Derivatives Pricing and Trading in Incomplete Markets: A Tutorial on Concepts

Dennis Yang
ATMIF LLC
dennis.yang@atmif.com

September 27, 2006

Published in Wilmott Magazine, January 2007, P52–62

Abstract

In this tutorial, various derivative pricing notions in incomplete markets are illustrated using a simple example, with emphasis on how to use these pricing concepts to make *systematic* trading decisions.

KEYWORDS: DERIVATIVES, INCOMPLETE MARKET

I. Introduction

This tutorial uses a computer game called *trading risky bonds* to illustrate derivatives pricing and trading in incomplete markets. In the game, a unit face value zero-coupon bond has a probability d of default before its maturity T. You, the player, are shown the current bond price p, and you then are asked to make a trading decision, *i.e.*, whether to buy or sell, and how many. To make matters as simple as possible, other idealized assumptions also apply, such as allowing short selling and leverage (borrowing), no transaction costs and no interest charges (zero interest rate).

This game is one of the simplest financial models that I can think of. In this model, the primary modeling object is the default probability d, which is specified exogenously. A risky bond is treated as a derivative in this setting, whose fair value can be "derived", which will be made more precise later on. In a financial model, it is important to distinguish what can and cannot be traded continuously. Here nothing is traded continuously. After the initial trade, the game does not guarantee that you will be able to trade the same risky bond again before its maturity.

Let me now explain the notion of complete and incomplete markets, and their relevance to derivatives pricing. Assume that we have a derivative position, if we can find a trading strategy that eliminates all risks associated with the position, *i.e.*, make the final wealth distribution a delta function, then the market is complete; otherwise it is called an incomplete market. The Black-Scholes world is a complete market, because applying the well-known delta hedging scheme makes the portfolio riskless, which means that a derivative's payoff function can be replicated using a dynamic trading strategy. Naturally a derivative's price is the value of the replicating portfolio in complete markets. Since in incomplete markets, no trading strategies exist to eliminate all risks, *i.e.*, to replicate a derivative's payoff function, we must use other means to price derivatives. It will be shown later in this tutorial that derivatives should be priced in the context of portfolio optimization.

Why bother with incomplete markets? Because they are much better representations of the real world, where a derivative position always incurs risks. The simple risky bond trading model here clearly belongs to incomplete markets. Despite the level of mathematics involved in this tutorial does not go beyond that of freshman calculus, the analysis of this simple model will reveal the richness of derivatives pricing in incomplete markets.

II. Decisions under Uncertainty

Suppose the outcomes of game A and B are Gaussian random variables with means 1.0 and 0.6, and standard deviations 1.0 and 0.5, respectively, which game do you prefer to play? Note that this is a preference question, so there is no right or wrong answer.

Why is this relevant to derivatives pricing? Because in incomplete markets, by definition, no trading strategies can eliminate uncertainty, so the final wealth is a random variable; applying different trading strategies will result in different final wealth distributions. Therefore a systematic way to rank wealth distributions will induce a natural way to rank trading strategies.

It is standard in economics and game theory to use the expected utility theory approach to rank different wealth distributions. In the expected utility framework, a wealth level w is mapped to a utility value, or happiness, by a utility function $U(w)$. The expected utility, or average happiness, is computed as $E[U] = \int U(w)\rho(w)dw$, where $\rho(w)$ is the probability density function of the final wealth (of course $\int$ should be replaced by $\sum$ if w is discrete). Using a utility function, any wealth distribution is now mapped into a real number that can be ranked.

The utility function $U(w)$ should be an increasing strictly concave function. It is an increasing function because people prefer states with more wealth; it is a concave function to model risk aversion. Note that an expected utility maximizer will decline a game of flipping a coin with equal payoffs of opposite sign. Risk aversion means that investors demand compensation for uncertainties. There are many choices for utility functions; in this tutorial, I use the exponential utility function (see Fig. 1),

$$U(w) = -\frac{1}{\gamma}\exp(-\gamma w) \tag{1}$$

where γ is a positive parameter that controls risk aversion. A larger γ means more risk averse. Section 2.4 of [1] offers some suggestions on how to choose γ in reality. However, for the purpose of this tutorial, it is not important to know what γ is. It will be seen shortly that γ and position size n always appear together as a product, so only the dimensionless size γn matters.

It is clear that multiplying a utility function by a positive constant and adding any constant to it will not affect the rankings of wealth distributions. This is called the affine transformation freedom of utility functions, which means that utility values do not have a natural scale. For the exponential utility function (1), the current wealth level w_0 does not affect future decisions, as it can be factored out as a positive multiplier. This fact is used implicitly below on the notion of indifference, in the sense that if you

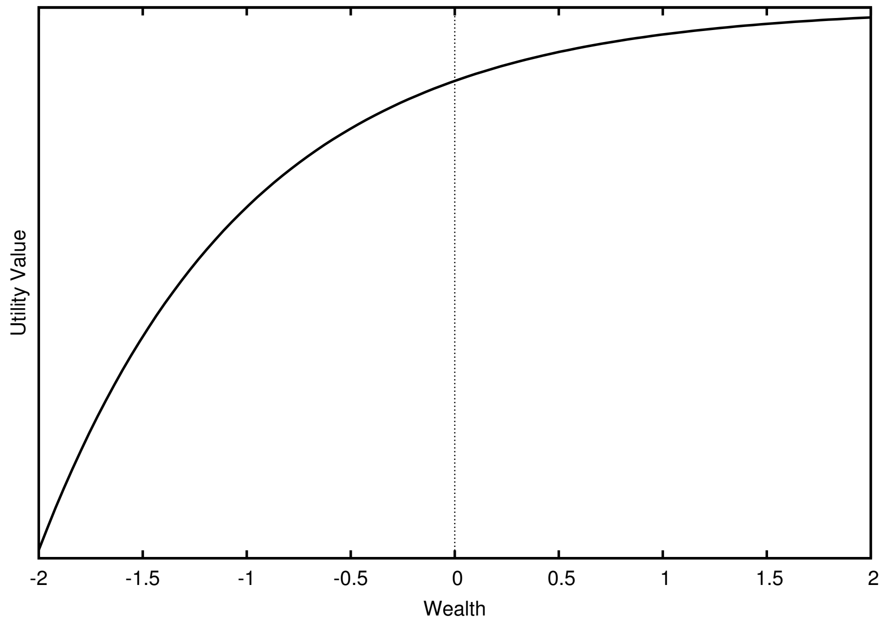

Figure 1: The exponential utility function with $\gamma = 1.0$. The vertical scale is left blank because utility values do not have a natural scale.

are indifferent between a certain payment and a random outcome, then you are also indifferent between paying that amount to receive the random outcome and doing nothing. The exponential utility function is memoryless because your past actions affect only your current wealth level, which has no bearing on your future decisions.

To summarize, once a utility function is chosen, the goal of playing any game with uncertain outcomes is to adopt a strategy that maximizes its expected value.

III. Portfolio Optimization

Let us now come back to the simple risky bond trading game. The portfolio optimization strategy here is to decide what the optimal position $\hat{n}$ is based on all the available information. Note that for mathematical convenience, a position size $\hat{n}$ is treated as a real number, instead of an integer.

Suppose that you have no initial position and the bond price is p, let w_0 be your initial wealth before paying $p\hat{n}$ for $\hat{n}$ units of the risky bond, then the final wealth without default is $w_0 + (1-p)\hat{n}$, and the final wealth with default is $w_0 - p\hat{n}$ (no recovery). Assuming your investment horizon is the same as the maturity of the risky bond, the expected utility of your final wealth is

$$E[U] = (1-d)\, U(w_0 + (1-p)\hat{n}) + d\, U(w_0 - p\hat{n})$$

where d is the default probability. Since by assumption $\hat{n}$ is optimal, the first order derivative of $E[U]$ with respect to $\hat{n}$ must be zero, which leads to

$$(1-d)(1-p)\, U'(w_0 + (1-p)\hat{n}) = dp\, U'(w_0 - p\hat{n}) \tag{2}$$

Using the exponential utility function (1), equation (2) now becomes

$$\gamma \hat{n} = \ln \frac{(1-d)(1-p)}{dp} \tag{3}$$

which gives the optimal position $\hat{n}$ to hold for a given trading price p. Notice that, as mentioned earlier, the current wealth level w_0 drops out, and that the risk aversion parameter γ and the position size $\hat{n}$ appear together as a product when using the exponential utility function. Because it is a zero interest environment, the bond maturity T does not enter the problem explicitly; however, the default probability d is implicitly T dependent.

