Timing Risk and Asset Prices

Indira Puri*
NYU Stern

March 25, 2026

Abstract

Understanding the valuation of uncertainty lies at the heart of finance. While foundational pricing frameworks focus extensively on uncertainty over payoff magnitudes (outcome risk), the timing at which payoffs arrive is also stochastic across a broad range of financial settings. I prove that the pricing of timing risk differs conceptually from that of outcome risk and characterize its pricing both when systematic and when idiosyncratic. Using only a no-arbitrage assumption, I show that timing risk's higher order moments affect price even when risk is purely idiosyncratic, a property not shared by outcome risk. I further show that for time, uncertainty can reasonably be valued higher or lower than certainty, unlike for outcomes; and timing risk features an asymmetry not present in outcome risk. These differences between outcome and timing risk survive diversification. I prove that the price of idiosyncratic timing risk is fully determined by the risk-free term structure. Finally, I decompose all priced risk into

*Email: indira.puri@stern.nyu.edu. I thank conference participants at the FMA-UC Davis Napa Conference (Best Paper Award), Institut Canadien des Dérivés, the LBS Summer Symposium (Best Paper Award), NBER Summer Institute: Asset Pricing, the Texas Finance Festival, SITE: The Macroeconomics of Uncertainty and Volatility, the Summer Symposium in Financial Markets (CEPR), and seminar participants at Booth, Brown, Harvard, MIT, Northwestern, Stanford, the University of Cambridge, the University of Chicago, the University of Maryland, the University of Michigan, the University of Wisconsin-Madison, and the Virtual Derivatives Seminar for helpful comments. I further thank Yakov Amihud, Aditya Chaudhry, Hui Chen, Ian Dew-Becker, Darrell Duffie, Steve Figlewski, Stefano Giglio, Ben Hebert, Philipp Illeditsch, Nikunj Kapadia, Peter Klibanoff, Theresa Kuchler, Andrew Lo, Sydney Ludvigson, Holger Mueller, Jonathan Parker, Thomas Philippon, Matt Richardson, Tony Saunders, Alexi Savov, Jeremy Stein, Johannes Stroebel, Marti Subrahmanyam, Quentin Vandeweyer, Jessica Wachter, and Jeff Wurgler for valuable feedback. This paper subsumes *Certain vs. Uncertain Timing: Financial Markets and Pricing Implications.*

outcome and timing components and formalize systematic timing risk. Three sample applications illustrate the the financial relevance of separating timing and outcome risk. First, unlike for outcome risk, the risk-adjusted discount rate for projects that include timing risk includes idiosyncratic risk, contrary to commonly taught intuition. Second, incorporating timing risk changes biotech valuations relative to standard industry benchmarks. Third, a simple application in option markets illustrates the theoretical co-movements. These results inform understanding of valuation and risk premia across financial settings.

1 Introduction

A central theme in finance is the valuation of future uncertain cash flows. In this paper, I consider uncertainty along two dimensions. One is uncertainty about outcomes. For example, consider a biotech company with a clinical trial underway. Here, the uncertainty is whether the drug will have a successful treatment effect relative to a placebo. Or consider the more mundane case of a European call option. Here, the uncertainty is about the future stock price on the option expiration date relative to the option strike.

Many economically important uncertainties are not only about outcomes but their *timing*. For example, in climate change (Giglio et al., 2021) uncertainty is not just about outcomes - the size of the consumption loss due to climate events - but also about timing: when a climate shock will hit. Similarly, corporate decisions in innovative firms (Eisfeldt and Papanikolaou, 2014) include outcome risk – whether the innovation succeeds – and timing risk – how long it will take to succeed and be adopted.

My study focuses on timing risk. Using a fairly minimal foundation of no-arbitrage, I analyze whether or not timing matters as a distinct dimension of uncertainty, and develop its implications for valuation. Because the analysis uses only no-arbitrage, it is similar in spirit to Hansen and Jagannathan (1991) or Modligliani and Miller (1958): baseline implications that arise from the law of one price. The analysis leads to several interesting results.

First, the effects of timing risk are distinct from those of outcome risk. Both risks are priced given non-zero covariance with the stochastic discount factor (SDF). However, the higher order moments of timing risk impact prices even when risk is purely idiosyncratic. The higher order moments of outcome risk, of course, do not matter in this case.

Second, the risk pecking order can change with timing risk. Certainty is always valued weakly higher than uncertainty for outcome risk. However, certainty can be valued higher, lower, or the same as, uncertainty for timing risk. I show that and which prevails depends both on the sign of the shock and the risk-free term structure.

Third, I show that the price of idiosyncratic timing risk can be fully characterized through the term structure of risk-free rates. Fourth and finally, I show that *all* priced risk can be fully decomposed into systematic outcome risk and systematic timing risk.

Three sample applications illustrate the the financial relevance of separating timing and outcome risk. First, I show that the risk-adjusted discount rate for timing risk changes with the project's idiosyncratic risk, unlike for outcomes; this runs counter to the typical intuition of risk-adjusted discount rates reflecting systematic risk (Myers and Turnbull, 1977; Berk and van Binsbergen, 2016). Second, as a practical example, incorporating timing risk changes benchmark industry biotech valuations; by 16% for typical numbers in the biotech space. Third, an application in option markets, considering two option classes that are contractually identical except that one includes timing risk, illustrates the predicted theoretical co-movements. Together, the theory and sample applications suggest that timing risk may deserve its own consideration in financial markets.

These results suggest a departure from the standard pedagogical treatment of time and states. The textbook convention often treats time as another state, either implicitly through recursive one-period-ahead techniques (Duffie, 2001) or explicitly by aggregating risks across states and periods (Cochrane, 2014). While efficient, this approach does not fully capture the wedge between timing and outcome risk. By disentangling these two forces, this paper identifies pricing dynamics that may not be readily apparent otherwise.

First, for assets with purely idiosyncratic outcome risk, one of the foundational results for pricing is that only the first moment matters for valuation: higher-order moments of the payoff do not affect price (Berk and DeMarzo, 2011). This is typically illustrated via diversification removing idiosyncratic risk. In contrast to this usual intuition, I show in Theorem 1 that higher-order moments of timing risk do affect price even when risk is purely idiosyncratic. This effect on price is not a risk premium but is a higher order effect unique to the dimension of timing relative to outcomes. Intuitively, timing "states" have embedded curvature not present in outcome states. Second, the pricing of certainty vs. uncertainty behaves differently for time than outcomes. In the outcome case, certainty is typically always priced weakly above uncertainty (Berk and DeMarzo, 2011; Maggiori, 2017). In the timing case, Theorem 2 proves that increasing timing risk can raise, lower, or leave unchanged the value relative to certain timing. Corollary 1 further proves a model-free asymmetry in the timing domain which need not be present for outcomes: for any stochastic discount factor, if uncertainty is valued higher than certainty for positive shocks, then it must be valued below certainty for negative timing shocks, and vice versa.

These conceptual differences – higher order moments, directionality, and asymmetry – survive diversification. Theorems 3 and Theorem 4 construct pseudo risk-free assets constructed from outcome risk vs. timing risk and prove that the pseudo risk-free timing portfolios continue to be priced with higher order moments, directionality, and asymmetry, in contrast to pseudo risk-free outcome portfolios.

Given these conceptual differences between timing and outcome risk, a natural question is what *does* characterize the pricing of uncertain timing? Theorem 5 proves that, for idiosyncratic timing risk, valuation is characterized entirely by the risk-free term structure. Precisely, the value of any claim with an idiosyncratic stopping time can be written as the value of a portfolio of zero-coupon bonds with weights equal to the stopping-time probabilities. Consequently, the effect of any SDF on timing dispersion is identical to its effect on zero-coupon bond pricing. This result might seem surprising because some models aim to change the pricing of uncertainty without moving risk-free rates, or vice versa (ex. Wachter (2013); Farhi and Maggiori (2018)). However, the theorem implies that – unlike for outcome risk – such distortions must operate on timing risk and risk-free rates in exactly the same way.

Many assets and financial settings have both outcome and timing risk. Theorem 6 provides a novel decomposition of priced risk – the covariance of the asset's payoffs with the SDF – into systematic outcome risk and systematic timing risk. This decomposition holds as an identity. Intuitively, systematic outcome risk is the covariance of the SDF with the asset payoffs conditional on the asset paying off, e.g. its stopping time being hit; this corresponds to the usual notion of systematic risk in macro-finance and asset pricing (Xu, 2021). Systematic timing risk, in contrast, measures whether the asset tends to pay off at times at which the SDF is higher or lower.

Three applications illustrate the theory. First, the standard taught intuition in corporate finance is that risk-adjusted discount rates adjust for systematic risk. This is certainly true for outcomes; but as I show, in the timing risk case, risk-adjusted discount rates reflect idiosyncratic risk as well.

Second, perhaps because of the taught intuition, industry norms which apply risk-adjusted net present value adjust for systematic risk and then use expected cash flows; biotech provides an illustrative example. However, this approach does not take idiosyncratic timing risk into account. I find that using the entire timing distribution of clinical trials for biotech, rather than the expected length of the trial, changes the

valuation of the average biotech firm by 16%.

Third, a simple application in options illustrates the directionality and co-movements predicted by the no-arbitrage theory. While American call options can be exercised at any time, European call options can be exercised only at expiry. As a result, the American options feature timing risk for the short side, and timing flexibility for the long side, which are not present in otherwise identical European options. An application of the theory to otherwise identical American vs. European options illustrates the differences between timing and outcome risk, and the characterization theorems.

The paper contributes to a number of literatures. First, the paper contributes to a deeper understanding of different types of uncertainty and their impact. While a large literature has studied uncertainty at different fixed durations (van Binsbergen, 2025), this paper studies the pricing when duration is uncertain.

Second, because the paper uses only a no-arbitrage assumption, it nests all pricing kernels. In particular, its assumptions mirror the so-called model-free results of Hansen and Jagannathan (1991) and Giglio et al. (2016) in assuming only the law of one price, thereby sidestepping an active debate over which stochastic discount factor to use (Maggiori, 2022; Kacperczyk, 2026). The reliance on minimal assumptions allows the paper's results to apply to a wide variety of settings. In particular, the results apply regardless of whether the SDF is friction-based (Krishnamurthy, 2010; Maggiori, 2017), production-based (Belo, 2010), preference-based (Doepke and Schneider, 2017), belief-based (Schneider and Piazzesi, 2015; Wachter, 2013), or otherwise – so long as no-arbitrage holds.

