Options and the Gamma Knife

Options and the Gamma Knife

Ian Martin*

April, 2018

Abstract

I survey work of Steve Ross (1976) and of Douglas Breeden and Robert Litzenberger (1978) that first showed how to use options to synthesize more complex securities. Their results made it possible to infer the risk-neutral measure associated with a traded asset, and underpinned the development of the VIX index. The other main result of Ross (1976), which shows how to infer *joint* risk-neutral distributions from option prices, has been much less influential. I explain why, and propose an alternative approach to the problem. This paper is dedicated to Steve Ross, and was written for a special issue of the *Journal of Portfolio Management* in memory of him.

Why do options have such a central place in finance? One reason is that options arise "in nature." A job offer letter gives its recipient the right, but not the obligation, to accept a job; the owner of an asset has the right, but not the obligation, to sell it; the owner of a plot of empty land may have the option to develop it; and so on. Another reason was provided in a classic paper of Ross (1976) and later elaborated by Breeden and Litzenberger (1978): options help to complete markets. This insight has become one of the most useful—and one of the few *robust*—tools in the financial economist's toolkit.

Suppose, for example, that the gold price is currently \$1000. How much would you pay to receive the cube root of the price of gold in a year? How much would you pay for the inverse of the price of gold in a year? If offered the opportunity to trade these contracts at \$10 and \$1/1000 respectively, should you buy or sell?

*London School of Economics. http://personal.lse.ac.uk/martiniw/.

Remarkably, it turns out that these questions have precise and unambiguous answers if you can observe the prices of European call and put options on gold. I explain why, in an exposition of the Ross–Breeden–Litzenberger papers. More generally, their results show that it is possible to infer the risk-neutral distribution of a random variable from the prices of European options on that random variable.

Equally remarkably, these and related questions—which at first sight might seem dryly academic—have, indirectly, had tremendous influence on financial markets. A leading example is the VIX index, whose definition is based[1] on the price of a claim to the logarithm of the level of the S&P 500 index. The Ross–Breeden–Litzenberger result is also conceptually important: as an example, I use it below to show (in a considerably more general setting than that of the Black and Scholes (1973) model) why volatility is central to option pricing.

Unfortunately the result only applies in one dimension: it shows how to determine the risk-neutral distribution of, say, the price of gold in a year's time. But it does not help to determine the *joint* distribution of the price of gold and the price of platinum; nor the joint distribution of the dollar-euro and dollar-yen exchange rates; nor the joint distribution of a given stock and the market index.

In the remainder of the paper, I address the question of whether there is a higher-dimensional version of the result that *would* reveal joint distributions such as these. On the face of it, the "main result," Theorem 4, of Ross (1976) does exactly this; and yet that result has had surprisingly little impact. I suggest an explanation for this fact, and propose an alternative approach that may be better suited for empirical implementation in practice. It is loosely inspired by the design of a device called a *gamma knife* that is used in neurosurgery to irradiate tumors while minimizing the damage to surrounding tissue.

The Ross–Breeden–Litzenberger result

The risk-neutral expectation operator, $\mathbb{E}_t^*$, has the property that the price (at time t) of a tradable payoff X_T (received at time T) is

$$\frac{1}{R_{f,t\to T}}\,\mathbb{E}_t^*\,X_T.$$

[1]The definition is based on the key contributions of Neuberger (1994) and Carr and Madan (2001), who built on the work of Ross (1976) and Breeden and Litzenberger (1978).

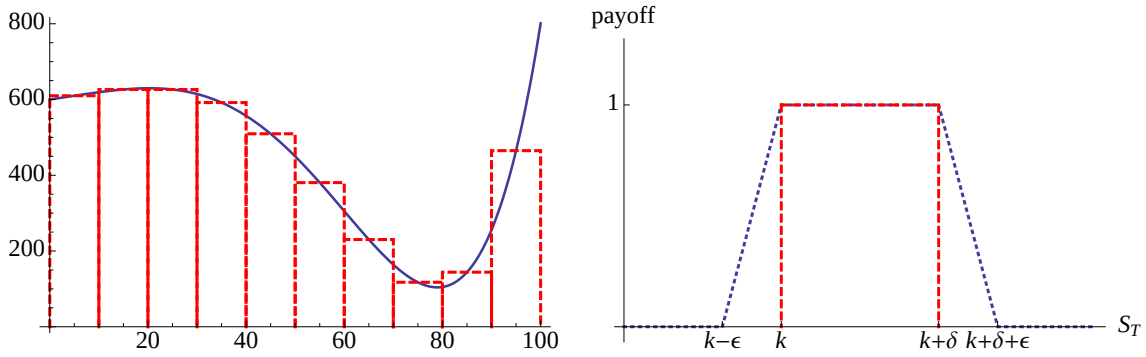

(a) Digitizing (b) Building a pixel from call options

Exhibit 1: The Breeden–Litzenberger logic. As $\varepsilon \to 0$ in panel (b), the blue dotted lines converge to the red dashed lines.

