

## RESEARCH STATEMENT

**Timothy C. Johnson**

April 2010

My research studies the determinants of asset prices and their dynamic properties. While my papers address a wide range of questions, they share a common motivation to try to simplify our understanding of market phenomena. Our field often rewards research that highlights the unexpected, anomalous, and difficult. As a result, we find ourselves awash with empirical puzzles and conflicting interpretations. I try to look for ways to understand complex patterns using familiar tools.

Beyond this, my choice of topics is also strongly motivated by current issues in the real world. I work on some purely theoretical and some purely empirical problems, although I always prefer to find ways to combine the two. In the following, I will divide my work into three categories:

- I. The stock market and the real economy,
- II. The cross-section of returns,
- III. Derivatives.

In each area, I describe my papers to date, and also discuss some on-going projects.

### I. The stock market and the real economy.

Several of my papers attempt to learn lessons about the real economy (technology, preferences, or information) from patterns in the stock market.

Recently, in **Inequality Risk Premia** (presented at the American Finance Association meetings in 2010), I proposed that market returns may be informative about the degree to which people are jealous. More technically, if utility depends on peer comparisons, then consumption externalities arise, which can be socially destructive. Understanding the strength of these effects, and conditions under which they are amplified, is important for many areas of economics and policy analysis. An important specific sub-question is whether people view relative and absolute consumption as substitutes or complements. Put simply (and somewhat loosely), does rising wealth make society more jealous or less?

The insight of the paper is that the stock market may contain information about this question. If relative well-being is subject to systematic shocks that investors care about, then asset risk premia will reflect the demand for insurance against

such shocks. Modeling this requires heterogenous agents with a general preference specification and an economy with incomplete markets. The model implies relationships between the expected return on assets that can hedge shifts in inequality and the strength of relative terms in utility. Further, the time-variation in these expected returns should be informative about the degree of complementarity. Using a long time series of inequality measures, the paper presents evidence consistent with the view that relative consumption concerns are significant, and perhaps countercyclical. The latter conclusion supports the conjecture that relative and absolute well-being may be substitutes, i.e., that rising wealth alleviates negative externalities.

In **Optimal Learning and New Technology Bubbles** (published in the *Journal of Monetary Economics* (2007)), I address some important issues that arose out of the experience of the bull market at the end of the 20th Century. Many commentators interpreted this episode as reinforcing two commonly held views of technological revolutions: first, that there is an unavoidable link between major innovations and speculative excess; and second, that the stock market is responsible for this link.

I argue, to the contrary, that what appears to be speculative excess may be due to *active learning* and that the market reflects, rather than causes, this pursuit. The paper views new, revolutionary technologies as characterized by uncertainty about returns to scale, and hence about the ultimate size of the new industry. Faced with this learning problem, overinvestment is the optimal way for a representative agent to acquire information. In effect, overshooting (relative to the full-information allocation) at the early stage of the industry enables the economy to converge to the long-run solution faster since more investment provides more information about returns to scale. An observer might think the overinvestment is stock market-driven since the price of installed capital is higher when the learning incentive is higher (because the average return on capital exceeds the marginal return), i.e., when the overinvestment is going on. The pattern of prices thus looks like a bubble, and *ex post* excess returns look predictably negative. The model suggest that the conventional view of such episodes as wasteful misallocation by irrational markets may be seriously incomplete. Taking some extra risk (by erring on the side of overinvesting) may be the socially efficient way to build a new industry.

Next, in a sequence of papers, I develop a theory of the nature and causes of fluctuations in marketwide liquidity. This line of inquiry was motivated by the growing sense, among both practitioners and academics, that fluctuation in market fragility – liquidity *risk* – is too little understood.

**In Dynamic Liquidity in Endowment Economies** (*Journal of Financial Economics* (2006)), I introduce a new quantification of market liquidity based on the slope of the representative agent's demand curve for risky assets. The advantage of this measure is that it allows one to deduce properties of liquidity which derive from primitive features of the economy and not from specific details of the market microstructure. In particular, agents will be more willing to accommodate exogenous trade demands (and hence markets will be more liquid) when such trades have less impact on the risk of their consumption stream or on their discount rates. These impacts themselves can vary for a variety of reasons in different models. For example, a small increase in risky asset holdings may matter less when that asset's dividends are less volatile or account for a smaller share of current consumption. In several example economies, I illustrate these mechanisms and show that my measure of liquidity is also reasonably large and time-varying in ways that may match important features of aggregate data. The point of the paper, though, is not to test particular model, but to show that a meaningful notion of liquidity exists irrespective of the institutional arrangements and frictions that describe the actual trading process. With this, one can view the empirical facts about liquidity dynamics as important moments for any model to match. Further, the proposed characterization of liquidity offers a tool for studying many specific questions concerning the interaction between market conditions and other state variables.

