

Taking Money Off the Table: Suboptimal Early Exercises, Risky Arbitrage, and American Put Returns*

Kevin Aretz[†]
Manchester Business School
kevin.aretz@manchester.ac.uk

Ian Garrett[‡]
Manchester Business School
ian.garrett@manchester.ac.uk

Adnan Gazi[§]
University of Liverpool
adnan.gazi@liverpool.ac.uk

This Draft: December 15, 2020

Abstract

Many studies report that American option investors often exercise their positions suboptimally late. Yet, when that can happen in case of puts, there is an arbitrage opportunity in perfect markets, exploitable by longing the asset-and-riskfree-asset portfolio replicating the put and shorting the put. Using early exercise data, we show that the arbitrage strategy also earns a highly significant mean return with low risk in real single-stock put markets, in which exactly replicating options is impossible. In line with theory, the strategy performs particularly well on high strike-price puts in high interest-rate regimes. It further performs well on short time-to-maturity puts on low volatility stocks, consistent with evidence that investors do not correctly incorporate those characteristics into their exercise decisions. The strategy survives accounting for trading and short-selling costs, at least when executed on liquid assets.

Keywords: Empirical asset pricing; cross-sectional option pricing; put options; early exercise.

JEL classification: G11, G12, G15.

[†]Accounting and Finance Division, Room 5.013 AMBS, Booth Street West, Manchester M15 6PB, UK, tel.: +44(0)161 275 0122, fax: +44(0)161 275 4023, URL address: <www.kevin-aretz.com>.

[‡]Accounting and Finance Division, Room 4.079 AMBS, Booth Street West, Manchester M15 6PB, UK, tel.: +44(0)161 275 4958.

[§]Accounting and Finance Group, ULMS, Chatham St, Liverpool L69 7ZH, UK, tel.: +44(0)151 795 3641.

*We are indebted to Michael Brennan, Olga Kolokolova, Alex Kostakis, Yifan Li, Christian Wagner and seminar participants at Alliance Manchester Business School, University of Liverpool Management School, the 6th Young Finance Scholars Conference (Brighton) and the 2019 Konstanz-Lancaster-Manchester-Warwick Workshop on Quantitative Finance and Econometrics (Manchester), the 2021 Southwestern Finance Association Annual Meeting (Virtual) for insightful and constructive feedback. We are also deeply grateful to Bob Whaley for sharing his manually-collected early exercise option data with us.

1 Introduction

Many studies suggest that American option investors do not always follow optimal early exercise policies, with them frequently exercising their positions too late. Pool et al. (2008), for example, estimate that the total profits lost from not optimally exercising single-stock American calls on ex-dividend dates are about \$500 million over a ten-year period, while Barraclough and Whaley (2012) estimate that those from not optimally exercising American puts are about \$1.9 billion over a twelve-year period. Yet, when American puts can be exercised too late, there is an arbitrage opportunity in perfect markets, exploitable by longing a dynamic underlying-asset-and-riskfree-asset portfolio replicating the put and shorting the traded put (Shreve (2004)). Intuitively, after the put should have but has not been exercised, the portfolio consists of one short underlying asset unit and an investment of the strike price into the money market, covering any obligations arising from the put but also earning interest. The earned interest represents the arbitrage profit.

In our paper, we study the profitability and total and systematic risk of the arbitrage strategy in real markets in which options can only be imperfectly replicated. While real investors are thus unable to earn a proper arbitrage profit, they may still be able to earn some profit with an only low risk (“risky arbitrage profit”). In accordance, we find that longing daily-rebalanced replication portfolios of single-stock American puts and shorting those puts earns us a highly significant mean return of 4.11% per month before transaction costs. Suggesting daily rebalancing creates efficient replication portfolios, the returns on the two legs of that strategy share a mean cross-sectional correlation of -0.85 ; the variance of the strategy return is more than 3-4 times lower than those on the returns of the legs; and the strategy only weakly loads on risk factors. In line with theory, the mean strategy return increases with the strike price and the interest rate. Also, it is higher for short time-to-maturity puts on low volatility stocks, aligning with further evidence that investors do not understand how those variables condition the optimal early exercise decision. Crucially, since the strategy requires frequent rebalancing and short-selling, we finally show that it survives accounting for trading and shorting costs, at least when executed on liquid assets.

We use options market data from Optionmetrics and early exercise data from Bob Whaley to study our arbitrage strategy in real markets. To derive the return on the long leg of the strategy, we invest the market price of the put into a portfolio consisting of the underlying stock and the riskfree asset at the start of the strategy return period, with the number of stocks equal to the (negative) delta of the put. At the end of each day, we then rebalance the stock holdings in the portfolio to its new delta. To derive the return on the short leg, we calculate the return of the traded put over the same period as the compounded early exercise payoff (if there is an early exercise) or the market price (if there is none) of the put at the end of that period to its market price at the start. To find out if and when the put is early exercised, we recognize that the Option Clearing Corporation randomly assigns exercise obligations to outstanding short positions (Pool et al. (2008)). We thus assume that the short leg is terminated over a day if a draw from the univariate distribution lies below the ratio of early exercised contracts of the put over that day to its open interest at the end of the prior day (“daily early exercise probability”).¹

We finally calculate the return of the risky arbitrage strategy as the spread in returns between the replication portfolio and the traded put over the strategy return period. Importantly, however, the strategy return period is not fixed, ranging from the start of a month to the earlier of the day over which the traded put is exercised and the end of the month. In other words, we always liquidate the long and the short leg of the risky arbitrage strategy on the same date.

Our evidence suggests that the risky arbitrage strategy is highly profitable with an only low total or systematic risk. Longing an equally-weighted portfolio of daily-rebalanced replication portfolios of all outstanding American puts and shorting an equally-weighted portfolio of those same puts yields a mean monthly return of 4.11% (t -statistic: 5.21). While the monthly variances of the returns on the long and the short leg of the strategy are, respectively, 0.32 and 0.26, the

¹Given that the put return depends on draws from a random distribution, it varies each time that we recalculate it, leading to concern that our empirical results cannot be replicated. Fortunately, however, those variations almost cancel out in the aggregate, likely due to a law of large numbers. Recalculating the mean monthly pooled-sample return of the risky arbitrage strategy three times, we, for example, obtain values of 4.11455%, 4.11022%, and 4.11812%. We plan to include a histogram of 1,000 mean strategy returns in future versions of our paper.

monthly variance of the strategy return is a much lower 0.09. Regressing the strategy's return on popular risk factors, such as the Fama-French (2015) five-factor model factors, we obtain alphas close to identical to the strategy's mean return but with higher t -statistics. In addition, we also obtain factor loadings markedly attenuated compared to those obtained from the two legs of the strategy. Calculating key performance evaluation statistic (e.g., the Sharpe (1966) ratio), we find yet more evidence that the strategy bodes satisfactory risk-adjusted performance.

We next condition the risky arbitrage strategy on put, stock, and macroeconomic characteristics including the strike price, the interest rate, moneyness, time-to-maturity, and underlying stock volatility. While theory suggests that, *ceteris paribus*, the strategy is more profitable for high strike-price puts in high interest-rate regimes, it is hard to generate further predictions since the strategy's profitability for some class of puts ultimately hinges on the extent to which investors early exercise puts too late within that class. Our evidence shows that the mean strategy return significantly increases with the strike price and the interest rate, but significantly decreases with moneyness, time-to-maturity, and stock volatility. Corroborating that the relations with time-to-maturity and stock volatility, but not that with moneyness, originate from variations in investors' tendency to exercise puts too late, we offer further evidence that investors do not seem to understand the negative effects of those characteristics on the optimal early exercise decision.

We finally turn to the transaction costs incurred by the risky arbitrage strategy. Assuming that an asset's trading costs are proportional to its bid-ask spread and that a stock can be borrowed at Markit's indicative rate, we show that trading and borrowing costs greatly eat into the profitability of that strategy. Considering in-the-money (ITM) puts with 30-60 days-to-maturity, the mean strategy return, for example, drops from 4.83% per month (t -statistic: 5.87) in the no-transaction-cost case to -2.17% (t -statistic: -2.05) in the 25% bid-ask-spread transaction-cost case. Importantly however, the strategy remains profitable even net of transaction costs when we restrict our attention to liquid puts written on liquid underlying assets. Excluding from the above ITM puts those with a bid-ask spread above the median and/or written on stocks with an Amihud (2002) illiquidity

value above the median, the mean strategy return is, for example, 3.92% (t -statistic: 3.20) even in the 25% bid-ask-spread transaction-cost case. We also establish that decreasing the rebalancing frequency of the replication portfolio can further help to raise the strategy's profitability net of transaction costs, without it greatly boosting the volatility of the strategy return.

Our work builds up on empirical studies suggesting that American option investors often do not follow optimal early exercise policies. While Overdahl and Martin (1994) show that most early exercises of single-stock American calls and puts fall within optimal boundaries on the underlying stock's price, Brennan and Schwartz (1977) document that the early exercises of such puts are typically inconsistent with the Black-Scholes (1973) framework. Finucane (1997) finds that investors often early exercise American calls on non-dividend stocks, conflicting with Merton's (1973) insight that it is never optimal to early exercise such calls. Digging deeper into Finucane's (1997) results, Poteshman and Serbin (2003) show that only individual but not institutional investors sometimes early exercise the former calls.² As already said, Pool et al. (2008) and Barraclough and Whaley (2012) find that the foregone profits from failing to optimally early exercise single-stock American calls and puts are economically large. More generally, Bauer et al. (2009) report that retail investors do not perform well in their option investments. While we offer further evidence that investors' early exercise strategies are often suboptimal, our main contribution to the above literature is to show how to make money from that suboptimality using a simple trading strategy.

We further add to empirical studies examining the performance of trading strategies. While stock strategies are dominant among those studies (see, e.g., Fama and French (1992), Lakonishok et al. (1994), and Carhart (1997)), recent studies have also evaluated option strategies. Coval and Shumway (2001), for example, report that writing zero-beta index straddles is profitable even after considering transaction costs. Vasquez (2017) documents that longing straddles with a high slope

²The recent studies of Jensen and Pedersen (2016), Battalio et al. (2020), and Figlewski (2019) highlight that trading costs can make it optimal to early exercise an American call on a non-dividend asset, implying that some of the early exercises classified as suboptimal by Finucane (1997) could be optimal. Notwithstanding, Pool et al.'s (2008) evidence that the vast majority of such calls are early exercised by retail (but not institutional) investors leads one to suspect that Finucane (1997) correctly classifies most of his early exercises.

of their implied volatility term structure and shorting those with a low slope yields positive mean cost-adjusted returns. Conversely, Goyal and Saretto (2009) find that longing options with a large difference between realized and implied volatility and shorting those with a low difference yields positive mean “delta-hedged returns” (i.e., the options’ returns neutralized with respect to their underlying stocks’ returns). Cao et al. (2020) show that sorting options based on characteristics used to predict stock returns also often yields significant mean delta-hedged returns, which can, however, differ in sign from the corresponding mean stock returns. We add to this literature by proposing a new risky arbitrage trading strategy involving single-stock American puts and their delta-replication portfolio firmly grounded in mathematical finance theory.

We finally add to studies looking into how option prices deviate from the values of their replication portfolios. While in a Black-Scholes (1973) perfect capital market the value of any option can be perfectly hedged/replicated using a dynamic asset-and-riskfree-asset portfolio, Leland (1985) proves that transaction costs can drive a wedge between the option’s value and the replication portfolio value. Adding underlying-asset dividends and stochastic volatility, Perrakis and Lefoll (2000) and Gondzio et al. (2003) come to the same conclusion. Relying on existing mathematical finance theory, we contribute to this literature by showing that suboptimally late early exercises of American puts can also make the values of those puts deviate from the values of their (standard/non-consuming) replication portfolios, not only in theory but also in practice.

We proceed as follows. In Section 2, we review the underlying theory. In Section 3, we discuss our methodology and data. In Sections 4, 5, and 6, we present the historical performance of the risky arbitrage strategy, condition the strategy on put, stock, and macroeconomic characteristics, and adjust it for trading and borrowing transaction costs, respectively. Section 7 concludes.