The asterisk means that this is the *risk-neutral* expectation; the subscript t indicates that it is the expectation *conditional on all information known at time t*. $R_{f,t\to T}$ is the gross riskless return from time t to T, which is determined by the corresponding zero-coupon bond yield. (Technically, therefore, $\mathbb{E}_t^*$ is the conditional expectation with respect to the $(T-t)$-*forward measure*.)

Breeden and Litzenberger (1978) built on Ross (1976) in showing how to compute risk-neutral expectations of the form

$$\mathbb{E}_t^* \, g(S_T)$$

for any random variable S_T—typically an asset price—on which European options are traded. So if we observe the prices of options on gold then we can calculate the risk-neutral expectation of some arbitrary *function* of the price of gold at some future time T. For instance, we can use option prices to calculate the risk-neutral distribution of S_T, as

$$\mathbb{P}_t^* \left(S_T \in [k, k+\delta] \right) = \mathbb{E}_t^* \left(\mathbf{1} \left\{ S_T \in [k, k+\delta] \right\} \right),$$

which we can evaluate by applying the Breeden–Litzenberger logic to the function $g(x) = \mathbf{1}\{x \in [k, k+\delta]\}$. Then divide by δ and let $\delta \to 0$ to compute the risk-neutral density of S_T at k.

The basic idea is illustrated in Exhibit 1. Panel a shows the function $g(S_T)$, viewed as an analog signal that can be digitized by being cut into "pixels" of width δ.

Next, we construct each pixel out of call options, as shown in panel b: to generate a payoff of 1 if $S_T \in [k, k+\delta]$ and 0 otherwise, buy $\frac{1}{\varepsilon}$ calls with strike $k-\varepsilon$, sell $\frac{1}{\varepsilon}$

calls with strike k and $\frac{1}{\varepsilon}$ calls with strike $k+\delta$, and buy $\frac{1}{\varepsilon}$ calls with strike $k+\delta+\varepsilon$. Letting $\varepsilon \to 0$, we have built the pixel.

The price of the pixel illustrated in panel b is therefore

$$\lim_{\varepsilon\to 0} \frac{\text{call}_{t,T}(k-\varepsilon) - \text{call}_{t,T}(k)}{\varepsilon} - \frac{\text{call}_{t,T}(k+\delta) - \text{call}_{t,T}(k+\delta+\varepsilon)}{\varepsilon} = \text{call}'_{t,T}(k+\delta) - \text{call}'_{t,T}(k)$$
$$\approx \text{call}''_{t,T}(k)\delta,$$

where the approximation becomes perfect as $\delta \to 0$, and where we write $\text{call}_{t,T}(k)$ for the time t price of a European call with strike k that expires at time T. The price of the digitized function in panel a is thus approximately

$$\sum_i g(k_i)\,\text{call}''_{t,T}(k_i)\delta.$$

Finally, we can increase the "resolution" of the resulting digitization by sending $\delta \to 0$ to find the exact price of the original (analog) function illustrated in panel a:

$$\text{price of a claim to } g(S_T) = \int_0^\infty g(K)\,\text{call}''_{t,T}(K)\,dK. \tag{1}$$

This is the Breeden and Litzenberger (1978) result.[2]

Exhibit 2 shows a hypothetical collection of call and put option prices; they intersect at the forward (to time T) price of the underlying asset, $F_{t,T}$. Equation (1) shows that the price of a claim to $g(S_T)$ depends on the second derivative (in the mathematical sense) of the call price as a function of strike, $\text{call}_{t,T}(K)$. Following Carr and Madan (2001), we can now integrate by parts twice[3] to find a more intuitive expression in terms of out-of-the-money option prices:

$$\text{price of a claim to } g(S_T) = \underbrace{\frac{g(F_{t,T})}{R_{f,t\to T}}}_{\text{naive guess}} + \underbrace{\int_{K=0}^{F_{t,T}} g''(K)\,\text{put}_{t,T}(K)\,dK + \int_{K=F_{t,T}}^{\infty} g''(K)\,\text{call}_{t,T}(K)\,dK}_{\text{convexity correction}}. \tag{2}$$

[2]Breeden and Litzenberger imposed a further assumption that $\text{call}_{t,T}$ should be twice differentiable. This is not needed, because Merton (1973) showed, in another classic paper, that call and put option prices (considered as a function of strike) are *convex*. It follows, by Alexandrov's theorem, that their second derivatives exist almost everywhere, which is all that is needed for (1) to make sense.