Third, the paper contributes to literatures on markets that include timing risk by formalizing it and developing results on precisely when and how timing risk differs from outcome risk. Such markets include settings with real options (Berk, 1999), investment decisions (Graham, 2022), and debt markets(Krishnamurthy, 2010).

Within the sample applications, innovation firms have been studied by Eisfeldt and Papanikolaou (2014); here, I find that incorporating timing risk is relevant to the valuation of biotech firms. Risk-adjusted discount rates are taught and used throughout finance (Graham et al., 2017), and I show that timing risk embodies a slightly different intuition than traditional outcome risk for these risk-adjusted rates. The options I study to illustrate the theoretical co-movements have previously been studied in Stein (1989) and Fleming and Whaley (1994), among others; the contribution is to empirically illustrate the difference between timing and outcome

valuation in derivatives markets (Atkeson et al., 2015).

While the paper studies time, it is conceptually distinct from other literatures on time-related concepts such as the elasticity of intertemporal substitution (Wachter, 2002) and resolution period effects (Epstein et al., 2014). One way to see this is that the results hold for any pricing approach, no matter its parametrization or resolution properties.

While decomposing risk has not historically been a pricing focus, the topic of uncertain timing has been studied in biology (Kacelnik and Bateson, 1996) and experimental and decision theory (Chesson and Viscusi (2003), Onay and Onculer (2007), DeJarnette et al. (2020) on expected utility, additive utility, and Epstein-Zin, respectively). The paper's no-arbitrage approach nests utility models as special cases, and studies pricing implications.

In addition to deepening the conceptual understanding of different components risk, there are many settings in which timing risk is relevant. These include but not limited to, corporate decisions; climate disasters; labor decisions; what makes an asset safe; distressed debt; policy implementation timelines; and so on. These varied economic settings provide ample avenues for applications of the pricing of timing vs. outcome risk in future work.

The paper proceeds as follows. Section 2 develops the theory, including conceptual differences between the pricing of timing vs. outcome risk, and a characterization of the pricing of idiosyncratic and systematic timing risk. Section 3.3 illustrates the findings of the theory in real-world market data. Section 3.4 discusses possible applications across economics. Section 4 concludes. All proofs are located in the appendix, with proof intuition discussed in the text.

2 Theory

I start by formalizing the set up to study timing vs. outcome risk, using the standard stochastic discount factor (SDF) framework (Duffie, 2001). I formalize outcome risk in the usual way, and timing risk as a stopping time. The only assumption used for the theoretical results is no arbitrage; no arbitrage in turn implies existence of a stochastic discount factor, regardless of whether the market is complete or incomplete (Berk and DeMarzo, 2011). For any given stochastic discount factor, all theorems derived in this paper must hold.

Definition 1 (The Stochastic Discount Factor Framework)**.**
There exists a probability space $(\Omega, \{\mathcal{F}_t\}_{t\in T}, P)$, where Ω is a set of states, $\{\mathcal{F}_t\}$ a filtration, $T \subseteq \mathbb{R}^+$ is a set of times, and P a probability measure. In this framework, assets are adapted processes: $X = (X_1, X_2, ...)$ where each X_t is a random variable measurable with respect to $\mathcal{F}_t$. We say that a stochastic discount factor $M_t = \{M_{t\to t'}\}_{t'>t}$ exists if, for all assets, the time t price of asset X is given by $Price_t(X) = E_{P_t}[\sum_{t'} M_{t\to t'} X_{t'}]$.

Definition 2 (Assets with Pure Uncertain Outcomes)**.**
An asset with pure uncertain outcomes X is one which has non-zero payoffs at only one time $t \in T$.

Definition 3 (Stopping Time)**.**
Given the filtered probability space $(\Omega, \{\mathcal{F}_t\}_{t\in T}, P)$ from Definition 1, a stopping time is a random variable $\tau : \Omega \to T$ such that for every $t \in T$, the event $\{\omega \in \Omega : \tau(\omega) \leq t\}$ is in the filtration F_t.

Say that τ is a non-trivial stopping time if there are at least two distinct times it takes on with positive probability: $\exists t_1 \neq t_2 \in T$ such that $P(\tau = t_1), P(\tau = t_2) > 0$.

The expected time of τ is given by $E_P[\tau]$.

Definition 4 (Assets with Pure Uncertain Timing)**.**
An asset with pure uncertain timing X is one whose payoff process is entirely characterized by some nontrivial stopping time τ. In particular,

$$X_t(\omega) = \begin{cases} c + a & \text{if } t = \tau(\omega) \\ c & \text{if } t \neq \tau(\omega) \end{cases}$$

for some $c, a \in \mathbb{R}$.

Example 1. *Suppose that the state space at time 0 is $\{G, B\}$, standing for good and bad. The state space at time 1 is then $\{(G,G),(G,B),(B,G),(B,B)\}$; at time 2 is $\{(G,G,G),(G,G,B),(G,B,G),(G,B,B),(B,G,G),(B,G,B),(B,B,G),(B,B,B)\}$; etc.*

An example of an asset with pure uncertain timing on this state space is one that pays out \$1 the first time the state space hits good only. This asset takes on a payoff value of 1 whenever the state is $(\underbrace{B, \ldots, B}_{\text{any number of times}}, G)$, and 0 in all other states.

An example of an asset with pure uncertain outcomes on this state space is one that pays out \$5 (in state B) or \$10 (in state G) at time 1, and 0 in all other times and states.

Definition 5 (Systematic vs. Idiosyncratic Risk)**.**
The time-t' systematic risk of an asset X from the perspective of period $t \leq t'$ is given by $Cov(M_{t\to t'}, X_{t\to t'})$.

An asset has purely idiosyncratic risk from the perspective of period t if its time-t' systematic risk is equal to zero at all times $t' \geq t$.

Definition 6 (Assets with Uncertain Timing and Outcomes)**.**
An asset with uncertain timing and outcomes is characterized by a baseline payoff $c \in \mathbb{R}$, a stopping time τ, and a random shock $Y : \Omega \to \mathbb{R}$. The payoff process is

$$X_t(\omega) = c + \mathbf{1}_{\{t=\tau(\omega)\}} \cdot Y(\omega), \qquad t \in T.$$

This nests:

- *Pure uncertain timing: $Y(\omega) = a \in \mathbb{R}$ (constant shock);*
- *Pure uncertain outcomes: $\tau(\omega) = t^*$ (fixed time) and $c = 0$.*

Remark 1 (Arbitrary Correlation Between Outcome and Timing Risk)**.** *For assets with uncertain timing and outcomes, any correlation structure between outcome risk Y and time risk τ is permitted by Definition 6; for example, it may be the case that the distribution of $Y|\tau = t_1$ differs from that of $Y|\tau = t_2$.*

Definition 7 (Valuation)**.**
Under Definition 1, the valuation of asset X at time t is given by $Price_t(X) = E_{P_t}[\sum_{t'} M_{t\to t'} X_{t'}]$.

Consider an asset with pure uncertain timing X which has associated stopping time τ. When $E_p[\tau] \in T$, define its certain-time equivalent,

$$X_t(\omega) = \begin{cases} c + a & \text{if } t = E_p[\tau] \\ c & \text{if } t \neq E_p[\tau] \end{cases}$$

Say that certain timing is valued higher (resp. lower) than uncertain timing if, for any asset with pure uncertain timing, its valuation is lower (resp. higher) than the valuation of its certain-time equivalent.

Definition 2 is the standard one period definition of an asset (Berk and DeMarzo, 2011). Definition 5 is the standard conceptualization of systematic risk and idiosyncratic risk across multiple periods (Berk and DeMarzo, 2011). Definitions 4 and 7 extend to the pricing setting definitions used in the decision theory literature of time lotteries and risk aversion over them (DeJarnette et al. (2020), also earlier studied in Chesson and Viscusi (2003) and Onay and Onculer (2007)). The present framework applies for any SDF; formalizes uncertainty as random variables through the use of stopping times; and allows assets to take on non-zero values away from these stopping times. Definition 6 combines uncertain outcomes and uncertain timing and will be used for the decomposition results.

All of the above definitions will play a role in analyzing the valuation of timing vs. outcome risk.

2.1 No-Arbitrage Assumption

The theoretical results of this paper rely only on the existence of a stochastic discount factor. From foundational pricing theory, an SDF exists if and only if no arbitrage holds. This short section reproduces this result for completeness to make transparent the reliance of the theory only on a no-arbitrage assumption.

To apply the First Fundamental Theorem of Asset Pricing (e.g. no arbitrage if and only if there exists a stochastic discount factor that prices all assets), pick an arbitrary time t; collapse future states relative to t, $\{\Omega_{t'} = S^{t'}\}_{t' \geq t}$ and time T into a single space $\mathcal{Z} = \cup_{t' \in T, t' \geq t} S^t$. When the cardinality of S and of T is discrete, so too is the cardinality of $\mathcal{Z}_t$. As above, assets are random variables defined on this space: $\mathcal{A}_t \subseteq \{f : \mathcal{Z}_t \to \mathbb{R}\}$. Assume standard portfolio formation: if asset a and asset b are in the set of assets, so too is the portfolio $\gamma a + (1-\gamma)b$ for any $\gamma \in (0,1)$.

Assumption (No-Arbitrage)**.** *There exists no portfolio whose price is below zero and whose payoffs are almost everywhere positive.*

Lemma 1 (Fundamental Theorem of Asset Pricing (Berk and DeMarzo, 2011))**.** *If no-arbitrage holds, then there exists a stochastic discount factor $M_t > 0$ that prices all assets: $Pr(X) = E[M_t(z)X_t(z)]$ for any asset $X_t \in \mathcal{A}_t$.*

Given the fundamental theorem, we can un-collapse back out to states over time simply by re-stating the expectation: $Pr(X) = E[M_t X_t] = E[\sum_{t' \in T, t' > t} M_{t \to t'} X_{t \to t'}]$.

Formally,

$$\begin{aligned} E[M_t X_t] &= \sum_{z \in \mathcal{Z}_t} p(z) M_t(z) X_t(z) \\ &= \sum_{z \in \cup_{t' \in T, t' \geq t} S^{t'}} p(z) M_t(z) X_t(z) \\ &= \sum_{t' \in T, t' \geq t} \sum_{\omega_{t'} \in S^{t'}} p(\omega_{t'}) M_t(\omega_{t'}) X_t(\omega_{t'}) \\ &= E[\sum_{t' \in T, t' \geq t} M_{t \to t'} X_{t \to t'}], \end{aligned}$$

exactly as in Definition 1.

