Online Appendix

Modifying Result 1 to accommodate power utility

Result 8 (Power utility version of Result 1)**.** *Let X_T be some random variable of interest whose value becomes known at time T. Then we can compute the expected value of X_T from the perspective of an investor with power utility over time-T wealth who chooses to invest fully in the market, written $\widetilde{\mathbb{E}}_t\, X_T$, as follows:*

$$\widetilde{\mathbb{E}}_t\, X_T = \frac{\mathbb{E}^*_t\left(X_T R_T^{\gamma}\right)}{\mathbb{E}^*_t\left(R_T^{\gamma}\right)}, \tag{1}$$

where γ is the investor's relative risk aversion.

Proof. The investor's portfolio choice problem is

$$\max_{\{w_i\}_{i=1,\dots,N}} \widetilde{\mathbb{E}}_t\, u\left(W_t \sum_{j=1}^{N} w_j R_{j,T}\right) \quad \text{s.t.} \quad \sum_{j=1}^{N} w_j = 1,$$

where w_i is the investor's portfolio weight on asset i, $u(x) = x^{1-\gamma}/(1-\gamma)$ is the power utility function, W_t is time-t wealth, and $R_{j,T}$ is the return on asset j. The first-order conditions for this problem, combined with the fact that the investor's optimally chosen portfolio return is (by assumption) the market return, imply that the SDF is proportional to $R_T^{-\gamma}$. The result follows as in the body of the paper. □

Result 9. *For any θ, we have*

$$\begin{aligned}\mathbb{E}^*_t\, R_T^{\theta} =& R_{f,t}^{\theta} + R_{f,t}\left\{\int_0^{F_{t,T}} \frac{\theta(\theta-1)}{S_t^{\theta}}\left(S_t R_{f,t} - F_{t,T} + K\right)^{\theta-2} \text{put}_{t,T}(K)\, dK + \right.\\ &\left. + \int_{F_{t,T}}^{\infty} \frac{\theta(\theta-1)}{S_t^{\theta}}\left(S_t R_{f,t} - F_{t,T} + K\right)^{\theta-2} \text{call}_{t,T}(K)\, dK\right\}.\end{aligned} \tag{2}$$

Proof. By the definition of R_T,

$$\mathbb{E}^*_t\, R_T^{\theta} = \mathbb{E}^*_t\left[\left(\frac{S_T + D_T}{S_t}\right)^{\theta}\right] = R_{f,t} \cdot \frac{1}{R_{f,t}}\, \mathbb{E}^*_t\left[\left(\frac{S_T + D_T}{S_t}\right)^{\theta}\right].$$

The term after the dot equals the price of a contract with payoff $[(S_T + D_T)/S_t]^{\theta}$ at time T. Breeden and Litzenberger (1978) show that given some function $f(\cdot)$, the time-0 price of the time-T payoff $f(S_T)$ is $\int_0^\infty f(K)\,\text{call}''_{t,T}(K)\,dK$. Thus,

$$\mathbb{E}_t^*\, R_T^\theta = R_{f,t} \cdot \int_0^\infty \left(\frac{K + D_T}{S_t}\right)^\theta \text{call}''_{t,T}(K)\,dK.$$

This expression can be simplified by integration by parts. Splitting the range of integration into two, and using the fact that $\text{call}''_{t,T}(K) = \text{put}''_{t,T}(K)$ (which follows by differentiating the put-call parity relationship twice with respect to K), we have

$$\mathbb{E}_t^*\, R_T^\theta = R_{f,t} \left\{ \int_0^{F_{t,T}} \left(\frac{K + D_T}{S_t}\right)^\theta \text{put}''_{t,T}(K)\,dK + \int_{F_{t,T}}^\infty \left(\frac{K + D_T}{S_t}\right)^\theta \text{call}''_{t,T}(K)\,dK \right\}.$$

The result follows by integrating by parts twice and using put-call parity to simplify the resulting expression. □

The above two results can be used to compute all moments of the return on the market (as perceived by an investor with power utility) from option prices. For example, Figure XV plots the equity premium perceived by investors with risk aversion $\gamma = 1, 2, 4$, and 8 at one-, three-, and twelve-month horizons. As another example, we can use the results to compute a measure of *true* (rather than risk-neutral) forward-looking volatility, as perceived by an investor with power utility. Figure XVI plots true forward-looking volatility (in the case $\gamma = 1$, i.e. taking the perspective of the log investor emphasized in Section VI of the main paper) at horizons of one month and one year. For comparison, the panels also show risk-neutral volatility, which is higher than true volatility, as one would expect.

