From Hyper Options to Options on Local Time

Peter Carr
Courant Institute,
New York University
251 Mercer Street
New York, NY 10012
(212) 260-3765
pcarr@nyc.rr.com

Roger Lee
Dept. of Mathematics
University of Chicago
5734 S. University Avenue
Chicago, IL 60637
(773) 702-5985
rl@math.uchicago.edu

Initial Version: September 22, 2002

Current Version: November 20, 2005

File reference: HyperOptions.tex

We thank Chris Heyde, Alex Levin, Mike Lipkin, Dilip Madan, Peter Maas, Carol Marquardt, Jan Vecer, Paul Wilmott, Liuren Wu, and participants at presentations at the Courant Institute of NYU and the 3rd Oxford-Princeton conference. The usual disclaimer applies.

From Hyper Options to Options on Local Time

Abstract

It is well known that a European option can only be exercised at maturity, while an American option can be exercised at any time prior to maturity. We introduce a third type of option called a hyper option. A hyper option is similar to an American option in that it can be exercised at any time prior to maturity, but it also differs from an American option in that it can be exercised many times. Exercising a hyper option locks in the exercise value while retaining optionality, albeit with a reversed polarity. Thus after a hyper call is first exercised, it can be exercised next as a put, then as a call, etc. A hyper option can be exercised an unlimited number of times, and all of the exercise proceeds are deferred without interest to maturity.

It is known that an American option written on a forward price always has positive probability of optimal early exercise. For a hyper option written on a forward price, we similarly show that there is always positive probability of multiple optimal early exercises. Despite this, we show that the arbitrage-free value of a hyper option on the forward is just the price of the corresponding European option. Our hedging and valuation results assume frictionless markets and no arbitrage, but are otherwise completely model-free. We also develop the arbitrage-free value and hedge of a hyper option written on the spot price of an asset paying dividends, under alternative assumptions on carrying costs: nonnegative carry, nonpositive carry, or zero carry. We again make no assumptions on the stochastic process generating uncertainty in the underlying asset price.

We also give conditions under which these results can be used to replicate a claim paying the number of upcrosses of the forward price of a given spatial interval. Under alternative assumptions on the symmetry of the stochastic process, we show how one can replicate the payoff to call options on the number of upcrosses. Finally, we show how the payoff to options on local time can be spanned.

From Hyper Options to Options on Local Time

I Introduction

In this paper, we introduce four new types of path-dependent options, whose payoffs are closely related. In general, the options which we introduce will be of interest to investors who believe that the price of an underlying asset will fluctuate around some pre-specified fixed level. By buying and holding one of our options to maturity, the investor's payoff increases each time that the underlying price crosses a fixed level, with the investor's maximum possible loss limited to the initial premium of the option. For the sellers of such options, we delineate robust or semi-robust hedges which just involve holding a semi-static position in a single standard option.

The first path-dependent option which we introduce is called a HYPER option (High Yielding Performance Enhancing Reversible option). Like any other option, a hyper option is issued as either a call or a put. A hyper option is similar to an American option in that it can be exercised at any time prior to maturity, but it also differs from an American option in that it can be exercised many times. Exercising a hyper option locks in the exercise value while retaining optionality, albeit with a reversed polarity. Thus after a hyper call is first exercised, it can be exercised next as a put, then as a call, etc. A hyper option can be exercised an unlimited number of times, and all of the exercise proceeds are deferred without interest to maturity. At maturity, the hyper option can be exercised a final time if it is in-the-money, or it can be left to expire worthless otherwise.

The second type of option which we introduce is called a super option (Structured UPside-Enhancing Re-activation). Like hyper options, super options also allow multiple early exercises, but in contrast to hyper options, they keep the same polarity throughout their life. After early exercise of a super option, the holder must pay the option's instrinsic value to re-open it. If this intrinsic value is negative, the super option can be re-opened at no cost.

The third type of option which we introduce is an option on the number of upcrosses of a given spatial interval. The user specifies a spatial corridor, a strike, a maturity, and a polarity. Over the life of the option, an upcross of the corridor is recorded each time that the underlying asset price transitions from below the lower end of the corridor to above the upper end of the corridor. At expiry, the call on upcrosses pays the difference, if any, between the number of upcrosses and its strike price.

The fourth and final type of path-dependent option that we introduce is an option on local time. Local time is typically only defined for continuous processes whose sample paths display unbounded variation, eg. diffusions. Local time is a stochastic process that increases only when the underlying reference process passes through some fixed level. Carr and Jarrow[1] show that the local time of the underlying asset at level K arises financially as the loss for an investor who prices and delta-hedges a straddle struck at K under the mistaken belief that the underlying asset has zero volatility. Imposing a symmetry condition, we show that a call on local time at K with strike $K_\ell \geq 0$ has twice the value of the out-of-the-money option struck K_ℓ dollars away from K. Setting $K_\ell = 0$ recovers the result in Carr Jarrow that the risk-neutral mean of the local time at K is twice the value of the out-of-the-money option of strike K.

In this paper, we transition through the four new types of path-dependent options in the order presented above. Hence, we first consider hyper options written on the forward price of some underlying asset, and then we subsequently consider hyper options written on the spot price. For a hyper option written on the forward price F, both the hyper option and the forward contract mature at some fixed date T. It is known that a standard American option on a forward price always has positive probability of optimal early exercise. In fact, if the option is exercised at all, it is optimally exercised early (assuming that interest rates are always positive). Let K denote the strike price of the hyper option. If a hyper call is exercised at any time $t \in [0, T]$, the owner locks in the payoff $F_t - K$, which is received at T. Exercising the hyper call converts it into a hyper put with the same underlying, strike, and maturity. We do not require that the hyper call be in-the-money (ITM) for it to be exercised. If the owner exercises his hyper call while $F < K$ to obtain the ITM hyper put, then $F - K$ is negative so the owner owes $K - F$ to the writer at maturity. If a hyper put is exercised at any time $t \in [0, T]$, the owner locks in the payoff $K - F_t$, which is received at T. Exercising the hyper put also reverses it into a hyper call with the same underlying, strike, and maturity.

If a hyper option is ITM at expiration, then we assume that the owner exercises it. We say that an exercise strategy with this property is *optimal* if it is value maximizing. Depending on the price path which is realized, we will show that it can be optimal for the owner of a hyper option to exercise early one or more times. In fact, in the zero carry case, at any time prior to maturity, there is always positive probability of multiple optimal early exercises. Thus, the writer of a hyper option must find a hedging strategy which defends against these multiple optimal early exercises, ideally without exposing the writer to model risk.

We show that there is a model-free strategy for replicating the payoff of a hyper option written on a forward price. This strategy implies the surprising result that this hyper option always has exactly the same value as the corresponding European option. This valuation result holds despite the fact that there is always positive probability of multiple optimal early exercises. The reason is that all possible exercise strategies are also optimal. Note that this differs from the result in Merton[2] for American calls on non-dividend paying stocks, for which the optimal exercise strategy is to wait to maturity and exercise if the call is ITM then.