No-arbitrage is the only assumption used throughout the paper. Given Lemma 1, the theorems can, and do, directly assume the existence of a stochastic discount factor that prices all assets. In the case that there are multiple admissible SDFs, as can occur for incomplete markets (Berk and DeMarzo, 2011), the theory holds for each SDF within the set. The theory therefore applies for both complete and incomplete markets: in the former case, for the unique SDF, and in the latter case, for any SDF within the permissible SDF set.

2.2 The Price of Timing Risk vs. Outcome Risk

2.2.1 Idiosyncratic Risk and Higher-Order Moments

Theorem 1 shows that for idiosyncratic risk, assets with uncertain timing are fundamentally different than assets with uncertain outcomes.

In particular, a foundational pricing result is that, for an asset with idiosyncratic risk and uncertain outcomes, its higher order moments are not priced. The intuition is straightforward and illustrated in Remark 2.

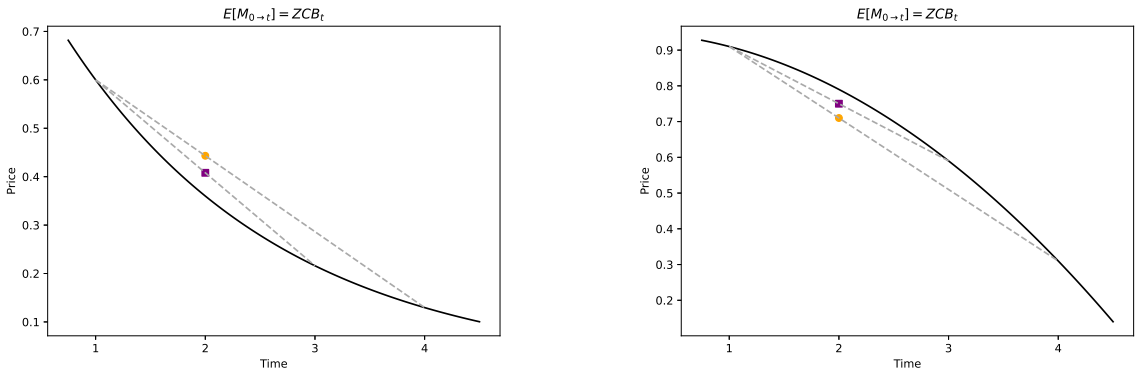
Remark 2 (Idiosyncratic Risk for Assets with Uncertain Outcomes (Berk and DeMarzo, 2011))**.** *With idiosyncratic risk, given SDF M and asset A with pure uncertain outcomes paying off at time t', $Price_t(A) = E_{P_t}[M_{t \to t'} A_{t'}] = E_{P_t}[M_{t \to t'}] E_{P_t}[A_{t'}]$. In particular the higher order moments of the asset such as $Var_{P_t}[A_{t'}]$ do not affect the price.*

Theorem 1, in contrast, shows that the higher order moments of the stopping time

for an asset with uncertain timing do affect the price, regardless of which SDF one uses.

Theorem 1 (Higher Order Moments for Uncertain Timing)**.** *Consider an asset with pure uncertain timing, with associated stopping time τ which has purely idiosyncratic risk. The higher order moments of τ affect the valuation of the asset.*

FIGURE 1: Pricing Assets with Uncertain Timing with Purely Idiosyncratic Risk



(A) Sample SDF 1 - Receiving \$1 at Time t

(B) Sample SDF 2 - Receiving \$1 at Time t

Note: These figures plot the price of two assets for two sample SDFs. In each figure, the curve corresponding to the price of zero-coupon bonds arising from the SDF is shown. In Figures 1a, this curve is convex, and in Figure 1b, it is concave. In each figure, the price of two assets is shown. The asset whose price is depicted by a purple square pays out \$1 at time 1 with 50% probability and at time 3 with 50% probability. The asset whose price is depicted by an orange circle pays out \$1 at time 1 with 66.6% probability and at time 4 with 33.3% probability. Both assets are assumed to have purely idiosyncratic risk. For Theorem 1, observe that both assets have an expected stopping time of 2 (thus lining up on the x axis), yet their prices differ regardless of the SDF used. For Theorem 5, observe that $E[M_{0\to t}]$ is exactly the price of a zero coupon bond maturing in t periods; then re-interpret the probabilities on the line as a portfolio.

The intuition is as follows. When an asset with uncertain outcomes has purely idiosyncratic risk, its price as discussed in Remark 2 becomes an expectation of the SDF times an expectation of the asset payoff. When the asset has uncertain timing, the price is similarly an expectation of the SDF, but that expectation is taken over time. If the SDF has any curvature - as shown for example in Figure 1 - this expectation embeds nonlinearity in it. Because of the nonlinearity, the higher order moments of the stopping time affect valuation.

To illustrate, Figure 1 plots the valuation of two assets with purely idiosyncratic risk. One pays out \$1 at Time 1 with 1/2 probability and at Time 3 with 1/2 probability. The other pays out \$1 at Time 1 with 2/3 probability and at Time 4 with 1/3 probability. Both have an expected stopping time value of 2. But they have

different higher order moments: the variance of stopping time of the former asset is 1, and of the latter asset 2. Because of this difference in the higher order moments of the stopping time, they are priced differently for any SDF, whenever the SDF has curvature in time. Figure 1 shows the different pricing for an SDFs with different types of curvature, illustrating that the valuation is different regardless of the SDF used, except in the special (and typically unused) case where the SDF is linear in time (Berk and DeMarzo, 2011; Maggiori, 2022).

2.2.2 Directionality and Asymmetry

With outcomes, certainty is always weakly better than uncertainty (Berk and DeMarzo, 2011), as illustrated formally in Remark 2. With timing, however, uncertainty can reasonably be valued higher, lower, or the same as, certain-time equivalents. Theorem 2 proves that which direction occurs is characterized by the risk-free term structure. Corollary 1 proves that any SDF must feature an asymmetry: if certainty is valued higher than uncertainty for positive shocks, then it must be valued lower than uncertainty for negative shocks, and vice versa.

The intuition is that when the yield curve is strongly convex, the value of possibly receiving a positive payment at some point later in the future is much less than the value of receiving it at the expected stopping time. The intuition flips for negative payments, since the corresponding pain becomes concentrated at near horizons. The sign intuition resembles that seen in utility functions for gains and losses (Onay and Onculer, 2007), but here applies for any stochastic discount factor, proven through the risk-free term structure. The proof intuition is similar to that used for convexity in bonds (Berk and DeMarzo, 2011), but here applies for a dimension of uncertainty.

Definition 8 (Risk-Free Term Structure)**.**
Given the discount function $m(T) := Price_t(ZCB_T)$ *representing the time-t price of a zero coupon bond with maturity* T*, the yield to maturity (or spot rate)* $y(T)$ *is defined by:* $m(T) = e^{-T \cdot y(T)}$*. Equivalently,* $y(T) = -\frac{\ln m(T)}{T}$*.*

Theorem 2 (Certainty and Uncertainty)**.** *Consider any asset with uncertain timing with associated stopping time* τ *and payoffs* $(c, c + a)$ *which has purely idiosyncratic risk. Let* $y(T)$ *denote the yield to maturity at time* T*.*

Then uncertain timing has a higher (resp. lower) valuation than certain timing

for $a > 0$ (resp. $a < 0$) if and only if

$$(y(T) + T \cdot y'(T))^2 > 2y'(T) + T \cdot y''(T) \quad \textit{for all } T \in \mathcal{T}$$

In particular, uncertainty can be valued higher, lower, or the same as, certainty.

The proof intuition is depicted in Figure 2. The directionality - whether the asset price lies above or below the black curve at the expected stopping time - depends on the curvature of the risk-free term structure.

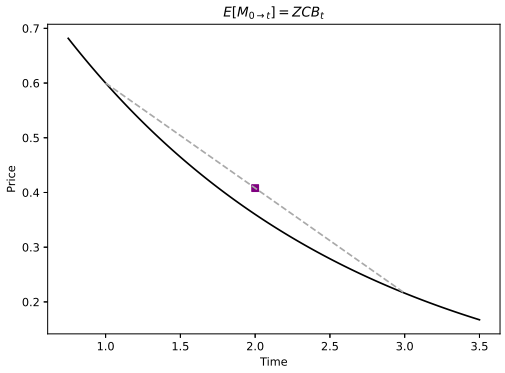
Corollary 1 (Asymmetry)**.** *Any SDF that has a higher valuation for uncertain timing than certain-time equivalents for positive shocks must have a higher valuation for certain-time equivalents than uncertain timing for negative shocks. Any SDF that has a higher valuation for certain-time equivalents than uncertain timing for positive shocks must have a higher valuation for uncertain timing than certain-time equivalents for negative shocks.*

The intuition is shown in Figure 2. For example, compare Figures 2a and 2c. These both depict the same SDF. In the former, the payoff is positive (e.g. receiving \$1 in t periods for each zero-coupon bond). Similarly the asset payoffs are positive: for example the purple square depicts receiving \$1 in period 1 with 50% probability and in period 3 with 50% probability. When considering an asset with uncertain timing which has negative payoffs, e.g. losing \$1 in period 1 with 50% probability and losing it in period 3 with 50% probability, the relevant zero coupon-bond comparison also becomes negative: losing \$1 at each possible maturity t. As a result, the SDF curve gets flipped over the x-axis. Because of this flipping, if the curve was convex when positive, it becomes concave when negative, and vice versa.