Pricing and hedging with $\Delta > 0$

The hedging strategy provided in Tables VIII and IX perfectly replicates the desired payoff when $\Delta > 0$, but requires positions in options at all expiry dates $\Delta, \ldots, T - \Delta$. Discretizing the continuous-time strategy provided in the statement of Result 6 (which is exactly valid in the limit as $\Delta \to 0$) is equivalent to ignoring all such positions in options with intermediate expiry dates. The cashflows in these rows contribute a term of size $O(\Delta)$ at time 0, and terms of size $O(\Delta^2)$ at dates between 1 and $T - \Delta$. Thus the overall replication error is of size $O(\Delta)$, so the limiting strike is a good approximation to the truth for sampling intervals $\Delta > 0$. The next result makes this formal.

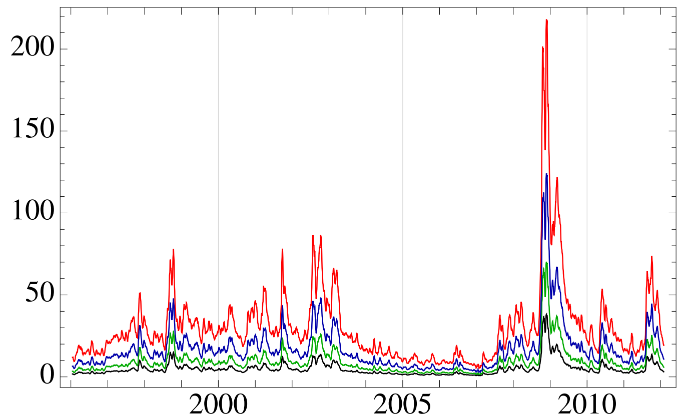

(a) 1-month horizon

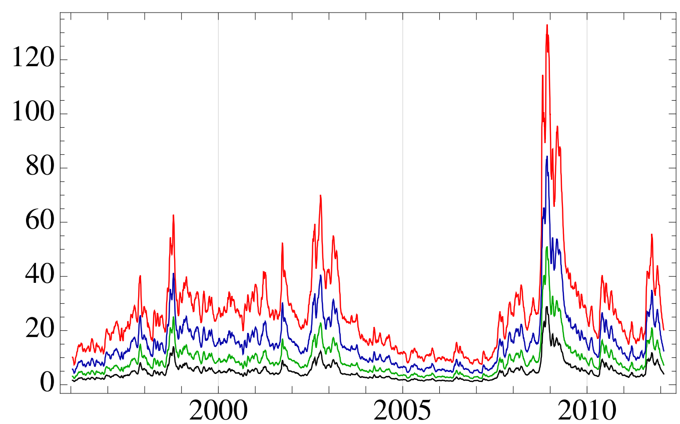

(b) 3-month horizon

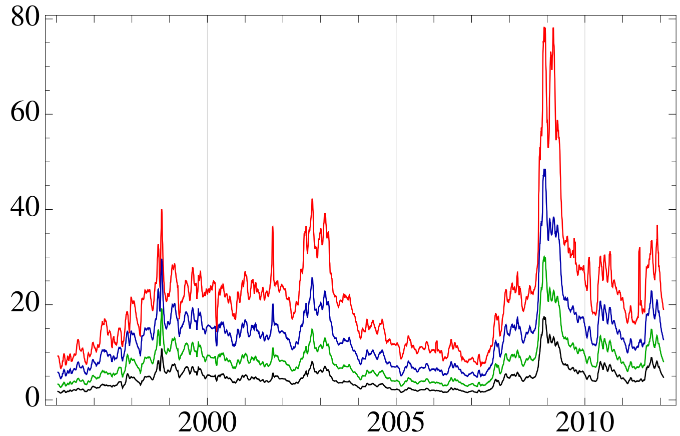

(c) 1-year horizon

FIGURE XV: The annualized equity premium at different horizons (in %), assuming risk aversion of 1 (black), 2 (green), 4 (blue) and 8 (red). The figures show 10-day moving averages.