Any market price p of the risky bond can be converted to a market implied default probability d^i by equating the winning average $(1-d^i)(1-p)$ with the losing average $d^i p$. A simple calculation shows that $d^i = 1-p$. It is clear that if the default probability d is the same as the market implied default probability d^i, then your optimal position is not to have a position ($\hat{n}=0$); you want to take a long position ($\hat{n}>0$) if $d<d^i$, and a short position ($\hat{n}<0$) if $d>d^i$. In addition to these qualitative conclusions, expression (3) provides a quantitative answer for the exact position.

For this simple trading game, expression (3) is the bottom line. So in some sense, the game is over. However, for the purpose of pricing derivatives in incomplete markets, I have just begun.

IV. Fair Value

I mentioned in Introduction that risky bonds are treated as derivatives in this simple trading game. Let me first provide a working definition for the meaning of a derivative's price, which applies to both complete and incomplete markets. Your model's fair value for a derivative is defined as follows: if the market price, or trading price, of the derivative is lower than your model's fair value, then you are a buyer, conversely you are a seller if it is higher. When the market price and the fair value agree, you stay put, *i.e.*, you are in equilibrium with the market. For this reason, the fair value is also called the equilibrium price (the two terms are used interchangeably). The fact that you stay put when the market price and the fair value agree means that the equilibrium state is optimal, otherwise you would have chosen to change your position.

What is the fair value f of the risky bond? It is given by the following formula

$$f = \frac{1-d}{(1-d) + d\exp(\gamma n)} \tag{4}$$

where n is your current position of the risky bond. Let us check whether f satisfied the working definition of the fair value. When the market price $p=f$, after substituting (4) into (3), the optimal position $\hat{n}$ equals your current position n; thus you stay put. Analogously, it is not difficult to show that if $p<f$, then the optimal position $\hat{n}$ is greater than n, hence you are a buyer; and vice versa. Therefore I have shown that the f given by formula (4) is indeed the fair value.

The following observation is *extremely* important: Formula (4) says that your fair value f depends not only on the model and its parameters, *i.e.*, the default probability d in this problem, but also on your risk preference γ and current position n. This position dependency offers a natural way to trade derivatives *systematically*, which I will address in a minute.

Where does formula (4) come from? It is obtained by inverting expression (3). The key to solve the problem is to ask yourself the following question: What should the market price be in order for my current position to be optimal? In other words, the necessary condition of the current portfolio being optimal is that the market price equals the current fair value. Hence derivatives can be priced naturally in the context of portfolio optimization.

In complete markets, derivatives can also be priced in the context of portfolio optimization. In this case, the position dependency happens to drop out of the option pricing equation—the Black-Scholes equation (see Section 3.2 of [1]). Therefore the position dependency feature of derivative pricing reveals that the underlying model market is incomplete.

In general different traders have different risk preferences and different positions. Due to risk preference and position dependency, derivative pricing in incomplete markets is only meaningful when you take the personal perspective. If you take the market perspective, then both concepts, *i.e.*, the aggregated market risk preference and the aggregated market position, are murky. However, from the personal perspective, you know, or should know, your own risk preference, and you definitely know your current position; thus there is no ambiguity.

V. How to Trade

The fair value f is position dependent and I now make the position dependency notationally explicit by writing it as $f(n)$. Suppose the current, or the pre-trade, position is n, and the trading size is m, then the post-trade position is $n + m$. As long as your fair value based on your post-trade position is not the market price, or the trading price, you are not in equilibrium with the market, you need to adjust your position, *i.e.*, trading with the market, to make them equal. Therefore for a given trading price p, the optimal trading size m is obtained by solving the local equilibrium equation

$$f(n + m) = p \tag{5}$$

which has the solution (cf. (3))

$$m = \frac{1}{\gamma} \ln \frac{(1 - d)(1 - p)}{dp} - n = \hat{n} - n \tag{6}$$

i.e., for a given price p, the optimal trading size is simply the optimal position $\hat{n}$ (post-trade) minus the current position n (pre-trade). Now you see how position dependency of the fair value offers a natural and systematic way to trade derivatives.

Let $q(m) := f(n + m)$ be the post-trade position based fair value as a function of the trading size m (cf. (4)),

$$q(m) = \frac{1 - d}{(1 - d) + d \exp[\gamma(n + m)]} \tag{7}$$

The function $q(m)$ is plotted in Fig. 2. The curve in the figure clearly demonstrates that you are a buyer ($m > 0$) when the trading price is less than the current fair value, or equilibrium price, 0.9202 (the horizontal dotted line); thus the right half of the curve can be thought as your personal demand curve. Conversely you are a seller ($m < 0$) when the trading price is greater than the current fair value, thus the left half of the curve can be regarded as your personal supply curve. Therefore the curve $q(m)$ is called the personal supply-demand curve. Note that the personal supply-demand curve depends on the pre-trade position n as well.

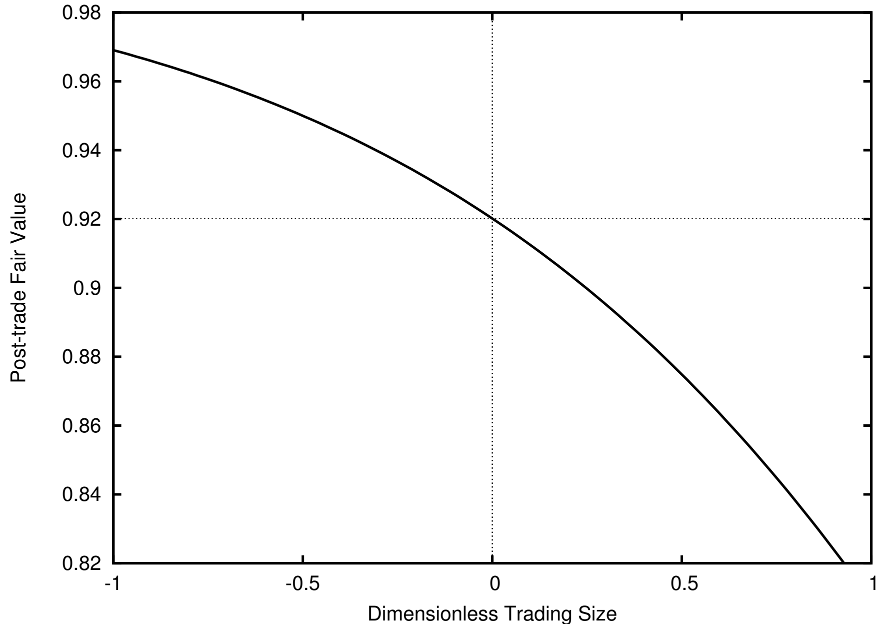

Figure 2: The personal supply-demand curve. The horizontal axis is γm. The parameters are: the default probability $d = 0.05$, and the pre-trade position $\gamma n = 0.5$.

From (7), it is easy to show that $q'(m) < 0$, *i.e.*, the personal supply-demand curve $q(m)$ is always downward sloping. This feature guarantees that trading with the market will eventually reach an equilibrium state. Because you are a buyer when the market price is lower than the fair value based on your current position; buying ($m > 0$) lowers your fair value based on the post-trade position, thus eventually your fair value will meet the market price. Similar logic works for the selling situation. The negative slope property of the personal supply-demand curve means that the inventory control mechanism, which is a natural consequence of risk aversion, is automatically built into the trading system, because buying lowers your fair value and selling raises it. This feature is important in reality, as all real-life traders use some sort of inventory control mechanism when making trading decisions.

The notion of a trading edge is defined to be the price difference between the trading price and the current fair value, *i.e.*, $p - f(n)$, with the convention that a positive trading edge is associated with selling, and the negative one with buying. The personal supply-demand curve shows that a bigger absolute trading edge $|p - f(n)|$ leads to a bigger absolute trading size $|m|$, which fits the intuition of all real-life traders.

There is another important use for the personal supply-demand curve—generating quotes. Suppose you are asked to make a bid on buying $m > 0$ risky bonds, how do you bid? You know that if you trade m units of the risky bond based on your post-trade fair value $q(m) := f(n + m)$, then you would be in equilibrium with the market after the trade, in the sense that the trading price, which can be viewed as the market price, equals your post-trade fair value. Thus the post-trade fair value $q(m)$ is a good candidate for the bid price. Similar logic applies to the ask price ($m < 0$) case. In summary, you use the fair value, or equilibrium price, based on the post-trade position as the quote price. For this reason, the personal supply-demand curve $q(m)$ is also known as the quote price curve. Pick a pair of numbers $m_b > 0$ and $m_a < 0$, at any given time, you are willing to post the following four numbers: $\{q(m_b), |m_b|\}$, $\{q(m_a), |m_a|\}$, which are the {bid price, bid size}, and the {ask price, ask size}. The action of posting these four numbers is called making a market. Therefore the quote price curve can be used to make markets.