2 Theory

In this section, we briefly review the mathematical finance theory suggesting that an ex-ante non-zero probability that an American put is exercised too late creates an arbitrage opportunity

in perfect capital markets. We keep our review as intuitive as possible, referring to other papers for more technical details. Consider an American put giving its owner the option to sell an underlying asset worth $S(t)$ at time t for a constant price of K (“strike price”) in each instant within the time period $t \in [0, T]$ (“maturity time”). To keep matters simple, the underlying asset does not pay out cash over the maturity time, and its value evolves according to the following geometric Brownian motion (GBM) under the risk-neutral (“martingale”) measure \mathbb{Q} :

$$dS(t) = rS(t)dt + \sigma S(t)d\widetilde{W}(t), \quad (1)$$

where r is the annualized risk-free rate of return, σ is the annualized volatility of the underlying asset’s return, and $d\widetilde{W}(t)$ is the differential of a Brownian motion.

Using risk-neutral pricing techniques, the value of the American put, $p(t, S(t))$, is:

$$p(t, S(t)) = \max_{\tau \in \Omega(t, T)} \tilde{E}[e^{-r(\tau-t)}(K - S(\tau)) | S(t)], \quad (2)$$

where \max is the maximum operator, \tilde{E} the expectation under the \mathbb{Q} measure, τ the (random) early exercise time associated with some *feasible* early exercise strategy (i.e., a strategy based only on information available at the current time), and $\Omega(t, T)$ the set of early exercise times associated with all feasible strategies. Karatzas (1988) and Jacka (1991) prove that the optimal feasible strategy is to early exercise the put as soon as the underlying asset value $S(t)$ drops below some time-variant boundary $L(t)$, which is bounded and increases convexly with time t . As long as $S(t) > L(t)$, we can thus replicate the put by investing $p(t, S(t))$ into a dynamically rebalanced portfolio containing the underlying asset and the riskfree asset and ensuring that the underlying asset investment is equal to $\Delta(t, S(t))S(t)$ in each instant, where $\Delta(t, S(t)) \equiv \partial p(t, S(t)) / \partial S(t)$ is the (negative) put delta. It follows from the replication strategy that:

$$\frac{1}{2}\sigma^2 S(t)^2 p_{SS}(t, S(t)) + rS(t)p_S(t, S(t)) + p_t(t, S(t)) - rp(t, S(t)) = 0, \quad (3)$$

where $p_S \equiv \partial p(t, S(t))/\partial S(t)$, $p_{SS} \equiv \partial^2 p(t, S(t))/\partial S(t)^2$, and $p_t \equiv \partial p(t, S(t))/\partial t$. Conversely, as soon as $S(t) \leq L(t)$, $p(t, S(t)) = K - S(t)$. A direct computation then yields:

$$\frac{1}{2}\sigma^2 S(t)^2 p_{SS}(t, S(t)) + rS(t)p_S(t, S(t)) + p_t(t, S(t)) - rp(t, S(t)) = rK. \quad (4)$$

We are now in a good position to repeat the insights in Shreve's (2004) Corollary 8.4.3:

COROLLARY 8.4.3 *Consider an agent with initial capital $X(0) = p(0, S(0))$. Suppose that, in each instant, this agent holds a portfolio consisting of $\Delta(t, S(t))$ units of the underlying asset and the residual value invested into the risk-free asset. Further, assume the agent consumes cash, $C(t)$, from that portfolio at rate rK per time unit if $S(t) \leq L(t)$ and else at rate zero. Then $X(t) = p(t, S(t))$ for all times $t \in [0, T]$. In particular, $X(t) \geq \max(K - S(t), 0)$ for all times t until T , so the agent can pay off a short position regardless of when the option is expired.*

To prove the corollary, Shreve (2004) starts with applying Itô's lemma to the differential of the discounted value of the American put, $d(e^{-rt}p(t, S(t)))$:

$$\begin{aligned} d(e^{-rt}p(t, S(t))) &= e^{-rt} \left[-rp(t, S(t))dt + p_t(t, S(t))dt + p_S(t, S(t))dS(t) \right. \\ &\quad \left. + \frac{1}{2}p_{SS}(t, S(t))dS(t)dS(t) \right] \end{aligned} \quad (5)$$

$$\begin{aligned} &= e^{-rt} \left[-rp(t, S(t)) + p_t(t, S(t)) + rS(t)p_S(t, S(t)) \right. \\ &\quad \left. + \frac{1}{2}\sigma^2 S(t)^2 p_{SS}(t, S(t)) \right] dt + e^{-rt}\sigma S(t)p_S(t, S(t))d\tilde{W}(t), \end{aligned} \quad (6)$$

noting that Equations (3) and (4) imply that the term in square parentheses in Equation (6) is zero if $S(t) > L(t)$ and else $-rK$. He then stresses that the differential of the value of the portfolio containing the underlying asset and the riskfree asset is equal to:

$$dX(t) = \Delta(t)dS(t) + r(X(t) - \Delta(t)S(t))dt - C(t)dt, \quad (7)$$

so that the differential of the discounted value of that portfolio is equal to:

$$d(e^{-rt}X(t)) = e^{-rt}(-rX(t)dt + dX(t)) \quad (8)$$

$$= e^{-rt}(\Delta(t)dS(t) - r\Delta(t)S(t)dt - C'(t)dt) \quad (9)$$

$$= e^{-rt}\sigma S(t)p_S(t, S(t))d\tilde{W}(t) - e^{-rt}rK\mathbb{I}_{S(t)\leq L(t)}dt, \quad (10)$$

where $\mathbb{I}_{S(t)\leq L(t)}$ is an indicator equal to one if $S(t) \leq L(t)$ and else zero. Comparing Equations (6) and (10), it is obvious that the values of the put and the portfolio are not only the same at the initial time $t = 0$, but also at any other time t over the time-to-maturity.

Corollary 8.4.3 implies that the arbitrage profit obtained from longing the asset-and-riskfree-asset portfolio and shorting the put over the period until the put is expired is $\max(t^E - \tau, 0)rK$ in perfect markets, where t^E is the actual expiration date of the put. As a result, the arbitrage profit increases with the length of the period over which the put should but has not been early exercised ($\max(t^E - \tau, 0)$), the risk-free rate of return (r), as well as the strike price (K). Intuitively, as the underlying asset value $S(t)$ reaches the early exercise boundary $L(t)$ from above, the replication portfolio consists of one short unit of the underlying asset and an investment of the strike price K into a money market account. If the put owner optimally early exercises at that point, we transfer the money market account to him/her in return for one long unit of the underlying asset, which we use to extinguish our short position in that asset. Conversely, if he/she does not early exercise, we hold onto the money market investment and the one short underlying asset unit, enabling us to earn interest on the money market investment equal to $rKdt$ in each instant.

Two remarks about the above analysis are in order. First, while we rely on the simplest possible assumptions in that analysis, Shreve (2004) highlights that the arbitrage opportunity generally exists in complete markets as long as investors' suboptimal exercise policies render the discounted American put value a *supermartingale* under the \mathbb{Q} measure. The arbitrage opportunity thus also exists in, for example, stochastic volatility and mixed jump-diffusion models. Second, while there

may be concern that the arbitrage payoff is negligible due to its dependence on the interest rate, that payoff can be large even in a low interest rate regime. To see that, consider an American put with $S(0) = K = 30$, $\sigma = 0.30$, $T = 0.25$, and $r = 0.01$, which ends up being early exercised six weeks too late (i.e., $\max(t^E - \tau, 0) = 0.125$). Scaling the arbitrage payoff, $0.01 \times 30 \times 0.125 = 0.0375$, by the initial put value, 1.7659, we obtain a three-month arbitrage return of 2.1%, translating into a monthly return of 70.80 basis points.^{3,4} Hence, even when $r = 0.01$, the arbitrage return compares well to the mean returns of popular risky stock strategies, such as the SMB, HML, and MOM strategies. Notwithstanding, the arbitrage return is obviously much larger in a higher interest rate regime, with it, for example, being 3.8% *per month* when $r = 0.05$.

In the remainder of our study, we evaluate the mean returns and the total and systematic risk of the arbitrage strategy outlined in this section in real markets. While many empirical studies suggest that real American put investors often early exercise their positions too late (a necessary condition for the strategy to be profitable), market imperfections such as transaction costs and discontinuous trading imply that we cannot perfectly replicate puts using their underlying assets and the riskfree asset in real markets. Given that, the arbitrage strategy in perfect markets, with a zero probability of a loss but a positive probability of a gain, becomes, at best, a risky arbitrage strategy in real markets, with a low probability of a loss but a high probability of a gain.

3 Methodology and Data

In this section, we review our methodology and data. We first explain how we calculate the returns of American puts and their replication portfolios with or without transaction costs, also elaborating on how we decide whether a short put is expired early. We next discuss how we compute optimal early exercise probabilities conditional on the Black-Scholes (1973) framework using the Longstaff-

³In these calculations, we follow the standard practice in the stock literature to compute the return of a long-short strategy as the difference in returns between the long and short leg. In our case, the long and short leg share the same initial value, so that the return of the strategy becomes the difference in payoffs between the two legs (which is the arbitrage payoff $\max(t^E - \tau, 0)rK$) scaled by the initial value of the American put.

⁴We use a binomial tree with 1,000 time steps to calculate the American put value.

Schwartz (2001) approach. We finally outline our data and data sources.

3.1 Calculating the American Put Return

We compute the holding-period gross return of an American put, $R^p(t_0, t_H)$, as its compounded-up early exercise payoff (if there is an early exercise over that period) or its market price at the end of that period (if there is none) to its market price at the start of that period:

$$R^p(t_0, t_H) = v(t_H)/p(t_0), \quad (11)$$

where t_0 and t_H are, respectively, the start and end day of the holding period, $v(t_x)$ is the put's value at the end of day t_x , and $p(t_x)$ is its market price at the end of the same day. If the put is early exercised on day $t_E < t_H$, its value at the end of the holding period, $v(t_H)$, is: $e^{r_f(t_E, t_H)} \max(K - S(t_E), 0)$, where $r_f(t_x, t_y)$ is the net riskfree rate of return from end of day x to end of day y , K is the strike price, and $S(t_x)$ is the underlying asset's value at the end of day x . Else, the value of the put on that date is its market price, $p(t_H)$, on the same date.⁵

Keeping in mind that the risky arbitrage strategy requires us to be short in the put, it is the put owners, not us, who determine if and when they early exercise their positions. When a put owner decides to early exercise a put, the Options Clearing Corporation randomly assigns the exercise to an outstanding short position. To mimic that procedure, we do the following. For each day within the holding period, we first calculate the daily early exercise probability for the entire put issue, defined as the number of early exercises of that issue over the day scaled by its open interest at the end of the prior day. Starting with the first day in that period, we draw a number from the univariate distribution with bounds zero and one, assuming that our short put position is expired if the drawn number lies below the daily early exercise probability. Unless the short

⁵Consistent with convention, Equation (11) gives the return on one long unit of the American put. Given that we are, however, short that put in our risky arbitrage strategy, what we mostly care about is $-R^p(t_0, t_H)$, and not $R^p(t_0, t_H)$. Recognizing that is especially important when we adjust returns for transaction costs since else one could erroneously gain the impression that our adjustments increase, and not decrease, returns.

position is expired, we move to the second day, again drawing a number from the univariate distribution and assuming the short position is expired if the number lies below the daily early exercise probability. We continue in that way until the end of the holding period.