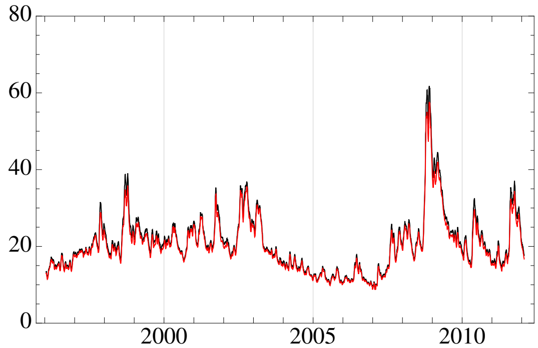

(a) 1-mo vol* (black) and vol (red)

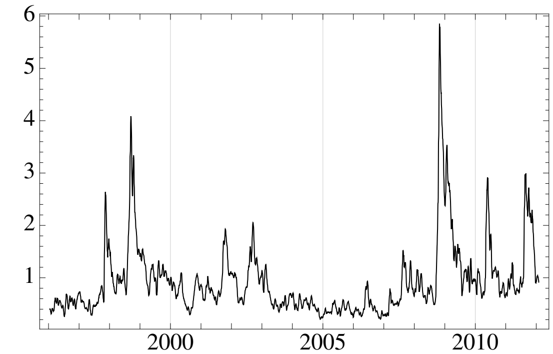

(b) 1-mo vol* minus vol

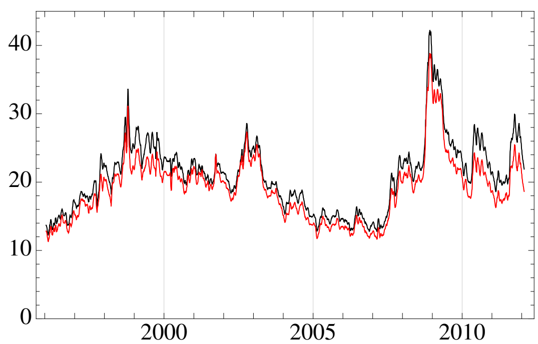

(c) 1-yr vol* (black) and vol (red)

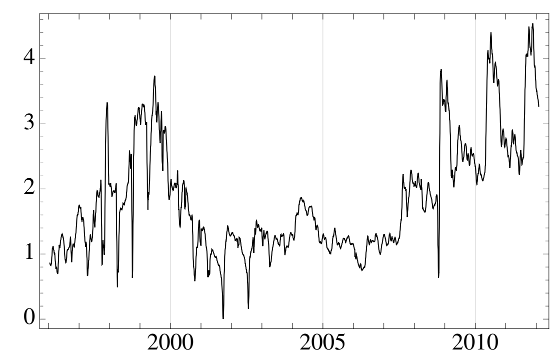

(d) 1-yr vol* minus vol

FIGURE XVI: Risk-neutral volatility (vol*) and the log investor's perceived true volatility (vol) over 1-month and 1-year horizons, annualized and in %. 10-day moving averages.

Result 10. *For $\Delta > 0$, the exact simple variance swap strike $V(\Delta)$ is well approximated by V, given in equation (32):*

$$|V(\Delta) - V| \leq \frac{T}{\Delta}\left(e^{(r-\delta)\Delta} - 1\right)^2 (1 + V) + \left|e^{2(r-\delta)\Delta} - 1\right| V \,.$$

If $T = 1$, $r - \delta = 0.02$, $V = 0.05$, then the right-hand side is less than 0.00001 with daily sampling ($\Delta = 1/252$), less than 0.00005 with weekly sampling ($\Delta = 1/52$), and less than 0.0002 with monthly sampling ($\Delta = 1/12$).

Proof. Result 6 implies that for $j < T/\Delta$,

$$\frac{e^{rj\Delta}P(j\Delta)}{F_{0,j\Delta}^2} = \lim_{\Delta\to 0} \mathbb{E}_0^* \sum_{i=1}^{j} \left[\frac{S_{i\Delta} - S_{(i-1)\Delta}}{F_{0,(i-1)\Delta}}\right]^2 \leq \lim_{\Delta\to 0} \mathbb{E}_0^* \sum_{i=1}^{T/\Delta} \left[\frac{S_{i\Delta} - S_{(i-1)\Delta}}{F_{0,(i-1)\Delta}}\right]^2 = \frac{e^{rT}P(T)}{F_{0,T}^2}.$$