Since many asset classes do not have transparent forward prices on which to base the payoff of a hyper option, we also consider hyper options written on the spot price of a dividend-paying asset. We again specify that all exercise proceeds be deferred without interest to maturity. Assuming essentially that the cost of carry is positive (negative), we find that a hyper call (put) on the spot has the same value as a European call, but a hyper put (call) has strictly greater value. In contrast to the case of hyper options written on the forward price, there is a unique optimal exercise strategy for hyper options written on the spot price, in the case of positive (negative) carry. This strategy is to exercise hyper puts (calls) immediately and to hold hyper calls (puts) to maturity and exercise if in-the-money then. We develop a hedging strategy for the hyper option writer which replicates perfectly if the owner follows this optimal strategy and super-replicates otherwise.

We also give conditions under which these results can be used to replicate a claim paying the number of upcrosses of the forward price of a given spatial interval. Under alternative assumptions on the symmetry of the stochastic process, we show how one can replicate the payoff to call options on the number of upcrosses. Finally, we show how the payoff to options on local time can be spanned.

The structure of this paper is as follows. We initially focus on hyper options written on a forward price. The next section presents our assumptions, which do not any include any process restrictions on the underlying assets. The following section demonstrates our key result that for hyper options written on a forward price, all exercise strategies are optimal. We consider some variations on hyper options in the fourth section such as hyper options written on the spot price of some underlying asset. In the fifth section, we present an exercise policy under which the payoff to a hyper option is the number of upcrosses of a given spatial interval. In the penultimate section, we also show that under some symmetry conditions, one can replicate the payaff to call options written on the number of upcrosses. Alternatively, the underlying of the call can be the product of the number of upcrosses with the width of the corridor. By shrinking the width of the corridor down to zero, we show how to span the payoff to options on local time. The final section summarizes the paper and presents various avenues for future research.

II Assumptions and Notation

We assume that there is a frictionless market in pure discount bonds paying one dollar at their maturity T. Let B_t be the price of this bond at time $t \in [0, T]$ with $B_T = 1$. For notational ease, we suppress the dependence of this price on the fixed maturity T. We also assume that there is a frictionless market in European options of maturity T. We will take the underlying risky asset to be a stock, but the same analysis works for many other underlyings. Note that we don't need limited liability, but we allow it. We assume no arbitrage, so the European call prices decrease with increasing strike, while the European put prices increase. The forward price F_t is defined as the unique strike price which equates European call and put prices at time t. Again, for notational ease, we suppress the dependence of this price on the fixed maturity T. We assume that one can always observe this forward price as a continuous time semi-martingale and hence write hyper options on it. While this process can have stochastic volatility or jumps, we assume that it is right continuous with left limits.

Let $C_t^e(K)$ and $P_t^e(K)$ denote the respective prices at time $t \in [0,T]$ of European calls and puts of strike $K \in \mathbb{R}$ and maturity $T > 0$. For notational ease, we have again suppressed the dependence of these prices on the maturity T. From no arbitrage, put call parity holds at all times:

$$C_t^e(K) - P_t^e(K) = B_t(F_t - K), \qquad t \in [0,T]. \tag{1}$$

III Hedging Hyper Options on a Forward Price

Let $C_t^h(K)$ and $P_t^h(K)$ denote the arbitrage-free value at time $t \in [0,T]$ of the hyper call and hyper put of strike $K \in \mathbb{R}$ and whose underlying is a forward price. As with every other security, the dependence on the fixed maturity T is understood. Our treatment of hyper options on a forward price assumes nothing about riskfree rates and dividends.

Theorem 1: *No arbitrage implies that for hyper options written on the forward price*:

$$C_t^h(K) = C_t^e(K) \text{ and } P_t^h(K) = P_t^e(K), \qquad t \in [0,T]. \tag{2}$$

A formal proof is given in the Appendix. Here, we content ourselves with giving the intuition behind the result. Let the adapted process $\gamma_t \in \{0,1\}$ be the call indicator at time $t \in [0,T]$. If the hyper option is a call at time t, then $\gamma_t = 1$, and $\gamma_t = 0$ otherwise. The process is to be left continuous with right limits, so if a hyper call is exercised at some t for $F_t - K$, then $\gamma_t = 1$ but $\gamma_{t+} = 0$, as the owner then has a hyper put. As a consequence of these assumptions, γ has finite variation. We leave to future investigation any extension to general predictable processes γ.

Let $\{N_t^{px}; t \in [0,T]\}$ and $\{N_t^{cx}; t \in [0,T]\}$ be continuous-time counting processes which keep track of the number of times at or before t that the hyper option has been exercised as a put and as a call. For $i = 1, 2, \ldots, N_T^{px}$, let $\sigma_i \in [0,T]$ be the stopping time at which the hyper option is exercised for the i-th time as a put. If no exercise as a put ever occurs, then the set $\{\sigma_1, \sigma_2, \ldots, \sigma_{N_T^{px}}\}$ is empty since $N_T^{px} = 0$. Similarly, for $i = 1, 2, \ldots, N_T^{cx}$, let $\tau_i \in [0, T]$ be the stopping time at which the hyper option is exercised for the i-th time as a call. If no exercise as a call ever occurs, then the set $\{\tau_1, \tau_2, \ldots, \tau_{N_T^{cx}}\}$ is empty since $N_T^{cx} = 0$.

Let N_t^c and N_t^p respectively denote the number of standard European calls and puts held at time $t \in [0, T]$. Then we claim the following dynamic trading strategy in European options is a perfect hedge for the sale of a hyper option:

$$N_t^c = \gamma_t \text{ and } N_t^p = 1 - \gamma_t, \qquad t \in [0, T]. \tag{3}$$

Since $\gamma_t \in \{0, 1\}$, the hedger only holds either a European call or a European put, but not both. In this strategy, the polarity of the European option held by the hedger (i.e. call or put) exactly mimics the polarity of the hyper option held by the owner. When the owner exercises, the hyper option changes polarity automatically and hence the hedger must also change the polarity of the European option held. Note that this dynamic option trading strategy is model-free in that it is independent of any parameters governing the forward price process. To implement it, all the hedger has to know at any time is whether or not the owner has exercised his hyper option at that time. In listed markets for American options, there is a lag between exercise by the owner and assignment to a writer. In an OTC market, one would presumably not record exercise proceeds or reverse the polarity until the hedger has executed his rebalancing trade. As the lag between owner exercise and writer notification is ignored in standard models of American option valuation, we also ignore it here. Also note that if one is concerned about transactions costs in options markets (which we have assumed to be zero), then one can use put call parity to change the strategy to one involving a static position in a single option combined with semi-static trading in a forward contract. More specifically, letting N_t^f be the holdings at t in forward contracts with strike K and maturity T, one can alternatively use either:

$$N_t^c = 1, N_t^p = 0, N_t^f = \gamma_t - 1 \tag{4}$$

or:

$$N_t^c = 0, N_t^p = 1, N_t^f = \gamma_t. \tag{5}$$

We now explain why the dynamic option trading strategy specified in (3) hedges the sale of a hyper option. Suppose for now that an investor has sold a hyper put. Since $\gamma_0 = 0$, the trading strategy (3) requires the purchase of the European put with the same underlying, strike, and maturity as the hyper put. If the hyper put is never exercised, then the European put expires worthless. If the hyper put is only exercised at maturity, then the European put is also exercised for the same amount. If the hyper put is exercised before maturity, then at $\sigma_1 \in [0, T)$, we have $\gamma_{\sigma_1+} = 1$. Hence the trading strategy (3) requires that the investor sell the European put and buy the corresponding European call at σ_1. From put call parity (1), this creates the cash flow $(K - F_{\sigma_1})B_{\sigma_1}$ immediately, which can be used to buy $K - F_{\sigma_1}$ pure discount bonds paying a total of $K - F_{\sigma_1}$ at maturity. This amount is exactly what is presently owed by the writer to the owner of the hyper option, since the hyper put exercise value of $K - F_{\sigma_1}$ is deferred to maturity. Thus, the bonds purchased by the hedger exactly cover his realized liability. Note that if the owner exercises his hyper put while it is OTM to obtain an ITM hyper call, then the above arguments still hold. The hedger's purchase of $K - F_{\sigma_1}$ bonds would actually be a sale of $F_{\sigma_1} - K$ bonds, so the only changes are semantic. We will continue to use words consistent with the owner exercising while ITM, but our arguments remain true in either case.

After this first exercise, the owner of the hyper option now has a hyper call. The hedger of this option is holding a European call with the same underlying, strike, and maturity. If after this first exercise, the only subsequent exercise is at maturity, then the hedger's European call provides the payoff. If after the first exercise, there is an early hyper call exercise, then at $\tau_1 \in (\sigma_1, T)$, we have $\gamma_{\tau_1+} = 0$. Hence, from (3), the hedger sells the European call and buys the European put at τ_1. From put call parity (1), this creates the payoff $F_{\tau_1} - K$ at maturity, which is precisely the increment to the hedger's total liability. After the first exercise as a hyper call, the hyper option owner now has a hyper put again. The hedger has a European put again. Thus, both parties are in the exact same position as at $t = 0$. As all liabilities were covered for every contingency, whether or not the round trip was completed, it follows that the hedger can continue to cover any additional liabilities realizing due to exercise of the hyper option. A formal argument proving this result can be found in the appendix. Intuitively, at every hyper put exercise prior to maturity, the ITM European put held in the hedge is reversed back into an out-of-the-money (OTM) European call and the money left over covers the increment in the exercise liability. Similarly, at every hyper call exercise prior to maturity, the ITM European call held in the hedge is converted into an OTM European put and the money left over again covers the increment in the exercise liability. At maturity, the European option held either expires worthless or in-the-money by exactly the same amount as the corresponding hyper option.

Now suppose that an investor initially writes a hyper call, rather than a hyper put. Then by a similar analysis, the mimicking strategy (3) also perfectly hedges the sale. Hence the hedger can meet all liabilities arising from the sale of the hyper call by charging a premium equal to the European call value.

As the cost at time $t \in [0, T]$ of starting a replicating strategy is just the price of the corresponding European option, no arbitrage implies:

$$C_t^h(K) = C_t^e(K) \text{ and } P_t^h(K) = P_t^e(K), \qquad t \in [0, T]. \tag{6}$$

Notice that we have been able to uniquely value a hyper option without specifying the owner's exercise strategy. The payoff for all possible strategies can be perfectly replicated by a dynamic option trading strategy, which when coupled with bonds is self-financing. It follows that all possible exercise strategies are optimal.

This differs from the case of a standard American call written on the spot price of a non-dividend paying stock. In that case, Merton[2] showed that the optimal exercise strategy is to wait to maturity and exercise if the call is ITM then. As a result, a hedge for the writer is to buy and hold the corresponding European call. If the owner of the American call exercises early, then the hedger would have money left over (assuming positive interest rates and randomness). In contrast, the hedger of a hyper call never has money left over for any possible exercise strategy.

If a stock pays sufficiently large dividends or if the American call is written on the forward, then there is an incentive to exercise early as is well known. If one exercises, one gets the cash flows thrown off by the reward function that would otherwise be lost by waiting. For an American call on stock, these cash flows are dividends on the stock less interest on the strike. For an American option on the forward, these cash flows are the interest on the immediate exercise proceeds. On the other hand, exercising an American option destroys its volatility value. Thus, the exercise decision for an American option involves a complicated comparison between the volatility value and the value of the cash flows generated by the exercise proceeds. In contrast, a hyper option exercise does not involve either a loss of volatility value or a relative gain in interest on the exercise proceeds. The volatility value is restored by the gift of another hyper option of the opposite parity. There is no gain of interest on the exercise proceeds because the immediate reward is also rising at the riskfree rate. Thus, the holder is indifferent to exercising or not.

To understand the source of the invariance of hyper option value to exercise strategy from another perspective, consider any dynamic trading strategy in European options which starts with a position in a single option and is not self-financing. If a claim is defined whose payoffs are the cash flows thrown off by the strategy, then the value of that claim is just the price of the initial option held. In general, this set of payoffs would be expressed in terms of the path of option prices, rather than in terms of the path of forward prices of the underlying asset. However, put call parity (1) implies that a conversion from a European call to a European put or a reversal back results in a cash flow depending only on the forward price. As a result, any trading strategy that just involves these conversions and reversals generates payoffs which just depend on the path of the underlying forward price. It is this implication of put call parity which makes the hyper option payoffs both potentially marketable and certainly hedgeable.

Note that the valuation results obtained are completely model-free. Regardless of whether the market is complete or not, the payoffs from hyper options lie in the semi-static span of European option trading and thus can be priced uniquely.

IV Variations on Hyper Options

This section examines several variations on the hyper option structure which still permit exact model-free valuation.

IV-A Hyper Options on Spot

In this section, we consider hyper options written on the spot price S of a dividend-paying asset. Without loss of generality we refer to the asset as a stock. We again allow exercise at any time and defer the receipt of all exercise proceeds to maturity without interest.

To distinguish the values of hyper options on spot prices from the values of hyper options on forward prices, we let C_t^H and P_t^H respectively denote the value at t of a hyper call and hyper put on the spot. Since European option prices on spot have the same value as European options on forward with the same maturity, we continue to use C_t^e and P_t^e to denote European option values.

Let $Z_t(T)$ be the price at t of a claim that pays the stock price at T, but receives no prior dividends. We will refer to this claim as the stripped stock, since the dividends to T have been stripped away. Equivalently, this claim can be viewed as a forward contract on S with delivery price zero, or as a zero-strike call on S. Its value is

$$Z_t(T) = B_t(T)F_t(T) \tag{7}$$

where $F_t(T)$ is the forward price at time t for T-delivery of the stock. To go long one stripped share, a trader would go long one forward contract (at zero cost), and buy F_t unit bonds.

We present results for three alternative assumptions:

- Nonnegative carry: With probability one, F_t/S_t is decreasing in $t \in [0, T]$.
- Nonpositive carry: With probability one, F_t/S_t is increasing in $t \in [0, T]$.
- Zero carry: With probability one, $F_t = S_t$ for all $t \in [0, T]$.

The zero carry case is the intersection of the other two, but we highlight it because it is equivalent to the case of a hyper option on a forward price, regardless of the dividend policy of the asset underlying the forward contract. Therefore these "spot" results actually generalize the forward results of section III.

