Timing Risk and Asset Prices

Indira Puri*
NYU Stern

February 9, 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

*Email: indira.puri@stern.nyu.edu. I thank seminar and conference participants at the the Canadian Derivatives Institute, CEPR ESSFM, the FMA-UC Davis Napa Conference (Best Paper Award), the LBS Summer Symposium (Best Paper Award), MIT, NBER SI: Asset Pricing, the Texas Finance Festival, SITE: The Macroeconomics of Uncertainty and Volatility, Stanford, the University of Cambridge, the University of Maryland, the University of Michigan, the University of Minnesota, and the Virtual Derivatives Seminar for helpful comments. I further thank Yakov Amihud, Jules van Binsbergen, 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, Lasse Pedersen, Thomas Philippon, Matt Richardson, Tony Saunders, Alexi Savov, Marti Subrahmanyam, Johannes Stroebel, Quentin Vandeweyer, Jessica Wachter, and Jeff Wurgler for valuable feedback. An earlier version of this paper was circulated under the title *Certain vs. Uncertain Timing: Financial Markets and Pricing Implications.*

into outcome and timing components and formalize systematic timing risk. An application of the theory in options markets illustrates the difference between timing and outcome risk, and the characterization theorems. 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 successful treatment effects 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, both technological revolutions (Pastor and Veronesi, 2009) and corporate decisions in innovative firms (Eisfeldt and Papanikolaou, 2014) include both 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. 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. That is, 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.

I illustrate the theory using two financial products which are contractually identical except for timing risk. An application of the decomposition theorem shows timing risk in this setting to be idiosyncratic. Given this, the theorems for idiosyncratic risk can be applied and illustrated. Consistent with the theory, in the data higher-order moments of timing risk affect prices while those of outcome risk do not, the sign of the timing risk coefficient varies with the direction of the shock, and the price of timing risk moves with the risk-free term structure. Together, the theory and empirical illustration suggest that timing risk may deserve its own treatment 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 (Berk and DeMarzo, 2011). 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; Gomes and Schmid, 2021). 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 asymme-

try – 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.

I illustrate the theory in as clean a financial setting as possible. In particular, I use the options market as a laboratory. While timing risk matters in other important markets, including real options (Berk, 1999), the valuation of private firms (Schmid and Steri, 2023), mortgages (Chen et al., 2020), and debt markets (Krishnamurthy, 2010; Kuehn and Schmid, 2014), these markets typically feature outcome risk intertwined with timing risk as well as unobserved individual heterogeneity. Options have the advantage of allowing for observability of timing risk and therefore a relatively clean illustration of the theory in market data. They have the additional advantage of being a sophisticated market which is large in its own right: my setting features options with a total notional value of $16 billion.

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. In equilibrium, the price of timing risk to the short side must equal the price of timing flexibility to the long side, and should be given by the pricing theorems introduced in this paper. I focus on call options without dividends in their lifetime because standard derivatives pricing theory predicts no price difference for these options, and observed deviations have typically been dismissed as irrational anomalies (Merton, 1973; Poteshman and Serbin, 2003) or as liquidity events (Figlewski, 2022). A theory of timing risk and how it affects asset prices says – regardless of the source of early exercise, its existence raises timing risk for the short side. The theory therefore makes predictions on how this timing risk should be priced.

Comparing real-world American and European options which are otherwise identical, I impute the realized early exercise distribution using open interest and volume to obtain realized outcome and timing risk. I apply Theorem 6, illustrating how to practically use the systematic outcome-timing risk decomposition in data, to test for whether early exercise constitutes idiosyncratic or systematic risk, both in outcomes and timing. A priori, the answer to this question is not obvious; while some papers find early exercise to be idiosyncratic (Poteshman and Serbin, 2003), others stress the role of liquidity (Figlewski, 2022) which may be impacted by intermediary frictions that tend to be systematic (Maggiori, 2017). Quantifying systematic timing risk and systematic outcome risk through the theorem, I find that both are statistically indistinguishable from zero – indicating that both timing and outcome risk in this case are idiosyncratic.

Given that the risk is idiosyncratic, the theorems for idiosyncratic timing risk can be applied. I next illustrate Theorem 1, which predicts that higher order moments of timing risk should affect price but of outcome risk should not. Illustrating the predictions of this theorem, and the difference between timing and outcomes, the option price difference between otherwise identical American and European options loads in a statistically significant way on timing risk variance ($p < 0.05$), but not statistically significantly on outcome risk variance.

In addition, consistent with Theorem 2, the sign of the coefficient on timing risk is positive when the shock to the short side is negative, and is negative when the shock

to the short side is positive. Thus, both in statistical significance and in sign, the pricing of timing risk is consistent with the theory and illustrates differences between timing and outcome risk.