**Liquid Capital and Market Liquidity**, the next paper in this series (*Economic Journal* (2009)), shows one such application. The work was specifically motivated by the apparently successful experience of central banks in affecting the resilience of the stock market through provision of liquidity to financial intermediaries.

Here the use of the word "liquidity" is different: it describes the availability of instruments (cash or negotiable assets) to agents in the economy. However, the question at issue is how this type of liquidity might affect *stock market liquidity* – the willingness of agents to accommodate the transaction demand of others.

One theory is that market liquidity is determined by intermediaries (market makers and arbitrageurs) with limited capital and financial constraints, whose constraints can be affected by central bank actions. While not disputing this possibility, my work points out another channel through which the two types of liquidity are linked. Specifically, when there is more transformable capital in the economy, it acts as a buffer stock that cushions consumption streams from external shocks. Thus discount rates (and marginal utility) are affected less by an (exogenous) trade of risky securities (for cash) when buffer stocks are plentiful. Hence the price impact of such a trade will be small. Thus asset markets will be liquid when the real liquidity of the economy is high.

The theory suggests that inefficient financing frictions may not be the sole, or even the most important, factor giving rise to the linkage between market fragility and the quantity of liquid capital. Beyond market liquidity, the theory presents a general equilibrium depiction of how available capital affects all aspects of the stock market. The degree of real liquidity in the economy can be a fundamental endogenous state variable. Decreases in real liquidity depress asset prices, raise risk premia, and increase volatility – for reasons that have nothing to do with the actions of monetary authorities or financial intermediaries.

A third paper in this series uses my theory of market liquidity to clarify the relation between liquidity and trading volume. **Volume, Liquidity, and Liquidity Risk** (*Journal of Financial Economics* (2008)) presents a simple model to illustrate the important economic distinction between these two market traits. In a nutshell, volume is high when there is *flux* in the characteristics (or just the number) of active participants, whereas liquidity is high when the risk-bearing capacity of those participants is high. This view contrasts with the conventional understanding of liquidity as a service provided by a segmented industry (intermediaries). However the conventional view has difficulty accounting for the fact that, in the data, there is *not* a strong positive association between volume and liquidity. My story thus resolves an empirical puzzle. Moreover it has a unique new implication: that volume should be positively associated with the *second moment* of liquidity. I verify this assertion using data for both the stock market and the government bond market.

A corollary of the liquidity theory developed in these papers is that net order flow (transaction demand) will affect market prices, not because of some temporary price pressure, but due to equilibrium responses to risk. Just how important this mechanism is at the aggregate level was unclear, however, since there have been few ways to measure either order flows or illiquidity for the stock market as a whole. This gap is addressed in **Market Liquidity and Flow-Driven Risk** with Prachi Deuskar (currently under second round review at *Review of Financial Studies*). This is an empirical project which exploits a unique data set provided to us by the Chicago Mercantile Exchange containing a virtually continuous record of orders and quotes for the most important stock market index product: the so-called e-mini futures on the S&P 500. The headline result of this study is that, not only does order flow affect prices, but that both the flow and its impact are huge. What we call *flow-driven risk* accounts for more than 40 percent of the variation in the stock market. Thus even buy-and-hold investors are exposed to “liquidity risk” in the sense that unpredictable exogenous order flow combined with high price impact move the market as much as information that is unrelated to order flow.

The paper also introduces new, directly observable measures of *ex ante* expected price impact (or market illiquidity), which allows us to study liquidity dynamics at high frequency during an extremely interesting period (2006-2008). We find that the most important determinant of illiquidity is the non-flow driven component of return volatility, consistent with aggregate risk driving the willingness to accommodate trade. Overall, we view the study as presenting a major challenge to asset pricing theory – how to quantitatively account for the extreme degree of flow-driven risk that we document.

While the recent financial crisis provides much motivation for continued study of market fragility, the period that immediately preceded the crisis also witnessed stock market sensitivity to a different type of stress. The extreme rise in the prices of many commodities raised significant questions about the exposure of aggregate wealth to the risk of increased scarcity.

In **Commodity Dependence and Aggregate Risk** (submitted to *American Economic Review*), I develop a simple model to try to understand some of the joint properties of commodity and stock market prices, and, at the same time, to delineate the real consequences of the economy’s need for raw materials. In the model, commodities are used for both production and consumption, and in each usage there is some scope for substitution with other inputs. Indeed, the description of the economy boils down to the elasticities of these two substitutions, along with the elasticity of intertemporal substitution. By appropriate choice of these numbers, the model, stylized though it is, can account for many dynamic patterns observed in the data. For example, commodity price bubbles arise as increased relative scarcity induces accelerated reallocation from production to consumption. At the same time, scarcity can induce the *sign* of the stock market’s exposure to commodity prices to reverse – a feature recently documented in work by Yilin Zhang, one of our doctoral students.