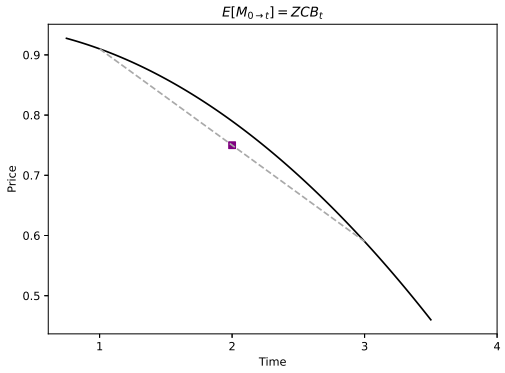
Remark 3 (Yield Curve Translation)**.** *Uncertain timing is always valued higher (resp. lower) than certain timing for positive (resp. negative) shocks when:*

- *The yield curve is flat and nonzero: $y' = 0, y'' = 0, y \neq 0$.*
- *The yield curve is downward sloping and weakly convex: $y' \leq 0, y'' \leq 0$.*
- *The yield curve does not rise too steeply: $2y'(t) \leq Ty''(t)$.*

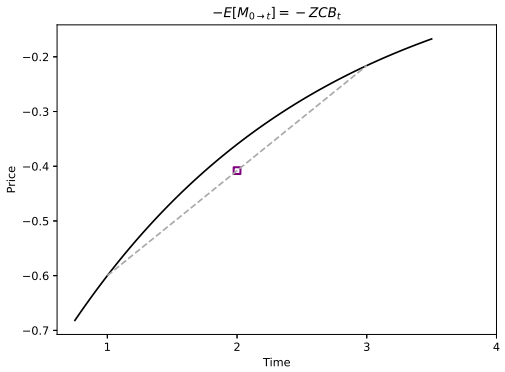
FIGURE 2: Certainty vs. Uncertainty and Asymmetry



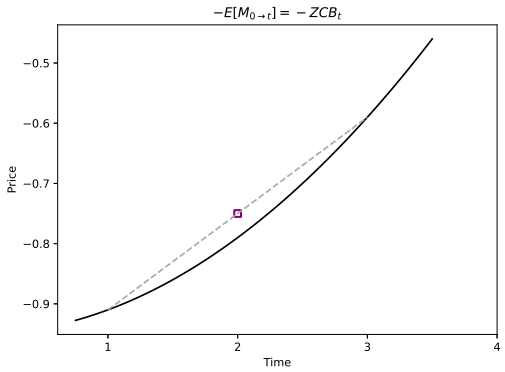
(A) Sample SDF 1 - Receiving $1 at Time t



(B) Sample SDF 2 - Receiving $1 at Time t



(C) Sample SDF 1 - Receiving (-$1) at Time t



(D) Sample SDF 2 - Receiving (-$1) at Time t

Note: These figures plot the price of gaining (subfigures 2a and 2b) and losing (subfigures 2c and 2d) $1 of purely idiosyncratic risk, for two sample SDFs (the same sample SDFs as in Figure 1). Observe that, for losing $1, the expected SDF curve flips over the x-axis (subfigures 2a and 2c show the flip for sample SDF 1; and subfigures 2b and 2d for sample SDF 2). The top two figures also depict, as a filled purple square, the price of an asset with uncertain timing with purely idiosyncratic risk which pays $1 in period 1 with 50% probability and in period 3 with 50% probability. The bottom two figures depict, as a hollow purple square, for each sample SDF, the price of an asset with uncertain timing with purely idiosyncratic risk which pays (-$1) in period 1 with 50% probability and in period 3 with 50% probability. For Theorem 2, observe that whether certainty is preferred to uncertainty is visually whether the purple square is below or above the black line; and that depending on the SDF, either direction can occur. For Corollary 1, observe that whenever the filled purple square is above the black line, the hollow purple square is below the black line, and vice versa; e.g., compare subfigures 2a and 2c. Similarly compare subfigures 2b and 2d. Because of the flipping over the x-axis, for *any* SDF, a higher valuation for certainty for positive shocks must be coupled with a higher valuation of uncertainty for negative shocks, and vice versa.

2.2.3 Diversification

For assets with uncertain outcomes, diversification pools assets with uncorrelated risks, improving the risk profile without affecting returns (Berk and DeMarzo, 2011). A natural question is whether the price impact of an asset with uncertain timing can be neutralized by holding a pool of such assets.

Theorems 3 and 4 seek to answer this question. Theorem 3 considers independent stopping times and shows that in the limit, flows become predictable, removing uncertainty. But pooling does not eliminate the impact of higher-order moments, even when each asset has purely idiosyncratic risk. The higher order impact is no longer due to risk – because flows become predictable – but rather due to the predictable flow's dispersion.

A natural next question is whether the time vs. outcome gap vanishes when the portfolio is chosen to smooth any shocks, for example as with perfectly negatively correlated assets. Theorem 4 shows that the difference in pricing time vs. pricing outcomes remains even with perfectly deterministic and smoothed payments through diversification. In particular, the higher-order moment and asymmetry results that applied for individual assets continue to hold even after a synthetic risk-free asset is created.

Independent Stopping Times. Theorem 3 treats the case with infinitely many independent stopping times and shows that while payments are now deterministic, the conceptual differences highlighted earlier in the paper for individual assets continue to hold.

Theorem 3 (Diversification of Idiosyncratic Uncertain Timing: Independent Stopping Times)**.** *Consider a set of N assets $A_1, .., A_n$ with uncertain timing which have payoffs $(c_i, a_i)_{i\in\{1,2,..,n\}}$ and associated stopping times $\tau_1, .., \tau_n$. Each asset A_i has purely idiosyncratic risk. The stopping times are assumed to be independent of each other. In addition, the limits $\bar{c} := \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} c_i$ and $\bar{a}(t') := \lim_{n\to\infty} \frac{1}{n} \sum_{i=1}^{n} a_i \, P_t(\tau_i = t')$ are assumed to be well-defined.*

Define the equally weighted portfolio $P^{(n)}$ as

$$P^{(n)}_{t'} \;:=\; \frac{1}{n}\sum_{i=1}^{n} A^{(i)}_{t'} \;=\; \underbrace{\frac{1}{n}\sum_{i=1}^{n} c_i}_{\textit{baseline}} \;+\; \underbrace{\frac{1}{n}\sum_{i=1}^{n} a_i\,\mathbf{1}_{\{\tau_i=t'\}}}_{\textit{timing shock at } t'}\,.$$

Then:

(a) *The payoff stream becomes deterministic in the limit.*

(b) *For any SDF $\{M_{0\to t}\}_{t\in T}$, the price of the portfolio $P^{(n)}_{t'}$ is equivalent to the price of an appropriately selected portfolio of zero-coupon bonds.*

(c) *The higher order moments of the dispersion across time, although now deterministic, affect valuation.*

Perfect Smoothing. We now study what happens with perfectly negatively correlated stopping times, so that payments are deterministic and smooth. For example, suppose that asset A_1 pays out \$1 at time 1 and at time 3 with 50% probability each, and asset A_2 does the same but is perfectly negatively correlated with asset A_1. Then holding a portfolio of $1/2$ A_1 and $1/2$ A_2 gives a payoff of \$0.5 at time 1 and \$0.5 at time 3 deterministically: a synthetic risk-free asset. Does this perfect negative correlation change the time vs. outcomes results of the earlier theorems?

It turns out that time and outcomes continue to differ. Whereas for outcomes, a synthetic risk-free asset created by combining two assets with identical payoffs and perfect negative correlation, is valued identically to the deterministic asset that pays the expected payoff, in the time dimension this is no longer the case. That is, the valuation of synthetic risk-free assets can differ from the valuation of a true risk-free asset paying at the expected stopping time.

Similarly, as in Theorem 3, the higher order moments of the temporal dispersion continue to affect valuation for time although higher order moments do not matter for diversified outcomes.

Theorem 4 proves the general case. Its statement is more general than perfect negative correlation: it says, instead, suppose there is a way to combine assets with uncertain timing to achieve identical and deterministically predictable payoffs at each date. Perfect negative correlation is one way to achieve this; there may be others. Even so, the above more nuanced time vs. outcomes differences continue to hold.

Theorem 4 (Diversification of Assets with Uncertain Timing: Smoothing Shocks)**.** *Let $A_1, \ldots, A_n$ be assets with uncertain timing with purely idiosyncratic risk, where asset i pays c_i at all times plus an additional a_i at a random stopping time τ_i. Assume further that each τ_i has the same marginal distribution.*

Suppose there exists a portfolio with weights $p = (p_1, \ldots, p_n)$ such that the portfolio deterministically pays amount c at all times plus an additional s at each shock time in S.

Then:

(a) The price of the portfolio is equivalent to the price of a portfolio of suitably chosen zero-coupon bonds.

(b) The temporal dispersion of the shock times affects valuation.

(c) The portfolio value can be higher, lower, or equal to its certain time equivalent, e.g. a payment of the same total amount at time $\mathbb{E}[\tau_1]$ (plus the baseline c at all times).

(d) If the portfolio value exceeds (resp. is below) the certain-time equivalent when the shock amount $s > 0$, it must be below (resp. exceed) it when $s < 0$.

Remark 4 (Contrast with Outcomes)**.** *Combining perfectly negatively correlated assets which have uncertain outcomes and pay at the same date t creates a risk-free asset. In particular, if A_1 and A_2 are assets with uncertain outcomes, and they have identical cumulative distribution functions but are perfectly negatively correlated, holding a portfolio with equal portfolio weights on each of them creates a risk-free asset whose valuation is identical to the asset that deterministically pays $E[A_1]$ at time t.*

By contrast, Theorem 4 shows that even when aggregate timing is made deterministic, for example via perfect negative correlation, the present value generally differs from paying the same total amount once at the average time. The wedge is a curvature effect of the SDF, not a risk premium.

In addition, the asymmetry implication of Corollary 1 remains, even though there is no uncertainty: if the wedge is positive (resp. negative) for positive shocks ($s > 0$), the wedge must be negative (resp. positive) for negative shocks ($s < 0$).

2.3 How is Timing Risk Priced?

2.3.1 Timing Risk and the Risk-Free Term Structure

So far, we have shown how timing risk differs from outcome risk. A natural question is what characterizes the pricing of uncertain timing, given these conceptual differences from uncertain outcomes.

I show in Theorem 5 that it is possible to completely characterize the price of an asset with uncertain timing for any SDF. The theorem proves that the price of an asset with uncertain timing is exactly equal to the price of a portfolio of zero-coupon bonds with portfolio weights equal to the asset's stopping time probabilities.

The theorem indicates a tight linkage between bond pricing results and timing risk. For example, by leveraging this result, known results for the pricing of bonds can be directly applied to price assets with uncertain timing which have purely idiosyncratic risk.