I then illustrate Theorem 5, which predicts that the pricing of uncertain timing is characterized by the risk-free term structure. Using contemporaneous short rates, I find correlational evidence consistent with the theorem: the option price difference between otherwise identical American and European options is higher per contract ($p < 0.05$) when short rates increase, even after accounting for liquidity and moneyness controls. Since the theorem says that the difference between an American and European option when early exercise occurs is equivalent to an appropriately constructed zero coupon bond portfolio – in this case of shorter-dated bonds less longer-dated bonds – and longer dates bonds are more sensitive to interest rates (Chen, 2010), the sign of the coefficient is also in the direction predicted by the theory.

Finally, as an additional illustration of Theorem 5, I split the sample into a 60-40 training set and a test set, where the test set chronologically follows the train set. The theorem predicts that the expected ZCB portfolio equivalent should move one for one with the pricing of uncertain timing. Because expectations are not observed but realizations are, I use the train set to regress realizations on contemporaneously observed variables for the option contracts (moneyness, days to expiry, and the short rate). Using the coefficients from these regressions obtained from the train set, I construct a proxy for expected ZCB portfolio equivalents in the test set. Because all variables are contemporaneously observed, and the training set chronologically precedes the test set, this method avoids look-ahead bias. In the test set, regressing the price difference between American and European options on the expected ZCB portfolio equivalent yields a coefficient that is statistically significant ($p < 0.05$) and indistinguishable from 1 across specifications, consistent with the prediction of Theorem 5.

These results in options illustrate various predictions of the theory. They also conceptually organize in a risk-based fashion a pricing difference that has historically been dismissed as an anomaly due to irrational behavior.

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 pric-

ing 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 (Papanikolaou, 2011; van Binsbergen, 2016; Gomes and Schmid, 2021; Maggiori, 2022). 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 (Papanikolaou, 2011), preference-based (van Binsbergen, 2016), belief-based (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 corporate bonds (Berk and DeMarzo, 2011). The options I study to illustrate the theory have previously been studied in Stein (1989) and Fleming and Whaley (1994), among others; the contribution is to empirically illustrate the novel theoretical results. As a supplemental contribution, these results provide a risk-based interpretation to a pricing deviation from standard theory.

While the paper studies time, it is conceptually distinct from other literatures on time-related concepts such as the elasticity of intertemporal substitution (Vissing-Jorgensen, 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 illustrates the findings of the theory in real-world market data. Section 4 discusses possible applications across economics. Section 5 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 $c = 0$.*

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 1.

Remark 1 (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.*

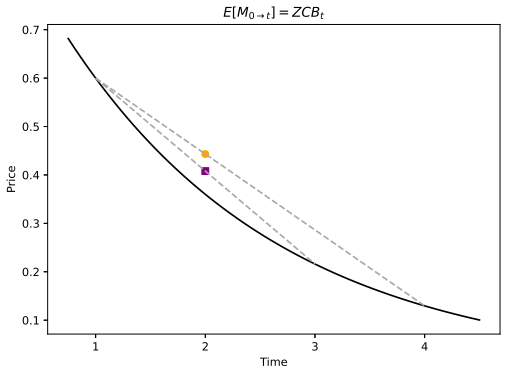
Theorem 1, in contrast, shows that the higher order moments of the stopping time for an asset with uncertain timing do affect the price, regardless of which SDF one uses.

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.*

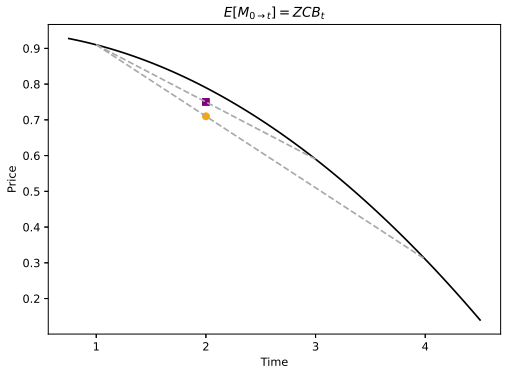
The intuition is as follows. When an asset with uncertain outcomes has purely idiosyncratic risk, its price as discussed in Remark 1 becomes an expectation of the SDF times an expectation of the asset payoff. When the asset has uncertain timing, the price is similarly an expectation of the SDF, but that 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

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.

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 1. 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

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

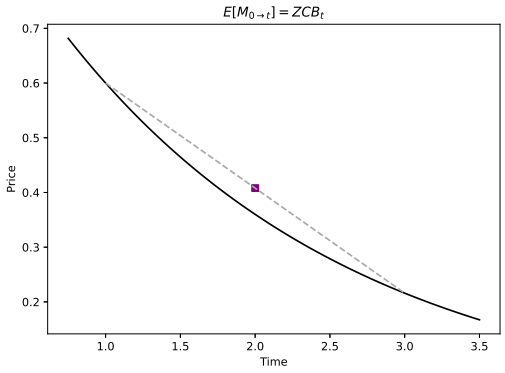
In particular, uncertainty can be valued higher, lower, or the same as, certainty.