Combining this observation with (37), we find that

$$\begin{aligned}\left|V(\Delta) - \frac{e^{rT}P(T)}{F_{0,T-\Delta}^2}\right| &= \sum_{j=1}^{T/\Delta-1} \left(e^{(r-\delta)\Delta} - 1\right)^2 \frac{e^{rj\Delta}P(j\Delta)}{F_{0,j\Delta}^2} + \frac{T}{\Delta}\left(e^{(r-\delta)\Delta} - 1\right)^2 \\ &\le \frac{T}{\Delta}\left(e^{(r-\delta)\Delta} - 1\right)^2 \frac{e^{rT}P(T)}{F_{0,T}^2} + \frac{T}{\Delta}\left(e^{(r-\delta)\Delta} - 1\right)^2.\end{aligned}$$

Now, by definition of V, we have $|e^{rT}P(T)/F_{0,T-\Delta}^2 - V| = \left|e^{2(r-\delta)\Delta} - 1\right| V$. Since $|V(\Delta) - V| \le |V(\Delta) - e^{rT}P(T)/F_{0,T-\Delta}^2| + |e^{rT}P(T)/F_{0,T-\Delta}^2 - V|$, by the triangle inequality, the result follows. □

Pricing and hedging when deep-out-of-the-money strikes are not tradable

Options that are sufficiently deep-out-of-the-money have prices so close to zero that they are not traded. Thus the idealized replicating portfolio, which comprises options of all strikes, is not attainable in practice. This issue affects both conventional variance swaps and simple variance swaps. Fortunately there is a practical solution to this problem. Suppose that, at time 0, options with strikes between A and B are tradable; the idealized scenario in which *all* strikes are tradable corresponds to $A = 0$, $B = \infty$. Then we can define the modified payoff

$$\left(\frac{S_\Delta - S_0}{F_{0,0}}\right)^2 + \left(\frac{S_{2\Delta} - S_\Delta}{F_{0,\Delta}}\right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{F_{0,T-\Delta}}\right)^2 - \phi(S_T),$$

where the correction term $\phi(S_T)$ is zero unless the underlying asset's price happens to end up outside the original strike range (A, B):

$$\phi(S_T) = \begin{cases} \left(\frac{A-S_T}{F_{0,T-\Delta}}\right)^2 & \text{if } S_T < A. \\ 0 & \text{if } A \le S_T \le B. \\ \left(\frac{S_T-B}{F_{0,T-\Delta}}\right)^2 & \text{if } S_T > B. \end{cases}$$

This modified payoff *can* be replicated without needing to trade options with strikes outside the range (A, B), by holding

(i) a static position in $2/F_{0,T}^2\, dK$ puts expiring at time T with strike K, for each $A < K \le F_{t,T}$,

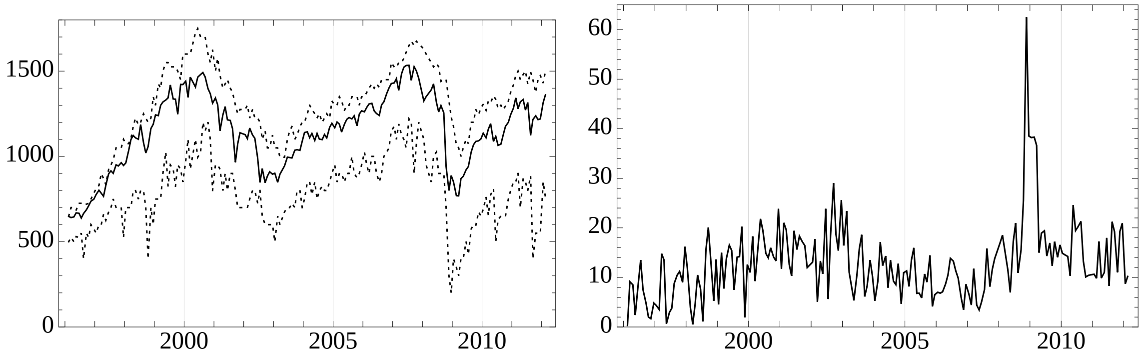

(a) S_T, A, and B (b) $\frac{\min\{S_T - A, B - S_T\}}{S_T}$

FIGURE XVII: Left: Upper and lower strike boundaries, A and B (dotted lines) and subsequent realized level of the market at expiry, S_T (solid line). Right: The distance at expiry from the edge of the strike range, expressed as a percentage of the terminal level S_T.