Theorem 5 (Risk-Free Term Structure Characterization)**.** *Consider an asset X with pure uncertain timing, associated stopping time τ, parameters (c, a), and purely idiosyncratic risk from time t. Define:*

- $ZCB_{t'}$ *= a zero coupon bond paying* \$1 *at time t' and* \$0 *otherwise*
- $Portfolio_\tau = \sum_{t' \geq t} P_t(\tau = t') \cdot ZCB_{t'}$

Then:

$$Price_t(X) = c \cdot \sum_{t' \geq t} Price_t(ZCB_{t'}) + a \cdot Price_t(Portfolio_\tau)$$

The intuition behind Theorem 5 is as follows. When writing out the price of an asset with uncertain timing which has purely idiosyncratic risk, the covariance between the SDF and the asset vanishes by virtue of the risk being idiosyncratic. What is left is a weighted sum over the expected value of the SDF at each point in time, where the weightage on the time is exactly the stopping time probability.

The crucial observation is that probabilities, like portfolio weights, sum to unity. This parallel admits a natural relabeling under which the pricing equation reduces to that of a portfolio of suitably chosen zero-coupon bonds. Figure 1 illustrates the intuition behind this equivalence.

2.3.2 Systematic Timing Risk

Many real-world assets have both uncertain timing and uncertain outcomes. For such assets, this section introduces a novel decomposition of systematic risk into systematic outcome risk and systematic timing risk.

These two components, formalized in Theorem 6 and Corollary 2, can be intuitively described as follows. Systematic outcome risk is what is typically thought of when discussing systematic risk (Berk and DeMarzo, 2011): conditional on the asset's stopping time being hit at time t, the covariance of the level of the SDF with the level of the asset payout. Systematic timing risk measures whether the asset is more likely to pay out at times when the SDF is higher or lower.

Systematic timing risk is computed at each date t by first computing the number of times at which the asset's stopping time takes on a non-zero probability; call this number N. Then create N mega-states: one for each time that the stopping time can take on with non-zero probability. Compute the conditional SDF in each of these mega-states. Compute the conditional asset payoff in each of these mega-states, then set the conditional amount to zero in all states except the one in which the asset pays out at time t. The probability-weighted covariance of these two series – the conditional SDFs against the conditional shock payoffs – is systematic timing risk. Intuitively, it measures whether, in the cases where the asset actually pays at t, the expected conditional discount factor at t tends to be higher or lower than its average across all possible paying dates.

The proof intuition of Theorem 6 is a direct application of the law of total covariance. One appealing property of this decomposition is that, unlike typical decompositions used in the literature, which rely on a first-order Taylor approximation (Giglio et al., 2016), the decomposition shown here is complete.

Theorem 6 (Time-t Systematic Risk Decomposition)**.** *Consider an asset with uncertain timing and outcomes as in Definition 6. At each time t, the systematic risk decomposes as:*

$$Cov(M_{0\to t}, X_t) = \underbrace{P(\tau = t) \cdot Cov(M_{0\to t}, Y \mid \tau = t)}_{\textit{Systematic Outcome Risk at } t} + \underbrace{Cov\left(\mathbb{E}[M_{0\to t} \mid \tau], \mathbf{1}_{\{t=\tau\}} \cdot \mathbb{E}[Y \mid \tau]\right)}_{\textit{Systematic Timing Risk at } t}$$

Corollary 2 (Overall Systematic Risk Decomposition)**.** *Let the asset pay the shock*

Y at the (random) date τ. Then

$$\operatorname{Cov}\left(\sum_t \mathbf{1}_{\{\tau=t\}} M_{0\to t}, Y\right) = \underbrace{\sum_t \mathbb{P}(\tau=t)\cdot \operatorname{Cov}(M_{0\to t}, Y \mid \tau=t)}_{\textit{Total systematic outcome risk}} + \underbrace{\operatorname{Cov}\left(\mathbb{E}\left[\sum_t \mathbf{1}_{\{\tau=t\}} M_{0\to t} \,\middle|\, \tau\right], \mathbb{E}[Y \mid \tau]\right)}_{\textit{Total systematic timing risk}},$$

where the first term exactly corresponds to the sum of similar terms in Theorem 6 and the second term is a probability weighted covariance over two random variables, where the probability weights are given by the stopping time τ's probabilities.

Remark 5 (Special Cases of the Systematic Risk Decomposition)**.** *Two special cases of the systematic risk decomposition are:*

1. *Pure systematic timing risk. If $Y = a$ (constant), then for every t, $\operatorname{Cov}(M_{0\to t}, Y \mid \tau = t) = \operatorname{Cov}(M_{0\to t}, a \mid \tau = t) = 0$, so the outcome component vanishes and all systematic risk is timing risk.*

2. *Pure systematic outcome risk. If $\mathbb{P}(\tau = t^*) = 1$ (deterministic timing), then the timing term $\operatorname{Cov}\big(\mathbb{E}[M_{0\to t} \mid \tau],\, \mathbf{1}_{\{\tau=t\}}\mathbb{E}[Y \mid \tau]\big)$ is zero, and only $\operatorname{Cov}(M_{0\to t^*}, Y \mid \tau = t^*) = \operatorname{Cov}(M_{0\to t^*}, Y)$ remains.*

Example 2 (Numerical Example: Systematic Risk Decomposition)**.** *Let $\Omega = \{GG, GB, BG, BB\}$, $T = \{1, 2\}$, and P be as in the table below. Consider an asset with baseline $c = 2$, stopping time τ, and shock Y, with payoff*

$$X_t(\omega) = c + \mathbf{1}_{\{t=\tau(\omega)\}} Y(\omega), \qquad t \in \{1, 2\},$$

and stochastic discount factors $M_t(\omega) = M_{0\to t}(\omega)$.

State ω	$P(\omega)$	$\tau(\omega)$	$Y(\omega)$	c	$X_1(\omega)$	$X_2(\omega)$	$M_1(\omega)$	$M_2(\omega)$
GG	*0.30*	*1*	*13.0*	*2.0*	*15.0*	*2.0*	*0.90*	*0.85*
GB	*0.30*	*1*	*5.0*	*2.0*	*7.0*	*2.0*	*0.94*	*0.91*
BG	*0.25*	*2*	*2.0*	*2.0*	*2.0*	*4.0*	*0.96*	*0.93*
BB	*0.15*	*2*	*0.0*	*2.0*	*2.0*	*2.0*	*0.97*	*0.95*

We compute the time-1 systematic risk decomposition.

For systematic outcome risk, we calculate the covariance of the SDF with the shock, conditional on paying at time 1. Since $P(GG \mid \tau = 1) = P(GB \mid \tau = 1) = 0.5$*:*

$$\mathbb{E}[M_1 \mid \tau = 1] = 0.92, \qquad \mathbb{E}[Y \mid \tau = 1] = 9.0,$$

$$\mathrm{Cov}(M_1, Y \mid \tau = 1) = \big(0.90 \cdot 13 \cdot 0.5 + 0.94 \cdot 5 \cdot 0.5\big) - 0.92 \cdot 9 = -0.08.$$

Thus systematic outcome risk at $t = 1$ *is* $P(\tau = 1)\,\mathrm{Cov}(M_1, Y \mid \tau = 1) = 0.6(-0.08) = -0.048$.

For systematic timing risk, partition states by τ*. The conditional means are:*

$$\mathbb{E}[M_1 \mid \tau = 1] = 0.92, \quad \mathbb{E}[M_1 \mid \tau = 2] = 0.96375, \quad \mathbb{E}[Y \mid \tau = 1] = 9.0, \quad \mathbb{E}[Y \mid \tau = 2] = 1.25.$$

Since the shock has not arrived at time 1 in $\tau = 2$ *states, these contribute zero to* X_1*:*

$$\mathrm{Cov}\big(\mathbb{E}[M_1 \mid \tau],\, \mathbf{1}_{\{\tau=1\}}\mathbb{E}[Y \mid \tau]\big) = 4.968 - 0.9375 \cdot 5.4 = -0.0945.$$

Total systematic risk at time 1 is $\mathrm{Cov}(M_1, X_1) = (-0.048) + (-0.0945) = -0.1425$.

2.3.3 Market-Clearing Prices

For financial products such as options, corporate bonds, and debt, one side often has time flexibility of payment while the other side holds the uncertain timing. How does being long uncertain timing versus short uncertain timing clear in equilibrium?

In these cases, the distribution of uncertain timing may arise probabilistically from an optimization problem on the side that has timing flexibility. But given that induced distribution, the value of flexibility to the long side must in equilibrium equal the cost of timing risk to the short side. If the long side has the flexibility to exercise at times when the SDF is systematically higher or lower (Theorem 6), then they must pay for that.

If the long side is using that flexibility idiosyncratically – for example due to idiosyncratic liquidity shocks or constraints – one of the takeaways of this paper's results is that such idiosyncracies do not zero out in their effect on price. Instead, the value of idiosyncratic timing risk co-varies with the risk-free term structure.

3 Sample Applications

3.1 Risk-Adjusted Valuation

Standard corporate finance teaching says that to adjust net present value for risk, adjust the discount rate. A typical intuition, through the lens of the capital asset pricing model, is that the discount rate reflects the project's systematic risk: the risk free rate plus a systematic risk adjustment (Berk and DeMarzo, 2011; Brealey et al., 2020). Consequently, textbooks teach that when two projects within the same firm have different systematic risk, or when the systematic risk changes over time, different adjusted discount rates should be used.

This is surely correct for outcome risk, but when the project has timing risk, the idiosyncratic risk, that is, the variance, skewness, etc. of the project's stopping time, also affects the discount rate. As a result, the widely taught intuition of Myers and Turnbull (1977) – that the discount rate used should only be changed if the systematic risk changes – while correct for outcome risk, is incomplete for timing risk. For projects with timing risk, the risk-adjusted discount rate can also change if the amount of idiosyncratic risk changes. Corollary 3 states and proves this formally. The proof is nearly a direct application of Theorems 1 and 5.

Corollary 3. *Consider any asset X and its risk-adjusted discount rate r_x, defined such that $\sum_{t=0}^{\infty} \frac{E[X_t]}{(1+r_x)^t} = E\left[\sum_{t=0}^{\infty} M_{0\to t} X_t\right]$.*

When asset X has pure uncertain outcomes, e.g. it pays off at some time t', r_x is a function of the risk-free rate from 0 to t' plus an adjustment for the asset's systematic risk.

However, when asset X has pure uncertain timing, r_x is a function of the zero-coupon bond portfolio with maturities where the asset pays out with positive probability with portfolio weights equal to the physical stopping time probabilities, plus an adjustment for the asset's systematic risk. In particular, the higher order moments of the stopping time, even when risk is purely idiosyncratic, affect the risk-adjusted discount rate.