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.

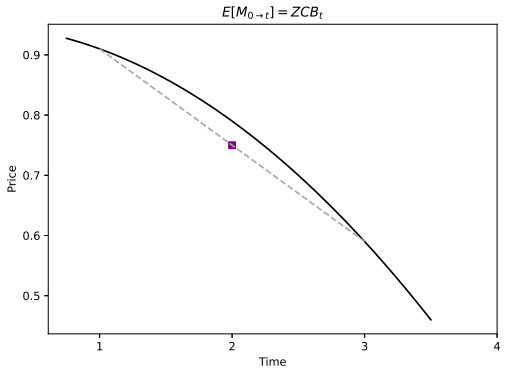
Corollary 1 (Asymmetry)**.** *Any SDF that has a higher valuation for uncertain timing than certain-time equivalents for positive shocks must have a higher valuation for certain-time equivalents than uncertain timing for negative shocks. Any SDF that has a higher valuation for certain-time equivalents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

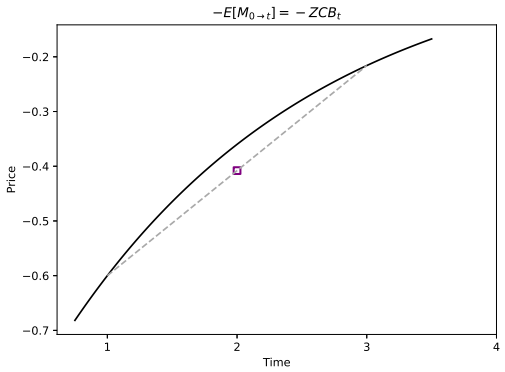
FIGURE 2: Certainty vs. Uncertainty and Asymmetry



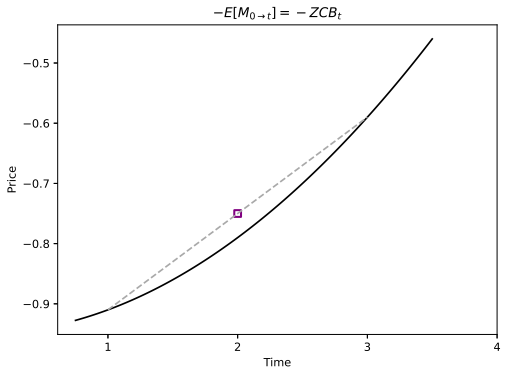
(A) Sample SDF 1 - Receiving $1 at Time t



(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.

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 2 (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:* $2y'(t) \leq Ty''(t)$.

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

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

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 $P_{t'}^{(n)}$ 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 3 (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*

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

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 4 (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, $\mathrm{Cov}(M_{0\to t}, Y \mid \tau = t) = \mathrm{Cov}(M_{0\to t}, a \mid \tau = t) = 0$, 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 $\mathrm{Cov}\big(\mathbb{E}[M_{0\to t} \mid \tau],\, \mathbf{1}_{\{\tau=t\}}\mathbb{E}[Y \mid \tau]\big)$ is zero, and only $\mathrm{Cov}(M_{0\to t^*}, Y \mid \tau = t^*) = \mathrm{Cov}(M_{0\to t^*}, Y)$ 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 Empirical Illustration in Market Prices

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.

The ideal test, 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.

I seek to implement this ideal illustration as closely as possible. 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] I impute the realized

[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

timing risk for the American option, and directly apply the theory of the paper to illustrate its findings in market prices.

3.1 Data

For best precision, I obtain data directly from CBOE, the exchange that offers the American (OEX) and European (XEO) contracts. I use the last quote within normal trading hours. The CBOE data is available starting in 2012. I therefore use data from January 2012 - December 2019.[2]

To focus cleanly on the theoretical predictions of this paper only, rather than derivatives pricing theory, I select a setting where standard pricing theory would predict no price difference between the American and European, but Theorems 1-6 would allow for a pricing of uncertain timing, even if idiosyncratic. I therefore focus only on call options without dividends in their lifetime; textbook derivatives pricing theory predicts no price difference between the American and European for these options (Merton, 1973). To the extent that early exercise has been found in such options, it has been dismissed as idiosyncratic or suboptimal (Poteshman and Serbin, 2003). The theory of this paper says that, even so, such timing risk should affect prices; and more generally, the components of an early exercise premium can be decomposed into outcome and timing risk, and the timing risk components interpreted and predicted. The theory thus provides a risk-based interpretation of deviations from standard theory.