(ii) a static position in $2/F_{0,T}^2\, dK$ calls expiring at time T with strike K, for each $F_{t,T} \leq K < B$, and

(iii) a dynamic position of $2e^{-\delta(T-t)}(1 - S_t/F_{0,t})/F_{0,T}$ units of the underlying at time t,

financed by borrowing. To see this, simply note that the payoff $\phi(S_T)$ is precisely the payoff on the "missing" options with strikes less than A and greater than B that are not included in the above position.

In the limit as $\Delta \to 0$, the fair strike for the modified payoff is

$$\widehat{V} \equiv \frac{2e^{rT}}{F_{0,T}^2} \left\{ \int_A^{F_{0,T}} \text{put}_{0,T}(K)\, dK + \int_{F_{0,T}}^B \text{call}_{0,T}(K)\, dK \right\}.$$

To explore how large the adjustment term $\phi(S_T)$ is in practice in the case of the S&P 500 index, I looked at every day in the sample on which *OptionMetrics* had data for options expiring in 30 days. On each such day, I recorded the lowest tradable strike (i.e. the strike of the most deep-out-of-the-money put option) and the highest tradable strike (i.e. the strike of the most deep-out-of-the-money call option), together with the subsequently realized level of the market at expiry time T.

The results are shown in Figure XVII. Over the sample period, the underlying asset's price *never* ended up outside the range of tradable strikes. In other words, the correction term $\phi(S_T)$ was zero in every case: in Figure XVIIa, the value of S_T at expiry is within the range of strikes that were tradable at initiation on every day in sample. Figure XVIIb shows how far the underlying ended from the closer of the two boundaries, expressing the result as a percentage of S_T; the graph is always positive, reflecting the fact that the strike boundary was never crossed over the sample period. The low point in the figure occurred at the very beginning of the sample, on January 18, 1996, when the S&P 500 closed at 608.24. On that day, the highest strike tradable on options expiring in 30 days—on Saturday, February 17, 1996—was 650; in the event, the S&P 500 closed just two points lower, at 647.98, on Friday, February 16.

As is apparent from Figure XVIIa, the width of the range of tradable strikes has tended to increase over time. The mean value of the percentage distance to the edge of the strike range, as illustrated in Figure XVIIb, is 12.9%; the median value is 11.9%. In other words, on the median day in sample, the S&P 500 would have had to move a *further* 11.9% in the appropriate direction in order to exit the range of tradable strikes.

Pricing and hedging under different assumptions on dividends

This section shows what happens to pricing and hedging of simple variance swaps under various different assumptions about dividend payout policies.

Completely unanticipated dividend payouts. Result 6 continues to hold if the asset makes unanticipated dividend payouts. Consider an extreme case in which the simple variance swap is priced and hedged, at time zero, as though $\delta = 0$; but immediately after inception of the trade, at time $t = \Delta$, the underlying asset is suddenly liquidated via an extraordinary dividend, causing its (ex-dividend) price to equal 0 from time Δ onwards. The payout that must be made by the counterparty who is short variance equals 1, in this example. Meanwhile, the hedge portfolio given in the above result will generate a positive payoff due to the put options going in-the-money. (The dynamic position will have zero payoff: it was neither long nor short at time 0, and subsequently the asset's price never moved from zero.) Since $S_T = 0$, the total payoff will be

$$\frac{2}{F_{0,T}^2} \int_0^{F_{0,T}} \max\{0, K - S_T\}\, dK = \frac{2}{F_{0,T}^2} \int_0^{F_{0,T}} K\, dK = 1.$$

In other words, the strategy perfectly replicates the desired payoff. This applies more generally: once the strike V is set and the replicating portfolio is in place, it does not matter why the price path moves around subsequently, whether due to the payment of unanticipated dividends or not.

Perfectly anticipated dividend payouts. Consider the case in which the asset pays a dividend $D_{k\Delta}$ at time $k\Delta$ for some k, and no dividends at any other time up to and including the expiry date, T. The price of a portfolio whose payoff is S_i^2 at time i continues to equal $\Pi(i)$, given by equation (35).