VI. Certainty Equivalent Profit and Loss

Once a trade occurs, a natural question is how good the trade is. One obvious measure is the amount of expected utility a trader gains through the trade. Since utility values do not have a natural scale due to the affine transformation freedom, a better measure for a trade is the certainty-equivalent profit and loss (CEPL), as I now explain.

The reason to participate in a trade is that it increases your expected utility. Receiving a lump sum of money (but not doing the trade) also increases your expected utility. If you are indifferent between the trade and the amount Υ in lieu of the trade, then Υ is defined as the CEPL of the trade. To be more specific, assuming that w_0 is your current wealth excluding the n risky bonds in your portfolio, furthermore you just traded m bonds at the price of p per bond. Then your final wealth without default is $w_0 - pm + (n + m)$, and your final wealth with default is $w_0 - pm$. The final expected utility in this case is

$$E_1[U] = (1-d)\,U(w_0 - pm + n + m) + d\,U(w_0 - pm)$$

If you receive a lump sum Υ, but forfeit the trade, then your final wealth without default is $w_0 + \Upsilon + n$, and your final wealth with default is $w_0 + \Upsilon$. The final expected utility in this case is

$$E_2[U] = (1-d)\,U(w_0 + \Upsilon + n) + d\,U(w_0 + \Upsilon)$$

By definition, $E_1[U] = E_2[U]$, after substituting the exponential utility function (1) for U, the expression for Υ is

$$\Upsilon(m,p) = -\frac{1}{\gamma}\ln\frac{d + (1-d)\exp[-\gamma(m+n)]}{d + (1-d)\exp(-\gamma n)} - mp \tag{8}$$

where I have made it notationally explicit that the CEPL of a trade is a function of both the trading size m and the trading price p.

I first investigate CEPL Υ as a function of the trading price p, or equivalently the trading edge $p - f(n)$, where $f(n)$ is the current fair value (based on the pre-trade position n). There are two cases to consider: Case (i), the trading size m is fixed. Expression (8) says that Υ is a linear function of p. In the case of buying ($m > 0$), the line has a negative slope, which means the lower the price, the higher the CEPL. Conversely, in the case of selling ($m < 0$), the line has a positive slope, which means the higher the price, the higher the CEPL. These conclusions are intuitive, as shown by the two dotted straight lines in Fig. 3. Case (ii), m is optimally chosen (cf. (6)). The CEPL in this case is defined to be the optimal CEPL Υ_o, which is plotted in Fig. 3 as the parabola-like curve. The minimum of the optimal CEPL curve is zero, which occurs when the trading price p agrees with your current fair value; there is no trading in this case because the current position is optimal. As long as the nonzero trading size is determined optimally, the corresponding Υ_o is always positive, otherwise you would choose not to trade. It is clear from the optimal CEPL curve that a larger absolute trading edge also means a larger optimal CEPL, which again fits the intuition of real-life traders. The parabola-like optimal CEPL curve always lies above the two straight lines, because by definition, the optimal trading size maximizes the CEPL at a given trading price.

Notice that when $|m|$ is fixed, the CEPL is negative when the trading price is too close to the current fair value, or when the absolute trading edge is too small. This is because when the trading price equals

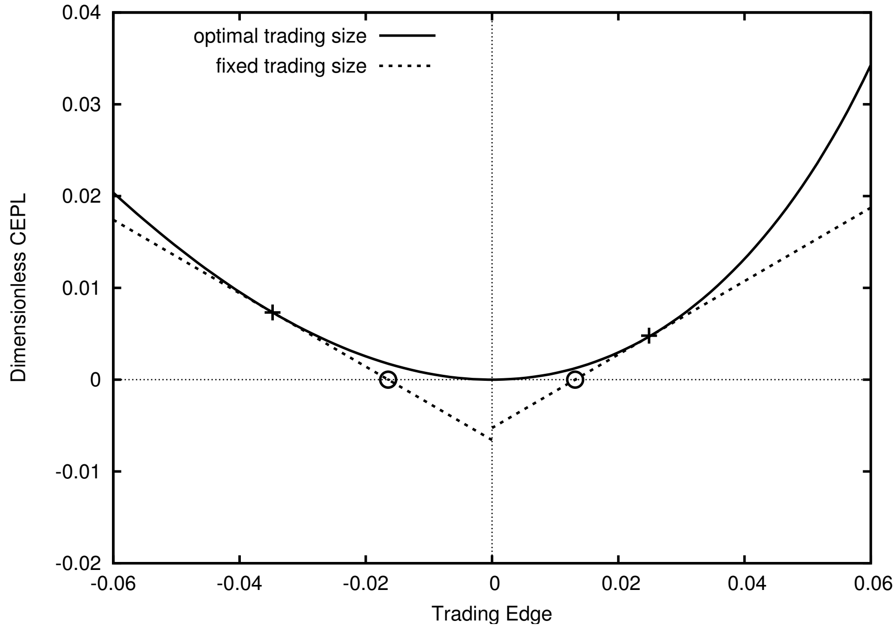

Figure 3: The dimensionless CEPL $\gamma\Upsilon$ is plotted against the trading edge $p - f(n)$. The solid curve corresponds to the case where the trading size is optimally chosen; whereas the two dotted straight lines correspond to the cases of the absolute value of the dimensionless trading size $\gamma|m|$ being fixed at 0.4. The parameters are: the default probability $d = 0.05$, and the pre-trade position $\gamma n = 0.5$.

the current fair value, your current position is optimal, changing the current equilibrium position without being compensated enough on the trading price decreases your expected utility, or equivalently produces a negative CEPL. It is evident from Fig. 3 that for the same absolute trading size, buying and selling at the current fair value may cause different damages. In the case of buying at the fixed size $\gamma m = 0.4$, as the trading price decreases from your current fair value (0.9202), your CEPL improves, until at a special price $r^b = 0.9038$ (trading edge -0.01645), which is called the buy reserve price (left small circle in Fig. 3), the CEPL is zero; *i.e.*, you break even utility-wise. Further discussion on the reserve price will be presented later. Lowering the trading price further makes your CEPL positive. At another special price $q^b = 0.8854$, which is called the buy quote price (left plus sign in Fig. 3), the line and the optimal CEPL curve touch. If trading takes place at this price, then the optimal trading size is precisely 0.4, because the trading price here is just the equilibrium price based on the post-trade position. Similar price points $r^s = 0.9334$ (right small circle) and $q^s = 0.9451$ (right plus sign) exist when selling at the fixed size $\gamma m = -0.4$.

I now consider CEPL Υ as a function of the trading size m. Again there are two cases to study: Case (i), the trading price p is fixed (at 0.8854). This corresponds to the dotted curve in Fig. 4. It is clear that selling ($m < 0$) below the current fair value (0.9202) produces a negative CEPL, which you will not do voluntarily. Instead of selling, buying a small amount results in a positive CEPL as expected; your CEPL increases with respect to your trading size, until it reaches its maximum at the optimal trading size ($\gamma m = 0.4$), after which it decreases as the trading size increases. Therefore an important message from Fig. 4 is that the optimal trading size problem under a given trading price has a unique answer. The CEPL will eventually dip below zero as the trading size becomes too large, which means that you are better off utility-wise not doing the trade. The crossover point is given a special name—the reserve trading size. In

