

Transaction Costs and the Pricing of Financial Assets

George M. Constantinides
University of Chicago and NBER

I would like to thank the officers of the Multinational Finance Society and the organizers of its 4th annual conference for bringing us together in the historic city of Thessaloniki to discuss research developments in finance. Specifically, I would like to recognize the President of the Society, Geoffrey Booth, President-elect, George Philippatos, Chairman of the Board of Trustees, Panayiotis Theodossiou, Program Chair, Nickolaos Travlos, and Program Co-chair, Angelos Tsaklanganos. They richly deserve a round of applause.

The theme of my address is *transaction costs and their impact on the pricing of financial assets*. In various forms, this has been a theme of my research over the course of the past several years. Oftentimes, in models of the financial markets we abstract from market imperfections in order to keep the model tractable and hope that this abstraction does not seriously impair the realism of the model. Thus, we abstract from bid-asked spreads, brokerage fees, execution costs, illiquidity (that is, lack of market depth), borrowing and short-selling restrictions and fees, and from the absence of certain markets which are needed to insure financial risk and labor income risk. How crucial is this abstraction? This is a very broad question and I do not intend to address it here in its full generality. I will focus on just one form of market imperfection: bid-asked spreads, brokerage fees and execution costs, collectively referred to as *transaction costs*. Other forms of market imperfections will be considered only indirectly and to the extent that they contribute to transaction costs.

* Leo Melamed Professor of Finance, Graduate School of Business, University of Chicago, 1101 East 58th Street, Chicago, Illinois 60637, e-mail <gmc@gsb.uchicago.edu>. I thank Yacine Ait-Sahalia and John Cochrane for helpful comments.

(*Multinational Finance Journal*, 1997, vol. 1, no. 2, pp. 93–99)
© by *Multinational Finance Society*, a nonprofit corporation. All rights reserved.

Transaction costs affect equilibrium asset prices and, therefore, affect the mean return of assets, gross of transaction costs. I define the *liquidity premium* as the difference of the mean rate of return of two assets, one of which is subjected to transaction costs, and the other not, but which are, otherwise, substitutable assets. How big is the liquidity premium?

Transaction costs also affect the prices of derivatives such as call and put options. With transaction costs, the concept of the *no arbitrage* price of a derivative is replaced by a *range of prices* such that one would be able to exploit a mis-pricing, net of transaction costs, only if the price of the derivative were to fall outside this range. How wide is the range of unexploitable price deviations in the presence of transaction costs?

Right from the start, we need to distinguish between the concepts of *endogenous* trading and *exogenous* trading. In endogenous trading the frequency and volume of trade are determined by the investor's optimal trading policy in the presence of transaction costs, that is to say, the policy which balances the costs against the benefits of trading in order to maximize the investor's *expected utility*.

It is hard to give a general definition of exogenous trading, but an example illustrates the point. Suppose that the one-way proportional transaction costs rate is 1%. If a portfolio manager is forced, as a matter of fund policy, to turn over a stock in the portfolio once every quarter, the annual transaction cost is 8%. Furthermore, if this portfolio manager is the *marginal investor* in this stock, it commands an 8% liquidity premium. If, on the other hand, the portfolio manager is forced to turn over this stock only once every 20 years, the annual transaction cost is only one-tenth of 1% and is negligible. Furthermore, if this portfolio manager is the *marginal investor* in this stock, it commands a negligible liquidity premium.

To focus the discussion, I put forth the following hypothesis: *In those cases where the trading of the marginal investor is endogenous, being the outcome of an optimal trading policy in the presence of transaction costs, transaction costs have a trivial effect on the liquidity premium and on the permissible deviation of a derivative's price from its no-arbitrage value.* This is not to say that transaction costs have a trivial effect on the trading activity. On the contrary, *investors deflect the impact of transaction costs by modifying drastically the frequency and size of their trades.*

I examine the hypothesis through a series of models in which an

investor considers the decision to trade in the presence of proportional transaction costs. None of the models allows for asymmetric information. In fact, the impact of transaction costs in cases where the agents are asymmetrically informed is largely uncharted territory and I will have very little to say about this case.

As my first example with endogenous trading, I consider an economy populated by investors, each maximizing expected utility and trading in order to re-balance the portfolio thrown out of balance by capital gains on its component assets. For simplicity, I assume that there are only two assets, a riskless one, such as a money-market account, and a risky one, such as an index fund. Transactions in the riskless asset incur no transaction costs while transactions in the risky one incur proportional transaction costs. I also assume that the borrowing and lending rates are equal, and that there are no limits to borrowing and selling the fund short. I considered a model along these lines in my paper titled “Capital Market Equilibrium with Transaction Costs” and found that the liquidity premium on the risky asset is about one-tenth of the proportional transaction costs rate. Essentially, investors accommodate even large transaction costs by drastically reducing the frequency and volume of trade because *the expected utility loss in holding an imbalanced portfolio is small*. I conclude that transaction costs have a negligible effect on the liquidity premium if the motive for trade is the rebalancing of a portfolio thrown out of balance by capital gains.