In the case of a deterministic interest rate $r(t)$ and dividend yield $q(t)$, the nonnegative carry assumption is equivalent to $r(t) \geq q(t)$ for all t; and the nonpositive carry assumption is equivalent to $r(t) \leq q(t)$ for all t.

Consider first the case of nonnegative carry:

Theorem 2a: *For hyper options written on the spot price, no arbitrage implies:*

$$\begin{aligned} C_t^H &= C_t^e \\ P_t^H &= P_t^e + Z_t g(t), \qquad t \in [0,T], \end{aligned} \tag{8}$$

for $r(t) \geq q(t)$, where $g(t) \equiv 1 - \frac{S_t(T)}{F_t(T)} \in [0,1]$ is decreasing in t with $g(T) = 0$.

Thus, Theorem 2a indicates that for nonnegative carry, hyper calls must have the same value as European calls, but hyper puts now must have the same or greater value than European puts.

Proof: Consider initially the situation of the hyper call writer. The initial hedge is to hold a European call. If the hyper call is never exercised early, then the European call covers any exercise at maturity. If the hyper call is exercised early at some time τ, then the hedger converts his European call into a European put. From put call parity,

$$\begin{aligned} C_\tau^e(K,T) - P_\tau^e(K,T) &= B_\tau(T)(F_\tau(T) - K) \\ &= B_\tau(T)(F_\tau(T) - S_\tau) + B_\tau(T)(S_\tau - K) \\ &= Z_\tau(T)g(\tau) + (S_\tau - K)B_\tau(T). \end{aligned} \tag{9}$$

Hence, the hedger can purchase $g(\tau)$ stripped shares and $S_\tau - K$ bonds with the money left over. The bonds cover the hedger's liability due to the hyper call exercise and the owner now has a hyper put. If the hyper put is held to maturity, then the hedger's European put covers any liability arising from terminal exercise. If the hyper put is not held to maturity, let $\sigma \in [\tau, T)$ be the time at which the owner exercises his hyper put. The hedger responds by selling his European put and buying a European call. From negating (9) and replacing τ with σ:

$$P_\sigma^e(K,T) - C_\sigma^e(K,T) = -Z_\sigma(T)g(\sigma) + (K - S_\sigma)B_\sigma(T). \tag{10}$$

Hence, the hedger must sell $g(\sigma)$ stripped shares, but can buy $S_\tau - K$ bonds with the money generated. The number of stripped shares in inventory is thus $g(\tau) - g(\sigma)$. Since $\sigma \geq \tau$ and g is decreasing, the hedger has a nonnegative number of stripped shares in inventory. The bonds cover the hedger's liability due to the hyper put exercise and the owner now has a hyper call. At this point, the hedger and the owner are in the same position as at initiation, except that the hedger is possibly richer. Thus the hedger can super-replicate the payoff from a hyper call by initially holding a European call.

If a hyper put is written initially instead, then the hedger can only super-replicate the payoff of the hyper put by holding a European put plus $g(0)$ stripped shares. The stripped shares are needed to protect against an immediate hyper put exercise followed by the owner holding his hyper call to maturity. Hence for any exercise strategy, we have:

$$\begin{aligned} C_t^H &\leq C_t^e \\ P_t^H &\leq P_t^e + Z_t g(t), \end{aligned} \tag{11}$$

since time 0 is arbitrary.

Suppose that the owner pursues the following exercise strategy going forward from any time t:

- Exercise the hyper put immediately, regardless of moneyness.
- Exercise the hyper call only at maturity, if in-the-money then.

This strategy removes all profit for the hedger and hence attains the upper bounds in (11). As a result, the arbitrage-free values of a hyper call and hyper put are given in (8). **Q.E.D.**

If instead we take the case of *nonpositive* carry, then similar arguments apply. Note that now $g(t) \leq 0$ and g is increasing. The conclusion is that:

Theorem 2b: *For hyper options written on the spot price, no arbitrage implies:*

$$\begin{aligned} C_t^H &= C_t^e + Z_t|g(t)| \\ P_t^H &= P_t^e, \qquad t \in [0, T], \end{aligned} \tag{12}$$

for $r(t) \leq q(t)$.

An optimal exercise strategy for the owner is to:

- Exercise the hyper call immediately, regardless of moneyness.
- Exercise the hyper put only at maturity, if in-the-money then.

If we take the case of *zero* carry, then for both kinds of hyper options, all exercise strategies are optimal. As a result, both kinds of hyper options on the spot price have the same value as the corresponding European option. Consolidating the three conclusions into a single statement, we have:

Theorem 2: *For nonnegative carry or nonpositive carry or zero carry, no arbitrage implies that for hyper options on the spot:*

$$C_t^H = C_t^e + B_t(T)(S_t - F_t(T))^+ \quad (13)$$

$$P_t^H = P_t^e + B_t(T)(F_t(T) - S_t)^+, \qquad t \in [0, T]. \quad (14)$$

Finally, we take special notice of the "strict" subcases of nonnegative and nonpositive carry, namely:

- Positive carry: F_t/S_t is strictly decreasing in $t \in [0, T]$.
- Negative carry: F_t/S_t is strictly increasing in $t \in [0, T]$.

The arguments given above show that in the positive carry case, hyper puts have *strictly* greater value than American puts, and the optimal exercise strategy is unique. In the negative carry case, hyper calls have *strictly* greater value than American calls, and the optimal exercise strategy is unique.

IV-B Super Options

Like hyper options, SUPER (Structured UPside-Enhancing Reactivation) options allow multiple exercises, with proceeds deferred to maturity. Unlike hyper options, no changes of polarity occur; instead, the holder of a super option retains the right to pay intrinsic value at any later date and receive a super option of the same polarity as the one exercised. For a super call on a spot price S, after the holder exercises at τ for a payoff of $S_\tau - K$ to be received at T, the holder can re-activate at any date $\sigma \in (\tau, T]$ to obtain a live super call, by paying $(S_\sigma - K)^+$ to be delivered at T.

It is clear that the value of a super option is bounded below by the value of a European option of the same polarity, because the holder of a super option can simply choose not to exercise until maturity. On the other hand, we claim that the super option value is bounded above by the same-polarity hyper option value. This holds because a hedger short a super option and long a same-polarity hyper option can super-replicate by exercising the hyper option whenever the super option is either exercised or re-activated. On each super option exercise, the payout is matched by the hyper option; and on each super option re-activation, the (possibly negative) payout to the super option holder is either equal to or less than the payout to the hyper option holder.

As a corollary, we observe that in the case of nonnegative carry (including zero carry), a super call has the same value as a European call, because the upper and lower bounds coincide. Likewise, in the case of nonpositive carry (including zero carry), a super put has the same value as a European put.

The hyper-option upper bound on super options, and hence the exact-price corollary, still hold if the super option payoff is defined instead to be the positive part $(S_\tau - K)^+$. If the super option holder exercises when $S_\tau < K$, then the hedger incurs no liability and should just hold the hyper option; if the super option holder later re-activates, the hedger should simply pocket the reactivation payment, if any.

The hyper-option upper bound is also robust to changing the definition of the reactivation payment to be $(\max(S_\sigma, F_\sigma) - K)^+$, because this can only increase the holder's liabilities.