Data on dividends paid by companies in the S&P 100, and end-of-day S&P 100 levels, comes from Compustat. As explained above, for as clean an illustration of the theory as possible, I consider only options in whose lifetime none of the 100 companies

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.

[2]The S&P 100 option contracts OEX and XEO also have the most traded volume in this period; volume drops starting in 2020, perhaps due to the popularity of other index options such as the S&P 500.

in the S&P 100 has dividend-relevant date. This reduces the sample substantially and makes most options short dated (no more than seven days to expiry), but makes for a clean test; given the size of the dataset, the cleaning still yields over 4,000 observations. To focus only on options in which timing risk may reasonably play a role, I remove any options quotes which expire on the same day as the quote (there is no timing risk in such options because all exercise occurs on the expiry day). Following standard data cleaning procedure for options data (Dew-Becker and Giglio, 2023), I drop any quotes with a bid less than 0 or where the bid-ask spread is negative. Summary statistics for the cleaned data are reported in Appendix B.1. The notional value (strike $\times$ multiplier $\times$ open interest) of all options in the final cleaned dataset is \$16 billion.

3.2 Empirical Early Exercise

To impute early exercise from the data, I use the change in open interest less the total traded volume for each day, clipped below by zero. The reason that this is a proxy for early exercise is as follows. When an option is exercised early, the volume traded does not change, but open interest decrements by one. As a result, $(-\Delta OI) - V$ changes by 1 when early exercise occurs. Formally, open interest denotes the number of derivatives contracts outstanding. A change in open interest occurs only due to early exercise, to an existing derivatives contract being closed, or to a new derivatives contract being opened. Formally, $\Delta OI = O - C - EE$ (ETrade from Morgan Stanley, 2021). As a result, $(-\Delta OI) = EE + C - O$. Volume, in turn increases with either an open or a close, but does not change with early exercise: $V = O + C$. Together, $(-\Delta OI) - V = EE + C - O - C - O = EE - 2O \leq EE$, so that the measure $max((-\Delta OI) - V, 0)$ is a lower bound on the amount of early exercise that occurs.

To sanity check the measure, I implement it for both American and European options. For the European options, we expect the measure to be exactly zero for all contracts because early exercise is not allowed. Zero is indeed what the implementation on the European options (XEO) empirically produces, indicating that the measure is not contaminated by other data artifacts.

This lower bound is a precise estimate of early exercise when volume is equal to zero. This is because volume $V = O + C$, the number of openings plus the number of closings, which in turn must always be weakly positive. When volume is exactly zero,

the number of openings O is therefore equal to zero, and so $(-\Delta OI)-V = EE-2O = EE$. In the options data, after all other data cleaning measures are implemented (e.g., no dividends in the lifetime of the option; non-zero bids; and asks at least as large as bids), 74% of the time when the early exercise lower bound $max((-\Delta OI) - V, 0)$ is strictly positive, total volume for that contract that day is zero, indicating a point estimation of early exercise. For the remaining 26% of cases where American volume is positive, I use European volume and open interest for the identical contract to back out a precise fraction of open and closes in the European counterpart, and apply this fraction to the imputation (formally, $-(\Delta \text{OI}_{European}) - V_{European} = -2O_{European}$; then impute Early Exercise$_{American} = (-\Delta OI_{American}) - V_{American} + 2O_{European}$ whenever $V_{American} > 0$). In the remaining minority of cases where American volume is nonzero and European volume is zero, I assume that opens equal closes (50% each).[3] This is similar to the 45% close fraction observed empirically in the European volume data for this sample. As a sanity check, running the entire procedure (both the precise bound and the later imputation) on European options correctly gives zero early exercise everywhere. Given the high degree of precision in this sample, this is the measure I use as a proxy for empirical early exercise. Using open interest to obtain an imputation of early exercise has precedent in the finance literature (for example, Barraclough and Whaley (2012)).

3.2.1 Determinants of Early Exercise

This paper is agnostic as to why early exercise occurs. It says only: given that early exercise occurs, it must be priced; and a theory of timing vs. outcome risk helps price it. However, for completeness, I note that in line with prior work that finds liquidity drivers of early exercise (Figlewski, 2022), 71% of the time, conditional on empirical early exercise occurring, the bid is less than the option's intrinsic value. Such instances violate the frictionless setting of Merton (1973). While this may help rationalize the source of early exercise, the contribution of the theory is to understand how American options are priced from a risk-based perspective both on timing and on options, irrespective of the source of early exercise.

[3]For robustness, Appendix B.2 reproduces the results using an alternative default open fraction of 75%. Because these are options, the fraction of opens must always be weakly greater than the fraction of closes (since a position can expire without being closed). The appendix finds qualitatively similar results, illustrating Theorem 6 and Corollary 2 and the predictions of Theorem 1.