Our return calculations above ignore transaction costs arising from trading the put. Following Goyal and Saretto (2009) and Cao and Han (2013), we assume that these costs are proportional to the put's bid-ask spread, $\varphi BAS^p(t_x)$, where φ is a constant and $BAS^p(t_x)$ that spread at the end of day t_x . Again keeping in mind that we are short the put, we compute the transaction-cost-adjusted gross put return, $R^{p,tca}(t_0, t_H)$, in the no-early-exercise case by *adding* $\varphi BAS^p(t_H)$ to the numerator of Equation (11) and *subtracting* $\varphi BAS^p(t_0)$ from its denominator. Conversely, in the early-exercise case, we add $(1 - \text{abs}(\Delta(t_E)))\varphi BAS^s(t_E)$ to the early exercise payoff in the numerator, where $\varphi BAS^s(t_x)$ and $\Delta(t_x)$ are, respectively, the underlying asset's bid-ask spread and the put's delta at the end of day t_x and abs is the absolute-value operator, while we again subtract $\varphi BAS^p(t_0)$ from the denominator. Thus, $R^{p,tca}(t_0, t_H)$ equals:

$$R^{p,tca}(t_0, t_H) = \begin{cases} \frac{e^{rf(t_E, t_H)} (\max(K - S(t_E), 0) + (1 - \text{abs}(\Delta(t_E)))\varphi BAS^s(t_E))}{p(t_0) - \varphi BAS^p(t_0)}; & t_E \leq t_H, \\ \frac{p(t_H) + \varphi BAS^p(t_H)}{p(t_0) - \varphi BAS^p(t_0)}; & t_E > t_H. \end{cases} \quad (12)$$

The $(1 - \text{abs}(\Delta(t_E)))\varphi BAS^s(t_E)$ adjustment in the early-exercise case arises since the long leg of our risky arbitrage strategy (the put replication portfolio to be discussed in the next subsection) is generally short delta underlying asset units on the early exercise date. As a result, when the put is early exercised against us, we use the underlying asset obtained from the put owner to extinguish our short delta position in that asset, implying that we only need to sell $(1 - \text{abs}(\Delta(t_E)))$ of that asset in the market at a unit transaction cost of $\varphi BAS^s(t_E)$.⁶

Using the daily early exercise probabilities above, we also calculate the real-world early exercise probability of a put within an issue over some other period. To do so, note that one minus a daily

⁶In case the put owner exercises his/her position on the optimal date or later, the put replication portfolio contains one unit of the underlying asset, implying that: $(1 - \text{abs}(\Delta(t_E)))\varphi BAS^s(t_E) = 0$. The adjustment thus deals with the unlikely case in which the put owner exercises his/her position too early.

early exercise probability yields the corresponding daily no-exercise (“survival”) probability. Computing one minus the product of the daily survival probabilities over the period, we obtain the real-world early exercise probability over that same period.

3.2 Calculating the Replication Portfolio Return

We next consider a dynamic portfolio noisily replicating an American put by shorting a number of underlying assets close to the put’s delta and investing the remaining portfolio value into the riskfree asset. Assuming we invest the market price of the put into that portfolio at the end of day t_0 and rebalance at the end of each day in $\mathbb{B} \in [t_1, t_2, \dots, t_N]$, where $t_0 < t_1 < \dots < t_N$ and t_N is the last date before the earlier of the put expiration date (t_E) and the holding period end date (t_H), we calculate the portfolio’s return, $R^X(t_0, t_H)$, over the holding period as follows. We first calculate the value of the portfolio at the end of the first rebalancing day, $X(t_1)$, as:

$$X(t_1) = \Delta(t_0)S(t_0)R(t_0, t_1) + (p(t_0) - \Delta(t_0)S(t_0))R_f(t_0, t_1), \quad (13)$$

where $R(t_x, t_{x+1})$ and $R_f(t_x, t_{x+1})$ are the gross underlying asset return and the riskfree rate of return from end of day t_x to end of day t_{x+1} , respectively. The portfolio’s value at the end of any other rebalancing date t_k is then:

$$X(t_k) = \Delta(t_{k-1})S(t_{k-1})R(t_{k-1}, t_k) + (X(t_{k-1}) - \Delta(t_{k-1})S(t_{k-1}))R_f(t_{k-1}, t_k), \quad (14)$$

while its value at the end of the holding period, t_H , can be written as:

$$X(t_H) = e^{r_f(t_E, t_H)} (\Delta(t_N)S(t_N)R(t_N, t_E) + (X(t_N) - \Delta(t_N)S(t_N))R_f(t_N, t_E)) \quad (15)$$

if the put is expired before the end of the holding period and as:

$$X(t_H) = \Delta(t_N)S(t_N)R(t_N, t_H) + (X(t_N) - \Delta(t_N)S(t_N))R_f(t_N, t_H) \quad (16)$$

if it is not. Notice that the compounding in Equation (15) ensures that the portfolio's value is measured at time t_H even when the put is expired earlier. We finally calculate the portfolio's gross return over the holding period by scaling by the initial investment:

$$R^X(t_0, t_H) = X(t_H)/X(t_0) = X(t_H)/p(t_0). \quad (17)$$

While we again abstract from transaction costs in our initial return calculations, such costs are likely to be even more important for the replication portfolio than the put due to the potentially frequent underlying-asset buys and sales necessary to ensure that the replication portfolio and put have similar deltas and due to the portfolio being short in the underlying asset. Again assuming that an asset's trading costs are proportional to its bid-ask spread, the total trading costs of the replication portfolio at the end of the holding period, $C^{TC}(t_0, t_H)$, are:

$$\begin{aligned} C^{TC}(t_0, t_H) &= \varphi \left(e^{r_f(t_0, t_H)} \text{abs}(\Delta(t_0)) \text{BAS}^s(t_0) \right. \\ &+ \sum_{i=1}^N e^{r_f(t_i, t_H)} \text{abs}(\Delta(t_i) - \Delta(t_{i-1})) \text{BAS}^s(t_i) \quad (18) \\ &\left. + \mathbb{I}_{\{t_E > t_H\}} \text{abs}(\Delta(t_H)) \text{BAS}^s(t_H) \right), \quad (19) \end{aligned}$$

where $\mathbb{I}_{\{t_E > t_H\}}$ is a dummy variable equal to one if $t_E > t_H$ and else zero. Intuitively, we only need to buy back the shorted underlying asset if the put is not exercised against us in the holding period. If it is, we obtain the underlying asset from the put owner, saving us $\varphi \text{abs}(\Delta(t_H)) \text{BAS}^s(t_H)$ in bid-ask trading costs. Conversely, the total costs originating from short-selling the

underlying asset at the end of the holding period, $C^{BC}(t_0, t_H)$, are equal to:

$$C^{BC}(t_0, t_H) = \sum_{i=0}^D e^{r_f(i+1, D)} r^{bc}(i+1) \text{abs}(\Delta(i)) S(i), \quad (20)$$

where the sum is taken over all days within the holding period, D is the number of days in that period, and $r^{bc}(i+1)$ is the daily stock-borrowing rate over day $i+1$. Subtracting the trading and stock-borrowing costs measured at the end of the holding period from the numerator of the unadjusted portfolio return, the adjusted portfolio return, $R^{X, tca}(t_0, t_H)$, is equal to:

$$R^{X, tca}(t_0, t_H) = \frac{X(t_H) - C^{TC}(t_0, t_H) - C^{BC}(t_0, t_H)}{p(t_0)}. \quad (21)$$

3.3 Calculating Theoretical Early Exercise Probabilities

To see whether real investors follow differentially suboptimal early exercise policies across different types of puts, we also contrast real-world early exercise probabilities with theoretical probabilities deduced from the optimal policies implied by the Black-Scholes (1973) model. To calculate the latter probabilities, we use Longstaff and Schwartz's (2001) least-squares approach. To be specific, we use a GBM to simulate q underlying-asset-value paths under the \mathbb{Q} measure (see Equation (1)), sampling the asset's value at times $t_0 < t_1 < \dots < t_k = T$. We next move backward through the paths, starting with calculating the path-specific maturity payoff of the put, $\max(K - S(t_k), 0)$. Moving to time t_{k-1} , we compare each path's early exercise payoff at that time, $K - S(t_{k-1})$, with the put's continuation value, which we define as the fitted value from a regression of the put's maturity payoffs discounted to time t_{k-1} on a function of the underlying asset value at time t_{k-1} .⁷ If the early exercise payoff exceeds the continuation value, we assume that an early exercise occurs for that path and at that time, replacing the put's value with the early exercise payoff.

Moving back to time t_{k-2} , we again regress the put's future payoff discounted to time t_{k-2} on

⁷To avoid bias, we run the regression on only observations for which the put is in-the-money at time t_{k-1} .

the function of the underlying asset's value at time t_{k-2} . This time, however, the future payoff is either the earliest early exercise payoff (if the put is early exercised in the future) or the maturity payoff (if it is not). As before, we assume that an early exercise occurs for a path at that time if the early exercise payoff exceeds the continuation value. We continue in that way until we reach time t_0 . We finally compute the *unconditional* theoretical early exercise probability over the put's time-to-maturity and imputed from the Black-Scholes (1973) model as the ratio of the number of paths over which the put is early exercised to the total number of paths q . Similarly, we compute the corresponding *conditional* theoretical probability over some window within the time-to-maturity as the ratio of the number of paths over which the put is first early exercised in that window to the total number of paths without an early exercise at the start of the window.

To implement the above methodology, we use 100,000 simulated underlying-asset-value paths for each put-month observation, with a number of time steps equal to the days-to-maturity of the put. We use a third-order polynomial to estimate the put's continuation value. While we directly observe the stock and strike price, the days-to-maturity, and the riskfree rate for each put-month observation, we estimate the underlying stock's volatility using monthly returns over the prior 60 months, allowing us to compute a forward-looking theoretical early exercise probability.

3.4 Data and Data Sources

We obtain daily data on American puts written on dividend-paying and non-dividend-paying single stocks and on those stocks from Optionmetrics. We retrieve riskfree rates of return from the zero-coupon yield curves also provided by Optionmetrics, and we always use that riskfree rate in our empirical work whose maturity date is closest to the date to which we want to compound or discount a cash flow. We obtain unique early exercise data manually collected from the Option Clearing Corporation's archives and containing the daily number of contracts exercised by put issue and owner (customers, market makers, and firms) from Bob

Whaley.⁸ Due to the availability of the early exercise data, our sample period is July 2001 to June 2014. We exclude the months November 2001, January 2002, July 2002, and January 2006 from that sample period because the early exercise data are consistently missing in these months. We further omit puts with a strike price-to-stock price ratio (moneyness) below 0.975 from our sample since such puts are hardly ever early exercised in real markets. We extract stock short-selling fees from Markit. We finally obtain the Fama-French benchmark factors, the VIX index, the TED spread, and a stock liquidity factor from Kenneth French’s website, the CBOE website, the Fred Database, and Lubos Pastor’s website, respectively.⁹

We apply standard filters to our data (see Goyal and Saretto (2009) and Cao and Han (2013)). To be specific, we exclude put-day observations for which the put price violates standard arbitrage bounds (as, e.g., that the put’s price must lie below its strike price). We further omit observations (i) for which the put price is below \$1 or one-half the bid-ask spread; (ii) for which the bid-ask spread is negative; or (iii) for which the underlying stock’s price is missing.

4 The Performance of the Risky Arbitrage Strategy

In this section, we examine the historical performance of the risky arbitrage strategy in the absence of transaction costs using the raw returns in Equations (11) and (17). We first look at the mean returns and return volatilities of the strategy and its two legs. We next adjust the mean returns for popular systematic factors, including firm-characteristic and macroeconomic factors.

⁸We are grateful to Bob Whaley for sharing these data with us. The shared data are an updated version of the data also used in Pool et al. (2008), Barraclough and Whaley (2012), and Jensen and Pedersen (2016).

⁹The benchmark factors are obtainable from: <https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/> and the stock liquidity factor from: <https://faculty.chicagobooth.edu/lubos.pastor/research/>. The URL for the VIX data are: <http://www.cboe.com/products/vix-index-volatility/vix-options-and-futures/vix-index/vix-historical-data>, and for the TED spread: <https://fred.stlouisfed.org/series/TEDRATE>.

4.1 Mean Return and Volatility of the Strategy

In Table 1, we offer descriptive statistics on the returns of our sample American put replication portfolios and traded American puts (columns (1) to (2), respectively), the spread across their returns (column (1)–(2)), and the moneyness and days-to-maturity of the puts (columns (3) to (4), respectively). We do not adjust returns for transaction costs, rebalance the replication portfolios daily, and define the strategy period to be one calendar month long. Each replication portfolio observation in column (1) corresponds to exactly one put observation in column (2). The descriptive statistics include the mean, the standard deviation (StDev), the Sharpe (1966) ratio, the mean’s t -statistic (Mean/StError), several percentiles, and the number of observations. With the exception of the t -statistic and the Sharpe ratio, we calculate the statistics by sample month and then average over time. Given that, we can interpret the means in columns (1) and (2) as the mean returns of equally-weighted portfolios of, respectively, the replication portfolios and of the puts. We calculate the t -statistic as the mean divided by the product of the standard deviation and the square root of the number of observations and the Sharpe (1966) ratio as the mean excess return (i.e., the mean return minus the riskfree rate of return) divided by the standard deviation. We finally compute moneyness as the strike-to-stock price ratio and measure both moneyness as well as time-to-maturity at the start of the strategy period.