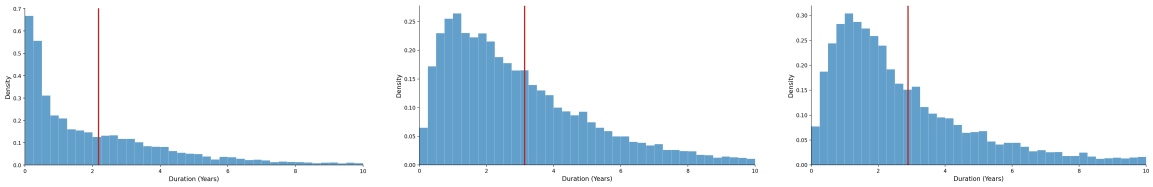
Example 3. *Consider Brealey et al. (2020) Chapter 9's Project A, which has cash flows of 100 million in years 1, 2, and 3, and a CAPM-computed risk-adjusted discount rate of 12%.*

Add now idiosyncratic timing risk: suppose that the cash flows in years 2 and 3 can arrive two years early or two years late, with equal probability. Solve for the new risk-adjusted rate r as:

$$\frac{100}{1.12} + \frac{1}{2}\left(\frac{100}{1.12^0} + \frac{100}{1.12^4}\right) + \frac{1}{2}\left(\frac{100}{1.12^1} + \frac{100}{1.12^5}\right) = \sum_{i=1}^{3} \frac{100}{(1+r)^i}.$$

Solving gives $r = 11.07\%$. Adding idiosyncratic timing risk changed the risk-adjusted discount rate by 93 basis points, or $0.93/12 * 100 = 7.8\%$ relative to the original systematic risk-based discount rate.

3.2 Innovation Firms



(A) Phase I (B) Phase II (C) Phase III

FIGURE 3: Empirical Phase Completion Times for Clinical Trials

Note: This figure plots, in years, the density of empirical completion times of each phase for all trials registered with clinicaltrials.gov. Data is taken from the AACT registry publicly provided by clinicaltrials.gov. To match the date range in DiMasi et al. (2016), data includes all trials from 1990 to 2013. The mean completion time of each phase is depicted as a vertical brown line.

Follow the risk-adjusted net present value method and in particular the usual taught intuition to adjust for risk using systematic risk through the CAPM (Berk and DeMarzo, 2011; Brealey et al., 2020), the industry standard in biotech is to obtain a risk-adjusted discount rate through the CAPM, then apply expected values to obtain a baseline valuation (Chandra and Mazumdar, 2024). DiMasi et al. (2016) describes the average biotech startup and Chandra and Mazumdar (2024), using these numbers, illustrate a typical industry benchmark valuation for this average biotech startup; the numbers are reproduced in Table 1, and the baseline valuation obtained is \$45.8M. Observe, however, that this takes expected cash flows and expected time of each phase into account.

However, what the baseline valuation does not include is uncertainty in how long it will take to complete each phase. Instead, it uses mean times. From a prior where outcome risk is similar to timing risk, and risk is treated as idiosyncratic, this approach

	Real Revenue ($M)	Time to Next Stage (Yrs)	Discount Rate	Probability of Occurrence
Phase I	-20.8	2	10.5%	100.0%
Phase II	-53.8	3	10.5%	59.5%
Phase III	-240	3	10.5%	21.1%
NDA Submission	-2.4	2	10.5%	13.1%
Approval	2838	-	10.5%	11.8%

TABLE 1: Biotech Valuation Illustration

Note: The first four rows (Phase I - NDA Submission) are taken directly from Chandra and Mazumdar (2024) Table 3, which in turn uses the average actual time and cost of biotech trials from DiMasi et al. (2016). The final row comes from Equation (1) of Chandra and Mazumdar (2024), computing the value at approval of the biotech startup given their assumptions on revenue and sales. Note that the time to next stage after approval is not applicable, beyond Chandra and Mazumdar (2024) assuming 10 years of sales.

may seem reasonable (Section 2.2); however, as Theorem 1 shows, the dispersion of timing affects value even when risk is purely idiosyncratic.

Augmenting the the valuation table using the empirical range obtained from clinicaltrials.gov (which approximately matches DiMasi et al. (2016) on the mean time for each phase) gives the empirical distribution over phase times reported in Figure 3. Notice that the empirical distributions at all phases are quite dispersed. Theorem 2 and Remark 3 therefore predict that taking timing risk into account should amplify both revenue and costs, since uncertainty in time raises value for positive shocks (revenue) and lowers value for negative shocks (costs).

Holding all else fixed and recomputing with timing risk changes the valuation from \$45.8M to \$53.0M, a difference of 16%. This decomposes as a +\$8.24 adjustment from revenues being more timing dispersed and a −\$1.03 adjustment from costs being more timing dispersed, amplifying both revenues and costs as the theory predicts; and since the present discounted values of revenues are higher in the industry standard example than the costs (as would typically be the case for viable biotech startups), amplifying revenues by more than costs.

3.3 Theoretical Co-Movements Illustrated

This next section provides an empirical illustration of the theoretical co-movements predicted by the theory. While using the relatively clean setting of option markets, the options that are cleanest in which to demonstrate the theoretical co-movements end up having short maturities, expiring in no more than seven days. As a result, while the application illustrates the directionality and timing vs. outcome differences

predicted by the theory, the magnitudes, because they take a week into account rather than months as in the prior examples, will naturally be lower.

In a cash-settled option with contractual strike price K, if the long side chooses to exercise, they receive $S - K$ dollars, where S is the value of the underlying. In European options, the long side can exercise only at the expiry of the option. In American options, the long side can exercise at any time up to and including expiry. As a result, there is an uncertain timing inherent in American options – timing flexibility for the long side, timing risk for the short side, which must be priced identically in equilibrium – which is not present in European options.

A clean illustration, as a laboratory, for the pricing of uncertain timing, is therefore to compare otherwise identical American and European options, observe the amount of timing risk in the American option, and check whether this timing risk affects price as per the theoretical predictions of the paper. This exercise is not a repudiation of canonical option pricing theory, but rather as an interpretation of the pricing of an American option that conceptually separates out timing from outcome risk.

The theory developed in the paper should apply for all options. However, it is cleanest to see in call options without dividends. The reason is that in such cases, frictionless markets (Merton, 1973) predict no early exercise and therefore no price difference. An empirical early exercise distribution, arising for example from wide bid-ask spread (Figlewski, 2022) or idiosyncratic behavior (Poteshman and Serbin, 2003), can truly be attributed to timing risk. From the perspective of the short side, whether or not that early exercise is optimal, it is timing risk they have to bear. On the other hand, in settings where early exercise is optimal, the empirical distribution of early exercise may be less cleanly interpreted as a timing risk distribution, since 'timing risk' would have to be measured relative to an optimality baseline which may be fragile and depend on expectations about dividends, etc.

To implement the clean laboratory, I use traded options that are contract-wise identical except for that one is American and one European. While most assets have only one option type associated with them, the Chicago Board Options Exchange (CBOE), one of the world's largest options exchanges, offers both American and European options on the S&P 100, which are identical on all contractual terms including multiplier, settlement details, strike price intervals, and platform rules; in their filing with the Securities and Exchange commission, CBOE explicitly says that the contracts are identical except for one being American and one being European (Securities

and Exchange Commission, 2007; CBOE, 2024*b*,*c*).[1]

Appendix C.2 details the data construction and early exercise imputation; as specified, I restrict to call options without dividend-relevant dates in their lifetimes, and otherwise implement standard data-cleaning measures for options. Early exercise is imputed using a combination of open interest and volume, following the literature (Barraclough and Whaley, 2012).

3.3.1 Systematic Timing and Outcome Risk Using Theorem 6

We first determine whether early exercise constitutes systematic outcome risk, systematic timing risk, both, or neither, using Theorem 6. The answer is not obvious. On the one hand, prior work finds that early exercise outside of the frictionless prediction can be due to person- or institution-specific factors (Poteshman and Serbin, 2003), which may be idiosyncratic. On the other hand, other work connects early exercise to bid-ask spreads (Figlewski, 2022), which, if liquidity correlates with intermediary frictions (Krishnamurthy, 2003; Maggiori, 2017), may be systematic. Theorem 6 allows us to answer this question precisely and to categorize each type of risk. Appendix C.3 details how to apply Theorem 6 empirically in this setting – it is a straightforward procedure – which may also be useful for empirically measuring systematic timing risk in other markets where timing is relevant.

Table 2 presents the results. Systematic outcome risk is statistically indistinguishable from zero regardless of whether the CAPM or Fama-French SDFs is used. Systematic timing risk is similarly statistically indistinguishable from zero regardless of whether the CAPM or Fama-French SDF is used. The results formalize the intuition from the options literature which finds that early exercise empirically is dependent on liquidity, short sale constraints, or irrationality; and provides quantitative evidence

[1]There also exist both American and European options on the S&P 500; in particular, American-style contracts on SPDR and European-style contracts on the S&P 500 index. However, these differ in at least three ways that OEX and XEO do not. First, the most liquid of the S&P 500 options, SPY and SPX, have different contract sizes. In contrast, OEX and XEO have the same contract size. Second, all S&P 500 index options have different CBOE trading fees than SPY, meaning that even after matching on contract size, one would not be able to differentiate fees from pricing premia (CBOE, 2024*a*). This is because it is not possible to see participant identity in the data, and the fees differ by participant type. In contrast, OEX and XEO have the same CBOE trading fees for all participants. Third, SPY versus SPX (and other S&P 500 index options) settle 15 minutes apart; that difference in settlement time can contribute to what is known as a wildcard premium (Fleming and Whaley, 1994), further affecting pricing. In contrast, OEX and XEO have the same settlement time.

against the hypothesis of early exercise due to systematic intermediary constraints.

TABLE 2: Early Exercise Constitutes Idiosyncratic Risk, Both in Outcomes and Timing: Applying Theorem 6 and Corollary 2

	Systematic Outcome Risk	**Systematic Timing Risk**
CAPM SDF	1.01	-1.66
	(1.10)	(1.42)
Fama-French SDF	-0.77	-1.49
	(2.32)	(1.97)

*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$.

