
Size	-0.08	-0.05	-0.08	-0.09	-0.08
	(-3.49)	(-2.51)	(-3.61)	(-3.47)	(-3.45)
Leverage	-0.01	-0.01	-0.01	-0.01	-0.02
	(-1.62)	(-1.40)	(-1.55)	(-1.49)	(-1.81)
Profitability	0.01	-0.02	-0.01	0.01	0.01
	(2.18)	(-3.09)	(-0.67)	(1.95)	(2.29)
M2B	-0.05		-0.04	-0.05	-0.05
	(-7.42)		(-6.50)	(-6.60)	(-7.56)
GO		-0.04	-0.02		
		(-5.89)	(-3.60)		
Return6				-0.02	
				(-5.63)	
Operating Leverage					0.00
					(0.29)
Adjusted-R^2	0.73	0.73	0.73	0.74	0.72
FE R^2	0.65	0.65	0.65	0.66	0.65