TABLE 1 ABOUT HERE

Table 1 suggests that the risky arbitrage strategy is highly profitable. While both legs of the strategy yield significantly negative mean monthly returns in columns (1) and (2), the mean return of the portfolio of replication portfolios (the long leg) is a less negative -9.08% (t -statistic: -4.09) compared to the -13.20% (t -statistic: -4.88) mean return of the put portfolio (the short leg). In turn, the mean spread return across them is 4.11% (t -statistic: 5.21) in column (1)–(2). Crucially, the spread return is much less volatile than the returns of the legs, as can be seen from the standard

deviations and percentiles. While the annualized standard deviation of the spread portfolio is, for example, only 30.18%, those of the portfolio of replication portfolios and of the puts are 56.70% and 51.16%, respectively (compare columns (1), (2), and (1)–(2)). Using the three standard deviations, we can easily calculate the correlation between the two legs to be -0.85 .¹⁰ The low volatility of the spread portfolio implies that it has a higher t -statistic than the two legs in absolute terms, despite it having a far less extreme mean return. It further implies that the spread portfolio has an impressive annualized Sharpe ratio of 1.64. The moneyness and days-to-maturity statistics in columns (3) and (4) suggest that the average put in our risky arbitrage strategy is in-the-money (moneyness: 1.07) and has slightly more than two months to maturity.

4.2 The Factor Model Alphas of the Strategy

In Table 2, we study the performance of the risky arbitrage strategy after taking its systematic risk into account. To do so, we use Black et al.’s (1972) time-series methodology and regress the monthly return on the spread portfolio long the equally-weighted portfolio of replication portfolios and short the equally-weighted put portfolio on risk factors, interpreting the constant (“alpha”) from that regression as the risk-adjusted mean return of the strategy. In column (1), we use only the excess market return and a constant as exogenous variables. Column (2) adds Fama and French’s (1993) benchmark factors SMB and HML, while column (3) also adds Carhart’s (1997) MOM factor. Column (4) adds Fama and French’s (2015) additional benchmark factors PRF and INV. Column (5) finally adds the change in the VIX index, the TED funding spread, and Pastor and Stambaugh’s (2003) liquidity factor.¹¹ Plain numbers

¹⁰To wit, $0.5670^2 + 0.5116^2 - 2 \times 0.8482 \times 0.5670 \times 0.5116 = 0.3018^2$.

¹¹The SMB factor is the return of a portfolio long small and short big stocks controlling for book-to-market, while the HML factor is the return of a portfolio long high book-to-market (“value”) and short low book-to-market (“growth”) stocks controlling for size. The MOM factor is the return of a portfolio long stocks with high returns over the recent past and short stocks with low returns over that period. The PRF factor is the return of a portfolio long more profitable and short less profitable stocks, while INV is the return of a portfolio long low-investment and short high-investment stocks, with both factors controlling for size. See Kenneth French’s website for more details. The VIX index is a portfolio of options mimicking option-implied volatility, the TED spread is the difference between the interest rate on short-term U.S. government debt and on interbank loans, and the systematic liquidity factor is the return of a portfolio long stocks with a high

in the table are monthly premium estimates, while numbers in parentheses are t -statistics derived from Newey and West's (1987) formula with a twelve-month lag length.

TABLE 2 ABOUT HERE

Table 2 suggests that, independent of the risk factors used in the regressions, the alpha of the risky arbitrage strategy is always significantly positive, with it, in fact, hardly differing from the mean strategy return. Considering the most comprehensive model in column (5), the alpha is, for example, 4.05% per month (t -statistic: 8.70) — virtually identical to the mean strategy return of 4.11% (t -statistic: 5.21) in Table 1. Notwithstanding, the strategy does load significantly on several risk factors due to us being unable to perfectly replicate puts in real markets. Columns (1) to (4), for example, suggest that the strategy loads positively and significantly on MKT, SMB, and MOM. Interestingly, however, the MKT, SMB, and MOM loadings seem mostly attributable to volatility and liquidity risk, as can be seen from column (5). Not only do VIX and LIQ command significant loadings of, respectively, -0.16 (t -statistic: -4.67) and 0.44 (t -statistic: 3.65) in that column, but they also drive out the MKT, SMB, and MOM loadings. Although the risky arbitrage strategy thus loads on some risks, its loadings are markedly attenuated compared to those of its two legs. While its univariate MKT loading is, for example, 0.76 in column (1), the portfolio of replication portfolios (the long leg) and the put portfolio (the short leg) attract corresponding MKT loadings of -5.52 and -6.27 , respectively (unreported to conserve space).

Taken together, the empirical results in this section suggest that the return of the risky arbitrage strategy is not spanned by those of other well-known trading strategies and that the strategy thus represents a novel trading opportunity for investors.

liquidity exposure and short stocks with a low exposure. See Lubos Pastor's website for more details.

5 Conditioning the Strategy Return

In this section, we condition the performance of the risky arbitrage strategy on put, stock, and macroeconomic factors. We start with those factors for which we can deduce the sign of their conditioning effect from the theory in Section 2, which are the strike price and the interest rate. We then turn to factors which may condition the strategy’s performance through them capturing variations in investors’ tendency to early exercise puts too late, namely, moneyness, time-to-maturity, and stock volatility. We finally offer evidence suggesting that the conditioning ability of the latter factors indeed comes through suboptimally late early exercises, with us showing that investors do not seem to correctly condition their early exercise decisions on time-to-maturity and stock volatility. As in Section 4, we continue to ignore transaction costs.

5.1 Conditioning on the Strike Price and the Interest Rate

Assuming that American put investors can sometimes exercise their positions too late, our theory in Section 2 suggests that the performance of the risky arbitrage strategy is higher for high strike-price puts in high interest-rate regimes. To test those predictions, we start with sorting our sample replication portfolios, their associated puts, and the spread portfolios long a replication portfolio and short its associated put into portfolios according to the quintile breakpoints of the strike price in month $t - 1$. We label these portfolios our “unconditional strike price portfolios.” Since the strike price may, however, relate to moneyness (defined as the strike-to-stock price ratio), possibly leading moneyness to confound the unconditional portfolio results, we also construct “strike-price portfolios controlling for moneyness” in the spirit of An et al. (2014). To do so, we further split each unconditional portfolio into portfolios based on the quintile breakpoints of moneyness in month $t - 1$ and then form equally-weighted portfolios of those portfolios within the same strike-price classification. For both sets of portfolios, we equally-weight the portfolios, set up a spread portfolio long the top and short the bottom quintile, and hold the portfolios over month t .

Table 3 presents the results from the univariate portfolio exercise, with Panel A focusing on

the unconditional portfolios and Panel B on those controlling for moneyness. Plain numbers are mean monthly returns, while those in square parentheses are t -statistics calculated from Newey and West’s (1987) formula with a twelve-month lag length. In accordance with theory, the table strongly suggests that the profitability of the risky arbitrage strategy improves with the strike price. Looking at the unconditional portfolios in Panel A, the mean return of the strategy is a mere 2.34% (t -statistic: 3.59) in the lowest strike-price quintile but a more impressive 9.02% (t -statistic: 5.22) in the highest quintile. The difference in those numbers is equal to a significant 6.68% (t -statistic: 4.57). Remarkably, Panel B suggests that controlling for moneyness in the strike-price sorts hardly alters results. The reason is that the strike price and moneyness share an average cross-sectional correlation of only -0.03 , indicating that the two variables are almost orthogonal and mitigating confounding effects potentially arising through moneyness.

TABLE 3 ABOUT HERE

We next condition on the interest rate on top of the strike price. To do so, we first remember that our theory in Section 2 suggests that the arbitrage profit is proportional to the interest rate times the strike price (rK) in a perfect market. We then recognize that the mean interest rate drops significantly from 2.83% per annum over the first half of our sample period (2001-2007) to only 0.27% over the second half (2008-2014). Combining these two observations, it becomes obvious that we can use a difference-in-difference (DID) approach to investigate the effect of the interest rate on the profitability of the risky arbitrage strategy, relying on the drop in the interest rate as shock variable and the strike price as treatment variable. To facilitate that approach, we first calculate the mean returns of the unconditional strike-price spread portfolios from Panel A of Table 3 separately for the 2001-2007 and 2008-2014 subsample periods. We then calculate the difference in mean returns for each portfolio over the two periods and finally the difference in those differences across the highest and lowest strike-price quintile (“DID estimate”).

Table 4 presents the results from our DID tests studying how the interest rate conditions the

profitability of the risky arbitrage strategy. As before, the plain numbers in the table are mean monthly returns, while those in square parentheses are t -statistics calculated from Newey and West’s (1987) formula with a twelve-month lag length. Again in accordance with theory, the table strongly suggests that a higher interest rate positively conditions the mean strategy return, with the positive effect, however, being amplified by the strike price. While the mean monthly return of the strategy is, for example, only an insignificant 0.75 percentage points (t -statistic: 0.79) higher in the high relative to the low interest rate subsample period in the lowest strike-price quintile, it is a significant 7.94 points (t -statistic: 4.87) higher in the highest quintile. The difference in those two percentage-point increases is a highly significant 7.18% (t -statistic: 5.38).

TABLE 4 ABOUT HERE

5.2 Conditioning on Moneyness, Maturity Time, and Volatility

While the former subsection supports our theory by establishing that the risky arbitrage strategy is more profitable on higher strike-price puts in higher interest-rate regimes, it is conceivable that other factors condition the strategy’s profitability if they capture variations in investor’s tendency to exercise puts too late (i.e., if they predict $\max(t^E - \tau, 0)$). To study that possibility, we now condition the strategy’s return on a parsimonious set of put and stock characteristics potentially capturing such variations, including moneyness, time-to-maturity, and stock volatility. We start with moneyness and time-to-maturity. At the end of each month $t - 1$, we first sort the replication portfolios, the puts, and the spread portfolios into a deep in-the-money (DITM; strike-to-stock price above 1.10), in-the-money (ITM; 1.025-1.10), and at-the-money (ATM; 0.975-1.025) portfolio according to the relevant put’s moneyness. We next independently sort them into a short (below 60 days), medium (60-90 days), and long (above 90 days) time-to-maturity portfolio according to the relevant put’s time-to-maturity. The intersection of the two univariate portfolio sorts then yields 3×3 portfolios double-sorted on moneyness and time-to-maturity. We equally-weight the

constituents of those portfolios. We finally hold the portfolios over month t .

Table 5 presents the results from the double-sorted portfolio exercise, with columns (1), (2), and (1)–(2) focusing on the put replication portfolios, the puts, and the spread portfolios long a replication portfolio and short the associated put, respectively. In turn, Panels A, B, and C look into DITM, ITM, and ATM put strategies, respectively. As before, plain numbers are mean monthly portfolio returns, whereas the numbers in parentheses are t -statistics calculated from Newey and West’s (1987) formula with a twelve-month lag length. Column (1)–(2) in the table offers strong evidence that the profitability of the risky arbitrage strategy decreases in both moneyness and time-to-maturity. Looking at strategies based on 30-60 day puts, the mean strategy return, for example, decreases from 9.66% (t -statistic: 6.53) in the ATM-put portfolio in Panel C to 2.49% (t -statistic: 5.28) in the DITM-put portfolio in Panel A. Conversely, looking at strategies based on ITM puts in Panel B, that same return decreases from 4.94% (t -statistic: 6.03) in the 30-60 day put portfolio to 2.58% (t -statistics: 4.16) in the 90-120 day put portfolio. Turning to the underlying replication portfolios and puts in columns (1) and (2), their mean returns drop with moneyness but rise with time-to-maturity, in line with the put results in Aretz and Gazi (2020).