Note: This table computes total systematic outcome risk and systematic timing risk per Theorem 6 and Corollary 2, using the explicit step-by-step procedure of detailed in Appendix C.3. The SDFs, computed based on a one year rolling basis, are derived from either the CAPM or the Fama-French three factor model; results are reported for both. Point estimates and standard errors are computed using 1000 bootstrap iterations where each draw is taken at the date level. Since both the SDF and the early exercise outcome distribution are computed in levels, the units in the table represent dollars per contract.

3.3.2 Higher Order Timing Moments and Pricing: An Empirical Illustration of Theorem 1

Because Section 3.3.1 shows that early exercise risk is purely idiosyncratic, the pricing results for idiosyncratic risk apply here. Standard theory (Berk and DeMarzo, 2011) predicts that higher moments of idiosyncratic outcome risk are not priced, whereas Theorem 1 predicts that higher moments of idiosyncratic timing risk are priced.

Since the American and European options are otherwise identical – differing only in the fact that the American can be exercised early – to compute a price measure of early exercise, I use $\text{Price}(American)_{i,t} - \text{Price}(European)_{i,t}$, where the price is computed as the midpoint quote, at the same minute, for the same strike price and expiry of an American option vs. an otherwise identical European one.

To compute timing and outcome moments, I use the realized mean and standard deviation of early-exercise timing and, conditional on exercise, the realized mean and standard deviation of payoff to the short side, respectively.

Table 3 regresses this price difference on timing- and outcome-risk moments, including liquidity and moneyness controls, with standard errors are clustered by day. The textbook prediction for outcome risk holds: higher moments are not priced, as outcome variance is insignificant throughout. Average outcome risk is marginally significant ($p < 0.1$) and economically small (about -0.002 per contract), implying that outcome risk plays little role in the American premium for these short-dated options.

By contrast, timing risk does correlate with price. Both the mean and variance of early-exercise timing are statistically significant ($p < 0.05$) predictors of the American-European price gap. A one-day earlier expected exercise lowers the premium by about \$6 to \$7.67 per contract, consistent with earlier exercise making the American more like a shorter-dated European and therefore less valuable.

Timing dispersion is also priced in the direction predicted by theory (Theorem 2). For contracts where early exercise is always bad for the short side, greater timing uncertainty raises the American premium: a one-percentage-point increase in timing standard deviation increases the price gap by about \$0.19 ($p < 0.05$). For contracts where early exercise is always good for the short side, theory predicts the opposite sign; there are only a handful of these cases in the data, which makes sense because one would not typically expect the long side to exercise for a negative payoff, and so the test is not sufficiently powered to detect significance. Nevertheless, the sign goes in the direction predicted by the theory: it is always negative across specifications.

Overall, Table 3 supports Theorem 1: higher moments of outcome risk are not priced, but higher moments of timing risk are. Quantitatively, timing uncertainty matters about half as much as expected timing. Using conditional sample averages, timing standard deviation implies roughly a \$5.13 higher American premium for the average standard deviation of early exercise timing conditional on early exercise occurring being 27 percentage points; versus \$9.60 to \$12.27 from expected exercise occurring 1.6 days early.

3.3.3 Risk-Free Term Structure and Pricing: An Empirical Illustration of Theorem 5

Thus far, Theorem 6 suggests that early exercise risk is purely idiosyncratic in both timing and outcomes. Taking that result as given, Theorem 1 provides a framework in which higher moments of timing risk, but not outcome risk, may be priced, while Theorem 2 helps interpret the signs of the estimated coefficients.

I next turn to Theorem 5, which suggests that the value of idiosyncratic timing risk should vary with the short rate. A simple way to illustrate this prediction is to regress the contract-level American-European price difference on the log short rate. The intuition is that, under the theorem, the American-European spread can be represented as the difference between a portfolio of shorter-duration zero-coupon bonds and a portfolio of longer-duration zero-coupon bonds (see Appendix C.4 for

TABLE 3: Theorem 1 Empirical Illustration: Higher Order Moments of Time Affect Price for Purely Idiosyncratic Risk

	Price(Am) − Price(Eur)		
	(1)	(2)	(3)
Timing Risk Mean	-7.67**	-6.24**	-6.00**
	(1.90)	(2.93)	(2.95)
Timing Risk Standard Deviation × 𝟙{All Negative Shocks}	0.19**	0.19**	0.18**
	(0.08)	(0.08)	(0.08)
Timing Risk Standard Deviation × 𝟙{All Positive Shocks}	-0.09	-0.09	-0.09
	(0.09)	(0.09)	(0.09)
Outcome Risk Mean	-0.00*	-0.00*	-0.00*
	(0.00)	(0.00)	(0.00)
Outcome Risk Standard Deviation	-0.02	-0.02	-0.02
	(0.02)	(0.02)	(0.02)
Moneyness Controls	N	Y	Y
Liquidity Controls	N	N	Y
R^2	0.02	0.02	0.04
N	4181	4181	4181

*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$

Note: Standard errors, in parentheses, are clustered at the date level. The dependent variable is the midpoint quote of an American option minus the midpoint quote of a European option observed at the same time, with the same strike price, underlying, and all other contract specifications, as described in Section C.2.1. This midpoint difference is multiplied by 100 to convert from index points to dollar units. The early exercise timing distribution for each contract is imputed as described in Section C.2.2. The timing risk mean is calculated as follows. We first calculate the average number of days to early exercise conditional on early exercise occurring. We then compute the timing risk mean as the difference between the number of days to expiry and this calculation: 𝟙{early exercise occurs with positive probability} (days to expiry - avg. days to early exercise conditional on early exercise occurring. Therefore, a zero indicates no difference in days to expiry and average days to early exercise, while the more positive timing mean gets, the earlier that early exercise (conditional on its occurring) happens on average. The timing risk standard deviation is the standard deviation of the number of trading days until early exercise, conditional on early exercise. The indicator for all negative shocks takes on a value of 1 if, for the contract, at all days where the American contract is exercised with positive probability, the payoff to the short side on that day of early exercise is negative (e.g. if the strike price of the option is lower than the value of the underlying). The indicator for all positive shocks takes on a value of 1 if, for the contract, at all days where the American contract is exercised with positive probability, the payoff to the short side on that day of early exercise is positive (e.g. if the strike price of the option is higher than the value of the underlying). Timing standard deviation is interacted with each of these indicators to study the sign predictions of Theorem 2. The outcome risk mean is the realized average dollar payoff to the short side conditional on early exercise, formed as a weighted average across possible early exercise days using the imputed exercise-timing weights. The outcome risk standard deviation is formed analogously: it is the empirical standard deviation of the dollar payoff to the short side conditional on early exercise. All regressions include a control for days to expiry and an indicator for whether a timing distribution (e.g at least two dates with early exercise mass) exists for the contract at hand. The moneyness control is the amount of dollars the option is in the money. The liquidity control is percentage difference in bid-ask spreads between the two contracts: the bid-ask spread (in dollars) of the American, minus the bid-ask spread of the European, all over the bid ask spread of the American.

the equations). Since longer-duration bonds are more sensitive to interest rates (Berk and DeMarzo, 2011), an increase in the short rate should, all else equal, increase this spread. Table B.2 in Appendix C.5 is consistent with that prediction: among contracts with empirical early exercise, a one-unit increase in the log three-month

Treasury rate is associated with a \$4.15-\$4.69 increase in the American-European price difference ($p < 0.1$ across specifications). Replacing the three-month Treasury with the one-month Treasury yields very similar results (Appendix C.7).

I then try for a more direct illustration of the theorem by constructing the realized zero-coupon-bond equivalent implied by observed early exercise. Regressing the American-European price difference on this realized ZCB-equivalent produces the predicted sign, although not the one-for-one magnitude implied by the theorem: a \$1 increase in realized present-value loss to the short side is associated with only about a \$0.01 increase in the spread. One possible reason is that the theorem is stated in expectations, whereas the realized ZCB-equivalent is an ex post object.

To approximate the ex ante quantity in the theorem, Table 4 uses a 60%/40% train-test split. In the training sample, realized ZCB-equivalents are projected on time-t observables only – moneyness, the short rate, and days to expiry – and the fitted coefficients are then used to form an out-of-sample proxy for expected ZCB-equivalents in the test sample. Regressions of the American-European price difference on this expectation proxy yield coefficients of about 0.7 ($p < 0.1$ across specifications), both with and without separate controls for liquidity and moneyness. These estimates are statistically indistinguishable from -1, the sign and magnitude suggested by the theorem.[2]

The economic magnitudes are also suggestive. The expectation proxy alone explains about 9% of out-of-sample variation in the American premium, while liquidity and moneyness controls add only 3 percentage points of explainatory power. Taken together, these results suggestively illustrate the co-movement of idiosyncratic timing risk with expected zero-coupon-bond equivalents.

3.4 Potential Avenues for Future Work

Understanding the components risk and its pricing is important across financial and economic settings. Below is a non-exhaustive set of economic settings in which understanding the price of timing risk may be relevant.

Climate Disasters. Both the impact as well as the timing of climate disasters are uncertain. It may be helpful to understand the timing component of climate risk as

[2]The relevant benchmark is -1, rather than 1, because a more negative expected ZCB-equivalent corresponds to a larger expected loss to the short side.

TABLE 4: Out-of-Sample Expected ZCB-Portfolio-Equivalent Coefficients

	Price(Am) − Price(Eur)		
	(1)	(2)	(3)
Expected ZCB-Portfolio Equivalent	-0.71**	-0.72**	-0.71**
	(0.36)	(0.37)	(0.37)
Moneyness Controls	N	Y	Y
Liquidity Controls	N	N	Y
R^2	0.09	0.11	0.12
N	1673	1673	1673

*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$.

Note: Standard errors, in parentheses, are computed via 1000 bootstrap iterations, with draws taken at the date level. The dependent variable is the midpoint quote of an American option minus the midpoint quote of a European option observed at the same time, with the same strike price, underlying, and all other contract specifications, as described in Section C.2.1. This midpoint difference is multiplied by 100 to convert from index points to dollars per contract. The full sample is split into 60% train and 40% test, by dates: the train sample contains all data until 3/15/2017 and the test sample is all data after that. Within the train sample, the realized ZCB portfolio equivalent difference between the American and European is regressed on short rate, the moneyness ratio (underlying/strike) times an indicator for being in the money, the moneyness ratio times an indicator for being out of the money, an indicator for being in the money, and dummies for the number of days to expiry. The coefficients on these variables obtained from this regression in the training set are used to construct contemporaneous expectations for the zero-coupon-bond portfolio equivalent in the training set. This contemporaneous expectation is the main regressor. The moneyness control is the amount of dollars the option is in the money. The liquidity control is percentage difference in bid-ask spreads between the two contracts: the bid-ask spread (in dollars) of the American, minus the bid-ask spread of the European, all over the bid ask spread of the American. The reported R^2 is the out-of-sample R^2 on the entire test set.

distinct from the outcome component. For example, climate policy shocks can have both timing and outcome components.