Using the parameters suggested by the observed price dynamics, the paper goes on to quantify the consequences of commodity dependence for the real economy. Perhaps surprisingly, the volatility of the effective consumption basket does not increase with increased scarcity. Nor does scarcity choke off growth. In fact, on average, outcomes with *lower* commodity prices pose more perceived risk (in terms of marginal utility) than ones with higher prices. Welfare calculations show that simply lowering the usage of commodities in production has little impact. Yet, crucially, there are potentially large benefits to increasing the economy’s ability to substitute away from commodities in response to future scarcity.

## II. The Cross-Section of Returns

A number of my papers have grappled with the wide variety of “anomalies” reported in the empirical literature on the cross-section of stock returns. My approach has been to take the findings seriously (even though some may be spurious), and to look for ways to understand or, at least, simplify them.

**Unifying Underreaction Anomalies** (co-authored with Andrew Jackson and published in the *Journal of Business* (2006)) illustrates this motivation by merging two whole classes of anomalies. The paper’s basic claim is that momentum effects and long-horizon post-event drift are fundamentally the same phenomenon. Momentum refers to the tendency of winner stocks to outperform loser stocks, and post-event drift refers to the finding of long-run outperformance or underperformance of firms following a variety of corporate actions (buybacks, takeovers, etc). If the paper is right, there is only *one* anomaly to be explained here, and not several. The starting point is the observation that, almost without exception, event studies examining price performance following good corporate news (when the initial stock reaction was positive) tend to find positive long-run drift, while those examining bad news find the reverse. That suggests the hypothesis that return continuation is a generic phenomenon, following *all* news. And *that* suggests that, lumping all types of news together, one would find momentum: the best performing firms in a given event-window will be the ones which have had *some* good corporate news, and hence ones that will be drifting upwards. Our empirical approach is to identify generic good news for a given stock by increases in its expected earnings, as measured by consensus forecasts of Wall Street analysts. Things get a bit tricky econometrically because these “forecasts” themselves are predictable and do not represent true expectations. But once their innovation component is correctly extracted, and once the return continuation attributable to this component is subtracted from stock returns, the resulting residuals show no signs of momentum or post-event drift.

Some anomalies can be accounted for simply by the failure to take into account the role of leverage. In **Forecast Dispersion and the Cross-Section of Expected Returns** (*Journal of Finance* (2004)), I point out that, for a given level of debt and systematic risk, firms with greater idiosyncratic risk will have *lower* expected excess returns to equity. This is just because adding idiosyncratic volatility raises the option component of equity value, hence lowering its expected returns. The immediate application of this observation is to the finding that firms about whom analysts disagree more do worse subsequently than firms about whom analysts agree (Diether, Malloy, and Scherbina, (2002)). My paper postulates that analyst disagreement is a measure of the noisiness of signals about a firm’s fundamental risk,

and that the noise is not itself systematic. That theoretical observation then implies that the resulting anomaly should be entirely accounted for by the interaction of this risk with leverage. Empirical tests bear out this prediction. This paper was nominated for the Smith-Breeden prize in 2004.

A similar framework is employed in **Rational Momentum Effects** (*Journal of Finance*, (2002)). This paper considers the implication of stochastic changes in expected growth rates of firms' dividends. Stock prices are actually more sensitive to these growth rates than to dividends themselves, in the sense that the price function is super-convex (its log is convex) in the former but only linear in the latter. This extreme sensitivity means that a firm's expected growth rate exposure rises with the level of its growth rate. If, in addition, expected growth rate risk is priced, then firms which have had positive shocks to expected growth rate will also have higher subsequent expected returns. The result is momentum: sorting stocks on past performance also sorts them on recent shocks to expected growth rate, which thus sorts on expected return. The paper goes on to explain the sensitivity of momentum profits to measurement horizon and holding period by generalizing the basic model to allow shocks of differing persistence to expected growth rates. The generalized version is able to match the empirical magnitudes of momentum profits *and* the sensitivity of these to the time-scales. This paper received a Smith-Breeden Distinguished Paper award.

A new project with Dirk Hackbarth attempts to clarify the role that real flexibility plays in determining firms' risk exposure. A number of recent papers in the literature have proposed that several cross-sectional anomalies (in particular, the "value effect") may ultimately stem from irreversibility of firm investments. The idea is sensible: when profitability declines, firms that cannot adjust their scale will subject stockholders to greater risk of future losses. Thus expected returns will be highest when firms are most constrained by irreversibility, and this is when valuations are lowest (and book-to-market is highest). On the other hand, the reasoning suggests that similar return patterns might be associated with measures like profitability or other measures of operating leverage, which they are not. We argue that this may be due to firm heterogeneity. In a simple model, we show that some firm investment technologies can lead to *decreasing* risk as profitability declines and operating leverage increases. We propose industry-based measures of operational flexibility. Empirical tests indicate that, consistent with the model, the sensitivity of expected returns to expected profitability decreases in more adjustable business sectors.