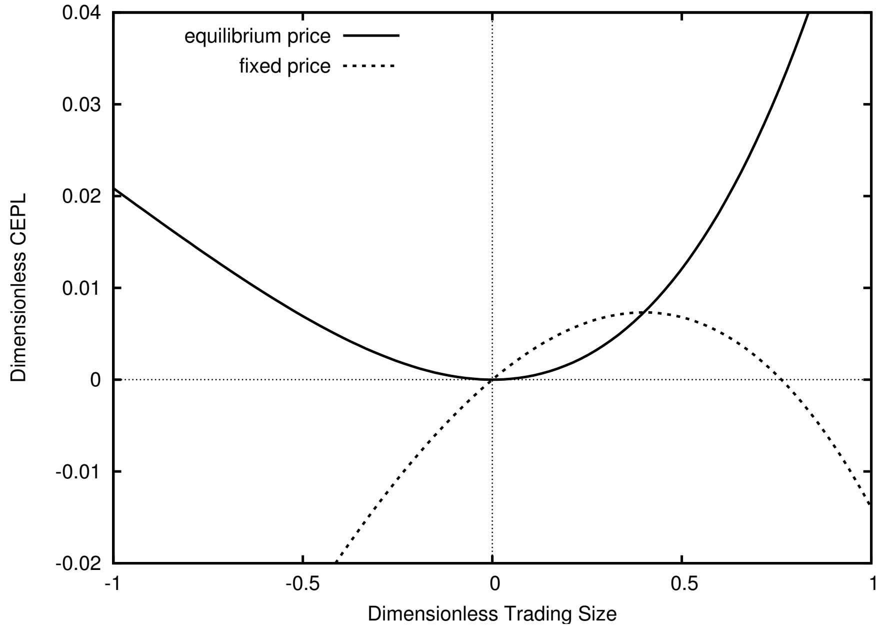

Figure 4: The dimensionless CEPL $\gamma\Upsilon$ is plotted against the dimensionless trading size γm. The solid curve corresponds to the case where the trading price is the post-trade based equilibrium price, whereas the dotted curve corresponds to the case of the trading price being fixed at 0.8854. The parameters are: the default probability $d = 0.05$, and the pre-trade position $\gamma n = 0.5$.

other words, as long as the trading size is below the reserve trading size, the CEPL is positive. The moral of the story is that you should never trade too big at a given trading price. Case (ii), p is the equilibrium price based on the post-trade position (cf. (7)). This is plotted as the solid curve in Fig. 4, which again corresponds to the optimal CEPL Υ_o, so it is nonnegative. Other than the origin, the two curves intersect at the optimal trading size (0.4). For a given trading size, the trading price for the two curves are different. When $m < 0$, the selling price on the solid curve is higher than that of the dotted curve, so the solid curve has a better CEPL. When $\gamma m > 0.4$, the buying price on the solid curve is lower than that of the dotted curve, so again the solid curve has a better CEPL. However, when $0 < \gamma m < 0.4$, the buying price on the dotted curve is lower than that of the solid curve, so that segment of the dotted curve lies above the solid one.

VII. Portfolio Indifference Price

I now pose a different question in this section. What is the most you are willing to pay to take over a position of n units of the risky bond from the initial empty portfolio? Assuming that w_0 is your initial wealth before paying h for the n risky bonds, then the final wealth without default is $w_0 + n - h$, the final wealth with default is $w_0 - h$. The final expected utility is

$$E[U] = (1 - d)\, U(w_0 + n - h) + d\, U(w_0 - h)$$

Since the goal is to maximize the expected utility, you would not do something voluntarily that reduces the expected utility, which implies that the most amount h you are willing to pay is to make the two expected utilities (before and after putting on the position) equal, *i.e.*,

$$U(w_0) = (1-d)\, U(w_0 - h + n) + d\, U(w_0 - h) \tag{9}$$

Substituting the exponential utility function into (9) leads to

$$h(n) = -\frac{1}{\gamma} \ln\left[d + (1-d)\exp(-\gamma n)\right] \tag{10}$$

The quantity $h(n)$ given by expression (10) is called the portfolio indifference price, which is an important concept, as I will show in a minute that other quantities are directly related to it. Of course, $h = 0$ when the portfolio is empty (*i.e.*, $n = 0$).

Based on expression (10), the CEPL formula (8) can be written as

$$\Upsilon(m, \mathcal{S}) = h(m+n) - h(n) - \mathcal{S} \tag{11}$$

where $\mathcal{S} := mp$ is the total amount paid (when $\mathcal{S} > 0$) or being paid (when $\mathcal{S} < 0$) to go from the pre-trade position n to the post-trade position $n + m$. This formula provides a general way to compute CEPL of a trade in incomplete markets. The reason formula (11) is valid for all incomplete market models is that it can be derived simply based on the notion of indifference. By definition, the CEPL of a trade with size m means that you are indifferent between the choices of (i) the pre-trade position n with a lump sum Υ, and (ii) paying $\mathcal{S}$ to do the trade, *i.e.*, to establish the post-trade position $n + m$. Since by the definition of portfolio indifference price, a position l and a lump sum $h(l)$ are equivalent, formula (11) immediately follows.

There is an important relation between the fair value $f(n)$ and the portfolio indifference price $h(n)$. Taking the first order derivative of $h(n)$ with respect to n, it is easy to verify from (4) and (10) that

$$f(n) = h'(n) \tag{12}$$

which is called the tangent relation. The tangent relation can also be derived from financial reasoning. Suppose the current market price of the risky bond p equals your current fair value $f(n)$, then your current position is optimal. Changing your current and optimal position by an infinitesimal amount $\epsilon \ll 1$ should leave your expected utility unchanged (to the first order), which means zero CEPL for the infinitesimal size trade, *i.e.*, $0 = h(\epsilon + n) - h(n) - \epsilon p$ (cf. (11)). This immediately leads to the tangent relation, bearing in mind $p = f(n)$.

There is another financial angle to view the tangent relation. Suppose your initial portfolio is empty, and you pay the market price p for each of the n risky bonds. The CEPL for this trade is $h(n) - pn$ (cf. (11) with $h(0) = 0$ and $\mathcal{S} = np$). Since the goal of doing a trade is to maximize the expected utility, or equivalently the CEPL, the optimal position n is determined through the equation $h'(n) - p = 0$. When your position is optimal, you are in equilibrium with the market. The local equilibrium equation says that your fair value $f = p$, which immediately leads to the tangent relation $f = h'$.

Other properties of the portfolio indifference price are discussed in [2, 3]. One especially important property is that $h(n)$ is always a concave function. The concavity of h coupled with the tangent relation immediately leads to $f'(n) < 0$, which is the negative slope result of the personal supply-demand curve. It guarantees that an equilibrium state can always be reached by trading with the market.

The concavity of h can also be used to show that the optimal CEPL is nonnegative. Recall that if the trading price equals the equilibrium price based on the post-trade position, or the trading size is optimally chosen, then the CEPL of the trade is called optimal, which can be computed as follows (cf. (11))

$$\Upsilon_o(m) = h(n+m) - h(n) - mf(n+m) \tag{13}$$

Draw a picture of a concave function and use the tangent relation $f = h'$, it is easy to see geometrically that $\Upsilon_o \geq 0$, with the equality holding when $m = 0$.

Since the quantity $h(n) - nf(n)$ is the optimal CEPL Υ_o of establishing the position n from the empty portfolio ($h(0) = 0$), the following option portfolio inequality is established

$$h(n) \geq nf(n) \tag{14}$$

In incomplete markets, which is the case here, the strict inequality holds as long as $n \neq 0$. In complete markets, the portfolio indifference price is a linear sum of fair values that are not position dependent, so the equality in (14) holds.

I now use formula (11) to show that CEPL is an additive quantity. Assuming $m = m_1 + m_2$ and $\mathcal{S} = \mathcal{S}_1 + \mathcal{S}_2$, then

$$\begin{aligned} \Upsilon(m, \mathcal{S}) &= [h(m_2 + m_1 + n) - h(m_1 + n) - \mathcal{S}_2] + [h(m_1 + n) - h(n) - \mathcal{S}_1] \\ &= \Upsilon(m_2, \mathcal{S}_2) + \Upsilon(m_1, \mathcal{S}_1) \end{aligned} \tag{15}$$

Thus the CEPL of a two-step trade is the sum of each step. The first two terms on the right hand side of (11) only depend on the final and initial positions; the last term $\mathcal{S}$ depends on the trading path, *i.e.*, intermediate trading prices. Therefore the overall CEPL is trading path dependent. A natural question is what are the best and worst paths given that the trading size is always chosen optimally at each step. Let the total trading size be m, the best path takes a single trade to reach the final state, the overall CEPL in this case is given by (13); the worst path takes infinite number of steps, of which the overall CEPL $h(n+m) - h(n) - \int_0^m f(n+y)\, dy$ becomes *zero* due to the tangent relation. This explains why in real life a dealer always wants to take the counter party out in a single big trade at the post-trade based fair value (final equilibrium price), rather than many small trades with trading prices progressing towards the final equilibrium price.

VIII. Reserve Price

I now expand on the reserve price concept mentioned earlier. The optimal trading size problem is to determine what size to trade for a given trading price. But the trading size is given in many real-life circumstances, especially for over-the-counter derivatives contracts, where there is only one trading size, either trade the whole contract or none at all. To decide whether to participate in a trade when the trading size m is prespecified, the notion of reserve price is needed.