3.3 Illustrating the Theory: Pricing Outcome vs. Timing Risk

3.3.1 Testing for Systematic Timing and Systematic 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 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 liquidity (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.

Following standard practice in empirical asset pricing, assume that the stochastic discount factor (SDF) process and early-exercise is stationary (or, more generally, satisfies conditions ensuring a law of large numbers for our panel) (Xu, 2017; Maggiori, 2017). These assumptions justify estimating conditional moments by pooling across contracts and dates.

Let $\tau_{i,t}$ denote the exercise stopping time for contract i observed at date t. The stopping time takes values in $\{t, t+1, \ldots, T_i\}$, where T_i is the expiry date, encompassing both early exercise (at any date strictly before T_i) and exercise at expiry. For American options, exercise may occur at any date; for European options, exercise occurs only at expiry, so $\tau_{i,t} = T_i$ with probability one.

Define days before expiry $u \equiv T_i - \tau_{i,t}$, so that $u = 0$ corresponds to exercise at expiry and $u > 0$ to early exercise u days before expiry. The stopping time distribution is then characterized by

$$W_{i,t}(u) \equiv \mathbb{P}\big(\tau_{i,t} = T_i - u\big), \qquad u = 0, 1, \ldots, T_i - t,$$

with $\sum_{u=0}^{T_i - t} W_{i,t}(u) = 1$ for each (i, t).

To estimate the systematic risk decomposition from Corollary 2, proceed as follows.

1. Consideration set. Retain all observations (i, t) for which the contract has not yet expired and the stopping time distribution can be estimated over the relevant horizon window $u = 0, 1, \ldots, H$.

2. SDF realizations. For each retained (i,t) and each $u = 0, 1, \ldots, H$, compute the realized SDF from t to $t+u$, denoted $M_{t,t+u}$, using the chosen factor model (e.g., CAPM or Fama–French).

3. Stopping-time probabilities. For each retained (i,t), compute the stopping-time distribution $\{W_{i,t}(u)\}_{u=0}^{H}$ and the conditional shares

$$q_{i,t}(u) \equiv \frac{W_{i,t}(u)}{\sum_{v=0}^{H} W_{i,t}(v)}.$$

4. Mega-states. Each value of u defines a mega-state $S^{(u)} = \bigcup_{i,t}\{\tau_{i,t} = T_i - u\}$, representing the event that exercise occurs u days before expiry.

5. Stopping-time distribution across mega-states. Estimate the probability that exercise occurs u days before expiry:

$$\widehat{P}(\tau = u) \equiv \frac{\sum_{i,t} W_{i,t}(u)}{\sum_{i,t}\sum_{v=0}^{H} W_{i,t}(v)}.$$

6. Conditional moments within each mega-state. For each u, estimate the conditional expected SDF and payoff using $q_{i,t}(u)$ as weights:

$$\hat{m}_u \equiv \frac{\sum_{i,t} q_{i,t}(u)\, M_{t,t+u}}{\sum_{i,t} q_{i,t}(u)}, \qquad \hat{Y}_u \equiv \frac{\sum_{i,t} q_{i,t}(u)\, Y_{i,T_i-u}}{\sum_{i,t} q_{i,t}(u)},$$

where Y_{i,T_i-u} is the payoff from exercise at u days before expiry. From the perspective of the short side, this is $Y_{i,T_i-u} = -(S_{T_i-u} - K_i)$.

7. Conditional covariance within each mega-state. For each u, estimate

$$\widehat{\mathrm{Cov}}_u(M,Y) \equiv \frac{\sum_{i,t} q_{i,t}(u)\left(M_{t,t+u} - \hat{m}_u\right)\left(Y_{i,T_i-u} - \hat{Y}_u\right)}{\sum_{i,t} q_{i,t}(u)}.$$

8. Total systematic outcome risk. The overall outcome-risk contribution is the probability-weighted sum of within-mega-state covariances:

$$\widehat{\text{Systematic Outcome Risk}} \equiv \sum_{u=0}^{H} \widehat{P}(\tau = u)\, \widehat{\mathrm{Cov}}_u(M,Y).$$

9. Total systematic timing risk. The timing-risk contribution is the covariance across mega-states between the conditional expected SDF and the conditional expected payoff:

$$\widehat{\text{Systematic Timing}} \text{ Risk} \equiv \sum_{u=0}^{H} \widehat{P}(\tau = u) \left(\hat{m}_u - \bar{m}\right) \left(\hat{Y}_u - \bar{Y}\right),$$

where the probability-weighted means are

$$\bar{m} \equiv \sum_{u=0}^{H} \widehat{P}(\tau = u)\, \hat{m}_u, \qquad \bar{Y} \equiv \sum_{u=0}^{H} \widehat{P}(\tau = u)\, \hat{Y}_u.$$

Because the test is for whether systematic risk is statistically zero, the sign convention (short vs. long side) does not affect inference. Standard errors are computed via bootstrap with draws taken at the date level.