V Restricted Hyper Options and Upcrosses

Suppose that we have two constant barriers set at K and at $H > K$. We say that an upcross of (K, H) is completed when the forward price F hits or crosses the barrier H from below after having been at or below K previously. A partial upcross of (K, H) is completed at maturity if after the last completed upcross prior to maturity, if any, the forward price touches or crosses K, and then finishes between K and H.

Restricted hyper options are hyper options struck at K that are initially issued out of the money. The restricted hyper option must be exercised as a call whenever its call moneyness $F - K$ equals or exceeds $H - K$, i.e. upon the completion of each upcross of (K, H). A restricted hyper option must also be exercised as a put whenever its put moneyness $K - F$ is nonnegative. If $F_0 > K$, this occurs at the first passage time to K and in general, this occurs upon the first subsequent return to K after each upcross. We will show that restricted hyper options are related to Doob's inequalities for the number of upcrossings of the interval (K, H).

Let $\tau_0 = 0$ and for $i = 1, 2, \ldots$, let us recursively re-define σ_i and τ_i by:

$$\sigma_i \equiv \inf\{t \in [\tau_{i-1}, T] : F_t \leq K\} \qquad \tau_i \equiv \inf\{t \in [\sigma_i, T] : F_t \geq H\}, \tag{15}$$

where we adopt the usual convention that the infimum of the empty set is infinity. Note that for each i for which σ_i is finite, we have $\sigma_i < \tau_i$.

Suppose first that $F_0 \leq K$ so that $\sigma_1 = 0$. Consider a restricted hyper call. The first exercise before maturity, if any, occurs the first time that the forward price reaches or crosses H. If there is such a completed upcross, then just after the first one, exercise results in the positive payoff $F_{\tau_1} - K$ received at T, and the hyper call becomes a hyper put. Since $F_{\tau_1} \geq H > K$, the payoff on this completed upcross of (K, H) is at least $H - K$ received at maturity. Hence, if we rescale our initial holdings to $\frac{1}{H-K}$ restricted hyper calls, then this payoff is at least one at maturity.

Now suppose that $F_0 > K$ so that $\sigma_1 > 0$. Consider an initial holding of $\frac{1}{H-K}$ hyper puts. These $\frac{1}{H-K}$ hyper puts are held until the first time that K is touched or crossed from above, if any. If this event never occurs, then the restricted hyper puts expire worthless. If this event does occur, then just after the first time that it does, exercise results in the nonnegative payoff $\frac{K-F_{\sigma_1}}{H-K}$ received at T, and the restricted hyper puts become restricted hyper calls. The $\frac{1}{H-K}$ restricted hyper calls are only exercised early once the forward price reaches or crosses $H > K$ from below, after having been below K previously. The payoff on each such completed upcross of K is at least one dollar received at maturity. Hence, if we hold $\frac{1}{H-K}$ OTM restricted hyper options initially, then the payoff is an upper bound on the number of upcrosses of the interval (K, H) over $[0, T]$.

Note that at maturity, the $\frac{1}{H-K}$ restricted hyper options may result in a payoff of $\frac{F_T-K}{H-K}$ if an upcross of K completes then. This payoff only arises if the number of put exercises just before maturity exceeds the number of call exercises just before maturity and if $F_T \in (K, H)$. In this case, the payoff $\frac{F_T-K}{H-K} \in (0, 1)$ and is considered to be a fractional upcross of the interval (K, H).

Let $n_t^u \equiv N_t^{cx}$ be the number of completed upcrosses of (K, H) in $[0, t]$. Likewise, let $n_t^d \equiv N_t^{px}$, but we warn the reader that n^d only represents the number of completed downcrosses if $F_0 \leq K$. In general, n^d represents the number of times that the stage has been set for a subsequent upcross. Hence, at times $t \in [0, T]$ when $n_t^d > n_t^u$, then an upcross can occur when F crosses H from below, but at times t when $n_t^d = n_t^u$, then F must return to K before an upcross can occur.

Let $U_t(K, H)$ be the arbitrage-free value at time $t \in [0, T]$ of a claim paying the total number of upcrosses of the interval (K, H) over $[0, T]$, including a possible fractional upcross at maturity. We will refer to this claim as an upcrosser. The terminal payoff to an upcrosser is:

$$U_T(K, H) = n_{T-}^u + 1(n_{T-}^d > n_{T-}^u, F_T < H)\frac{(F_T - K)^+}{H - K}. \tag{16}$$

The last term gives credit for a partially completed upcross, if any.

There exists a model-free upper bound on the initial value of an upcrosser, which is implied by the absence of arbitrage. Let $\theta_t^e(K)$ denote the time value of a standard European option with strike K and maturity T at time $t \in [0, T]$. Hence, if $F_t \leq K$, then $\theta_t^e(K) = C_t^e(K)$, while if $F_t > K$, then

$\theta_t^e(K) = P_t^e(K)$. By the foregoing argument, we have:

$$U_t(K,H) \leq \frac{\theta_t^e(K)}{H-K}, \qquad t \in [0,T]. \tag{17}$$

This upper bound is achieved on every path if the process is skip-free in the sense that it can't cross K or H without first touching there. For such a process, the payoff to an initially OTM restricted hyper option is just $U_T(K,H)$, i.e. the the number of completed upcrosses in $[0,T]$, (including any fractional one completing at T). As a result, we have:

$$U_t(K,H) = \frac{\theta_t^e(K)}{H-K}, \qquad t \in [0,T]. \tag{18}$$

In fact, it is sufficient to just proscribe up jumps over H when $n^d > n^u$ and down jumps over K when $n^d = n^u$.

Note that as $H \downarrow K$, the model-free upper bound on the initial value of an upcrosser rises as the reciprocal of $H-K$ towards infinity. If the price process is continuous (but not necessarily Markov), then it is skipfree and hence from (18), the arbitrage-free value explodes as $H \downarrow K$. This makes sense because it is already known that the number of upcrosses of an infinitessimally small interval by a continuous process would realize at either zero or infinity.

VI Options on Upcrosses

Let $C_t^{nu}(k,T)$ denote the arbitrage-free value at time $t \in [0,T]$ of a call on the number of upcrosses of the interval (K,H). The terminal payoff is:

$$C_T^{nu}(k,T) = \left[n_{T-}^u + 1(n_{T-}^d > n_{T-}^u, F_T < H) \frac{(F_T - K)^+}{H-K} - k \right]^+, \tag{19}$$

where for simplicity, we assume that the call strike is a positive integer k. In this section, we work with a skipfree forward price process F. We also impose alternative forms of symmetry on this forward price process which allow exact replication of the payoff of a call on the number of upcrosses.

VI-A Calls on Upcrosses under APCS

We say that Arithmetic Put Call Symmetry (APCS) holds at a given forward price G for options of maturity T if:

$$C_t^e(K_c) = P_t^e(K_p), \qquad t \in [0,T], \tag{20}$$

where the simple arithmetic mean of the two strikes is G, i.e. $\frac{K_c+K_p}{2} = G$. APCS implies in particular that equally out-of-the-money options have the same value.