As shown by the left dotted line in Fig. 3, which corresponds to the case of buying with a given size, the CEPL is positive only if the unit-trading price is lower than the threshold $r^b(m)$ (left small circle). Therefore the most you are willing to pay for the m bonds is $r^b(m)$ per bond. When the trading price is precisely $r^b(m)$, the CEPL of the trade is zero, thus you are indifferent between whether to buy or not.

The sell reserve price $r^s(|m|)$ is defined in a similar way. The reserve price $r(m)$ is clearly a function of the trading size m, with $m > 0$ meaning buy and $m < 0$ meaning sell. For a given absolute trading size of $|m|$ units, your corresponding buy reserve price $r(|m|)$ and sell reserve price $r(-|m|)$ form a range (see the two small circles in Fig. 3). If the unit trading price p falls within this range, then you do not wish to participate in that trade, buy or sell, because doing so lowers your expected utility.

The question now becomes how to compute $r(m)$ in general. By the definition of reserve price, after paying the amount $\mathcal{S} = m\, r(m)$ (or being paid if $\mathcal{S} < 0$) to go from the pre-trade position n to the post-trade position $n + m$, you break even utility wise, *i.e.*, zero CEPL. The CEPL formula (11) immediately leads to

$$r(m) = \frac{1}{m}\,[h(n+m) - h(n)] \tag{16}$$

In the case of trading risky bonds, the explicit expression for the reserve price is (cf. (10))

$$r(m) = \frac{1}{\gamma m} \ln \frac{d + (1-d)\exp(-\gamma n)}{d + (1-d)\exp[-\gamma(n+m)]} \tag{17}$$

Based on the notion of reserve price, the CEPL formula (11) can be rewritten as

$$\Upsilon(m, p) = m\,[r(m) - p] \tag{18}$$

which means that in order for the CEPL of the trade to be positive, you must buy ($m > 0$) below your buy reserve price, or sell ($m < 0$) above your sell reserve price. The financial interpretation of (18) is simple: If you view $r(m)$ as your break even cost per unit, then your profit of the trade is the difference of the unit prices times the total number of units being traded. Recall that if the trading price is the quote price, then the CEPL is called optimal, which can be computed with the following formula

$$\Upsilon_o(m) = m\,[r(m) - q(m)] \tag{19}$$

Since the optimal CEPL is always positive when $m \neq 0$, the buy quote price $q^b(m)$ is always lower than the corresponding buy reserve price $r^b(m)$, whereas the sell quote price $q^s(m)$ is always higher than the corresponding sell reserve price $r^s(m)$. This situation is shown clearly in Fig. 5. Notice that the two curves in the figure intersect at the current fair value with coordinate (0.0, 0.9202), which is the result of reserve price formula (16) and the tangent relation (12). The two dotted lines in Fig. 5 will be commented on later.

After buying m units of the risky bond at price p per bond, there are five possible situations: (i) If $p = q^b$, then you maintain equilibrium with the market, the CEPL of this trade is called optimal. (ii) If $p < q^b$, then your new fair value (based on the post-trade position) is still higher than p, which means you want to buy more; the CEPL of this trade is greater than the optimal value, I call it the super-optimal CEPL. (iii) If $q^b < p < r^b$, then your new fair value is lower than p, which implies you want to sell out some bonds if the trading size were divisible; the CEPL of this trade is called the sub-optimal CEPL. (iv) If $p = r^b$, then the CEPL of this trade is zero, which means you are indifferent whether to trade or not, so probably you are doing someone a favor by buying at r^b. (v) If $p > r^b$, then the trade has negative CEPL, which you do not want to do voluntarily. Obviously similar concepts are applicable to the sell situation as well.

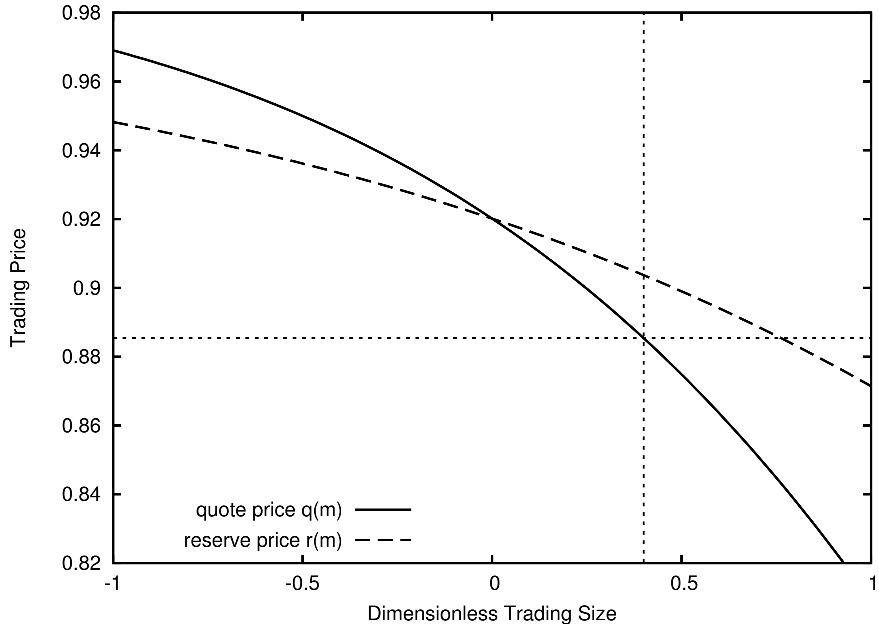

Figure 5: The trading price is plotted against the dimensionless trading size. The dashed curve is the reserve price $r(m)$ (cf. (17)); the solid curve is the quote price $q(m)$ (cf. (7)), which is the same as the one shown in Fig. 2. The two dotted straight lines represent, respectively, the situations of trading with a fixed price $p = 0.8854$ (horizontal) and trading with a fixed size $\gamma m = 0.4$ (vertical). The parameters are: the default probability $d = 0.05$, and the pre-trade position $\gamma n = 0.5$.

IX. Arbitrage Price

Recall that the personal supply-demand curve is downward sloping (see Fig. 2), which means that $q(m)$ is a monotone decreasing function of the trading size m. Now define two limits (assuming they exist) as follows:

$$a^b := \lim_{m \to +\infty} q(m) \qquad (20)$$

$$a^s := \lim_{m \to -\infty} q(m) \qquad (21)$$

The number a^b is called the buy arbitrage price, because if the risky bond were to trade at a price p that is less or equal to a^b, then you would want to buy an infinite amount. Since you are risk averse, the only possible way for this to occur is for the model to admit an arbitrage when $p \leq a^b$. Similarly, a^s is called the sell arbitrage price.

For the risky bond, it is easy to check that the buy arbitrage price is 0, and the sell arbitrage price is 1, which does not convey much useful information. In a general incomplete market model, there may not be an easy way to compute these two arbitrage prices. Fortunately for real-life traders, the buy and sell arbitrage prices are of limited use, because usually (a^b, a^s) form a wide range, and the market price p is not anywhere near a^b or a^s in normal circumstances.

X. Connecting the Dots

I now comment on the two dotted lines in Fig. 5, and their relationship with Fig. 3 and Fig. 4.

I first discuss the connection to Fig. 3, in which the left straight line is for the case of buying with the fixed dimensionless trading size $\gamma m = 0.4$. This corresponds to moving along the vertical dotted line in Fig. 5 from bottom to top. Note that the vertical axis of Fig. 5 corresponds to the horizontal one of Fig. 3. The CEPL of the trade decreases along the way because your purchasing price is continuously getting worse; recall that your CEPL is called super-optimal before you hit the quote price curve; when you are on the quote price curve, which corresponds to the left plus sign in Fig. 3, your post-trade position is optimal for that trading price. When you hit the reserve price curve, which corresponds to the left circle in Fig. 3, your CEPL crosses zero into the negative territory. Of course moving further up the line gives worse negative CEPLs. The solid curve in Fig. 3 corresponds to moving from bottom to top along the quote price curve in Fig. 5. If you wonder what the CEPL curve (such as the ones in Fig. 3) would look like when you move along the reserve price curve in Fig. 5 (the dashed one), the answer is a flat line at zero.