I enrich the above example by introducing the life-cycle motive for trade, where investors save during the years of their peak wage earnings and dis-save upon their retirement. Vayanos (1998) considered a model along these lines in a general equilibrium context. He found that for a realistic transaction costs rate the liquidity premium on stocks is negligible. This is because the round-trip transaction costs in trading the stocks are amortized over a long holding period and the equilibrium liquidity premium is negligible. I conclude that proportional transaction costs have a negligible impact on the liquidity premium even when we introduce the life-cycle motive to trade.

I broaden the scope of the example by introducing random labor income fluctuations as a third motive for portfolio rebalancing. Aiyagari and Gertler (1991) and Heaton and Lucas (1996) considered models along these lines and showed, as before, that transaction costs have a negligible effect on the liquidity premium of the risky asset. This is because the investor smooths consumption by borrowing or

lending and tolerates large deviations of the portfolio from its optimal composition in order to deflect the impact of transaction costs. Can we generate a substantial liquidity premium by introducing additional market imperfections which force the investor to trade the risky asset frequently and incur in the process substantial transaction costs? We can make it hard for the investor to borrow by either forbidding borrowing outright or by introducing a wedge between the borrowing and lending rates. In this case we find that investors evade the need to trade the risky asset (and incur transaction costs) by keeping a positive balance in their money-market account and smooth their consumption by drawing down this balance. As before, transaction costs induce only a minuscule liquidity premium on the risky asset unless one assumes an unrealistically low supply of the riskless asset in the economy. In all of these examples the first part of the hypothesis holds up—transaction costs have a trivial effect on the liquidity premium when trading is endogenous.

I now examine the impact of transaction costs when the motive for trade is the dynamic hedging of a derivative. Recall that in the Black-Scholes setting, if the *market price* of a derivative differs from its Black-Scholes *theoretical value*, the investor buys the underpriced derivative or writes the overpriced one and perfectly hedges the position by dynamic trading, thereby realizing as an arbitrage profit the difference between the market price and the theoretical value. The dynamic trading policy incurs an *infinite volume of trade* over the lifetime of the derivative. This is just fine in the Black-Scholes model because transaction costs are assumed away. But if there are proportional transaction costs, *however small the proportional transaction costs rate may be*, the total transaction costs of the dynamic trading policy are infinite. Thus, it appears that transaction costs may have a major impact on derivative prices, no matter how small the transaction costs rate may be.

Here we need to draw a major distinction between *exchange-traded* derivatives, *over-the-counter*, *plain vanilla* derivatives, and *client-customized* derivatives. Examples of exchange-traded derivatives are the European calls and puts on the S&P 500 index which are traded on the Chicago Board Options Exchange. Their relevant characteristics are that they are standardized, are traded by numerous buyers and sellers, and are traded in small and large orders. Examples of over-the-counter, plain vanilla derivatives are the FX and interest-rate, plain vanilla swaps. Their relevant characteristics are that they are standardized and

are traded by numerous buyers and sellers. Examples of client-customized derivatives are exotics and structured notes issued by dealers such as investment banks. Their relevant characteristics are that they are nonstandard, are traded between two institutional parties, have prices often in the millions of dollars, and have payoffs often in the hundreds of millions of dollars.

I will argue that a sizable fraction of the trades in exchange-traded derivatives is *endogenous* and, therefore, the permissible deviation of a derivative's price from its no-arbitrage value is small. By contrast, trades in client-customized derivatives are primarily *exogenous* and, therefore, the permissible deviation is large.

Consider first exchange-traded derivatives. A sizable fraction of the buyers and writers of these derivatives *have exposure in the payoff of the option which is relatively small compared to their net worth*. Therefore, it is plausible to ascribe to them a utility function and view their trades as endogenous trades. With endogenous trading, it turns out that the investor accommodates transaction costs by drastically reducing the frequency and volume of trade. Naturally, this results in imperfect hedging and exposes the investor to risk. How much does the option price deviate from its theoretical, no-arbitrage value before an investor can increase utility by trading the option, given the risk imposed by the imperfect hedging policy?