TABLE 5 ABOUT HERE

As a next step, we condition the performance of the risky arbitrage strategy on the idiosyncratic volatility of the underlying stocks. At the end of each month $t - 1$, we thus split the replication portfolios, puts, and the spread portfolios long a replication portfolio and short the associated put into portfolios according to the quintile breakpoints of that volatility estimated using the market model or the Fama-French-Carhart (FFC; 1997) model. We can write the market model as:

$$R_{i,\tau} = \alpha_i + \beta_i^{mkt}(R_\tau^{mkt} - R_{f\tau}) + \epsilon_{i,\tau}, \quad (22)$$

where $R_{i,\tau}$ is stock i ’s return over month τ , $R_\tau^{mkt} - R_{f\tau}$ is the excess market return, α_i and β_i^{mkt}

are parameters, and $\epsilon_{i,\tau}$ is the residual. We can write the FFC model as:

$$R_{i,\tau} = \alpha_i + \beta_i^{mkt}(R_\tau^{mkt} - R_{f\tau}) + \beta_i^{smb}R_\tau^{smb} + \beta_i^{hml}R_\tau^{hml} + \beta_i^{mom}R_\tau^{mom} + \epsilon_{i,\tau}, \quad (23)$$

where R_τ^{smb} , R_τ^{hml} , and R_τ^{mom} are the returns of spread portfolios on size, the book-to-market ratio, and the eleven-month (momentum) past return, respectively, and β_i^{smb} , β_i^{hml} , and β_i^{mom} are additional parameters. We estimate both models over the prior 60 months of monthly data, calculating idiosyncratic volatility as the standard deviation of the residual, $\epsilon_{i,\tau}$. We equally-weight the constituents of the quintile portfolios and hold them over month t .

Table 6 presents the results from the univariate portfolio exercise, with Panels A and B using market- and FFC-model estimates to proxy for idiosyncratic stock volatility, respectively. As before, plain numbers are mean monthly returns, whereas numbers in parentheses are Newey-West (1987) t -statistics. The table offers strong evidence that the profitability of the risky arbitrage strategy deteriorates with idiosyncratic stock volatility. Looking at the market-model sorts in Panel A, the mean strategy return in the third row, for example, drops from 6.07% (t -statistic: 4.70) in the low-volatility portfolio to 2.81% (t -statistic: 5.36) in the high-volatility portfolio. The difference in those two numbers is a highly significant -3.26% (t -statistic: -3.19). Turning to the replication portfolios and puts in the first and second rows of each panel, their mean returns increase with idiosyncratic stock volatility, aligning with Aretz and Gazi’s (2020) put results.

TABLE 6 ABOUT HERE

5.3 Jointly Conditioning on All Factors

We next investigate how the strike price, moneyness, time-to-maturity, and stock volatility jointly condition the performance of the risky arbitrage strategy and verify that our portfolio results are robust to variations in methodology. To do so, we run Fama-MacBeth (FM; 1973) regressions of

the returns on the replication portfolios, the puts, and the spread portfolios long a replication portfolio and short its associated put over month t on combinations of those conditioning factors calculated using only data until the end of month $t - 1$. To mitigate outlier effects, we use the log strike price (instead of the strike price) and time-to-maturity stated as fraction of a year in the regressions. As before, moneyness is the strike-to-stock price ratio, while stock volatility is the annualized FFC idiosyncratic volatility estimate, introduced in Section 5.2.

Table 7 gives the regression results, with Panels A, B, and C using the spread portfolio return, replication portfolio return, and put return as dependent variable, respectively. Plain numbers are monthly premium estimates, whereas the numbers in parentheses are Newey and West (1987) t -statistics with a twelve-month lag length. The table shows that the regressions yield results almost exactly identical with those from the portfolio exercises. To be specific, column (1) in Panel A reveals that the unconditional mean spread return is 4.11% (t -statistic: 7.33) per month, aligning with column (1)–(2) in Table 1. Moreover, columns (2)–(5) in that panel confirm that the mean spread return significantly rises in the strike price but significantly drops in moneyness, time-to-maturity, and stock volatility, at least when strike price and stock volatility are not jointly included as independent variables. Interestingly, when we jointly include the strike price and stock volatility, as we do in column (6) of Panel A, the strike-price premium hardly changes, while the stock volatility premium switches from being significantly negative to significantly positive. The reason for this unexpected outcome is that the strike price and stock volatility are negatively correlated, with higher volatility stocks often having many low strike-price puts written on them.

TABLE 7 ABOUT HERE

Turning to the replication portfolio and put results in Panels B and C of Table 7, we notice that their monthly premiums on the strike price, moneyness, time-to-maturity, and stock volatility are of exactly the same sign as those obtained in the portfolio exercises.

5.4 Why Do Moneyness, Time-to-Maturity, and Stock Volatility Condition the Success of the Risky Arbitrage Strategy?

While we argue that the ability of moneyness, time-to-maturity, and stock volatility to condition the risky arbitrage strategy in Sections 5.2 and 5.3 is due to those factors capturing variations in investors' tendency to exercise puts too late, there may be other reasons for that ability. To offer more support for the hypothesis that suboptimally late exercises do indeed lie behind that ability, we next take a closer look at real investors' early exercise behavior. We start with benchmarking the real-world early exercise probabilities of our sample puts against their corresponding optimal probabilities deduced from the Black and Scholes (1973) model. We calculate the probabilities as described in Sections 3.1 and 3.3, respectively. While we acknowledge that the shortcomings of the Black and Scholes (1973) model imply that the optimal probabilities deduced from it differ from the true optimal probabilities, we nonetheless hope to learn some broader lessons about the optimality of real investor's early exercise strategies from the comparisons.¹²

To take a first look at how the real-world and Black-Scholes (1973) early exercise probabilities of our sample puts over their entire times-to-maturity are related, Table 8 reports their mean values over portfolios formed according to the decile breakpoints of the theoretical probabilities at the start of the strategy return period. We calculate the mean values first by cross-section and then average over time. The table reveals that the real-world and the Black-Scholes (1973) probabilities are strongly positively correlated. While the monotonic increase in the mean Black-Scholes (1973) probabilities over the portfolios from 16.35% to 76.74% is by construction, the mean real-world probabilities, remarkably, also monotonically increase over them, from 6.27% to 27.05%, sharing an average cross-sectional correlation of 0.23 between the two probabilities. Notwithstanding, the mean real-world probabilities are consistently 2-3 times lower than the

¹²It is not obvious to us how to calculate more accurate theoretical early exercise probabilities. Switching to a more sophisticated option pricing model, as, for example, the Heston (1993) model, would require us to estimate additional parameters governing, for example, the mean reversion in volatility, the long-run volatility, the correlation between asset value and volatility, etc. Unfortunately, these additional parameters are tremendously difficult to estimate, especially at the stock level at which we only have limited amounts of data.

mean Black-Scholes (1973) probabilities, suggesting that real-world put investors often wait too long with exercising their positions and that the necessary condition for a risky arbitrage profit to exist is fulfilled. While relaxing certain Black-Scholes (1973) assumptions (as, e.g., the constant volatility and/or no asset-value jumps assumptions) could help to close the gaps between the probabilities, we deem it unlikely that it can account for the entire gaps.

TABLE 8 ABOUT HERE

We next investigate how the differences between the real-world and Black-Scholes (1973) early exercise probabilities relate to moneyness, time-to-maturity, and stock volatility. To do so, Table 9 reports the mean differences in those probabilities calculated over the strategy return period for portfolios triple-sorted according to those characteristics. We construct the portfolios as follows. At the end of month $t - 1$, we independently sort our puts according to the same moneyness, time-to-maturity, and FFC idiosyncratic stock volatility breakpoints as in Tables 5 and 6, using the intersections of the univariate portfolios to create the triple-sorted portfolios. As before, we calculate the mean probabilities first by cross-section and then average over time.

TABLE 9 ABOUT HERE

The table suggests that the difference between the probabilities tends to be more pronounced for shorter time-to-maturity puts written on lower volatility stocks. Looking at ITM puts on stocks with a third-quintile volatility in Panel B, the mean difference is, for example, 24.20% for 30-60 day puts but only 3.58% for 90-120 day puts. Conversely, looking at ITM puts with 60-90 days-to-maturity in the same panel, the mean difference is 22.96% for puts written on the lowest-volatility stocks but only 4.22% for those written on the highest-volatility stocks. The upshot is that real investors seem particularly bad in timely exercising short time-to-maturity puts on low-volatility stocks, which may explain why the risky arbitrage strategy works better on such puts. Interestingly, however, the

table further suggests that the difference in the probabilities also tends to be more pronounced for higher moneyness puts. Looking at 30-60 day puts on third-quintile volatility stocks, the mean difference is, for example, 36.66% on DITM puts but only 12.07% on ATM puts. The implication is that variations in real investors' tendency to early exercise puts too late does not help to explain why the risky arbitrage strategy works better on lower-moneyness puts.

In Table 10, we switch to FM regressions to find out how real investors' tendency to exercise puts too late varies with moneyness, time-to-maturity, and stock volatility. While the dependent variable is the difference between the Black-Scholes (1973) and real-world early exercise probabilities, the independent variables are the characteristics defined as in Table 7. The regressions deliver results in alignment with the portfolio exercise results, showing that the difference in the probabilities tends to rise with moneyness and to drop with time-to-maturity and stock volatility.

TABLE 10 ABOUT HERE

6 Adjusting for Transaction Costs

In this section, we adjust the profitability of the risky arbitrage strategy for bid-ask and short-selling transaction costs. As already said, it is crucial to adjust for those costs since the strategy involves potentially frequent stock purchases and sales and shorting the stock. In addition, Goyal and Saretto (2009) and Cao and Han (2013) show that high bid-ask costs in the options market greatly eat into the profitability of option trading strategies, often rendering profits insignificant or even negative. To account for transaction costs, we switch from studying the raw returns in Equations (11) and (17) to studying the transaction cost adjusted returns in Equations (12) and (21), setting φ , the proportion of the bid-ask spread representing trading costs, to either zero, 0.10, 0.25, or 0.50. When $\varphi = 0.50$, investors buy at the ask price and sell at the bid price.

Table 10 reevaluates the profitability of the double-sorted moneyness and time-to-maturity portfolios originally studied in Table 5 under transaction costs, reporting, however, only the mean

strategy return (and not the mean leg returns). While Panel A looks into our full sample, Panel B focuses only on strategies executed on liquid assets. A liquid stock (put) is defined as one with an Amihud (2002) stock illiquidity estimate (bid-ask spread scaled by put price) below the median, with the liquidity proxies measured at the start of the strategy period. Conversely, columns (1) to (4) consider the $\varphi =$ zero, 0.10, 0.25, and 0.50 case, with the final three columns (but not the first) also adjusting for shorting costs. As in Table 5, plain numbers are mean monthly returns, while those in parentheses are Newey-West (1987) t -statistics. Since the Markit short-selling data are available from only January 2002, our empirical work adjusting for transaction costs relies on the sample period from that date to June 2014 (leading us to lose five sample months).

TABLE 11 ABOUT HERE

In line with expectations, Panel A of Table 11 confirms that adjusting for transaction costs greatly reduces the profitability of our risky arbitrage strategy. Starting with column (1) in that panel, it is obvious that the slight change in our sample period does not materially affect mean strategy returns (compare the column with column (1)–(2) in Table 5). In contrast, assuming that investors incur trading costs φ equal to 10% of an asset’s bid-ask spread plus stock short-selling costs, column (2) reveals that the mean strategy return remains significantly positive only for strategies based on 30-60 day ITM or ATM puts (see Panels A.2 and A.3). In all other cases, it is either insignificant or, in one case, even significantly negative. Assuming that investors incur even higher trading costs (i.e., $\varphi = 0.25$ or 0.50), columns (3) and (4) finally show that the mean strategy return becomes highly significantly negative in the vast majority of cases.

Fortunately, Panel B of Table 11 suggests that the picture markedly improves once we restrict our attention to strategies executed on liquid assets. Interestingly, column (1) in that panel shows that the mean strategy return is generally higher in case of such strategies, even in the absence of transaction cost adjustments. More importantly, however, columns (2) to (4) indicate that the mean strategy return remains positive and significant even under the assumption that investors

incur a 10% (25%) [50%] bid-ask trading cost plus stock short-selling costs in case of nine (six) [one] out of the nine strategies conditioned on moneyness and time-to-maturity.