I now discuss the connection to Fig. 4. The dotted curve in Fig. 4 is for the case of trading at the fixed price $p = 0.8854$. This corresponds to moving along the horizontal dotted line in Fig. 5 from left to right. The CEPL of the trade is negative for $m < 0$, since selling below your current fair value cannot possibly be good. As you start to buy ($m > 0$), the CEPL crosses into positive territory. Increasing the trading size increases the CEPL, until it reaches its maximum value (optimal CEPL) when you hit the quote price curve, which gives you the optimal trading size for that trading price. Moving further towards right decreases the CEPL, until it crosses zero when you hit the reserve price curve, which gives you the reserve trading size for that trading price. The solid curve in Fig. 4 corresponds to moving from left to right along the quote price curve in Fig. 5 (the solid one).

Other than the instrument being traded, you need two quantities to describe a trade, namely trading size and trading price. The quote price curve and the reserve price curve shown in the size–price plot Fig. 5 are *extremely* important, because they are the basis for making rational trading decisions, which are summarized as follows: (i) When the trading price or the trading edge is fixed, corresponding to a horizontal line in the graph, you want the trade to occur at the optimal trading size, which is the intersection of the horizontal line and the quote price curve. You never want to trade at a size that is bigger than the reserve trading size, which is the intersection of the horizontal line and the reserve price curve. (ii) When the trading size is predetermined, corresponding to a vertical line in the graph, you aim for the trade to occur at the quote price, which is the intersection of the vertical line and the quote price curve. However if you can trade at a better price, *i.e.*, buying below your buy quote price or selling above your sell quote price, then jump on it. Never participate in a trade with a price that is worse than your reserve price, which is the intersection of the vertical line and the reserve price curve, *i.e.*, do not buy above your buy reserve price or sell below your sell reserve price.

Having explained all the relevant prices, I present a schematic drawing in Fig. 6 to illustrate their relative relationship. Here is a quick recap of various concepts involved: The current fair value, or equilibrium price, is associated with an infinitesimal trading size; both the (buy/sell) reserve prices and the (buy/sell) quote prices are associated with a *given* finite absolute trading size; and the (buy/sell) arbitrage prices are associated with an infinite absolute trading size.

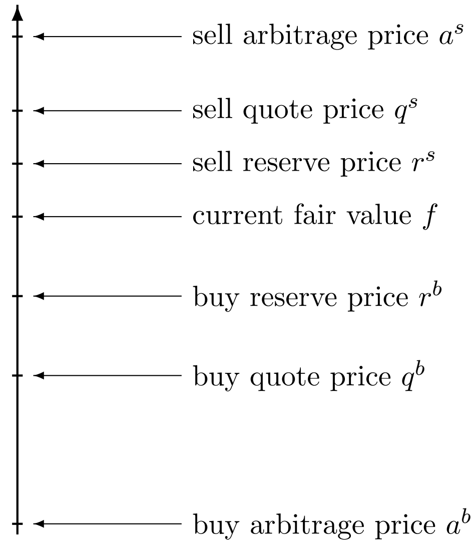

Figure 6: Schematic drawing for various prices

I now point out two features in Fig. 6. The first is that the buy and sell prices (both quote and reserve) are in general *asymmetric* around the current fair value. The second is that the reserve price $r(m)$ always lies between the current fair value $f(n)$ and the quote price $q(m)$. The ranking of various prices is intuitive with respect to their financial meanings.

XI. Mutually Beneficial Trading

In complete markets, derivatives can be priced using the arbitrage argument. The crux of the arbitrage argument is perfect replication, which means derivatives are redundant. Therefore there is no economical reason for their existence. Since derivatives do exist in the real world, we need to use an incomplete market model, in which derivatives are no longer redundant, to approximate the reality. In this section, I show that trading derivatives in incomplete markets is mutually beneficial.

Let me provide an example first. Suppose you think that a risky bond has 5% chance of default before maturity, and it is offer on the market at 0.9367. You know it is cheap, your computation (cf. (6)) shows that the optimal trading size at this price is $\gamma m = 0.25$ (assuming no pre-trade position). Naturally you go ahead and do the trade, and have gained a CEPL of $\gamma \Upsilon = 1.724 \times 10^{-3}$ (cf. (13)).

Why would anyone sell you the bond at that price? A possible rational explanation is that he thought the bond had a higher default rate. What if you are told that he shared your risk aversion parameter γ and that he also believed the bond had 5% chance of default, would you then regard him as an idiot to sell the bond below 0.95? Well, I might sell it to you below 0.95. Let me explain. My bond fair value was only 0.9202 before I traded with you, because I had an inventory (pre-trade position) of $\gamma n = 0.5$ bonds. At price 0.9367, my optimal trading size was $\gamma m = -0.25$; in fact 0.9367 was my sell quote price for selling 0.25 bonds. I have gained a CEPL of $\gamma \Upsilon = 1.994 \times 10^{-3}$ by trading with you. Therefore, I might be an idiot overall, but not on this trade.

In this example, the trade occurred because we had different initial pre-trade positions. Clearly we have both benefited from the trading (positive CEPLs). What is the economical reason behind it? The answer is risk transfer. In incomplete markets, having positions means taking risks. But any risk has a price associated with it. At price level 0.9202, I was willing to take the risk associated with the position $\gamma n = 0.5$; but that position became too risky at price level 0.9367, hence I was happy to unload some of my position to you. Note that to establish my original position, it must be that I was able to buy some of the bond at 0.9202 earlier. However, under the exponential utility function, how I established my pre-trade position had no bearing on whether or how to trade with you. After doing this single trade, the two of us have established a local equilibrium, as both of us share the same fair value for the bond (0.9367).

So far I have assumed that we both share the same risk aversion parameter γ, and the model parameter default probability. What if this is not the case? The fair value of a risky bond is position and model dependent; in general we will have different fair values for that bond based on our own model and position, *i.e.*, $f(n) \neq \tilde{f}(\tilde{n})$, where n and $\tilde{n}$ are our respective current position. I now prove that we can always establish a local equilibrium state by trading with each other; *i.e.*, there exist a trading size m_* such that

$$f(n + m_*) = \tilde{f}(\tilde{n} - m_*) \tag{22}$$

Note one of us is buying and the other is selling depending on the sign of m_*. To prove this result, define an auxiliary function g as follows

$$g(m) := h(n + m) + \tilde{h}(\tilde{n} - m) \tag{23}$$

where h and $\tilde{h}$ are portfolio indifference prices for the two of us. Note that $g(m)$ is simply the sum of the two CEPLs (omitting terms that are independent of m). Because both h and $\tilde{h}$ are concave functions, the auxiliary function $g(m)$ is also a concave function, which has a unique global maximum point m_*. Using the tangent relation $f = h'$, it is easy to see that m_* satisfies the equilibrium relation (22). Although this argument does not exclude the possibility that the unique maximum point m_* might be at plus or minus infinity, it cannot happen for this problem because (22) would have been violated (one side is the buy arbitrage price zero, while the other is the sell arbitrage price one).

Therefore whenever we disagree on a risky bond's fair value, we can always eliminate the valuation disparities by trading with each other. Since the equilibrium state is optimal, both of us gain expected utilities through trading, thus trading is a mutually beneficial, *i.e.*, it is a win-win game utility wise. In conclusion, derivatives trading in incomplete markets increases happiness of all parties involved.

XII. Static Hedging

So far there is only one type of risky bond in this trading game. Since the concept of hedging is associated with more than one financial instrument, the trading universe is enlarged from one bond to two bonds with different maturities T_1 and T_2 ($T_1 < T_2$). Of course the two bonds share the same default risk factor, with the default probability during the first period $[0, T_1]$ being d_1, and the one during the second period $[T_1, T_2]$ being d_2.

Dynamic hedging is allowed if one of the primary modeling objects can be continuously traded, *e.g.*, continuous-time geometric Brownian motion model for a stock price. Since nothing is traded continuously in this trading game, only static hedging is possible. In the sense that the game provides you an opportunity to trade now, if you do not take the opportunity, then the game may not let you trade the same instrument at later time.

Let me present a few basic formulas first before discussing static hedging. Assuming your current position has n_1 and n_2 of the T_1 and T_2 bonds, respectively, there are three possible scenarios: (i) Default occurs in the first period. The probability is d_1, and the final wealth is 0. (ii) Default occurs in the second period. The probability is $(1-d_1)d_2$, and the final wealth is n_1. (iii) Default does not occur in either period. The probability is $(1-d_1)(1-d_2)$, and the final wealth is n_1+n_2. Thus the final expected utility is

$$E[U] = d_1\,U(0) + (1-d_1)d_2\,U(n_1) + (1-d_1)(1-d_2)\,U(n_1+n_2)$$

By the definition of portfolio indifference price h, we have $U(h) = E[U]$. Substituting the exponential utility function (1) leads to

$$h(n_1,n_2) = -\frac{1}{\gamma}\ln D(n_1,n_2) \tag{24}$$

where D is defined as

$$D(n_1,n_2) := d_1 + (1-d_1)\exp(-\gamma n_1)[d_2 + (1-d_2)\exp(-\gamma n_2)] \tag{25}$$

Notice that $h(n_1,n_2) \neq h(n_1,0) + h(0,n_2)$, so h is a quantity associated with the whole portfolio, which is not a simple sum of its parts in incomplete markets.