Table 1 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 against the hypothesis of early exercise due to systematic intermediary constraints.

TABLE 1: 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 Section 3.3.1. 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

Given that Section 3.3.1 found early exercise to constitute purely idiosyncratic risk (e.g., systematic timing and outcome risk were not statistically different from zero), we may apply the theorems for pricing idiosyncratic risk.

Classic pricing theory (Berk and DeMarzo, 2011) says that higher order moments of idiosyncratic outcome risk should not be priced. On the other hand, Theorem 1 says that higher order moments of timing risk, even when purely idiosyncratic, are priced. Do these predictions appear in the data?

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, 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.

I then compute mean and standard deviation of both timing and outcome risk at the contract level. For each contract, the realized timing risk moments are simply the realized mean and standard deviation of the early exercise stopping time probability distribution. Similarly, the realized outcome mean and standard deviation are computed as, conditional on early exercise, the average payoff amount and standard deviation of $-(S-K)$ to the short side (as earlier, I take the perspective of the short side, though either the perspective of the long or the short side is permissible here; the test is for statistical significance).

Table 2 regresses the price difference between American and European options on timing risk mean and standard deviation and outcome risk mean and standard deviation, with and without liquidity controls and time fixed effects. In each specification, standard errors are clustered by day.

As a sanity check, observe in Table 2 that the textbook result for outcome risk holds. Higher order moments of outcome risk do not affect price (e.g. outcome variance is not statistically significant). Mean outcome risk slightly affects price ($p < 0.10$ across specifications). Though quite small (the precise estimate is -\$0.002 per contract), indicating that outcome risk plays a minimal role in pricing American vs. European options without dividends, the sign of the coefficient is economically sensible: the less painful is the early exercise payoff to the short side (equivalently, the less upside for the long side when exercising early), the lower is the premium for

the American option.

The predictions made by Theorem 1 are illustrated in Table 2's regressions. While higher order outcome moments do not affect price, higher order timing risk moments do. In particular, both the mean and the variance in the time of early exercise are statistically significant ($p < 0.05$) in affecting the price difference between otherwise identical American and European options.

As early exercise becomes one day earlier, the price difference reduces by \$6-\$7.67 per contract on average. The negative sign on timing mean is reasonable: an earlier mean time of early exercise effectively makes the American resemble a European option with an earlier expiration date. Since an early expiration kills option value, it makes sense that, all else equal, sooner early exercise reduces the premium of the American relative to the European.

Timing standard deviation also affects the American-European price difference in a statistically significant way, consistent with Theorem 1. To interpret the sign of its coefficient, per Theorem 2, we need to know whether the shock is positive or negative. Table 2 therefore considers timing standard deviation for options where all early exercise payoffs to the short side are negative, separately from timing standard deviation for option contracts where all early exercise payoffs to the short side are positive. Per Theorem 2 and Remark 2, since these are short-dated options, the term structure is locally flat. Therefore, the compensation for uncertain timing should increase in timing standard deviation when the shock to the short side is negative (equivalently, when the shock to the long side is positive). Intuitively, in this case, the long side is paying for timing flexibility and the short side is being compensated for timing risk. The coefficients on timing standard deviation for this shock case (negative shocks only to the short side) are indeed positive: a one percent increase in timing standard deviation increases the American-European price difference by \$0.19 on average ($p < 0.05$, both with and without moneyness and liquidity controls).

In the case where all shocks to the short side are positive – e.g. the long side is exercising early and losing money – the same theorem and remark predict that more timing uncertainty should reduce the value of uncertain timing. 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.

Finally, in addition to illustrating and validating Theorem 1, the coefficients of Table 2 indicate that uncertainty in the time of early exercise is about half as meaningful to the market as its expectation when the time of early exercise is quite uncertain. For example, the conditional average empirical timing standard deviation, conditional on its being greater than zero, is 27 percentage points; indicating that timing standard deviation alone would predict an approximately \$5.13 more expensive American than European option for the average contract which has a timing distribution. In comparison, the average empirical timing mean, conditional on its being greater than zero, is being exercised 1.6 days early: indicating an average contribution to the American-European price difference of \$9.60-\$12.27.

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

Thus far, we have used Theorem 6 to show that early exercise constitutes purely idiosyncratic risk in both outcomes and timing; given that it is idiosyncratic, we applied Theorem 1 to illustrate the pricing of higher order moments for time but not for outcomes, and used Theorem 2 to interpret the coefficient signs .