In Table 12, we aim to decrease transaction costs further by reducing the frequency with which we rebalance the put replication portfolio, saving us costs arising from stock purchases and sales over the strategy return period. The downside is that a lower rebalancing frequency decreases the ability of a replication portfolio to track its associated put, boosting the volatility of the strategy return and making the strategy diverge even more from a textbook arbitrage strategy. To study the effects of a lower rebalancing frequency, Panels A, B, and C of Table 12 present the mean transaction cost adjusted returns of strategies executed on only liquid assets and using daily, weekly, and no rebalancing, respectively. The no-rebalancing case is identical to the “buy-and-hold delta-hedging strategy” of Goyal and Saretto (2009), who form replication portfolios at the start of the strategy period and hold those over the entire period. Within each panel, we consider strategies involving only puts with a price above \$1, \$2, and \$5 at the start of the strategy return period.

TABLE 12 ABOUT HERE

The table suggests that a lower rebalancing frequency helps to further increase the profitability of the risky arbitrage strategy. While, surprisingly, a lower rebalancing frequency often also boosts the mean strategy return and its t -statistic in the absence of transaction cost adjustments (see column (1)), the improvements are far more pronounced in their presence. Looking at the case in which investors incur a 50% bid-ask trading cost plus stock short-selling costs in column (4), Panel A shows that only a single mean strategy return out of three is positive and weakly significant (with a t -statistic of 2.04) under daily rebalancing. In contrast, Panel B suggests that the corresponding mean returns under no rebalancing are significantly positive with t -statistics above 3.88.

Overall, this section shows that the risky arbitrage strategy does not only exist in theory but can profitably be exploited by real investors, at least when executed on liquid assets.

7 Concluding Remarks

We evaluate the profitability of a risky arbitrage strategy exploiting evidence that real put investors often exercise their positions far too late. Grounded in mathematical finance theory, the strategy consists of a long position in a stock-and-riskfree-asset portfolio replicating a put and a short position in the put. Our empirical work suggests that the strategy yields a highly significant positive mean return, with a low total as well as systematic risk. Consistent with theory, the mean strategy return rises with both the strike price and the interest rate. Interestingly, however, it drops with moneyness, time-to-maturity, and stock volatility, aligning with further evidence that investors do not appear to understand the implications of those characteristics for the optimal early exercise decision. Crucially, we finally show that the strategy survives accounting for both bid-ask trading costs and shorting costs, at least when it is executed on liquid assets. The upshot is that the strategy does not only exist on paper but can be pursued by real investors.

References

- Amihud, Y., 2002. Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1), pp.31-56.
- An, B.J., Ang, A., Bali, T.G. and Cakici, N., 2014. The joint cross section of stocks and options. *The Journal of Finance*, 69(5), pp.2279-2337.
- Aretz, K. and Gazi, A., 2020. The early exercise risk premium. *Working Paper*.
- Barraclough, K. and Whaley, R.E., 2012. Early exercise of put options on stocks. *The Journal of Finance*, 67(4), pp.1423-1456.
- Battalio, R., Figlewski, S. and Neal, R., 2020. Option investor rationality revisited: The role of exercise boundary violations. *Financial Analysts Journal*, 76(1), pp.82-99.
- Bauer, R., Cosemans, M. and Eichholtz, P., 2009. Option trading and individual investor performance. *Journal of Banking & Finance*, 33(4), pp.731-746.
- Black, F., Jensen, M.C., and Scholes, M.S., 1972. The capital asset pricing model: Some empirical tests. *Studies in the Theory of Capital Markets*, New York: Praeger, pp.79-121.
- Black, F. and Scholes, M.S., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), pp.637-654.
- Brennan, M.J. and Schwartz, E.S., 1977. The valuation of American put options. *The Journal of Finance*, 32(2), pp.449-462.
- Cao, J. and Han, B., 2013. Cross section of option returns and idiosyncratic stock volatility. *Journal of Financial Economics*, 108(1), pp.231-249.
- Cao, J., Han, B., Tong, Q. and Zhan, X., 2020. Option return predictability. *Working Paper*.
- Carhart, M.M., 1997. On persistence in mutual fund performance. *The Journal of Finance*, 52(1), pp.57-82.
- Coval, J.D. and Shumway, T., 2001. Expected option returns. *The Journal of Finance*, 56(3), pp.983-1009.
- Fama, E.F. and French, K.R., 1992. The cross-section of expected stock returns. *The Journal of Finance*, 47(2), pp.427-465.
- Fama, E.F., and French, K.R., 1993. Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), pp.3-56.
- Fama, E.F. and French, K.R., 2015. A five-factor asset pricing model. *Journal of Financial Economics*, 116(1), pp.1-22.
- Fama, E.F. and MacBeth, J.D., 1973. Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3), pp.607-636.
- Figlewski, S., 2019. An American call is worth more than a European call: The value of American exercise when the market is not perfectly liquid. *Working Paper*.
- Finucane, T.J., 1997. An empirical analysis of common stock call exercise: A note. *Journal of Banking & Finance*, 21(4), pp.563-571.

- Gondzio, J., Kouwenberg, R. and Vorst, T., 2003. Hedging options under transaction costs and stochastic volatility. *Journal of Economic Dynamics and Control*, 27(6), pp.1045-1068.
- Goyal, A., and Saretto, A., 2009. Cross-section of option returns and volatility. *Journal of Financial Economics*, 94(2), pp.310-326.
- Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2), pp.327-343.
- Jacka, S.D., 1991. Optimal stopping and the American put. *Mathematical Finance*, 1(2), pp.1-14.
- Jensen, M.V., and Pedersen, L.H., 2016. Early option exercise: Never say never. *Journal of Financial Economics*, 121(2), pp.278-299.
- Karatzas, I., 1988. On the pricing of American options. *Applied Mathematics and Optimization*, 17(1), pp.37-60.
- Lakonishok, J., Shleifer, A. and Vishny, R.W., 1994. Contrarian investment, extrapolation, and risk. *The Journal of Finance*, 49(5), pp.1541-1578.
- Leland, H.E., 1985. Option pricing and replication with transactions costs. *The Journal of Finance*, 40(5), pp.1283-1301.
- Longstaff, F.A. and Schwartz, E.S., 2001. Valuing American options by simulation: A simple least-squares approach. *The Review of Financial Studies*, 14(1), pp.113-147.
- Merton, R.C., 1973. Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1), pp.141-183.
- Newey, W.K., and West, K.D., 1987. A simple positive-definite heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55, pp.703-708.
- Overdahl, J.A. and Martin, P.G., 1994. The exercise of equity options: Theory and empirical tests. *The Journal of Derivatives*, 2(1), pp.38-51.
- Pástor, L., and Stambaugh, R.F., 2003. Liquidity risk and expected stock returns. *Journal of Political Economy*, 111(3), pp.642-685.
- Perrakis, S. and Lefoll, J., 2000. Option pricing and replication with transaction costs and dividends. *Journal of Economic Dynamics and Control*, 24(11-12), pp.1527-1561.
- Pool, V.K., Stoll, H.R. and Whaley, R.E., 2008. Failure to exercise call options: An anomaly and a trading game. *Journal of Financial Markets*, 11(1), pp.1-35.
- Poteshman, A.M. and Serbin, V., 2003. Clearly irrational financial market behavior: Evidence from the early exercise of exchange traded stock options. *The Journal of Finance*, 58(1), pp.37-70.
- Sharpe, W.F., 1966. Mutual fund performance. *The Journal of Business*, 39(1), pp.119-138.
- Shreve, S.E., 2004. *Stochastic calculus for finance II: Continuous-time models* (Vol. 11). Springer Science & Business Media.
- Vasquez, A., 2017. Equity volatility term structures and the cross section of option returns. *Journal of Financial and Quantitative Analysis*, 52(6), pp.2727-2754.

Table 1: The Performance of the Risky Arbitrage Strategy

The table presents descriptive statistics for the monthly returns on daily-rebalanced stock-and-riskfree-asset portfolios replicating American puts (column (1)), on the puts (column (2)), and on spread portfolios long a replication portfolio and short the corresponding put (column (1)–(2)). It further presents descriptive statistics on the moneyness (column (3)) and time-to-maturity (column (4)) of the puts. The descriptive statistics are the mean, the standard deviation (StDev), the t -statistic for the mean ($Mean/StError$), the monthly and annual Sharpe (1966) ratio (Monthly and Annual Sharpe Ratio, respectively), several percentiles, and the number of observations. Each observation used in column (1) corresponds to one observation used in column (2). We calculate moneyness as the strike-to-stock price ratio and time-to-maturity as the number of calendar days until maturity, both at the start of the return period. With the exception of the t -statistic and the Sharpe ratio, we calculate each statistic as the time-series mean of the cross-sectional statistic. The t -statistic is the mean scaled by the standard error of the mean. The Sharpe ratio is the difference between mean return and riskfree rate scaled by the standard deviation of the mean in column (1) and (2), and the mean scaled by the standard deviation of the mean in column (1)–(2).

	Monthly Replication Portfolio Return (in %)	Monthly American Put Option Return (in %)	Monthly Spread Portfolio Return (in %)	Money- ness	Days to Maturity
	(1)	(2)	(1)–(2)	(3)	(4)
Mean	−9.08	−13.20	4.11	1.07	73
StDev	56.70	51.16	30.18		
Mean/StError	[−4.09]	[−4.88]	[5.21]		
Monthly Sharpe Ratio			0.47		
Annual Sharpe Ratio			1.64		
Percentile 1	−105.50	−90.63	−39.39	0.98	49
Percentile 5	−80.26	−79.94	−17.27	0.99	49
Quartile 1	−43.35	−48.79	−2.87	1.02	50
Median	−14.99	−20.04	2.72	1.07	62
Quartile 3	15.89	12.17	8.73	1.12	96
Percentile 95	78.65	76.95	21.66	1.18	111
Percentile 99	169.58	152.01	67.66	1.20	111
Observations	5,612	5,612	5,612	5,612	5,612

Table 2: Adjusting the Strategy's Performance for Systematic Risk

The table shows the results from time-series regressions of the month- t return of a spread portfolio long a stock-and-riskfree-asset portfolio replicating a put and short the put on several sets of benchmark factors measured over the same month and a constant. The factors include the market return minus the risk-free rate of return (MKT); the return of a spread portfolio long small and short large stocks (SMB); the return of a spread portfolio long value and short growth stocks (HML); the return of a spread portfolio long winner and short loser stocks (MOM); the return of a spread portfolio long profitable and short unprofitable stocks (PRF); the return of a spread portfolio long non-investing and short investing stocks (INV); the change in the VIX option-implied volatility index (VIX); the three-month LIBOR rate minus the Treasury bill rate (TED); and the return of a spread portfolio long high-liquidity and short low-liquidity stocks (LIQ). Plain numbers are estimates, while the numbers in square parentheses are t -statistics calculated using Newey and West's (1987) formula with a twelve-month lag length.

	Time-Series Regression Model:				
	(1)	(2)	(3)	(4)	(5)
MKT	0.76 [6.89]	0.67 [5.67]	0.82 [6.31]	0.81 [5.63]	0.03 [0.17]
SMB		0.48 [2.19]	0.45 [2.09]	0.48 [2.21]	0.34 [1.74]
HML		-0.03 [-0.12]	0.03 [0.14]	0.15 [0.65]	0.64 [3.04]
MOM			0.27 [2.54]	0.27 [2.41]	0.16 [1.58]
PRF				0.08 [0.26]	-0.25 [-0.96]
INV				-0.46 [-1.42]	-0.52 [-1.85]
VIX					-0.16 [-4.67]
TED					-3.03 [-1.72]
LIQ					0.44 [3.65]
Constant	0.04 [7.51]	0.04 [7.25]	0.04 [7.18]	0.04 [6.96]	0.04 [8.70]

Table 3: Conditioning the Strategy on the Strike Price

The table presents the mean returns of portfolios of stock-and-riskfree-asset portfolios replicating a put and of the associated puts sorted on the strike price. It further presents the mean returns of the corresponding spread portfolios long the portfolios of put replication portfolios and short the put portfolios. At the end of each sample month $t - 1$, we first sort each of those assets into portfolios according to the quintile breakpoints of the associated strike price, without controlling for the strike-to-stock price ratio (“moneyness;” Panel A). Within each strike price portfolio, we next sort them into further portfolios according to the quintile breakpoints of moneyness in month $t - 1$. We then form equally-weighted portfolios of those portfolios within the same strike price classification, averaging out the effect of moneyness (Panel B). We also form a spread portfolio long the top and short the bottom quintile (“High–Low”). We equally-weight the portfolios and hold them over month t . The put and replication portfolio observations are matched, so each put observation corresponds to one replication portfolio observation. Plain numbers are mean monthly portfolio returns (in %) and the numbers in square parentheses are t -statistics calculated using Newey and West’s (1987) formula with a twelve-month lag length.