Using the tangent relation (12), the fair-value formulas for the T_1 and T_2 bonds are

$$\begin{aligned} f_1(n_1,n_2) = \frac{\partial h}{\partial n_1} &= \frac{1}{D(n_1,n_2)}(1-d_1)\exp(-\gamma n_1)[d_2+(1-d_2)\exp(-\gamma n_2)] && (26)\\ f_2(n_1,n_2) = \frac{\partial h}{\partial n_2} &= \frac{1}{D(n_1,n_2)}(1-d_1)(1-d_2)\exp[-\gamma(n_1+n_2)] && (27)\end{aligned}$$

where $D(n_1,n_2)$ is given by (25). Let p_1 and p_2 be the trading prices for the T_1 and T_2 bonds, respectively. The local equilibrium equation says that the optimal post-trade positions $\hat{n}_1$ and $\hat{n}_2$ must satisfy

$$\begin{aligned} f_1(\hat{n}_1,\hat{n}_2) &= p_1 && (28)\\ f_2(\hat{n}_1,\hat{n}_2) &= p_2 && (29)\end{aligned}$$

These two equations are easy to solve, the solution is

$$\begin{aligned} \gamma\hat{n}_1 &= \ln\frac{d_2(1-d_1)(1-p_1)}{d_1(p_1-p_2)} && (30)\\ \gamma\hat{n}_2 &= \ln\frac{(1-d_2)(p_1-p_2)}{d_2 p_2} && (31)\end{aligned}$$

The optimal trading sizes for the two bonds are $m_1 = \hat{n}_1 - n_1$ and $m_2 = \hat{n}_2 - n_2$. Notice that because $T_2 > T_1$, the T_2 bond is more likely to default, so we must have $p_1 > p_2$, otherwise there is a trivial arbitrage by buying the T_1 bond and shorting the T_2 bond.

I now investigate the scenario of making trading decisions on the T_1 bond, while static hedging with the T_2 bond. This could be the case because the T_1 bond is illiquid, whereas the T_2 bond is more liquid. I will compare the following three situations: (i) no hedging with the T_2 bond; (ii) the T_2 bond has zero bid-ask spread; (iii) the T_2 bond has finite bid-ask spread. It is assumed that you can trade sufficient numbers of the T_2 bond at the bid or ask price. The following parameters are used to generate the figures

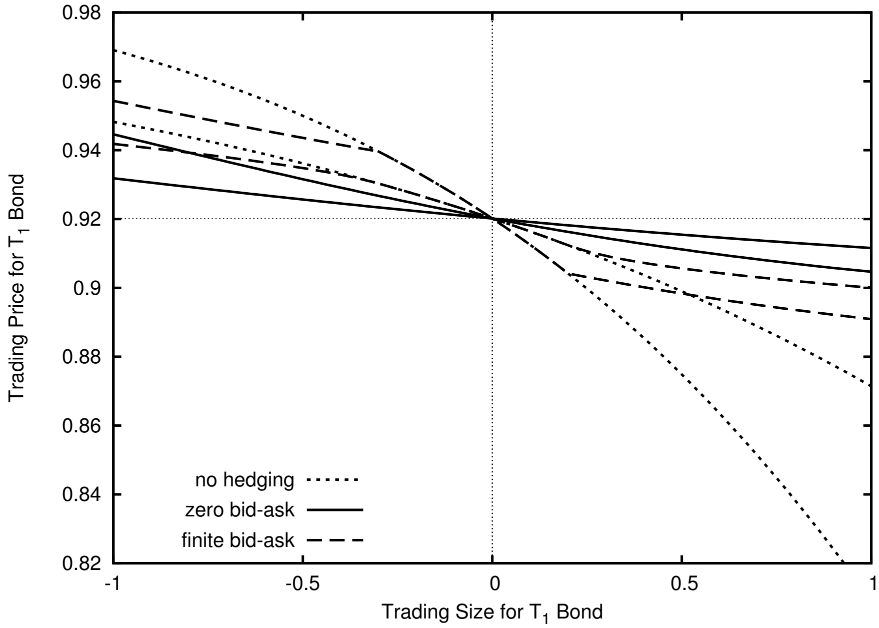

Figure 7: The trading price is plotted against the dimensionless trading size for the T_1 bond. There are three sets of quote and reserve price curves. The dotted ones are for the situation of no hedging; the solid ones are for static hedging with zero bid-ask spread; the dashed ones are for static hedging with finite bid-ask spread. Within each set, the quote price curve always has a larger negative slope.

of this section: the default probabilities are $d_1 = 0.05$ and $d_2 = 0.03$; the pre-trade position is $\gamma n_1 = 0.5$ and $\gamma n_2 = 0$; in the case of zero bid-ask spread, the T_2 bond is trading at $p_2 = 0.8925$, which is the current fair value of the T_2 bond based on the pre-trade position; in the case of finite bid-ask spread, the bid and ask prices of the T_2 bond are $p_{2b} = 0.8770$ and $p_{2a} = 0.9114$, respectively.

The most important things to know with respect to making rational trading decisions on the T_1 bond are its quote and reserve price curves. In situation (i) of no hedging, the quote price curve is $q_1(m_1) = f_1(n_1 + m_1, n_2)$; the reserve price curve is $r_1(m_1) = [h(n_1 + m_1, n_2) - h(n_1, n_2)]/m_1$. This set of curves is plotted in Fig. 7 as the dotted ones, which are the same as the ones in Fig. 5. In situation (ii) of zero bid-ask spread, the quote price curve is $\bar{q}_1(m_1) = f_1(n_1 + m_1, n_2 + m_2)$, where the optimal hedging size m_2 is solved from the local equilibrium equation $f_2(n_1 + m_1, n_2 + m_2) = p_2$, with p_2 being the hedging price for the T_2 bond. By definition, if a trade occurs at the reserve price, the CEPL of the trade is zero, thus $0 = h(n_1 + m_1, n_2 + m_2) - h(n_1, n_2) - m_1\bar{r}_1(m_1) - m_2 p_2$, which means that the reserve price curve is $\bar{r}_1(m_1) = [h(n_1 + m_1, n_2 + m_2) - h(n_1, n_2) - m_2 p_2]/m_1$. This set of curves is plotted in Fig. 7 as the solid ones. It is clear that the slopes of the solid curves are much smaller than the dotted ones, which means for a given trading edge, the optimal trading size is much bigger when static hedging is allowed; alternatively for a predetermined trading size, you demand much less price concession under static hedging. These conclusions are intuitive.

I now study situation (iii) of static hedging with finite bid-ask spread. I will focus on the case of buying ($m_1 > 0$) T_1 bond first. Buying the T_1 bond lowers the fair value (based on the post-trade position) of the T_2 bond (cf. (27)). As long as the fair value of the T_2 bond is higher than its bid price, no static hedging takes place, because doing so (selling below the current fair value) lowers the expected utility. At the critical trading size $\gamma m_{1b} = 0.2$ of the T_1 bond with the corresponding critical trading edge $p_{1c} - f_1(n_1, n_2) = -0.01598$, the fair value of the T_2 bond equals its bid price, *i.e.*, $f_2(n_1 + m_{1b}, n_2) = p_{2b} = 0.8770$. Increasing the buying size of the T_1 bond further causes the static hedging to occur by triggering the selling of the T_2 bond at its bid price. Thus the quote and reserve price curves of the T_1 bond with $m_1 > 0$ consist of two segments—no hedging first and then static hedging. The selling case is completely analogous: at the critical trading size $\gamma m_{1a} = -0.3$ of the T_1 bond with the corresponding critical trading edge $p_{1c} - f_1(n_1, n_2) = 0.01944$, the fair value of the T_2 bond equals its ask price, *i.e.*, $f_2(n_1 + m_{1a}, n_2) = p_{2a} = 0.9114$. Therefore the dashed curves in Fig. 7 corresponding to the finite bid-ask spread static hedging case is a mixture of no hedging and then static hedging, with the middle segments of the dashed curves coincide with those of the no hedging case. When the bid-ask spread narrows, the no-hedging region becomes smaller, so the dashed curves approach the solid ones; when the bid-ask spread widens, the no-hedging region becomes larger, so the dashed curves approach the dotted ones.

