Markus K. Brunnermeier
December 2007

MAJOR THEME OF RESEARCH

My research studies financial crises and significant mispricings due to institutional frictions, strategic considerations, and behavioral trading. My current, past and future work can be grouped into three closely related lines of research: (i) asset price bubbles, (ii) liquidity crises and (iii) behavioral economics.

Asset Price Bubbles

A bubble is said to exist if the price of an asset substantially exceeds its fundamental value. Notorious examples of bubbles include the Dutch tulip mania of the 1630's, the South Sea bubble of 1719-1720, the Japanese real estate bubble of late 1980's, and more recently the Internet bubble which peaked in March 2000.

My research on bubbles started with my book **"Asymmetric Information: Bubbles, Crashes, Technical Analysis and Herding"** (Oxford University Press, 2001). The book surveys and critiques the literature on bubbles and herding in settings in which different investors hold different information. One conclusion of the book is that it is difficult to obtain bubbles in models in which all market participants are fully rational.

While even proponents of the "efficient market hypothesis"[1] admit that not all market participants are fully rational, they argue that bubbles can not persist since well-informed sophisticated investors will undo the price impact of behavioral traders. That is, rational investors should have an incentive to go against the bubble even before it emerges. The work on limits to arbitrage challenges this view. My theoretical research **"Bubbles and Crashes"** (Econometrica, 2003) with Dilip Abreu offers a new mechanism for why rational traders might prefer to ride the bubble rather than attack it. We argue that timing the bubble is a difficult task. Since a single trader alone cannot burst the bubble, he faces the following trade-off: if he attacks the bubble too early, he forgoes profits from the subsequent run-up caused by momentum traders; if he attacks too late, he will suffer from the subsequent crash. Thus, rational traders have to coordinate their actions over time. Each trader tries to forecast when other rational traders will go against the bubble. Timing other traders' moves is difficult in our model because traders become sequentially aware of the bubble and they do not know where in the queue they are. Because of this "sequential awareness", it is never common knowledge that a bubble has emerged. That is, it might be the case that everybody knows of it and that everybody knows that everybody knows it and so on but this "knowledge iteration" can never be made infinitely often. It is this lack of common knowledge which removes the bite of the standard backwards induction argument. Hence, in equilibrium, all traders ride the bubble for a number of periods, even though they know that the price exceeds the fundamental value.

The second important message of our work is that relatively insignificant news events can trigger large price movements. While this phenomenon is empirically well documented, it was up to now not well understood. This phenomenon occurs quite naturally in our model because even unimportant news events allow traders to synchronize their sell strategies. Of course, large price drops are themselves significant synchronizing events and hence can trigger a whole price cascade.

[1] The semi-strong form of the efficient market hypothesis asserts that asset prices fully and correctly reflect all pertinent information in the public domain.

My article **"Hedge Funds and the Technology Bubble"** (Journal of Finance, 2004) empirically examines the response of top hedge funds to the growth of the technology bubble using previously unexplored data on hedge fund stock holdings. We find that hedge funds, which probably come closest to the ideal of unconstrained rational traders, did not exert a correcting force on stock prices during the technology bubble. Instead, they were heavily invested in technology stocks and were riding the bubble rather than shorting these stocks. Hedge funds understood that prices would eventually deflate: they captured the upturn, but skillfully avoided much of the downturn by reducing their long exposure on a stock-by-stock basis when prices were about to decline. These empirical findings question the efficient markets notion that rational speculators always immediately stabilize prices by shorting overpriced assets. The study also reveals that short-sales constraints, emphasized in recent work on the technology bubble, are not sufficient to explain the failure of rational speculative activity to contain the technology bubble. Short-sale constraints and arbitrage risks alone can rationalize reluctance to take short positions, but do not explain why sophisticated investors would buy into the overpriced technology sector.