	Strike Price, K					High–Low
	1 (Low)	2	3	4	5 (High)	
Panel A: Univariate Strike-Price Portfolios						
Replication Portfolio	−8.52 [−4.19]	−10.89 [−4.95]	−11.83 [−5.28]	−11.60 [−4.89]	−6.32 [−2.26]	2.21 [1.26]
American Put Option	−10.86 [−4.31]	−12.87 [−4.66]	−14.22 [−5.06]	−14.85 [−5.07]	−15.33 [−5.06]	−4.47 [−2.80]
Spread Portfolio	2.34 [3.59]	1.99 [2.98]	2.40 [3.68]	3.25 [4.70]	9.02 [5.22]	6.68 [4.57]
Panel B: Strike-Price Portfolios Controlling for Moneyness						
Replication Portfolio	−8.49 [−4.20]	−10.72 [−4.90]	−11.73 [−5.26]	−11.51 [−4.86]	−6.23 [−2.23]	2.26 [1.27]
American Put Option	−10.87 [−4.33]	−12.74 [−4.63]	−14.18 [−5.05]	−14.81 [−5.06]	−15.34 [−5.06]	−4.47 [−2.77]
Spread Portfolio	2.38 [3.62]	2.02 [3.01]	2.44 [3.72]	3.29 [4.72]	9.10 [5.24]	6.72 [4.58]

Table 4: Conditioning the Strategy on the Strike Price and Interest Rate

The table presents the mean returns of spread portfolios long an equally-weighted portfolio of put replication portfolios and short an identically-weighted portfolio of the associated puts sorted on the strike price and separately calculated over the July-2001 to December-2007 (column (1)) and the January-2008 to June-2014 (column (2)) subsample periods. The table also reports the differences in mean spread portfolio returns across the subsample periods (column (1)–(2)). At the end of each sample month $t - 1$ within a subsample period, we first sort the spread portfolios into portfolios according to the quintile breakpoints of the strike price of the associated put. We also form a spread portfolio long the top and short the bottom quintile (“High–Low”). We equally-weight the portfolios and hold them over month t . The replication portfolio and put observations are matched, so that each replication portfolio observation corresponds to one put observation. Plain numbers are mean monthly portfolio returns (in %), and the numbers in square parentheses are t -statistics calculated using Newey and West’s (1987) formula with a twelve-month lag length.

Strike Price, K	Until 2007	From 2008	Difference
	(1)	(2)	(1)–(2)
1 (Low)	2.75 [5.95]	1.99 [1.85]	0.75 [0.79]
2	2.00 [4.57]	1.98 [1.74]	0.03 [0.03]
3	2.57 [5.55]	2.24 [2.04]	0.33 [0.32]
4	3.72 [6.48]	2.85 [2.52]	0.87 [0.81]
5 (High)	13.32 [6.33]	5.38 [3.97]	7.94 [4.87]
High–Low	10.57 [5.39]	3.39 [4.90]	7.18 [5.38]

Table 5: Conditioning the Strategy on Moneyness and Maturity Time

The table presents the mean returns of portfolios of stock-and-riskfree-asset portfolios replicating a put (column (1)) and of the associated puts (column (2)) sorted on the puts' moneyness and time-to-maturity. It further presents the mean returns of the corresponding spread portfolios long the portfolios of put replication portfolios and short the put portfolios (column (1)–(2)). At the end of each sample month $t - 1$, we first sort each of these assets into portfolios according to whether the strike-to-stock price ratio (“moneyness”) of the associated put lies above 1.10 (Panel A), between 1.025 and 1.10 (Panel B), or between 0.975 and 1.025 (Panel C). Within each moneyness portfolio, we next sort them into portfolios according to whether their days-to-maturity are below 60, between 60 and 90, or above 90 days. We equally-weight the portfolios and hold them over month t . The observations used in columns (1) and (2) are matched, so that each observation in column (1) corresponds to one observation in column (2). Plain numbers are mean monthly portfolio returns (in %) and the numbers in square parentheses are t -statistics calculated using Newey and West’s (1987) formula with a lag length of twelve months.

Days-to-Maturity	Monthly Replication Portfolio Return (in %)	Monthly American Put Option Return (in %)	Monthly Spread Portfolio Return (in %)
	(1)	(2)	(1)–(2)
Panel A: Deep In-The-Money (Strike-to-Stock Price > 1.10)			
30-60	–15.31 [–7.61]	–17.80 [–7.94]	2.49 [5.28]
60-90	–8.67 [–5.17]	–10.46 [–5.31]	1.79 [3.72]
90-120	–5.80 [–3.61]	–7.28 [–3.85]	1.48 [3.90]
Panel B: In-The-Money (Strike-to-Stock Price 1.025 to 1.10)			
30-60	–13.33 [–5.23]	–18.27 [–6.08]	4.94 [6.03]
60-90	–7.10 [–3.29]	–10.42 [–3.92]	3.32 [3.95]
90-120	–4.58 [–2.30]	–7.16 [–2.90]	2.58 [4.16]

(continued on next page)

Table 5: Conditioning the Strategy on Moneyness and Maturity Time (Cont.)

Days-to-Maturity	Monthly Replication Portfolio Return (in %)	Monthly American Put Option Return (in %)	Monthly Spread Portfolio Return (in %)
	(1)	(2)	(1)-(2)
Panel C: At-The-Money (Strike-to-Stock Price 0.975 to 1.025)			
30-60	-7.36 [-2.49]	-17.01 [-4.55]	9.66 [6.53]
60-90	-5.83 [-2.17]	-11.19 [-3.34]	5.36 [3.92]
90-120	-4.43 [-1.86]	-8.03 [-2.65]	3.60 [4.12]

Table 6: Conditioning the Strategy on Idiosyncratic Stock Volatility

The table presents the mean returns of portfolios of stock-and-riskfree-asset portfolios replicating a put and of the associated puts sorted on the idiosyncratic volatility of the stock. It further presents the mean returns of the corresponding spread portfolios long the portfolios of put replication portfolios and short the put portfolios. At the end of each sample month $t - 1$, we sort each of these assets into portfolios according to the quintile breakpoints of a stock's market-model (Panel A) or Fama-French-Carhart-model (Panel B) idiosyncratic volatility. We estimate the models over the prior 60 months, defining idiosyncratic volatility as the volatility of the residual. We also form a spread portfolio long the top and short the bottom quintile ("High-Low"). We equally-weight the portfolios and hold them over month t . The replication portfolio observations and put observations are matched, so that each replication portfolio observation corresponds to one put observation. Plain numbers are mean monthly portfolio returns (in %), and the numbers in square parentheses are t -statistics calculated using Newey and West's (1987) formula with a twelve-month lag length.

	Idiosyncratic Stock Volatility					
	1 (Low)	2	3	4	5 (High)	High-Low
Panel A: Market Model Idiosyncratic Volatility						
Replication Portfolio	-8.99 [-3.27]	-9.70 [-4.43]	-9.92 [-4.57]	-9.50 [-4.36]	-7.30 [-3.38]	1.69 [1.07]
American Put Option	-15.06 [-4.71]	-14.37 [-5.10]	-13.99 [-5.26]	-12.47 [-4.77]	-10.11 [-4.02]	4.95 [3.25]
Spread Portfolio	6.07 [4.70]	4.67 [4.83]	4.07 [5.22]	2.96 [4.38]	2.81 [5.36]	-3.26 [-3.19]
Panel B: FFC Model Idiosyncratic Volatility						
Replication Portfolio	-9.04 [-3.29]	-9.78 [-4.46]	-9.63 [-4.24]	-9.56 [-4.45]	-7.41 [-3.57]	1.63 [1.05]
American Put Option	-14.94 [-4.63]	-14.56 [-5.16]	-13.63 [-5.03]	-12.67 [-4.86]	-10.21 [-4.21]	4.73 [3.03]
Spread Portfolio	5.90 [4.76]	4.77 [4.82]	4.00 [5.12]	3.11 [4.40]	2.80 [5.57]	-3.10 [-3.18]

Table 7: Fama-MacBeth (1973) Regressions

The table presents the results of Fama-MacBeth (1973) regressions of the month- t return of a spread portfolio long a stock-and-riskfree-asset portfolio replicating a put and short the put (Panel A), of the replication portfolio (Panel B), and of the put (Panel C) on subsets of stock and option characteristics plus a constant. The characteristics are the log strike price, the strike-to-stock price ratio (“moneyness”), time-to-maturity as fraction of a year, and idiosyncratic stock volatility, all measured at the start of month t . We calculate idiosyncratic volatility using the Fama-French-Carhart model estimated over the prior 60 months. The replication portfolio observations and put observations are matched, so that each replication portfolio observation corresponds to one put observation. The plain numbers are premium estimates, and the numbers in square parentheses are t -statistics calculated using Newey and West’s (1987) formula with a twelve-month lag length.

	Regression Model:					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Spread Portfolio Return						
Strike Price		0.04 [8.34]	0.04 [7.91]			0.04 [8.03]
Moneyness			-0.26 [-6.70]		-0.29 [-7.46]	-0.26 [-6.96]
Time-to-Maturity			-0.17 [-13.85]		-0.18 [-14.06]	-0.16 [-13.78]
Volatility				-0.03 [-3.81]	-0.02 [-2.59]	0.04 [5.02]
Constant	0.04 [7.33]	-0.10 [-6.26]	0.22 [4.63]	0.05 [7.25]	0.40 [8.07]	0.20 [4.04]
Panel B: Replication Portfolio Return						
Strike Price		0.02 [2.27]	0.02 [2.07]			0.03 [3.31]
Moneyness			-0.32 [-5.00]		-0.37 [-5.79]	-0.35 [-5.53]
Time-to-Maturity			0.49 [14.21]		0.49 [14.18]	0.50 [14.64]
Volatility				0.05 [2.17]	0.08 [3.24]	0.11 [4.11]
Constant	-0.09 [-4.38]	-0.16 [-4.91]	0.09 [1.02]	-0.12 [-4.89]	0.17 [1.86]	0.04 [0.38]

(continued on next page)

Table 7: Fama-MacBeth (1973) Regressions (Cont.)

	Regression Model:					
	(1)	(2)	(3)	(4)	(5)	(6)
Panel C: Put Return						
Strike Price		-0.02 [-3.67]	-0.02 [-3.45]			-0.01 [-1.99]
Moneyness			-0.06 [-0.77]		-0.08 [-0.93]	-0.08 [-1.00]
Time-to-Maturity			0.66 [18.68]		0.67 [18.92]	0.67 [19.00]
Volatility				0.08 [3.39]	0.09 [3.94]	0.07 [2.62]
Constant	-0.13 [-5.52]	-0.06 [-1.81]	-0.13 [-1.07]	-0.17 [-6.21]	-0.22 [-1.92]	-0.16 [-1.33]

Table 8: Mean Black-Scholes Vs. Real-World Early Exercise Probabilities

The table presents the mean optimal Black-Scholes (1973; column (1)) and real-world (column (2)) early exercise probabilities of our sample puts separately by moneyness. We calculate both probabilities over the remaining time-to-maturity of the puts. At the end of each sample month $t - 1$, we first sort the puts into moneyness portfolios according to the decile breakpoints of their Black-Scholes (1973) probabilities. We then calculate means first by portfolio and sample month and then average over our sample period by portfolio. We also report the mean differences between the two probabilities (column (1)–(2)) and their associated t -statistics calculated using Newey and West's (1987) formula with a twelve-month lag length (column (3)).