It is reasonable to question whether there exists a stochastic process for the underlying forward price which generates arbitrage-free option prices which always satisfy APCS. Fortunately, it is very easy to construct a wide class of such processes. Consider any stochastic process arising from Brownian motion by performing an *independent* stochastic time change. The resulting process is a symmetric martingale which jumps whenever the stochastic clock jumps. Options written on such processes satisfy APCS at all price levels.

Now suppose that we further require that the process generating APCS not jump over one or more price levels. The easiest way to ensure this skipfree condition is to require that the independent stochastic clock be a continuous increasing process. In this case, the resulting symmetric mattingale is also continuous. If the time change is absolutely continuous, then these continuous symmetric martingales are called Ocone martingales. Using the language of stochastic differential equations, Ocone martingales solve:

$$dF_t = a_t dW_t, \qquad t \in [0,T],$$

where the absolute volatility process a evolves independently of W.

With existence assured, we now define **Assumption set A1:** *The price process F is a $\mathbb{Q}$ martingale which never jumps over K or H and for which APCS holds at all times σ_i and $\tau_i, i = 1, 2, \ldots$.*

Theorem 3: Semi-Robust Pricing of Call on Upcrosses Under APCS

Under frictionless markets and **A1***, no arbitrage implies that for any positive integer* k*,* $t \in [0, T]$ *:*

$$C_t^{nu}(k) = (n_t^u - k)^+ B_t + \frac{1(n_t^d > n_t^u)}{H-K} C_t^e(K + 2(H-K)(k - n_t^u)^+) + \frac{1(n_t^d = n_t^u)}{H-K} P_t^e(K - 2(H-K)(k - n_t^u)^+). \quad (21)$$

Proof:

Suppose that an investor sells a call on upcrosses at $t = 0$. We will show that the nature of the hedging strategy depends on whether or not the call on upcrosses is in-the-money. At initiation, the call on upcrosses is out-of-the-money, and so long as this remains true, the hedger flips back and forth between OTM calls and puts of different strikes. If the call on upcrosses goes into the money before maturity, then the hedger flips back and forth between calls and puts both struck at K.

At $t = 0$, $n_t^u = 0$ and suppose that the investor purchases $\frac{1}{H-K}$ units of the OTM option struck $2(H-K)k$ dollars away from K, i.e.:

$$\begin{aligned} &\frac{1}{H-K} C_0^e(K + 2(H-K)k) \quad \text{if} \quad F_0 \leq K \\ &\frac{1}{H-K} P_0^e(K - 2(H-K)k) \quad \text{if} \quad F_0 > K. \end{aligned} \quad (22)$$

At each time that n^u or n^d increases while $n^u \leq k$, the hedger switches the polarity and strike of the option held. The new strike is chosen so that the trade is self-financing. Hence, at each time τ that n^u increases to catch up to n^d, we have $n_{\tau-}^u = n_\tau^u - 1$. The hedger sells his holding in $\frac{1}{H-K}$ calls struck at $K + 2(H-K)(k - n_\tau^u + 1)$ and buys a position in $\frac{1}{H-K}$ puts struck at $K - 2(H-K)(k - n_\tau^u)$. Notice that the average of the two strikes is H. At each time τ that n^u increases to catch up to n^d, we have $F_\tau = H$, so APCS (20) implies that this trade is self-financing.

While $n^u \leq k$, consider each time σ that n^d increases above n^u. We have $n_{\sigma-}^u = n_\sigma^u$, so the hedger sells his holding in $\frac{1}{H-K}$ puts struck at $K - 2(H-K)(k - n_\sigma^u)$ dollars and buys a position in $\frac{1}{H-K}$ calls struck at $K + 2(H-K)(k - n_\sigma^u)$ dollars. Notice that the average of the two strikes is K. At each time σ that n^d increases above n^u, we have $F_\sigma = K$, so APCS (20) implies that this trade is also self-financing.

Define a round trip as an upcross followed by a subsequent first return to K. At the end of each such round trip, the forward price is at K and the strikes being traded are each $2(H-K)$ dollars closer to K than they were at the beginning of the round trip. When a round trip ends with the number of upcrosses equal to k, the hedger is holding $\frac{1}{H-K}$ puts struck at K as the forward price returns to K. This ends the first regime. The second regime begins with the sale of these puts which from put call parity generates exactly enough cash to buy $\frac{1}{H-K}$ calls struck at K. After this point in time, the hedger can create the number of upcrosses beyond k, by exploiting put call parity as indicated in the last section. **Q.E.D.**

VI-B Calls on Local Time under APCS

In this subsection, we suppose that the underlying of the call is the product of the number of upcrosses with the width of the corridor, $H-K$. Let $C_t^\pi(K_\pi)$ be the arbitrage-free price at time $t \in [0,T]$ of a call written on this product, where $K_\pi \equiv (H-K)k$. At expiry, this call pays:

$$C_T^\pi(K_\pi) = [n_{T-}^u(H-K) + 1(n_{T-}^d > n_{T-}^u, F_T < H)(F_T-K)^+ - K_\pi]^+. \tag{23}$$

Since the payoff is simply $H-K$ times larger than the payoff (19) described in the last subsection, Theorem 3 implies that:

$$\begin{aligned} C_t^\pi(K_\pi) &= [(H-K)n_t^u - K_\pi]^+ B_t \quad +1(n_t^d > n_t^u) C_t^e(K + 2[K_\pi - (H-K)n_t^u]^+) \\ &\qquad\qquad\qquad\qquad\qquad\quad +1(n_t^d = n_t^u) P_t^e(K - 2[K_\pi - (H-K)n_t^u]^+), \qquad t \in [0,T]. \end{aligned} \tag{24}$$

We now define **Assumption set A2:** *The price process F is an Ocone martingale under $\mathbb{Q}$.* By letting $H \downarrow K$, we shrink the width of the corridor down to zero. Under **A2**, the random variable underlying the call payoff in (23) converges to half of the *local time* of F at K as of time T. Let $L_t^F(K)$ denote the local time of F at K, considered as an increasing stochastic process in t. This process is formally defined by the Tanaka Meyer formula, i.e.:

$$L_t^F(K) \equiv |F_t - K| - |F_0 - K| - \int_0^t \operatorname{sgn}(F_s - K) dF_s, \qquad t \in [0,T], \tag{25}$$

where:

$$\operatorname{sgn}(x) \equiv \begin{cases} 1 \text{ if } x > 0, \\ -1 \text{ if } x \leq 0. \end{cases}$$

Let $C_t^L(K_\ell)$ be the arbitrage-free price at time $t \in [0, T]$ of a call written on the local time realized over $[0, T]$. At expiry, this call pays:

$$C_T^L(K_\ell) = [L_t^F(K) - K_\ell]^+.$$

Letting $H \downarrow K$ in (24), doubling both sides of (24), and letting $K_\ell \equiv 2K_\pi$, we see that we can always value this call:

Theorem 4: Semi-Robust Pricing of Call on Local Time

Under frictionless markets and **A2***, no arbitrage implies that:*

$$\begin{aligned} C_t^L(K_\ell) \;&= [L_t^F(K) - K_\ell]^+ B_t \;\; +1(F_t \leq K) 2C_t^e(K + [K_\ell - L_t^F(K)]^+) \\ &\qquad\qquad\qquad\qquad\quad +1(F_t > K) 2P_t^e(K - [K_\ell - L_t^F(K)]^+), \qquad t \in [0, T]. \end{aligned} \tag{26}$$