Other than the quote and reserve price curves, the optimal CEPL plots in Fig. 8 are also of interest, which are obtained from formula (19). Since there is a one-to-one relationship between the optimal trading size m_1 and the trading edge $p_1 - f_1(n_1, n_2)$ by way of the quote price curve, the horizontal axis of the left panel can be mapped into the one of the right panel. The two panels in the figure seem to provide different conclusions, as the left one says that the optimal CEPL for the no hedging case is the highest, whereas the right one says the opposite, *i.e.*, static hedging with zero bid-ask spread case is the best. The apparent contradiction is the result of comparing apples with oranges, as the two horizontal axes represent different quantities. For a fixed trading size (left panel), the optimal CEPL for the no hedging case is the best simply because it requires the larges price concession in order for the trade to occur. It is intuitive that at a given trading price, static hedging with zero bid-ask spread gives the best optimal CEPL, otherwise you would have chosen not to do static hedging.

Another interesting quantity to know is the optimal hedging size γm_2 for the T_2 bond, which is plotted in Fig. 9. The curves in the left panel have negative slopes, which mean buying the T_1 bond requires the selling of the T_2 bond to hedge, and vice versa. The right panel says the same thing, since by convention, a negative trading edge is associated with buying, and a positive edge is for selling. This is obvious, but it explains why the trading size under static hedging is much bigger than that of no hedging, because you are essentially trading a spread in the former case. It is known to real life traders that spreads are usually traded with much bigger sizes. Note that as long as the trading size for the T_1 bond is given (left panel), the optimal hedging size for the T_2 bond does not depend on the trading price of the T_1 bond due to the memoryless property of the exponential utility function, but it does depend on the hedging price of the T_2 bond. The key point is that for a given hedging price, without any guesswork, the system automatically produces the optimal hedge ratio.

The opposite scenario of making trading decisions on the T_2 bond while static hedging with the T_1 bond can be investigated analogously. This will be left as an excise for the readers.

XIII. Conclusion

In this simple trading game, almost all quantities of interest have simple analytical expressions, which allow you to have a hands-on experience. But for more complicated incomplete market models, usually no analytical formulas are available. For example, in diffusion type of models, the fair value f and the

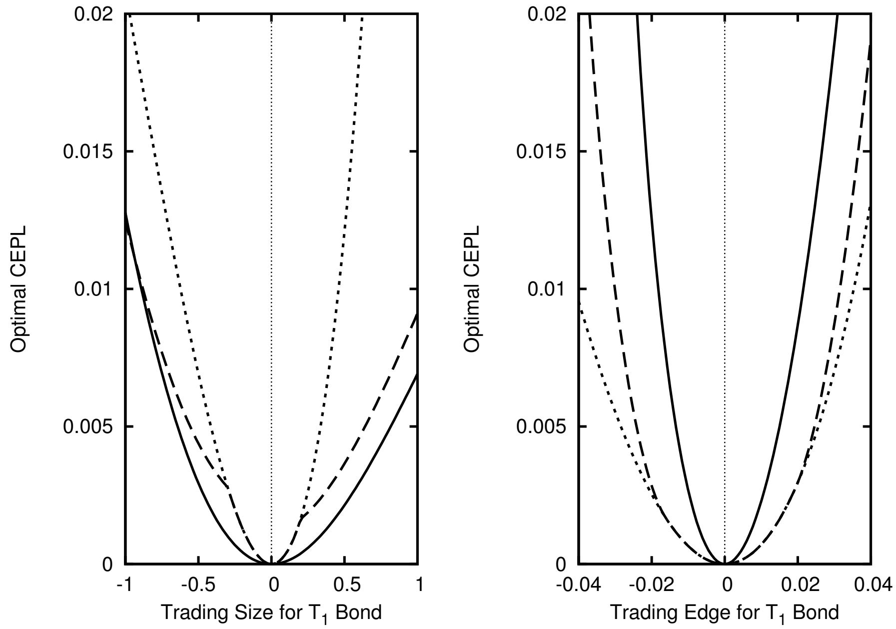

Figure 8: The dimensionless optimal CEPL $\gamma\Upsilon_o$ is plotted against the dimensionless trading size γm_1 (left panel), and against the trading edge $p_1 - f_1(n_1, n_2)$ (right panel). The dotted curve is for no hedging; the solid one is for static hedging with zero bid-ask spread; the dashed one is for static hedging with finite bid-ask spread.

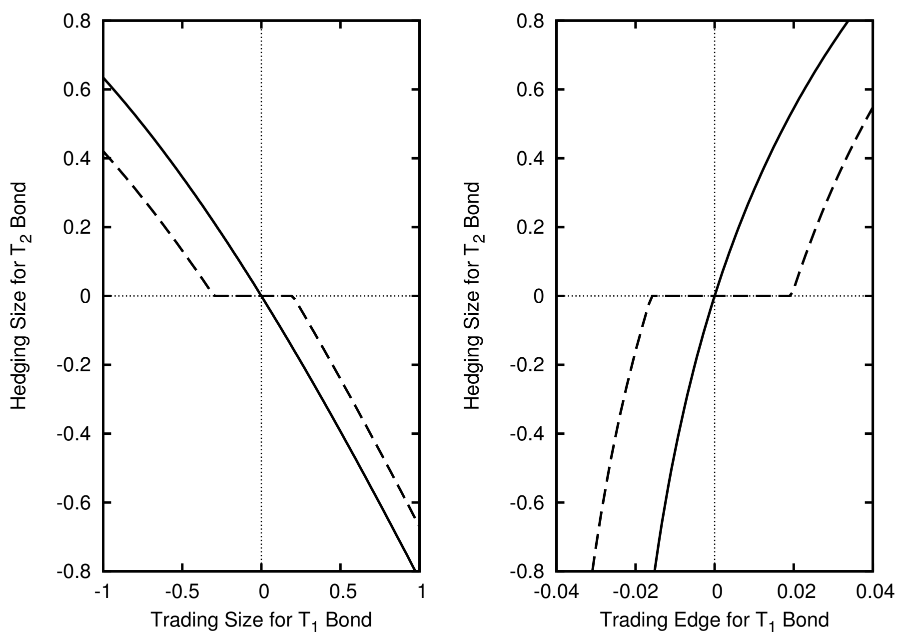

Figure 9: The dimensionless optimal hedging size γm_2 for the T_2 bond is plotted against the dimensionless trading size γm_1 (left panel), and against the trading edge $p_1 - f_1(n_1, n_2)$ (right panel). The solid curve is for static hedging with zero bid-ask spread; the dashed one is for static hedging with finite bid-ask spread.

portfolio indifference price h, which are needed to compute quote and reserve price curves, satisfy a pair of PDEs [1, 4]. This pair of PDEs needs to be solved numerically in general, as the one for h is nonlinear. Despite that the way to compute f and h may vary based on different incomplete market models, the basic concepts and *qualitative* features described in this tutorial remain valid. If you go behind the analytical expressions in this tutorial to think about their financial meanings, you will discover that many conclusions presented here are a natural consequence of risk aversion in incomplete markets.

Other than various derivative pricing concepts and their usages in incomplete markets, the main points covered in this tutorial are:

- derivatives should be priced in the context of portfolio optimization;
- derivatives pricing is preference and position dependent in incomplete markets, which is only meaningful from the personal perspective;
- the position dependent pricing offers a natural and systematic way to trade derivatives;
- derivatives trading in incomplete markets is mutually beneficial.

References

[1] Dennis Yang (2006): *Quantitative Strategies for Derivatives Trading*, ATMIF, New Jersey. Excerpt available at http://www.atmif.com/qsdt

[2] Vicky Henderson and David Hobson, *Utility Indifference Pricing–An Overview*, Volume on Indifference Pricing, ed. R. Carmona, forthcoming. Available at http://orfe.princeton.edu/˜vhenders/preprints.html

[3] Aytac İlhan, Mattias Jonsson and Ronnie Sircar, *Portfolio Optimization with Derivatives and Indifference Pricing*, Volume on Indifference Pricing, ed. R. Carmona, forthcoming. Available at http://www.princeton.edu/˜sircar/

[4] Dennis Yang, *A Simple Jump to Default Model*, Available at http://www.atmif.com/papers