[3]In more detail: split the range of integration into two pieces, $[0, F_{t,T})$ and $[F_{t,T}, \infty)$, and use put-call parity to write $\text{call}''_{t,T} = \text{put}''_{t,T}$ in the lower range. Then integrate by parts twice, using put-call parity to cancel some of the resulting terms (assuming that the function g behaves sufficiently nicely

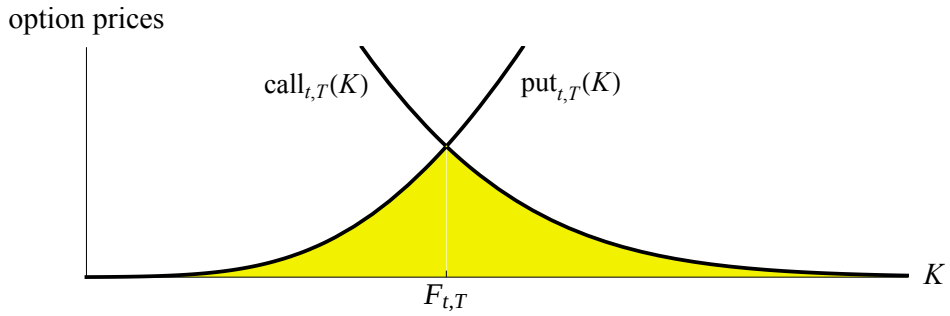

Exhibit 2: The prices, at time t, of call and put options expiring at time T.

Suppose, then, that the (1-year-forward) price of gold is \$1000, and that interest rates are zero (that is, for convenience, $R_{f,t\to T} = 1$). How much is a claim to the cube root of the price of gold worth? The first term in (2) is $\$\sqrt[3]{1000} = \$10$. The second and third terms supply the "convexity correction" to this naive first approximation. Without knowing anything at all about option prices—other than that they are nonnegative—we can be sure that the price of the cube root contract is less than \$10 because the function $g(K) = \sqrt[3]{K}$ is concave, so that $g''(K) < 0$. Conversely, the price of an "inverse contract" that pays $1/S_T$ must be more than \$1/1000, because the function $g(K) = 1/K$ is convex. In both cases, the convexity correction will be large if option prices are high.

Equation (2) reveals a general relationship between option prices and (some notion of) volatility. Setting $g(K) = K^2$ and neglecting dividends (so that $\mathbb{E}_t^* S_T = S_t R_{f,t\to T}$ and we can write $R_T = S_T/S_t$ for the return on the asset), we find

$$\operatorname{var}_t^* R_T = \frac{2R_{f,t\to T}}{S_t^2}\left\{\int_{K=0}^{F_{t,T}} \text{put}_{t,T}(K)\,dK + \int_{K=F_{t,T}}^{\infty} \text{call}_{t,T}(K)\,dK\right\}. \tag{3}$$

Option prices reveal the risk-neutral variance of the underlying asset's price at time T: the integrals inside the curly brackets are equal to the shaded area in Exhibit 2. Aside from its theoretical interest, this result has proved useful empirically. In Martin (2017), I defined the so-called SVIX index based on this calculation, and provided applications to forecasting the return on the S&P 500 index. Martin and Wagner (2018) defined SVIX indexes for individual stocks and show that they can be used to forecast returns on individual stocks.[4]

at zero and infinity). The result is (2).

[4]We also applied this idea in Martin and Ross (2018) to show that under the assumptions made by Ross (2015), put and call options on a (very) long-dated zero-coupon bond should, in principle, forecast

The name of the SVIX index is intended to recall the VIX index, which emerges from (2) on setting $g(K) = -\log K$:

$$2\left(\log \mathbb{E}_t^* R_T - \mathbb{E}_t^* \log R_T\right) = 2R_{f,t\to T}\left\{\int_0^{F_{t,T}} \frac{1}{K^2}\,\mathrm{put}_{t,T}(K)\,dK + \int_{F_{t,T}}^{\infty} \frac{1}{K^2}\,\mathrm{call}_{t,T}(K)\,dK\right\}. \tag{4}$$

If we take the underlying asset to be the S&P 500 index, then the right-hand side of (4) is the definition of the VIX index (squared), while the left-hand side is a measure of the variability of the return on the S&P 500 index that has been called its (risk-neutral) *entropy* (Alvarez and Jermann, 2005; Backus et al., 2011).

The original motivation for the definition of the VIX index was that if S_T follows a diffusion, we would have

$$\mathbb{E}_t^* \int_{\tau=t}^{T} \sigma_\tau^2\,d\tau = 2R_{f,t\to T}\left\{\int_0^{F_{t,T}} \frac{1}{K^2}\,\mathrm{put}_{t,T}(K)\,dK + \int_{F_{t,T}}^{\infty} \frac{1}{K^2}\,\mathrm{call}_{t,T}(K)\,dK\right\}, \tag{5}$$

where σ_τ is the instantaneous volatility of the underlying asset at time τ. That is, the formula that defines the VIX index (squared) would also provide the fair strike for a variance swap if S_T followed a diffusion.