Suppose that an investor sells a call on local time at time 0 and that the investor wishes to replicate the payoff via semi-static option trading. Since the local time is initially zero, (26) indicates that the replicating strategy starts by holding 2 OTM options struck K_ℓ dollars away from K. As the hedger moves through calendar time with $L_t^F(K) \leq K_\ell$, the hedger always holds two units of whichever standard option is out-of-the-money. The strike of the option held is $K_\ell - L_t^F(K) \geq 0$ dollars away from K. The reason that the strikes move in towards K is that all option trades are made just as F moves away from K in such a way that the two options held become ITM. Each such trade involves a disposition of the two ITM options and an acquisition of two OTM options with the same price. As a result, the strike of the two acquired options must necessarily be closer to K. This strategy is self-financing by design. If the running local time $L_t^F(K)$ never reaches K_ℓ before T, the hedger just ends up holding 2 worthless OTM options. On the other hand, if $L^F(K)$ reaches K_ℓ before T, then at the first passage time of $L^F(K)$ to K_ℓ, the strike of the 2 options held first reaches K. Between this first passage time and expiry, the hedger holds 2 puts with strike K if $F_t > K$, and 2 calls with strike K otherwise. This strategy generates a positive cash flow each time that F crosses K, whose magnitude at T accumulates to $L_T^F(K) - K_\ell$.

Evaluating (26) at $t = 0$ implies:

$$C_0^L(K_\ell) = 2C_0^e(K + K_\ell). \tag{27}$$

Hence at initiation, a call on the local time at K with strike K_ℓ has twice the value of a standard call struck K_ℓ dollars above K.

VI-C Calls on Upcrosses under GPCS

An unfortunate implication of APCS is that there is positive probability that prices can go negative whenever there is positive probability that prices can more than double. To remedy this, we will consider a family of underlying forward price processes which takes values on the whole positive half line.

It will be convenient to set $H = Ku$ with $u > 1$. Hence, $C_t^{nu}(k, T)$ denotes the arbitrage-free value at time $t \in [0, T]$ of a call on the number of upcrosses of the interval (K, Ku). For simplicity, we again assume that the call strike is a positive integer k.

We say that Geometric Put Call Symmetry (GPCS) holds at a given forward price G for options of maturity T if:

$$C_t^e(K_c) = \frac{K_c}{G} P_t^e(K_p), \qquad t \in [0, T], \tag{28}$$

where the geometric mean of the two strikes is G, i.e. $\sqrt{K_c K_p} = G$. To understand this result, suppose that we are comparing values of a call and a put which are equally out-of-the-money on a geometric basis. Then the call which is x percent out-of-the-money has x percent more value than the equally out-of-the-money put. Intuitively, the source of the greater call value is the higher normal volatility experienced when the forward price is at K_c compared to when it is at K_p.

Assumption set A3: *The price process F is a $\mathbb{Q}$ martingale which never jumps over K or Ku and for which GPCS holds at all times σ_i and $\tau_i, i = 1, 2, \ldots$.*

Theorem 5: Semi-Robust Pricing of Call on Upcrosses Under GPCS

Under frictionless markets and **A3**, *no arbitrage implies that for any positive integer k, $t \in [0,T]$:*

$$\begin{aligned} C_t^{nu}(k) \quad &= (n_t^u - k)^+ B_t \quad + \frac{1(n_t^d > n_t^u)}{K(u-1)} u^{-(k-n_t^u)^+} C_t^e(K u^{2(k-n_t^u)^+}) \\ &\qquad + \frac{1(n_t^d = n_t^u)}{K(u-1)} u^{(k-n_t^u)^+} P_t^e(K u^{-2(k-n_t^u)^+}), \end{aligned} \tag{29}$$

where n_t^u is the number of upcrosses of (K, Ku) completed over $[0,t]$.

Proof: Suppose again that an investor sells a call on upcrosses at $t = 0$. At $t = 0$, $n_t^u = 0$ and suppose that the investor initiates a hedging strategy by buying $\frac{1}{K(u-1)}$ units of the OTM option whose strike K^o is such that $|\ln(K^o/K)|$ is $2k \ln u$:

$$\begin{aligned} &\frac{1}{K(u-1)} C_0^e(K u^{2k}) \quad \text{if} \quad F_0 \leq K \\ &\frac{1}{K(u-1)} P_0^e(K u^{-2k}) \quad \text{if} \quad F_0 > K. \end{aligned} \tag{30}$$

As under APCS, the nature of the trading strategy depends on whether the call on upcrosses is in-the-money. In the first regime, this call is OTM and the hedger flips back and forth between OTM calls and puts of different strikes. At each time τ that a new upcross completes, we have $F_\tau = Ku$ and $n_{\tau-}^u = n_\tau^u - 1$. The hedger sells his holding in $\frac{1}{K(u-1)} u^{-(k-n_\tau^u+1)}$ calls struck at $K u^{2(k-n_\tau^u+1)}$ and buys a position in $\frac{1}{K(u-1)} u^{k-n_\tau^u}$ puts struck at $K u^{-2(k-n_\tau^u)}$. Notice that the geometric mean of the two strikes is Ku. Notice that the ratio of the call strike to this mean is:

$$\frac{K u^{2(k-n_\tau^u+1)}}{Ku} = u^{2(k-n_\tau^u)+1}.$$

Finally, notice that the ratio of the number of puts bought to the number of calls sold is also:

$$\frac{\frac{1}{K(u-1)} u^{k-n_\tau^u}}{\frac{1}{K(u-1)} u^{-(k-n_\tau^u+1)}} = u^{2(k-n_\tau^u)+1}.$$

As a result, GPCS (28) implies that this trade is self-financing.

While the call on upcrosses is still OTM, at each time σ that F returns to K so that n^d exceeds n^u, we have $F_\sigma = K$, and $n_{\sigma-}^u = n_\sigma^u$. The hedger sells his holding in $\frac{1}{K(u-1)} u^{(k-n_\sigma^u)}$ puts struck at $K u^{-2(k-n_\sigma^u)}$ dollars and buys a position in $\frac{1}{K(u-1)} u^{-(k-n_\sigma^u)}$ calls struck at $K u^{2(k-n_\sigma^u)}$ dollars. Notice that the geometric mean of the two strikes is K. Notice that the ratio of the call strike to this mean is:

$$\frac{K u^{2(k-n_\sigma^u)}}{K} = u^{2(k-n_t^u)}.$$

Finally, notice that the ratio of the number of puts sold to the number of calls bought is also:

$$\frac{\frac{1}{K(u-1)}u^{k-n_\sigma^u}}{\frac{1}{K(u-1)}u^{-(k-n_\sigma^u)}} = u^{2(k-n_t^u)}.$$

As a result, GPCS (28) implies that this trade is self-financing.