While in "Bubbles and Crashes" behavioral momentum traders continuously push up the price until the bubble bursts, both price and fundamentals follow a random walk in **"Synchronization Risk and Delayed Arbitrage"** (Journal of Financial Economics, 2002) with Dilip Abreu. Assets can be over- or underpriced. Price corrections are significantly delayed because of synchronization risk (each trader is uncertain when other traders will trade against the mispricing) and because traders incur holding costs for non-balanced positions. This framework has many applications. For example, it can explain the persistence of mispricing and sudden price adjustments in the bond market and the foreign exchange (FX) market. Currently, I am pursuing an empirical project with Gabriele Galati that verifies our theory using proprietary high frequency data on FX carry trades. Investors are willing to "carry" an overvalued currency if it yields a higher interest rate even though it might collapse as soon as a synchronizing event occurs. Going against this currency is costly since one foregoes the interest differential. We cover episodes like the previously unexplained 10% dollar-yen collapse on October 7th and 8th, 1998 that occurred without a corresponding shift in fundamentals.

The contributions of the working paper **"Clock Games: Theory and Experiments"** (with John Morgan) are twofold. First, it introduces the notion of clock games. Clock games are timing games in which the timing of other key players' moves is random and depends on de-synchronized clocks. We consider games with a finite number of strategic players, vary the observability of opponents' moves and consider a setting with information clustering. By doing so, we are able to study a rich class of situations in which players' timing decisions are crucial and each player neither wants to be the first nor the last to move. Potential cross-disciplinary applications range from asset price bubbles to the introduction of new products to the starting of revolutions. Second, the paper tests the empirical predictions of this model in the laboratory environment with a novel experimental design. Our main finding in the experiments is that the key comparative static properties of the model are largely borne out in the choice behavior of the subjects. Most notably, strategic delays increase with the uncertainty about other players' timing; and subjects immediately follow the first mover when moves are observable.

There was a sharp run-up in housing prices from the late 1990s through 2005 in almost all major countries. My paper **"Money Illusion and Housing Frenzies"** (Review of Financial Studies, 2007) with Christian Julliard provides a possible explanation for the link between inflation and housing prices. Unlike in the stock market, in the real estate market it is obvious that rational traders cannot undo the price impact of behavioral investors. In this paper we study whether the behavioral bias of "money illusion" can explain housing price movements. Agents that suffer from money illusion do

not properly take into account the fact that a decline in nominal interest rates due to a drop in inflation does not reduce future real mortgage payments. This leads to a non-linear increase in their evaluation of housing. After empirically decomposing the price-rent ratio into a rational component and an implied mispricing component, we find that inflation and the nominal interest rate explain a large share of the time-series variation of the mispricing. We also find that the run-up in the housing prices starting in the late 1990s is reconcilable with the contemporaneous reduction in inflation and nominal interest rates during that period. Moreover, we reject the currently predominant view that inflation effects are due to the so-called tilt effect.[2]

My recent working paper **"Carry Trades and Currency Crashes"** (joint with Stefan Nagel and Lasse Pedersen) documents that carry trades are subject to crash risk, i.e. exchange rates between investment and funding currencies are negatively skewed. We argue that this negative skewness is due to a sudden unwinding of carry trades and increases when global volatility, as measured by the VIX, rises. Using price data from FX risk-reversals (option) contracts – a measure of implied skewness – we document that after a crash the price of insurance against a further crash sharply increases, despite the fact that a subsequent crash is less likely. Finally, we empirically document that funding currencies (like the Japanese Yen or Swiss Franc) commove with each other. Similar comovement is also observed for investment currencies.

Liquidity Crises – Market Microstructure

While bubbles can persist because traders are reluctant to sell overvalued shares, liquidity crises occur when prices drop and nobody wants to step in to buy shares. A prominent example was the liquidity crisis during August and September 1998 which was triggered by the financial difficulties of the hedge fund Long Term Capital Management (LTCM). Market liquidity dried up exactly when LTCM was forced to liquidate its positions. In our article **"Predatory Trading"** (Journal of Finance, 2005), Lasse Pedersen and I argue that this lack of market liquidity was not a coincidence but was due to predatory trading by other strategic traders. If one trader needs to sell, others also sell and only subsequently buy back the asset. This induces endogenous time-variations in liquidity, leads to price overshooting, and a reduced liquidation value for the distressed trader. Hence, the market is illiquid when liquidity is most needed. Even worse, predatory trading can also trigger another trader's crisis, and the crisis can spill over across traders and across markets leading to systemic risk. Our model provides a new framework for studying the strategic interaction among large traders and has important implications for risk management and for financial regulators. It advocates for "dealer exit stress tests" and argues in favor of coordinated actions by regulators, bailouts and granularity in disclosure rules for asset holdings of large hedge funds.