Unfortunately the diffusion assumption is strong and demonstrably false; as, therefore, is equation (5).[5] By contrast, the results of Ross, Breeden, and Litzenberger—and in particular equations (1)–(4)—rely only on the logic of *static* replication and therefore allow the underlying asset prices to follow essentially any arbitrage-free process. In contrast, *dynamic* replication arguments of the type often invoked in the theory of option pricing, leading up to and beyond the famous formula of Black and Scholes (1973), require far stronger assumptions about asset price behavior and the ways in which uncertainty evolves.

---

the bond's returns. Testing this prediction is challenging, however, as tolerable approximations to the long bond option prices that the theory asks for are hard to come by.

[5]This can be seen directly by looking at financial asset prices, which jump both at predictable times (eg, when important economic numbers are released) and at unpredictable times (eg, when a plane hits the World Trade Center); or indirectly, from the fact that variance swap strikes diverge, in practice, from (5).

Ross–Breeden–Litzenberger in two dimensions

Ross, Breeden, and Litzenberger tell us how to calculate the risk-neutral distribution of dollar-yen in a month's time, or how to calculate the risk-neutral distribution of dollar-euro in a month's time. But how do the two interact? What is the *joint* risk-neutral distribution of dollar-yen and dollar-euro?

More generally, a practical method of computing quantities of the form $\mathbb{E}_t^* f(S_{1,T}, S_{2,T})$ would have many applications. Most obviously, joint risk-neutral distributions would be of direct interest in themselves, as they are in the one-dimensional case. A less obvious example is provided by Kremens and Martin (2018), who argue that the risk-neutral covariance of a given foreign currency with the S&P 500 index should (in theory) and does (in practice) forecast the currency's excess return. As it happens, precisely the assets whose prices must be observed to reveal this risk-neutral covariance—quanto forward contracts on the S&P 500 index—are traded. But this is something of a coincidence. Along similar lines, the approach of Martin and Wagner (2018) to forecasting individual stock returns would be greatly simplified if the risk-neutral covariances of individual stocks and the S&P 500 index were directly observable. This would be true if (in addition to stock and index options) outperformance options on stock i, relative to the index, were widely traded and liquid; as they are not, Martin and Wagner use a linearization to relate the desired risk-neutral covariances to risk-neutral variances, which are observable via (3).

The following result of Ross (1976) handles the case in which there are only finitely many states of the world. I provide a different proof that reveals that the result is fragile, in a sense that will become clear below. (For that reason, I specialize to the two-dimensional case for simplicity.)

Result 1 (Ross (1976), two-dimensional case)**.** *With finitely many states of the world, all Arrow–Debreu prices can be inferred from the prices of European calls and puts on a single, appropriately chosen, linear combination $S_{1,T} + \lambda S_{2,T}$ for some fixed $\lambda \in \mathbb{R}$.*

Proof. Suppose there are N states, which can be viewed as points in $\mathbb{R}^2$. There are at most $\binom{N}{2}$ lines joining these points in pairs; on each, $S_{1,T} + \mu_i S_{2,T}$ is constant for some μ_i, $i = 1, \ldots, \binom{N}{2}$. Let λ be a number *not* equal to any of the μ_i; then any line of the form $S_{1,T} + \lambda S_{2,T} = \text{constant}$ intersects at most one of the points. By trading options on the portfolio $S_{1,T} + \lambda S_{2,T}$, we can create an Arrow security for an arbitrary state

$(s_{1,T}, s_{2,T})$. Let $c = s_{1,T} + \lambda s_{2,T}$. The payoff on a "butterfly spread,"

$$\frac{1}{\varepsilon}\left[\max\left\{0, S_{1,T} + \lambda S_{2,T} - (c-\varepsilon)\right\} - 2\max\left\{0, S_{1,T} + \lambda S_{2,T} - c\right\} + \right.$$
$$\left. + \max\left\{0, S_{1,T} + \lambda S_{2,T} - (c+\varepsilon)\right\}\right], \qquad (6)$$

equals 1 if $(S_{1,T}, S_{2,T}) = (s_{1,T}, s_{2,T})$ by the definition of c. Furthermore, by taking $\varepsilon > 0$ sufficiently small, we can guarantee that $|S_{1,T} + \lambda S_{2,T} - c| > \varepsilon$ (and hence that the payoff (6) is zero) for all $(S_{1,T}, S_{2,T}) \neq (s_{1,T}, s_{2,T})$. The result follows. □