### III. Derivatives

The final group of papers consists of studies inspired mostly by real-world experience and interaction with practitioners and students. Both my previous employment and my current teaching center around derivatives.

**Insider Trading in Credit Derivatives** (written with Viral Acharya and published in *Journal of Financial Economics* (2007)) is an empirical attempt to isolate proxies for the likelihood of informed trade in this market, and then to assess whether this likelihood has any effect on market prices or liquidity. We argue that this setting is ideal for examining how markets cope with information asymmetries because there are identifiable insiders (lending banks) with both access to non-public information and a strong incentive to exploit it. To the extent that this market is serving an important risk-transfer function, moreover, protecting its orderly functioning may be a valid regulatory objective. We demonstrate that the number of relationship banks for a given entity is a significant determinant of the degree of information flow from credit markets to equity markets. This number then becomes our measure of the likelihood of insider activity. However we find no role for this measure in determining either the credit spread (i.e. the cost of bankruptcy insurance) or the associated bid/ask spread (i.e. liquidity). It appears, then, that the market is able to cope with potentially severe information asymmetries.

A follow-up paper is **More Insiders, More Insider Trading: Evidence from Private Equity Buyouts** (also with Viral Acharya and forthcoming in *Journal of Financial Economics*). The empirical strategy of the prior paper exploited the fact that insider trading was related to the number of insiders. We were tempted to use this observation to interpret what was reported anecdotally and in the business press to be an unprecedented increase in insider trading during the private equity buy-out wave of the middle 2000s. This period also saw changes in the financing structure of takeovers, which included “clubbing” of multiple private equity firms to bid for larger targets as well as greater use of broad syndicates of debt providers. However, before leaping to conclusions, we asked whether economic logic, did, in fact, support the idea that more widely held information would be more widely exploited. The answer, it turns out, depends crucially on the enforcement regime. Perhaps surprisingly, there is almost no literature analyzing insider trading under conditions in which it is explicitly penalized. So we developed a model of competitive insiders facing a regulator with limited enforcement powers. We describe conditions, which we argue are not implausible, under which the efficient strategy for the regulator is one in which more insiders do not necessarily trade more than fewer insiders. Intuitively, the regulator will want to make marginal trading by

each insider raise the probability that all of them are caught.

Returning to the market context, the theoretical logic is especially important in formulating the right hypothesis, because the markets in question (unlike the CDS market studied in the prior paper) *are* subject to active enforcement. We therefore adopt the null that the enforcement regime is efficient in the sense of the model, so that bids involving more insiders should *not* exhibit more insider trading. We then gather information on the number of debt and equity providers for a large sample of buy-outs of public companies and also develop a number of measures of suspicious pre-bid trading in the targets' securities (stock and bonds) and on derivatives referencing these (options and CDS). The data do confirm the intuitive conjecture (and thus reject the null) that deals with wider financing syndicates are associated with more suspicious pre-bid trade. We interpret this cautiously as evidence for less-than-optimal regulation. More broadly, the work suggests a methodology for analyzing other regulatory enforcement games, which my co-author and I continue to explore.

Finally, the two papers that grew out of my thesis were concerned with anomalies in options markets. As a general matter, the equilibrium implications of volatility patterns are under-researched. In the foreign exchange market, for example, a well-known empirical regularity is that, following a sustained strengthening of a currency, the tails of the risk-neutral distribution get fatter in the direction of further strengthening. In **Return Dynamics when Persistence is Unobservable** (*Mathematical Finance*, (2001)), I present a model of exchange-rate determination which accounts for this, without recourse to any time-varying risk premia. Instead, the model views agents as having to infer whether recent shocks to fundamentals (in this case, money deflated by consumption) are persistent or transitory. This involves characterizing the solution to a non-linear filtering problem, but the intuition is simple. A given shock is more likely to have been permanent when there are subsequent surprises in the same direction, i.e. a realized trend. Following a trend, then, the likelihood that current shocks are persistent is also higher, which means their impact on prices will be higher. Hence, the resulting price dynamic can be summarized by the simple rule that volatility increases when trends continue. This, in turn, implies that the distribution of future returns at any one time is likely to be skewed in the direction of recent trends. **Volatility, Momentum and Time-Varying Skewness in Foreign Exchange Returns** (*Journal of Business and Economic Statistics*, (2002)) is a companion empirical work. Here I fit the continuous time model to currency data using Bayesian diffusion estimation techniques. The forward-looking distributions implied by the fitted model do a good job of explaining the time-series of moments extracted from options distributions, even though no options data is used in the estimation.