	Mean Black-Scholes Early Exercise Probability (in %)	Mean Real-World Early Exercise Probability (in %)	Difference (in %)	t -statistic of the Difference
	(1)	(2)	(1)–(2)	(3)
1 (Low)	16.35	6.27	10.08	[4.62]
2	25.93	8.28	17.65	[6.82]
3	32.51	10.77	21.74	[8.52]
4	38.19	13.34	24.85	[10.09]
5	43.36	15.11	28.25	[10.33]
6	48.31	17.30	31.01	[11.78]
7	53.06	20.81	32.25	[11.75]
8	58.04	22.35	35.69	[11.30]
9	64.15	24.28	39.87	[11.45]
10 (High)	76.74	27.05	49.69	[13.50]

Table 9: Mean Difference Between Black-Scholes and Real-World Early Exercise Probabilities By Moneyness, Maturity Time, and Volatility

The table presents the mean difference between the optimal Black-Scholes (1973) and the real-world early exercise probability of our sample puts by moneyness, time-to-maturity, and idiosyncratic stock volatility. We calculate both probabilities over the strategy return period. At the end of each month $t - 1$, we first sort our sample puts into portfolios according to whether their strike-to-stock price ratio (“moneyness”) lies above 1.10 (Panel A), between 1.025 and 1.10 (Panel B), or between 0.975 and 1.025 (Panel C). Within each moneyness portfolio, we next sort into portfolios according to whether their days-to-maturity are below 60, between 60 and 90, or above 90 days. Within each moneyness-maturity sorted portfolio, we finally sort into portfolios according to the quintile breakpoints of the Fama-French-Carhart-model idiosyncratic stock volatility. See the captions of Tables 5 and 6 for details on the sorting variables. We then calculate means first by portfolio and sample month and then average over our sample period by portfolio. Plain numbers are the mean early exercise probability differences (in %), while the numbers in square parentheses are t -statistics calculated using Newey and West’s (1987) formula with a twelve-month lag length.

Days-to-Maturity	Idiosyncratic FFC Stock Volatility				
	1 (Low)	2	3	4	5 (High)
Panel A: Deep In-The-Money (Strike-to-Stock Price > 1.10)					
30-60	54.83 [23.64]	44.17 [19.23]	36.66 [17.07]	31.47 [16.12]	25.28 [16.96]
60-90	48.86 [13.25]	32.17 [10.67]	21.98 [10.15]	14.29 [8.75]	7.54 [8.72]
90-120	38.02 [8.09]	19.93 [5.75]	10.61 [4.92]	5.47 [4.56]	1.98 [4.87]
Panel B: In-The-Money (Strike-to-Stock Price 1.025 to 1.10)					
30-60	35.00 [20.18]	28.08 [29.06]	24.20 [35.83]	21.12 [30.90]	18.48 [37.09]
60-90	22.96 [7.54]	13.93 [8.69]	9.49 [9.34]	6.75 [11.74]	4.22 [16.37]
90-120	14.39 [4.27]	6.32 [4.07]	3.58 [3.84]	1.76 [3.84]	0.89 [6.73]

(continued on next page)

Table 9: Mean Difference Between Black-Scholes and Real-World Early Exercise Probabilities By Moneyness, Maturity Time, and Volatility (Cont.)

Days-to-Maturity	Idiosyncratic FFC Stock Volatility				
	1 (Low)	2	3	4	5 (High)
Panel C: At-The-Money (Strike-to-Stock Price 0.975 to 1.025)					
30-60	13.50 [35.12]	12.58 [47.26]	12.07 [43.46]	12.05 [58.22]	12.07 [50.32]
60-90	4.89 [7.36]	3.86 [10.54]	3.41 [10.93]	3.05 [14.06]	2.29 [14.05]
90-120	1.69 [3.04]	0.86 [2.59]	0.56 [2.61]	0.37 [2.49]	0.24 [6.69]

Table 10: Fama-MacBeth Regressions Explaining the Difference Between Black-Scholes and Real-World Early Exercise Probabilities

The table presents the results of Fama-MacBeth (1973) regressions of the difference between Black-Scholes (1973) and real-world early exercise probabilities for our sample puts on subsets of stock and option characteristics plus a constant. We calculate the two probabilities over the strategy return period. The characteristics include the strike-to-stock price ratio (“moneyness”), time-to-maturity as fraction of a year, and idiosyncratic stock volatility, and they are measured until the start of the strategy return period. We calculate idiosyncratic stock volatility from the Fama-French-Carhart model estimated over the prior 60 months. The plain numbers are the average estimates, while the numbers in square parentheses are t -statistics calculated using Newey and West’s (1987) formula with a twelve-month lag length.

	Regression Model:				
	(1)	(2)	(3)	(4)	(5)
Moneyiness	1.37 [30.97]			1.47 [33.68]	1.61 [35.75]
Time-to-Maturity		-1.05 [-35.32]		-1.15 [-39.52]	-1.18 [-41.43]
Volatility			-0.24 [-16.57]		-0.32 [-20.18]
Constant	-1.29 [-29.41]	0.39 [50.65]	0.24 [32.59]	-1.16 [-27.39]	-1.20 [-29.27]

Table 11: Adjusting the Strategy for Transaction Costs

The table presents the mean transaction-cost-adjusted returns of spread portfolios long stock-and-riskfree-asset portfolios replicating a put and short the puts sorted on the puts' moneyness and time-to-maturity. To adjust for transaction costs, we assume that investors always buy (sell) at the midpoint price plus (minus) φ times the bid-ask spread. We furthermore assume that investors can borrow stocks at Markit's indicative rate. Panel A considers the full sample, while Panel B restricts attention to only those strategies involving puts with a bid-ask spread below the median and stocks with an Amihud (2002) illiquidity value below the median, both measured at the start of the spread return period. At the end of each sample month $t - 1$, we first sort the spread portfolios into portfolios according to whether the strike-to-stock price ratio ("moneyness") of the associated put lies above 1.10 (Panels A.1 and B.1), between 1.025 and 1.10 (Panels A.2 and B.2), or between 0.975 and 1.025 (Panels A.3 and B.3). Within each moneyness portfolio, we next sort them into portfolios according to whether their days-to-maturity are below 60, between 60 and 90, or above 90 days. We equally-weight the portfolios and hold them over month t . Plain numbers are mean monthly portfolio returns (in %) and the numbers in square parentheses are t -statistics calculated using Newey and West's (1987) formula with a lag length of twelve months.

Days-to-Maturity	Borrowing Cost Adjustment/Bid-Ask Spread Fraction φ :			
	No/0.00	Yes/0.10	Yes/0.25	Yes/0.50
	(1)	(2)	(3)	(4)
Panel A: Full Sample				
<i>Panel A1: Deep In-The-Money (Strike-to-Stock Price > 1.10)</i>				
30-60	2.48 [4.92]	-0.08 [-0.16]	-3.30 [-4.42]	-10.05 [-6.26]
60-90	1.76 [3.78]	-0.70 [-1.37]	-3.94 [-5.11]	-10.45 [-6.62]
90-120	1.41 [3.48]	-1.02 [-2.15]	-4.38 [-5.81]	-11.39 [-6.67]
<i>Panel A2: In-The-Money (Strike-to-Stock Price 1.025 to 1.10)</i>				
30-60	4.83 [5.87]	1.68 [1.96]	-2.17 [-2.05]	-10.37 [-5.71]
60-90	3.15 [4.18]	0.19 [0.24]	-3.64 [-3.63]	-11.67 [-6.83]
90-120	2.47 [3.72]	-0.37 [-0.52]	-4.21 [-4.52]	-12.64 [-6.84]

(continued on next page)

Table 11: Adjusting the Strategy for Transaction Costs (Cont.)

Days-to-Maturity	Borrowing Cost Adjustment/Bid-Ask Spread Fraction φ :			
	No/0.00	Yes/0.10	Yes/0.25	Yes/0.50
	(1)	(2)	(3)	(4)
<i>Panel A3: At-The-Money (Strike-to-Stock Price 0.975 to 1.025)</i>				
30-60	9.49 [6.25]	5.56 [3.62]	1.04 [0.60]	-8.93 [-3.79]
60-90	4.72 [4.20]	1.36 [1.14]	-2.93 [-2.10]	-12.29 [-6.15]
90-120	3.32 [3.31]	0.10 [0.10]	-4.24 [-3.40]	-14.16 [-6.72]
Panel B: High-Liquidity Put and Stock Sample				
<i>Panel B1: Deep In-The-Money (Strike-to-Stock Price > 1.10)</i>				
30-60	3.93 [3.75]	2.65 [2.93]	1.71 [1.83]	0.14 [0.14]
60-90	3.27 [3.43]	1.93 [2.57]	1.02 [1.32]	-0.52 [-0.65]
90-120	2.57 [4.41]	1.61 [2.86]	0.69 [1.18]	-0.85 [-1.34]
<i>Panel B2: In-The-Money (Strike-to-Stock Price 1.025 to 1.10)</i>				
30-60	6.87 [4.93]	5.09 [4.24]	3.92 [3.20]	1.95 [1.51]
60-90	5.85 [3.96]	4.22 [3.35]	3.10 [2.44]	1.22 [0.93]
90-120	3.79 [4.01]	2.60 [2.76]	1.52 [1.56]	-0.29 [-0.28]
<i>Panel B3: At-The-Money (Strike-to-Stock Price 0.975 to 1.025)</i>				
30-60	12.15 [4.93]	9.37 [4.39]	7.98 [3.66]	5.66 [2.48]
60-90	8.32 [3.95]	5.79 [3.52]	4.52 [2.72]	2.40 [1.38]
90-120	5.70 [3.73]	4.35 [2.86]	3.14 [2.00]	1.10 [0.66]

Table 12: Adjusting the Strategy for Transaction Costs Under Different Replication Portfolio Rebalancing Schemes

The table presents the mean transaction-cost-adjusted returns of spread portfolios long stock-and-riskfree-asset portfolios replicating a put and short the puts using different rebalancing schemes. To adjust for transaction costs, we assume that investors always buy (sell) at the midpoint price plus (minus) φ times the bid-ask spread. We further assume investors can borrow stocks at Markit's indicative rate. In Panels A, B, and C, we consider the daily, weekly, and no rebalancing cases, respectively. The table restricts attention to only those strategies involving puts with a bid-ask spread below the median and stocks with an Amihud (2002) illiquidity value below the median, both measured at the start of the spread return period. Within each panel, we further separately look into strategies involving only puts with a price above \$1, \$2, and \$5 at the start of the strategy return period. We equally-weight the portfolios and hold them over month t . Plain numbers are mean monthly portfolio returns (in %) and the numbers in square parentheses are t -statistics calculated using Newey and West's (1987) formula with a lag length of twelve months.

Days-to-Maturity	Borrowing Cost Adjustment/Bid-Ask Spread Fraction φ :			
	No/0.00	Yes/0.10	Yes/0.25	Yes/0.50
	(1)	(2)	(3)	(4)
Panel A: Daily Rebalancing				
Put Price \geq \$1	5.00 [4.90]	3.51 [3.86]	2.40 [2.55]	0.53 [0.52]
Put Price \geq \$2	5.12 [4.65]	3.64 [3.79]	2.61 [2.64]	0.87 [0.83]
Put Price \geq \$5	6.54 [4.14]	5.06 [3.72]	4.25 [3.08]	2.89 [2.04]
Panel B: Weekly Rebalancing				
Put Price \geq \$1	5.29 [6.08]	4.14 [5.48]	3.11 [3.98]	1.39 [1.67]
Put Price \geq \$2	5.32 [5.67]	4.21 [5.19]	3.25 [3.89]	1.64 [1.86]
Put Price \geq \$5	6.12 [4.73]	4.97 [4.49]	4.21 [3.73]	2.93 [2.51]

(continued on next page)

Table 12: Adjusting the Strategy for Transaction Costs Under Different Replication Portfolio Rebalancing Schemes (Cont.)

Days-to-Maturity	Borrowing Cost Adjustment/Bid-Ask Spread Fraction φ :			
	No/0.00	Yes/0.10	Yes/0.25	Yes/0.50
	(1)	(2)	(3)	(4)
Panel C: No Rebalancing				
Put Price \geq \$1	6.80 [8.27]	5.61 [9.19]	4.72 [7.47]	3.21 [4.78]
Put Price \geq \$2	6.56 [7.37]	5.33 [8.39]	4.50 [6.88]	3.10 [4.49]
Put Price \geq \$5	6.74 [5.48]	5.35 [6.10]	4.69 [5.25]	3.59 [3.88]