We now illustrate the predictions of Theorem 5, which says that for purely idiosyncratic timing risk, there should be a relationship between the valuation of such risk and the short rate. A simple test regresses the American-European contract level price difference against the log short rate (Table 3). The theory predicts a positive coefficient for the log short rate. Intuitively, when early exercise occurs, from Theorem 5, the price difference between an American and European is equivalent to a portfolio of shorter zero-coupon bonds less a portfolio of longer zero-coupon bonds. Since longer-dated ZCBs are more sensitive to interest rates (Berk and DeMarzo, 2011), when the interest rate goes up, the value of the longer-dated ZCB portfolio falls by more than the value of the shorter-dated ZCB portfolio; as a result, the American-European price difference should increase as interest rates increase. Table 3 bears out this prediction: an increase of one in the log short rate when the American contract features empirical early exercise, increases the American-European price difference by \$4.15 - \$4.69 on average ($p < 0.01$ in all specifications). The short rate here is taken to be the three-month treasury; Appendix B.3 re-computes all tables in this section with the one month treasury rate instead and finds close to identical results.

We now try for an even more direct illustration of the theorem by computing the

TABLE 2: Theorem 1 Empirical Illustration: Higher Order Moments of Time Affect Price for Purely Idiosyncratic Risk

	(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 3.1. 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 3.2. 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.

realized zero-coupon bond equivalent given by realized early exercise. In particular, the value of an American when early exercise occurs, relative to if early exercise did

TABLE 3: Theorem 5: Short Rates and The Value of Timing Risk

	(1)	(2)	(3)
Log Short Rate × 𝟙 {P(early exercise) > 0}	4.69***	4.24***	4.15***
	(1.18)	(1.34)	(1.34)
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 3.1. This midpoint difference is multiplied by 100 to convert from index points to dollar units. The log short rate is the log of the daily analog of the annualized three-month treasury rate, multiplied by an indicator for whether the American contract features empirical early exercise. All regressions include a control for days to expiry and a control for the probability of early exercise. 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.

not occur, in ZCB terms is

$$V_{American} = E[\mathbb{1}\{\text{ee}\}(\sum_{t'=t}^{T_i-1} \frac{1}{(1+y_{t\to t'})^{t'-t}} \mathbb{1}\{\text{ee at } t'|\text{ee}\}(-(S_{t'}-K))) + \mathbb{1}\{\text{no ee}\}\frac{1}{(1+y_{t\to T_i})^{T_i-t}} \min((-(S_{T_i}-K),0)],$$

where ee denotes early exercise occurring, and T_i denotes the expiry time of the option. Since the value of the European in ZCB terms is

$$V_{European} = E[\frac{1}{(1+y_{t\to T_i})^{T_i-t}} \min((-(S_{T_i}-K),0)],$$

I compute the $V_{American} - V_{European}$ object from the empirical early exercise data and short rates.

There is one subtlety. Theorem 5 says that the price of uncertain timing should equal the expected value of the zero coupon bond portfolio equivalent. However, we have only the realized value. In general, the expectation need not be the same as the realization. As a first pass, I regresses the American-European price difference on realized ZCB portfolio equivalents. The coefficient is in the direction expected: for an increase in \$1 of present-value realized loss to the short side, the price difference increases by \$0.01 on average. However, the magnitude is less than the one-for-one relationship between expected present value and price that the theorem would

predict. A possible culprit is realization versus expectation; for example, on other dimensions of risk, options are known to be much less sensitive to realizations than to expectations (Dew-Becker and Giglio, 2024).

To proxy for expectations, Table 4 splits the sample into an 60% train set (1/1/2012 to 3/15/2017) and a 40% test set (3/16/2017 to 12/31/2019). In the training set, the realized ZCB values are regressed on time-t observable characteristics: the moneyness of the option, the short rate, and the number of days to expiry. Using the coefficients from this regression, I then compute expected ZCB values in the test set using only contemporaneously observable variables, so as to avoid look ahead bias. Formally the regression in the training set is:

$$ZCB_{i,t}^{(realized)} = \alpha + \beta \text{ contemporaneous}_{i,t} + \epsilon_{i,t}. \qquad \forall (i,t) \in \text{Train Set},$$

where $\text{contemporaneous}_{i,t}$ includes the 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.