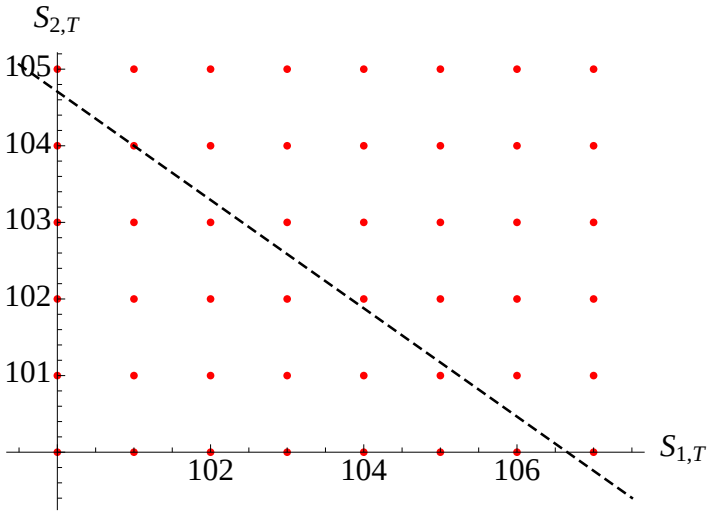

Exhibit 3: An illustration of Ross's (1976) result that with finitely many states of the world, options on $S_{1,T} + \lambda S_{2,T}$ will complete markets, if λ is chosen appropriately. The dashed line passes through $(101, 104)$ and avoids all other points on the grid; and any line parallel to the dashed line intersects *at most one* dot on the grid.

The proof is illustrated in Exhibit 3 with $(S_{1,T}, S_{2,T})$ lying on a grid. (The proof does not require the points to lie on a grid, but it is the relevant case in practice.) Along the dashed line, $S_{1,T} + \lambda S_{2,T} = c$ for some constants λ and c. As drawn, λ lies between 1 and 2, while c ensures that the dashed line passes through the point $(101, 104)$. The key to the proof is that we can arrange things (i.e., choose λ) so that any line parallel to the dashed line picks out *at most* one of the dots: that is, for any given c,

$$S_{1,T} + \lambda S_{2,T} = c \text{ in } \textit{at most} \text{ one of the states of the world.} \qquad (7)$$

Notice that if a line intersects two or more points, then its slope is a rational number.

Thus, if we choose λ to be any *irrational* number,[6] property (7) will hold; in Exhibit 3, λ is equal to $\sqrt{2}$. (At this point, the sense in which the result is fragile is starting to become clear.)

As in the one-dimensional case illustrated in Exhibit 1b, it is possible to buy a butterfly spread on $S_{1,T} + \lambda S_{2,T}$ that pays off precisely on the dashed line; and as the line only intersects the state $(101, 104)$, we have synthesized the Arrow–Debreu security for $(101, 104)$. By the same logic, options on this linear combination can be used to synthesize all the other Arrow–Debreu securities (by varying the strikes of the options in the butterfly spread, i.e., by varying c in the above construction).

Why, then, has the approach not been useful in practice? The construction fails if options are only traded on linear combinations for a limited set of λ, or if the state space is not finite. If we only observe options on combinations with, say, $\lambda = \pm 1$ or $\pm\frac{1}{2}$, then the butterfly spread construction will not "separate the dots": we cannot set things up so that the dashed lines only intersect a single dot. The larger and finer the grid, the more difficult it is to ensure that the dashed line threads between points and, correspondingly, the more extreme are the requirements on λ. If $S_{1,T}$ and $S_{2,T}$ take values on an infinite grid, then the construction will not work for any rational λ; and if $S_{1,T}$ and $S_{2,T}$ vary continuously, the result fails even if λ is allowed to be irrational.

The gamma knife

With the earlier results in mind, it is natural to try to synthesize a two-dimensional pixel, that is, a payoff of \$1 if both $S_{1,T} \in [s_1, s_1 + \delta]$ and $S_{2,T} \in [s_2, s_2 + \delta]$ for some small δ. It turns out that by taking a different approach, we will be able to work out the prices of individual pixels so long as we can observe options on multiple linear combinations of $S_{1,T}$ and $S_{2,T}$. Moreover, the construction will work even in the case in which $S_{1,T}$ and $S_{2,T}$ can vary continuously. The key idea is analogous to the *gamma knife* technology used by neurosurgeons, in which multiple beams of radiation are shone at a single target—perhaps a tumor—so that surrounding tissue receives relatively little harm; and similar ideas arise more generally in X-ray tomography.

Exhibit 4 illustrates how one could use this approach to find the price of a two-dimensional pixel (as shown, the price of an Arrow–Debreu security that pays off if $S_{1,T} = 3$ and $S_{2,T} = 5$) in an idealized example with perfect data. The line in Exhibit 4a

[6] This gives a direct way to see that *almost any* (in a measure-theoretic sense) choice of λ will do, as was shown by Arditti and John (1980).