In this section, it will be important to distinguish between $F_{0,t}$, the forward price of the dividend-paying asset to time t, and $\widetilde{F}_{0,t} \equiv S_0 e^{rt}$, the appropriate normalization for the definition of a simple variance swap in this case. A standard no-arbitrage argument implies that the forward price is $F_{0,t} = S_0 e^{rt}$ if $t < k\Delta$, and $F_{0,t} = S_0 e^{rt} - D_{k\Delta} e^{r(t-k\Delta)}$ if $t \geq k\Delta$, so $F_{0,t}$ and $\widetilde{F}_{0,t}$ coincide for times t before the payment of the dividend, but differ thereafter. It turns out that $\widetilde{F}_{0,t}$ is the appropriate normalization so that the intermediate option positions are negligibly small, as was the case in the main text.

The definition of the payoff on the simple variance swap must be modified to allow for the presence of the dividend. At time T, the counterparties to the simple variance swap now exchange V for

$$\left(\frac{S_\Delta - S_0}{\widetilde{F}_{0,0}}\right)^2 + \cdots + \left(\frac{S_{(k-1)\Delta} - S_{(k-2)\Delta}}{\widetilde{F}_{0,(k-2)\Delta}}\right)^2 + \left(\frac{S_{k\Delta} + D_{k\Delta} - S_{(k-1)\Delta}}{\widetilde{F}_{0,(k-1)\Delta}}\right)^2 + \\ + \left(\frac{S_{(k+1)\Delta} - S_{k\Delta}}{\widetilde{F}_{0,k\Delta}}\right)^2 + \cdots + \left(\frac{S_T - S_{T-\Delta}}{\widetilde{F}_{0,T-\Delta}}\right)^2.$$

If the stock price happens to track the forward price at all points in time, then this payoff will be zero in the $\Delta \to 0$ limit, as is the case with variance swaps and simple variance swaps in the absence of dividends.

The starting point of the replicating strategy will be to carry out precisely the trades listed in Tables VIII and IX of the main paper with δ set equal to zero (and replacing $F_{0,t}$ with $\widetilde{F}_{0,t}$ wherever it occurs in the tables). This replicating strategy generates the above payoff minus V, plus an extra payoff of $(D_{k\Delta}/\widetilde{F}_{0,(k-1)\Delta})^2 - 2D_{k\Delta}(S_{k\Delta} + D_{k\Delta})/\widetilde{F}^2_{0,(k-1)\Delta}$. To offset this extra payoff, two new positions are required: (i) a short position of $e^{-rT}(D_{k\Delta}/\widetilde{F}_{0,(k-1)\Delta})^2$ (measured in dollars) in bonds, and (ii) a long position of $2D_{k\Delta}e^{-r(T-k\Delta)}/\widetilde{F}^2_{0,(k-1)\Delta}$ units of the underlying held until time $k\Delta$, then rolled into bonds.

After some algebra (and up to terms of order Δ, as usual) this implies that the simple variance swap strike is given by

$$V = \frac{2e^{rT}}{\widetilde{F}_{0,T}^2} \left\{ \int_0^{F_{0,T}} \text{put}_{0,T}(K)\, dK + \int_{F_{0,T}}^{\infty} \text{call}_{0,T}(K)\, dK \right\},$$

and that the replicating portfolio is equivalent to holding

(i) a static position of $2/\widetilde{F}_{0,T}^2\, dK$ puts expiring at time T with strike K, for each $K \leq F_{0,T}$,

(ii) a static position of $2/\widetilde{F}_{0,T}^2\, dK$ calls expiring at time T with strike K, for each $K \geq F_{0,T}$, and

(iii) a dynamic position of $2(F_{0,t} - S_t)/(\widetilde{F}_{0,t}\widetilde{F}_{0,T})$ units of the underlying at time t,

financed by borrowing.

Imperfectly anticipated dividend payouts. In the general case in which dividends are of unknown size and timing, simple variance swaps can be straightforwardly priced and hedged if *total return* options can be traded: these are options on a claim to the underlying-with-dividends-reinvested (the latter being a tradable asset that does *not* pay dividends). We can price and hedge a simple variance swap on the underlying-with-dividends-reinvested by reinterpreting the inputs to Result 6. The price S_t corresponds to the value of the underlying-with-dividends-reinvested (so S_0 is the spot price of the underlying asset); the instantaneous dividend yield $\delta = 0$; $F_{0,t}$ is the forward price of the dividend-adjusted underlying, which equals $S_0 e^{rt}$ for all t by a static no-arbitrage argument; and $\text{put}_{0,T}(K)$ and $\text{call}_{0,T}(K)$ are the prices of total return options expiring at time T.