A proxy for the expected zero-coupon bond portfolio in the test set is then constructed using the above estimated coefficients:

$$\hat{ZCB}_{i,t}^{(expectation)} = \alpha + \beta \text{ contemporaneous}_{i,t}. \qquad \forall (i,t) \in \text{Test Set}$$

And the regressions in Table 4 are of the form:

$$Price(American)_{i,t} - Price(European)_{i,t} = \alpha + \gamma \hat{ZCB}_{i,t}^{(expectation)} + \eta X + \epsilon_{i,r} \qquad \forall (i,t) \in \text{Test Set}$$

where X are controls for liquidity and moneyness. To ensure that standard errors reflect the uncertainty inherent in the training set, I bootstrap the procedure, drawing with replacement (but keeping the temporal split fixed) for the train and test sets, and compute the standard errors over 1000 bootstrap iterations. To allow for correlations within a given day, I draw at the date rather than contract level (using dates before the temporal split for the bootstrapped train set; and after the temporal split for the bootstrapped test set).

Running this regression in the test set yields statistically significant coefficients on expected ZCB portfolio equivalents of approximately -0.7, both with and without

liquidity and moneyness controls ($p < 0.1$ in all specifications). These are statistically indistinguishable from -1 all specifications, including after controlling for liquidity and moneyness separately from the expectation proxy; consistent with the one-for-one prediction of the theorem.[4] In addition, this single, theoretically-derived proxy alone accounts for 9% of out-of-sample variation; controls for liquidity and moneyness add only an additional 3 percentage points.

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

	(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 3.1. 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.

4 Potential Applications Across Finance

Understanding the components risk and its pricing has application 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.

Corporate Decisions. Companies often have uncertainty not only over the magnitude of cash flows on an uncertain project, but also when they will occur. One puzzle in the corporate literature is that companies appear to use hurdle rates that

[4]Note that rather than expecting a coefficient of 1 in the table, we expect a -1 because the more negative the regressor (the expected-ZCB-equivalent), the greater the expected loss to the short side.

are higher than their weighted average cost of capital (Graham, 2022). It may be worth exploring whether part of this difference arises from variability in timing risk across projects and companies.

Bond Spreads. Callable and convertible bonds are subject to uncertain timing. Pricing models for these bonds typically conceptualize them as a standard bond plus an option (Berk and DeMarzo, 2011). Just as, in Section 3, options prices including a pricing of timing risk (Section 3), corporate bond spreads may differ as well.

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.

Mortgages. In mortgages, prepayment risk is at heart uncertainty over the time at which the mortgage will be paid back. In the literature, prepayment is often predicted using interest rate factors, 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 is priced.

Legal Risk. Work in law and economics has considered the impact of outcome uncertainty (Malmendier, 2009) on market entry and regulation. In practice, legal trials can be lengthy and have uncertain duration; it may be interesting to consider how timing risk affects market outcomes and therefore optimal regulation.

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

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 include both types of uncertainty. 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.

5 Conclusion

Using only a no-arbitrage assumption, I showed that 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. 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. Using market data from options as a relatively clean laboratory 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

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

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.

van Binsbergen, Jules. 2016. "Good-Specific Habit Formation and the Cross-Section of Expected Returns." *The Journal of Finance*, 71: 1699–1732.

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

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

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

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

Chen, Hui. 2010. "Macroeconomic Conditions and the Puzzles of Credit Spreads and Capital Structure." *The Journal of Finance*, 65: 2171–2212.

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.

Dew-Becker, Ian, and Stefano Giglio. 2023. "Cross-Sectional Uncertainty and the Business Cycle: Evidence from 40 Years of Options Data." *American Economic Journal: Macroeconomics*, 15: 65–96.

Dew-Becker, Ian, and Stefano Giglio. 2024. "Recent Developments in Financial Risk and the Real Economy." *Annual Review of Financial Economics*, 16: 39–60.

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.

ETrade from Morgan Stanley. 2021. "Open Interest."

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.

Gomes, Joao F., and Lukas Schmid. 2021. "Equilibrium Asset Pricing with Leverage and Default." *The Journal of Finance*, 76: 977–1018.

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

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.

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.

Kuehn, Lars Alexander, and Lukas Schmid. 2014. "Investment-Based Corporate Bond Pricing." *The Journal of Finance*, 69: 2741–2776.

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.

Malmendier, Ulrike. 2009. "Law and Finance 'at the Origin'." *Journal of Economic Literature*, 47: 1076–1108.

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

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.

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

Papanikolaou, Dimitris. 2011. "Investment Shocks and Asset Prices." *Journal of Political Economy*, 119: 639–685.

Pastor, Lubos, and Pietro Veronesi. 2009. "Technological Revolutions and Stock Prices." *American Economic Review*, 99: 1451–1483.

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*.

Schmid, Lukas, and Roberto Steri. 2023. "Credit market equivalents and the valuation of private firms." *Working Paper*.

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.

Vissing-Jorgensen, Annette. 2002. "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution." *Journal of Political Economy*, 110: 825–853.

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. 2017. "Global Risk Aversion and International Return Comovements." *Working Paper*.

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