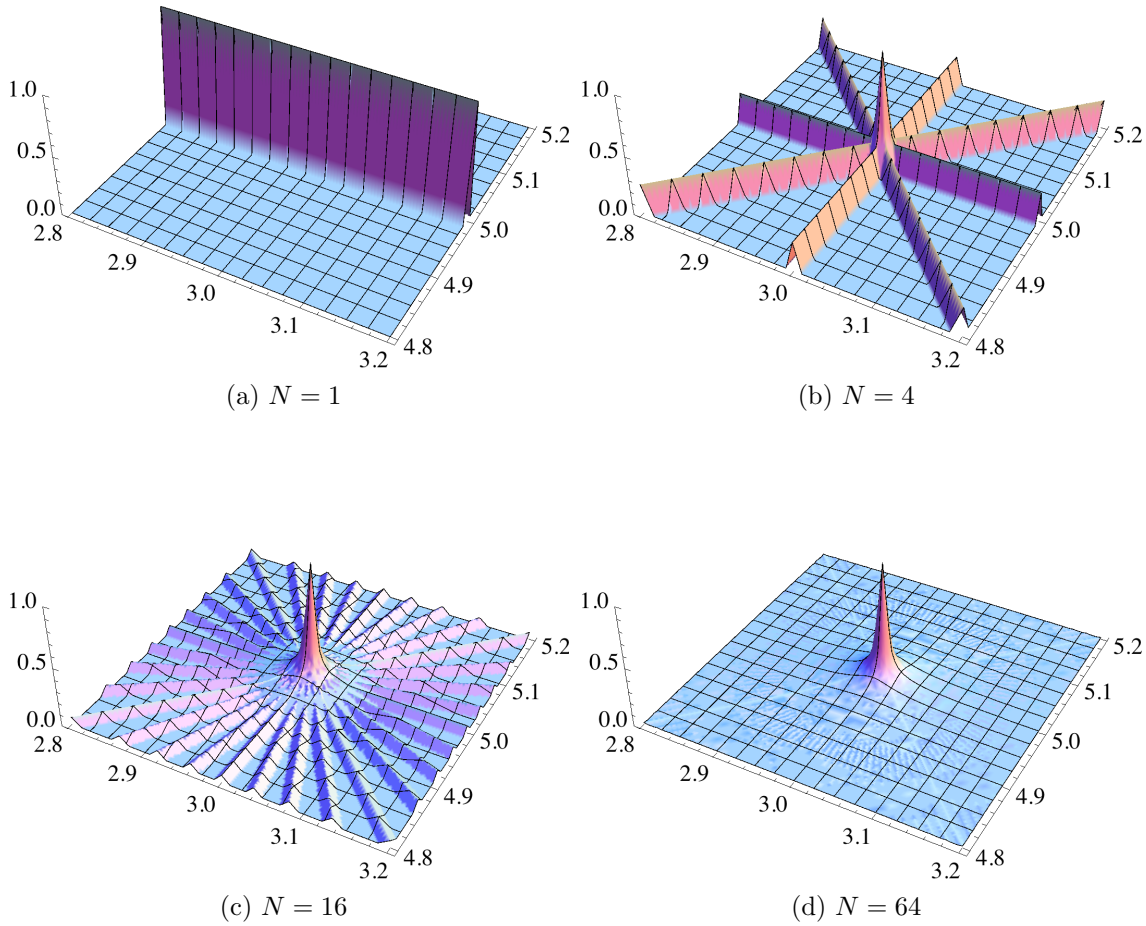

(a) $N = 1$ (b) $N = 4$

(c) $N = 16$ (d) $N = 64$

Exhibit 4: Using options on linear combinations of the two state variables to create approximate Arrow–Debreu securities.

corresponds to the dashed line of Exhibit 3: it illustrates the set of points at which a butterfly spread on $S_{2,T}$, constructed as in Exhibit 1b, has unit payoff. Exhibit 4b shows what happens when we also buy a butterfly spread on $S_{1,T}$, together with butterfly spreads on $S_{1,T} + S_{2,T}$ and $S_{1,T} - S_{2,T}$. Strikes are chosen so that all four lines intersect at the single point of interest. Here, the payoff on the portfolio of butterfly spreads is four times as great as the payoff on an individual butterfly spread, so we can scale back position sizes by a factor of 4. Continuing this process, Exhibits 4c and 4d use options on 16 and 64 different linear combinations of $S_{1,T}$ and $S_{2,T}$, respectively, to pick out the point of interest, where the payoff is 1. The resulting payoff approaches the desired pixel payoff as the number of linear combinations approaches infinity.