Mortgages. In mortgages, prepayment risk is uncertainty over the time at which the mortgage will be paid back. In the literature, prepayment is often predicted using projected interest rates, but part of the pricing of MBS appears to reflect something beyond interest rates (Scharfstein and Sundaram, 2016; Chen et al., 2020). It may be interesting to explore idiosyncratic risk at the aggregate level in MBS and the extent to which this affects prices.

Labor Search. Once unemployed, the time until finding employment is uncertain, in addition to the wage once employed. Canonical search and matching models in labor employ outcome uncertainty but on the time dimension, only consider the expected time rather than taking the full distribution of timing risk into account (Mortensen and Pissarides, 1994). The results in this paper suggest that the full distribution of timing risk may affect optimal search and valuation of unemployment.

Private Equity. The time of capital call for private equity is uncertain, as studied through an additive utility model in Gourier et al. (2024); however, this approach embeds strong directional assumptions. It may be interesting to study valuation of uncertain timing in private equity in a more general and model-free way.

Policy Decisions. A number of policy decisions involve uncertainty both over the outcome and over the timeline of implementation. For example, infrastructure projects often include delays, as does the passage of legislation. A pricing of uncertain timing may therefore be relevant in policy contexts.

Safe Assets. Safe assets are typically thought to be so because of their lack of outcome risk (Krishnamurthy, 2010). Timing risk may also play a role in the definition and pricing of safe assets.

4 Conclusion

Using only a no-arbitrage assumption, I showed how the valuation of timing risk and outcome risk conceptually differ. For example, higher order moments of timing risk affect prices, while higher order moments of outcome risk do not; and any SDF that prices all assets must feature a sign-based asymmetry for timing risk that need not be present for outcome risk. These differences survive diversification. I characterized the pricing of both idiosyncratic and systematic timing risk, including formalizing the systematic timing vs. outcome risk in any asset or economic setting. Three sample applications, ranging from practices used in industry to discount rate intuition, to market data from options, illustrated the findings of several theorems. The conceptual or empirical application of the theory developed here may be interesting to study across economic settings.

References

Atkeson, Andrew G., Andrea L. Eisfeldt, and Pierre-Olivier Weill. 2015. "Entry and Exit in OTC Derivatives Markets." *Econometrica*, 83: 2231–2292.

Barraclough, Kathryn, and Robert E. Whaley. 2012. "Early Exercise of Put Options on Stocks." *The Journal of Finance*, 67: 1423–1456.

Belo, Frederico. 2010. "Production-based measures of risk for asset pricing." *Journal of Monetary Economics*, 57: 146–163.

Berk, Jonathan B. 1999. "A Simple Approach for Deciding When to Invest." *American Economic Review*, 89: 1319–1326.

Berk, Jonathan B., and Peter M. DeMarzo. 2011. *Corporate Finance. The Prentice Hall series in finance.* 2., ed ed., Boston:Prentice Hall.

Berk, Jonathan B., and Jules van Binsbergen. 2016. "Assessing asset pricing models using revealed preference." *Journal of Financial Economics*, 119: 1–23.

van Binsbergen, Jules. 2025. "Duration-Based Stock Valuation." *Journal of Financial Economics*, Forthcoming.

Brealey, Richard A., Stewart C. Myers, and Franklin Allen. 2020. *Principles of corporate finance. The McGraw-Hill/Irwin series in finance, insurance, and real estate.* Thirteenth edition, international student edition ed., New York, NY:McGraw-Hill Education.

CBOE. 2024*a*. "Cboe U.S. Options Fee Schedules."

CBOE. 2024*b*. "OEX Options Product Specifications."

CBOE. 2024*c*. "XEO Options Product Specifications."

Chandra, Amitabh, and Sumon Mazumdar. 2024. "Biotech Asset Valuation Methods: A Practitioner's Guide." *Journal of Investment Management*, 22: 36–57.

Chen, Hui, Michael Michaux, and Nikolai Roussanov. 2020. "Houses as ATMs: Mortgage Refinancing and Macroeconomic Uncertainty." *The Journal of Finance*, 75: 323–375.

Chesson, Harrell W., and W. Kip Viscusi. 2003. "Commonalities in Time and Ambiguity Aversion for Long-Term Risks." *Theory and Decision*, 54: 57–71.

Cochrane, John H. 2014. "A Mean-Variance Benchmark for Intertemporal Portfolio Theory." *The Journal of Finance*, 69: 1–49.

DeJarnette, Patrick, David Dillenberger, Daniel Gottlieb, and Pietro Ortoleva. 2020. "Time Lotteries and Stochastic Impatience." *Econometrica*, 88: 619–656.

DiMasi, Joseph A., Henry G. Grabowski, and Ronald W. Hansen. 2016. "Innovation in the pharmaceutical industry: New estimates of R&D costs." *Journal of Health Economics*, 47: 20–33.

Doepke, Matthias, and Martin Schneider. 2017. "Money as a Unit of Account." *Econometrica*, 85: 1537–1574.

Duffie, Darrell. 2001. *Dynamic asset pricing theory.* . 3rd ed ed., Princeton (N.J.):Princeton university press.

Eisfeldt, Andrea L., and Dimitris Papanikolaou. 2014. "The Value and Ownership of Intangible Capital." *American Economic Review*, 104: 189–194.

Epstein, Larry G., Emmanuel Farhi, and Tomasz Strzalecki. 2014. "How Much Would You Pay to Resolve Long-Run Risk?" *American Economic Review*, 104: 2680–2697.

Farhi, Emmanuel, and Matteo Maggiori. 2018. "A Model of the International Monetary System." *The Quarterly Journal of Economics*, 133: 295–355.

Figlewski, Stephen. 2022. "An American Call Is Worth More Than a European Call: The Value of American Exercise When the Market Is Not Perfectly Liquid." *Journal of Financial and Quantitative Analysis*, 57: 1023–1057.

Fleming, Jeff, and Robert E. Whaley. 1994. "The Value of Wildcard Options." *The Journal of Finance*, 49: 215–236.

Giglio, Stefano, Matteo Maggiori, and Johannes Stroebel. 2016. "No-Bubble Condition: Model-Free Tests in Housing Markets." *Econometrica*, 84: 1047–1091.

Giglio, Stefano, Bryan Kelly, and Johannes Stroebel. 2021. "Climate Finance." *Annual Review of Financial Economics*, 13: 15–36.

Gourier, Elise, Ludovic Phalippou, and Mark M. Westerfield. 2024. "Capital Commitment." *The Journal of Finance*, 79: 3407–3457.

Graham, John R., Scott B. Smart, Christopher Adam, and Brindha Gunasingham. 2017. *Introduction to corporate finance.* . 2nd Ed. ed., Australia ; Mason, Ohio:South-Western/Cengage Learning.

Graham, John R. 2022. "Presidential Address: Corporate Finance and Reality." *The Journal of Finance*, 77: 1975–2049.

Hansen, Lars Peter, and Ravi Jagannathan. 1991. "Implications of Security Market Data for Models of Dynamic Economies." *Journal of Political Economy*, 99: 225–262.

Kacelnik, Alex, and Melissa Bateson. 1996. "Risky Theories-The Effects of Variance on Foraging Decisions." *American Zoologist*, 36: 402–434.

Kacperczyk, Marcin T. 2026. "What is the Carbon Premium a Premium on?"

Krishnamurthy, Arvind. 2003. "Collateral constraints and the amplification mechanism." *Journal of Economic Theory*, 111: 277–292.

Krishnamurthy, Arvind. 2010. "How Debt Markets Have Malfunctioned in the Crisis." *Journal of Economic Perspectives*, 24: 3–28.

Maggiori, Matteo. 2017. "Financial Intermediation, International Risk Sharing, and Reserve Currencies." *American Economic Review*, 107: 3038–3071.

Maggiori, Matteo. 2022. "International macroeconomics with imperfect financial markets." In *Handbook of International Economics*. Vol. 6, 199–236. Elsevier.

Merton, Robert C. 1973. "Theory of Rational Option Pricing." *The Bell Journal of Economics and Management Science*, 4: 141.

Modligliani, Franco, and Merton Miller. 1958. "The Cost of Capital, Corporation Finance and the Theory of Investment." *American Economic Review*, 48.

Mortensen, D. T., and C. A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *The Review of Economic Studies*, 61: 397–415.

Myers, Stewart C., and Stuart M. Turnbull. 1977. "Capital Budgeting and the Capital Asset Pricing Model: Good and Bad News." *The Journal of Finance*, 32: 321–333.

Onay, Selcuk, and Ayse Onculer. 2007. "Intertemporal choice under timing risk: An experimental approach." *Journal of Risk and Uncertainty*, 34: 99–121.

Poteshman, Allen M., and Vitaly Serbin. 2003. "Clearly Irrational Financial Market Behavior: Evidence from the Early Exercise of Exchange Traded Stock Options." *The Journal of Finance*, 58: 37–70.

Scharfstein, David, and Adi Sundaram. 2016. "Market Power in Mortgage Lending and the Transmission of Monetary Policy." *Harvard Business School Working Paper*.

Schneider, Martin, and Monika Piazzesi. 2015. "Payments, Credit and Asset Prices." *Meeting Papers, Society for Economic Dynamics*.

Securities and Exchange Commission. 2007. "Release No. 34-56191; File No. SR-CBOE-2007-79."

Stein, Jeremy. 1989. "Overreactions in the Options Market." *The Journal of Finance*, 44: 1011–1023.

Wachter, Jessica A. 2002. "Portfolio and Consumption Decisions under Mean-Reverting Returns: An Exact Solution for Complete Markets." *The Journal of Financial and Quantitative Analysis*, 37: 63.

Wachter, Jessica A. 2013. "Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility?" *The Journal of Finance*, 68: 987–1035.

Xu, Nancy. 2021. "Procyclicality of the comovement between dividend growth and consumption growth." *Journal of Financial Economics*, 139: 288–312.