I addressed this question in a series of papers.¹ In the paper entitled "Transaction Costs and the Volatility Implied by Option Prices," I find that when the proportional transaction costs rate is 1%, the permissible deviation of a European call price from its theoretical, no-arbitrage value is just a few percent of the theoretical value. The result holds even with relative risk aversion coefficient as high as six and irrespective of whether the option is in or out of the money. In "Bounds on Option Prices in an Intertemporal Setting with Proportional Transaction Costs and Multiple Securities," similar results are obtained for American and path-dependent options. These results provide a theoretical justification for the generally held view in the derivatives industry that observed deviations of the market prices of exchange-traded call and put options from their Black-Scholes theoretical price, commonly referred to as the volatility smile or skew, are not accounted for by transaction costs.

Thus it appears that the second part of the hypothesis holds up

1. See Constantinides (1998) and Constantinides and Zariphopoulou (1998a,b).

also—*transaction costs have a trivial effect on the permissible deviation of derivatives from their theoretical, no-arbitrage value.* I stress that the hypothesis works in this case because this is an example of *endogenous trading* where the investor chooses the frequency and size of the hedging trades to maximize expected utility in the presence of transaction costs. Essentially, the investor finds it optimal to deviate substantially from the Black-Scholes dynamic hedging policy in order to contain the transaction costs despite the loss of expected utility due to imperfect hedging.

Over-the-counter, plain vanilla derivatives are traded between two institutional parties, have prices often in the millions of dollars, and have payoffs often in the hundreds of millions of dollars. Because of the sheer size of their payoff, we can no longer argue that the dealer has sufficient reserves to afford to leave the book exposed to substantial risk through endogenous trading. However, *because these derivatives are standardized*, a dealer may oftentimes hedge a long position in such a derivative by taking a short position in a similar derivative in a transaction with a third party. Thus, the dealer may have to hedge only the residual exposure of the book to risk and this may be done with low transaction costs. Transaction costs play only a minor role in the pricing of these derivatives, but the reason is the ease of hedging the book rather than endogenous trading. Nowadays, the bid-asked spread on plain vanilla swaps is just a few basis points

Unlike the case with over-the-counter, plain vanilla derivatives, a long position in a *client-customized derivative* cannot typically be hedged by taking a short position in a similar derivative transaction with a third party. The reason is that these derivatives are not standardized and it is difficult to match them in pairs. The dealer in this case is obliged to hedge the derivative on a *stand alone* basis with tight exposure limits to the various sources of risk, such as delta and vega risk, and incurs substantial transaction costs in the process. This is an instance of *exogenous* rather than *endogenous* trading. I conclude that client-customized derivatives need not conform to the hypothesis and may trade at a price substantially above or below the theoretical value. In practice, client-customized derivatives do trade at prices substantially above or below the theoretical value.

So far I have assumed that investors have homogeneous information and are fully rational. A disturbing observation is that under these assumptions one comes nowhere close to explaining the observed volume of trade on the stock exchanges and the observed volume of

trade and open interest on the futures and options exchanges. Clearly something else must be going on. Therefore, we are led to examine models in which investors have or believe that they have different information than other investors. In such models, investors trade primarily to exploit their private information. What is the impact of transaction costs in this case? I speculate that by decreasing the frequency and volume of trade, transaction costs decrease the informativeness of prices and possibly raise the liquidity premium. I also speculate that fewer investors collect information on assets traded with large transaction costs, compounding the information asymmetry on these assets and possibly raising their liquidity premium. How important are these effects in practice? I do not know. But it seems appropriate to conclude my address with a question: *Do transaction costs have a major or trivial impact on the liquidity premium in situations where the motive to trade is private information?*

References

- Aiyagari, S. R. and Gertler, M. 1991. Asset returns with transactions costs and uninsured individual risk. *Journal of Monetary Economics* 27: 311–31.
- Constantinides, G. M. 1986. Capital market equilibrium with transaction costs. *Journal of Political Economy* 94: 842–862.
- Constantinides, G. M. 1998. Transaction costs and the volatility implied by option prices. Working paper. Chicago: University of Chicago.
- Constantinides, G. M. and Zarihopoulou, T. 1998a. Bounds on prices of contingent claims in an intertemporal economy with proportional transaction costs and general preferences. Working paper. Chicago: University of Chicago and University of Wisconsin, Madison.
- Constantinides, G. M. and Zarihopoulou, T. 1998b. Bounds on option prices in an intertemporal setting with proportional transaction costs and multiple securities. Working paper. Chicago: University of Chicago and University of Wisconsin, Madison.
- Heaton, J., and Lucas, D. J. 1996. Evaluating the effects of incomplete markets on risk sharing and asset pricing. *Journal of Political Economy* 104: 443–87.
- Vayanos, D. 1998. Transaction costs and asset prices: A dynamic equilibrium model. *The Review of Financial Studies* 11: 1–58.