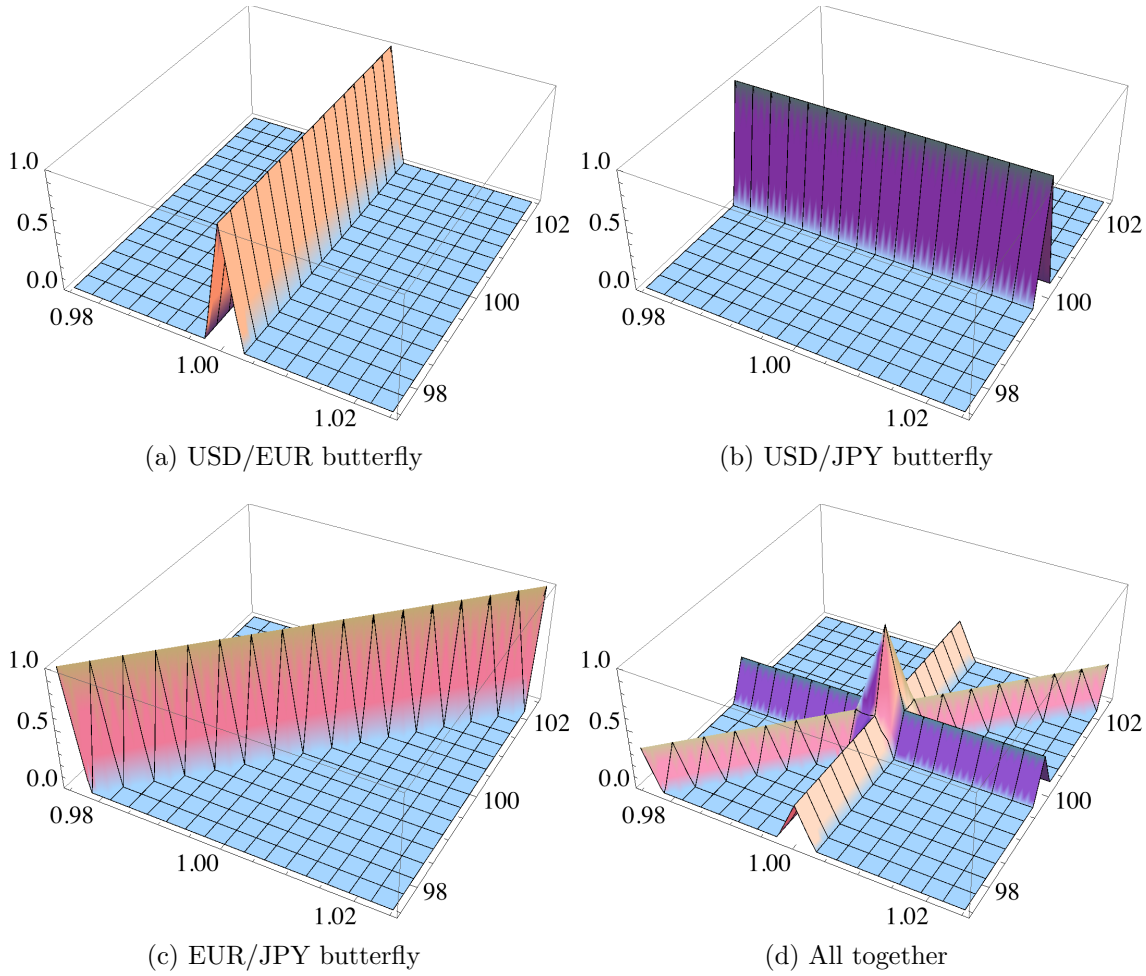

(a) USD/EUR butterfly

(b) USD/JPY butterfly

(c) EUR/JPY butterfly

(d) All together

Exhibit 5: Using options on linear combinations of the state variables to create approximate Arrow–Debreu securities.

Exhibit 5 shows a more realistic example, with $S_{1,T}$ corresponding to the USD/EUR exchange rate and $S_{2,T}$ to the USD/JPY exchange rate. Using options on USD/EUR, we can create a butterfly spread that pays one unit only if $S_{1,T}$ equals some pre-specified value, say 1; and using options on USD/JPY, we can create a butterfly spread that pays one unit only if $S_{2,T}$ equals, say, 100. (The exhibits show approximations to these securities.) Using options on EUR/JPY we can create a butterfly spread that pays one unit only if $S_{2,T}/S_{1,T} = 100$, i.e., $S_{2,T} = 100S_{1,T}$. Finally, adding the three together, and scaling the position sizes by $\frac{1}{3}$, we obtain the approximate Arrow–Debreu security illustrated in Exhibit 5d.[7]

[7]Basket options, which have payoffs of the form $\max\{0, a_1\text{USD/EUR} + a_2\text{USD/JPY} - K\}$ for a variety of a_1 and a_2, move us closer to the examples depicted in Exhibit 4.