At the end of each round trip, the forward price is at K and the strikes being traded are each a factor of u^2 closer to K than they were at the beginning of the round trip. When a round trip ends with the number of upcrosses equal to k, the hedger is holding $\frac{1}{K(u-1)}$ puts struck at K as the forward price returns to K. This ends the first regime. The second regime begins with the sale of these puts which from put call parity generates exactly enough cash to buy $\frac{1}{K(u-1)}$ calls struck at K. After this point in time, the hedger can create the number of upcrosses beyond k, by exploiting put call parity as indicated in the last section. **Q.E.D.**

VII Summary and Future Research

We introduced a new option type called hyper options, which can be exercised many times. We showed that prior to maturity, hyper options written on the forward price always have positive probability of multiple optimal early exercises, and yet their arbitrage-free value is given by the value of the corresponding European option. The resolution of this apparent paradox is that for such hyper options, all possible exercise strategies are also optimal. In contrast, for hyper options written on the spot price, there can be a unique optimal exercise strategy. If the underlying has positive cost of carry, then hyper put values exceed European put values, while hyper calls have the same value as European calls. The situation is reversed for negative carrying costs.

We considered the consequences of restricting the exercise strategy of a hyper option to one requiring exercise upon completing an upcross of (K, H) and when the forward price subsequently returns to K. We showed that the ratio of the hyper option value to the corridor width $H - K$ is an upper bound on the value of a claim paying the number of upcrosses of (K, H) over $[0, T]$. Furthermore, if the process is skipfree, then the payoff to this claim can be exactly replicated and the upper bound is attained. When the process is both skipfree and displays symmetry, the payoff to calls on the number of upcrosses can also be exactly replicated. Alternatively, the underlying of the call can be the product of the number of upcrosses with the width of the corridor. By shrinking the width of the corridor down to zero, we showed how to span the payoff to options on local time.

Future research can focus on alternative conditions under which calls on upcrosses can be replicated. For example, it should be possible to allow the underlying forward price to be the sum of a constant process and a geometric Brownian martingale. Replication of a call on upcrosses would still be possible even if this sum is subjected to an independent time change. The two cases discussed in this paper each arise as a special case.

Appendix: Proof of Theorem 1

This appendix provides a formal proof of Theorem 1. We first argue that a hyper put on a forward has the same value as a European put. Note that if X is a semimartingale and N is a left-continuous right-limits process with finite variation, then $N_t^+ := N_{t+}$ is a right-continuous left-limits semimartingale. We have

$$d(X_t N_t^+) = N_{t-}^+ dX_t + X_{t-} dN_t^+ + d[X, N^+]_t = N_t dX_t + X_{t-} dN_t^+ + \Delta X_t dN_t^+ = N_t dX_t + X_t dN_t^+, \quad (31)$$

where each dN_t^+ can be understood in the Lebesgue-Stieltjes sense.

Consider a position (N^p, N^c, N^b) in options and bonds, with N^p and N^c defined in (3), and

$$N_t^b := \sum_{i=1}^{N_{t-}^{px}} (K - F_{\sigma_i}) + \sum_{i=1}^{N_{t-}^{cx}} (F_{\tau_i} - K). \quad (32)$$

We will show that this portfolio self-finances and replicates the hyper put. The portfolio value in dollars is $V_t := N_t^p P_t^e + N_t^c C_t^e + N_t^b$. Taking B_t as numeraire, the value in bonds is

$$\frac{V_t}{B_t} = N_t^p \frac{P_t^e}{B_t} + N_t^c \frac{C_t^e}{B_t} + N_t^b \qquad t \in [0, T]. \quad (33)$$

By the integration by parts formula (31),

$$\frac{V_T}{B_T} = \frac{V_0}{B_0} + \int_0^T N_t^p d\frac{P_t^e}{B_t} + \int_0^T \frac{P_t^e}{B_t} dN_{t+}^p + \int_0^T N_t^c d\frac{C_t^e}{B_t} + \int_0^T \frac{C_t^e}{B_t} dN_{t+}^c + \int_0^T dN_{t+}^b. \quad (34)$$

where we have used the readily verifiable fact that $V_t = N_{t+}^p P_t^e + N_{t+}^c C_t^e + N_{t+}^b$ for $t = 0$ and $t = T$.

The integrals with respect to dN_{t+}^c and dN_{t+}^p sample the integrand at $t = \sigma_1, \tau_1, \sigma_2, \tau_2, \ldots, \sigma_{N_T^{px}}, \tau_{N_T^{cx}}$. If $t = \sigma_1, \sigma_2, \ldots, \sigma_{N_T^{px}}$, then $dN_{t+}^p = -1$ and $dN_{t+}^c = 1$; while if $t = \tau_1, \tau_2, \ldots, \tau_{N_T^{cx}}$, then $dN_{t+}^p = 1$ and $dN_{t+}^c = -1$. Therefore

$$\begin{aligned} \frac{V_T}{B_T} &= \frac{V_0}{B_0} + \int_0^T (1-\gamma_t) d\frac{P_t^e}{B_t} + \sum_{i=1}^{N_T^{px}} \frac{C_{\sigma_i}^e - P_{\sigma_i}^e}{B_{\sigma_i}} + \int_0^T \gamma_t d\frac{C_t^e}{B_t} + \sum_{i=1}^{N_T^{cx}} \frac{P_{\tau_i}^e - C_{\tau_i}^e}{B_{\tau_i}} + N_{T+}^b \\ &= \frac{V_0}{B_0} + \int_0^T (1-\gamma_t) d\frac{P_t^e}{B_t} + \sum_{i=1}^{N_{T-}^{px}} (F_{\sigma_i} - K) + \int_0^T \gamma_t d\frac{C_t^e}{B_t} + \sum_{i=1}^{N_{T-}^{cx}} (K - F_{\tau_i}) + N_T^b, \\ &= \frac{V_0}{B_0} + \int_0^T (1-\gamma_t) d\frac{P_t^e}{B_t} + \int_0^T \gamma_t d\frac{C_t^e}{B_t} + \int_0^T N_t^b d\frac{B_t}{B_t}, \end{aligned} \quad (35)$$

by (1) and (32). This shows that the trading strategy is self-financing.

Also, since $B_T = 1, P^e_T = (K - F_T)^+$, and $C^e_T = (F_T - K)^+$, we have

$$V_T = (1 - \gamma_T)(K - F_T)^+ + \gamma_T(F_T - K)^+ + \sum_{i=1}^{N^{px}_{T-}} (K - F_{\sigma_i}) + \sum_{i=1}^{N^{cx}_{T-}} (F_{\tau_i} - K), \quad (36)$$

which is the terminal payoff P^h_T to the holder of a hyper put, for any exercise strategy.

By no-arbitrage, the initial hyper put value must equal the initial value of the replicating strategy:

$$P^h_0 = V_0 = (1 - \gamma_0)P^e_0 + \gamma_0 C^e_0 + 0 \cdot B_0 = P^e_0.$$

Hence the hyper put has the same initial value as a European put.

A similar argument holds for hyper calls. **Q.E.D.**

References

[1] Carr, P. and R. Jarrow, 1990, "The Stop-Loss Start Gain Paradox and Option Valuation: A New Decomposition into Intrinsic and Time Value", *Review of Financial Studies*, **3**, 469–492.

[2] Merton, R., 1973, "Theory of Rational Option Pricing", *Bell Journal of Economics and Management Science*, **4**, 141–183.