In the interest of completeness, we have the following result. I write $f(x, y)$ for the joint risk-neutral density of $S_{1,T}$ and $S_{2,T}$; and assume that we can observe the prices of options on arbitrary linear combinations of $S_{1,T}$ and $S_{2,T}$.[8] We can then construct a butterfly spread, as in (6), that pays off along an arbitrary line L parametrized by p, its distance from the origin, and $\alpha \in [0, \pi)$, the angle of the normal to L. The butterfly spread associated with L reveals

$$\mathscr{R}f(L) = \mathscr{R}f(p, \alpha) = \int_{-\infty}^{\infty} f(p\cos\alpha - t\sin\alpha, p\sin\alpha + t\cos\alpha)\, dt.$$

Result 2 (The gamma knife, two-dimensional[9] case)**.** *If the prices of options on arbitrary linear combinations $aS_{1,T} + bS_{2,T}$ are observable then we can treat $\mathscr{R}f(L)$ as observable for all lines L and reconstruct the joint risk-neutral density of $S_{1,T}$ and $S_{2,T}$ via the formula*

$$f(x, y) = \lim_{\varepsilon \to 0} \frac{1}{\pi} \int_0^{\pi} \int_{-\infty}^{\infty} \mathscr{R}f(p - x\cos\alpha - y\sin\alpha, \alpha) G_{\varepsilon}(p)\, dp\, d\alpha, \tag{8}$$

where

$$G_{\varepsilon}(p) = \frac{1}{\pi\varepsilon^2}\left(1 - \frac{1}{\sqrt{1 - \varepsilon^2/p^2}}\right) \quad \text{if } |p| > \varepsilon$$

and $G_{\varepsilon}(p) = 1/(\pi\varepsilon^2)$ otherwise.

Proof. Observability of $\mathscr{R}f(L)$ follows as in the one-dimensional case. The resulting butterfly spread prices are Radon transforms of the risk-neutral distribution, which can be recovered by inverting the Radon transform as in equation (8) (see, for example, Nievergelt, 1986a,b). □

Result 2 is an idealization, of course. Practical implementations will have to deal with a finite (and small) number of "slices," along the lines illustrated in Exhibit 5. The analogy with medical practice may be helpful here, as X-ray tomography confronts precisely the same issue, which has, therefore, received considerable attention in the literature: see, for example, Shepp and Kruskal (1978).

---

[8] This is a strong assumption, though arguably closer to the contracts one observes in reality than that of Nachman (1989), who proposed an alternative that requires observability of options on portfolios of options on $S_{1,T}$ and $S_{2,T}$. While Nachman's approach is theoretically interesting, prices of the relevant options-on-options are not observable, even approximately, in practice.

[9] Result 2 can be extended to N dimensions using higher-dimensional Radon transforms (or the closely related X-ray transforms), but this extension is unlikely to be implementable in practice.

Conclusion

David Foster Wallace once told this story:

> There are these two young fish swimming along and they happen to meet an older fish swimming the other way, who nods at them and says, "Morning, boys. How's the water?" And the two young fish swim on for a bit, and then eventually one of them looks over at the other and goes, "What the hell is water?"

As it happens, he was not ruminating on financial economics at the time.[10] But the quotation captures something of the influence of Steve Ross on economists of my generation. Steve's way of looking at the world is now enshrined in thousands of papers and textbooks, and continues to organize the thinking of financial economists. His work lives on in them, and us.

[10]The story was part of his 2005 commencement speech to the graduating class at Kenyon College.

References

Alvarez, F. and Jermann, U. J. (2005). Using asset prices to measure the persistence of the marginal utility of wealth. *Econometrica*, 73(6):1977–2016.

Arditti, F. D. and John, K. (1980). Spanning the state space with options. *Journal of Financial and Quantitative Analysis*, 15(1):1–9.

Backus, D., Chernov, M., and Martin, I. (2011). Disasters implied by equity index options. *Journal of Finance*, 66:1969–2012.

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81:637–659.

Breeden, D. T. and Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651.

Carr, P. and Madan, D. (2001). Towards a theory of volatility trading. In Jouini, E., Cvitanic, J., and Musiela, M., editors, *Option Pricing, Interest Rates and Risk Management*, Handbooks in Mathematical Finance, pages 458–476. Cambridge University Press.

Kremens, L. and Martin, I. (2018). The quanto theory of exchange rates. Working paper.

Martin, I. (2017). What is the expected return on the market? *Quarterly Journal of Economics*, 132(1):367–433.

Martin, I. and Ross, S. A. (2018). Notes on the yield curve. Working paper.

Martin, I. and Wagner, C. (2018). What is the expected return on a stock? Working paper.

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183.

Nachman, D. C. (1989). Spanning and completeness with options. *Review of Financial Studies*, 1(3):311–328.

Neuberger, A. (1994). The log contract. *Journal of Portfolio Management*, 20(2):74–80.

Nievergelt, Y. (1986a). Elementary inversion of Radon's transform. *SIAM Review*, 28(1):79–84.

Nievergelt, Y. (1986b). Exact reconstruction filters to invert Radon transforms with finite elements. *Journal of Mathematical Analysis and Applications*, 120:288–314.

Ross, S. A. (1976). Options and efficiency. *Quarterly Journal of Economics*, 90(1):75–89.

Ross, S. A. (2015). The recovery theorem. *Journal of Finance*, 70(2):615–648.

Shepp, L. A. and Kruskal, J. B. (1978). Computerized tomography: The new medical X-ray technology. *American Mathematical Monthly*, 85(6):420–439.