Our article **"Market Liquidity and Funding Liquidity"** (Review of Financial Studies, forthcoming) shows the interplay between a security's market liquidity – i.e., the ease of trading it – and traders' funding liquidity – i.e., their availability of funds. Traders provide market liquidity and their ability to do so depends on their funding, that is, their capital and the margins charged by their financiers. In times of crisis, reductions in market liquidity and funding liquidity are mutually reinforcing, leading to a liquidity spiral. Our model provides a natural and cohesive explanation for several empirically well documented features of market liquidity, such as why it (i) can suddenly dry up and trigger a liquidity crisis, (ii) has commonality across securities, (iii) is related to volatility, (iv) experiences

[2] The tilt effect arises for mortgages with a fixed nominal repayment schedule. In a high inflation environment the repayment schedule is, in real terms, tilted towards early repayments. This constrains the size of mortgages and hence, reduces housing demand.

"flight to liquidity" events, and (v) comoves with the market. Finally, the model shows how the Fed can improve current market liquidity by committing to improve funding in the event of a future crisis.

Recently, I worked on a piece titled **"Deciphering the 2007 Liquidity Crunch"**. While the estimated losses in subprime mortgages of US$ 200 to 300 billion seem to be very large, they are relatively modest when put into perspective. For example, it roughly corresponds to a not so uncommon drop between 1% and 2% of the U.S. stock market. In this light it seems surprising that the mortgage crisis has caused such turmoil in financial markets. To understand this, I study various amplification mechanisms that explain how modest shocks can cause large dislocations. A crucial concept in this research is the emergence of liquidity spirals. At times of crisis when asset prices and market liquidity drops, funding requirements for financial institutions increase. This is the case since the collateral value of the underlying assets erodes and margins rise. Higher margins force financial institutions to cut back on leverage, exacerbating the initial price decline and so forth. Another amplification mechanism is the risk of financial gridlock. This can emerge when financial institutions are lenders and borrowers at the same time and each individual institution is not able to pay its obligation only because the others are not paying theirs. While coordination might resolve this, this may be difficult in today's complex and interweaved financial system.

My recent draft paper **"Hedge Fund Tail Risk"** (joint with Tobias Adrian) studies the tail dependency of hedge fund returns and (predictable) spill-over effects to the banking sector using quantile regressions.

Two of my earlier papers touch upon liquidity, in the sense that asymmetric information and market regulation affect liquidity and market efficiency. Figuring out the correct price impact of a public news announcement is difficult. This is specially the case if some of the news might have leaked beforehand to some insiders. In this case, an unknown part of the information content is already reflected in the pre-announcement price. My paper **"Information Leakage and Market Efficiency"** (Review of Financial Studies, 2005) models such a situation (with sequential signal arrival) and derives the optimal trading strategy for traders who receive some information early. It also shows that even though information leakage enhances informational efficiency in the very short-run, it lowers it in the long run. The analysis provides support for Regulation Fair Disclosure which was recently promulgated by the Securities and Exchange Commission.

My paper **"Disclosure Requirements and Stock Exchange Listing Choice in an International Context"** with Steven Huddart and John Hughes (Journal of Accounting and Economics, 1999) studies competition among exchanges in listing standards. Exchanges require listed firms to regularly disclose company information. Some argue that competition among exchanges leads to a "race to the bottom" in disclosure standards. Our paper argues to the contrary. Exchanges have no incentive to undercut listing standards since this would scare uninformed liquidity traders away and with it corporate insiders who benefit from trading with them.

Behavioral Finance – Economics and Psychology

My research on bubbles primarily focused on mechanisms that explain why rational sophisticated agents do not undo the price impact of behavioral agents. The next logical step – taken in my third line of research – is to model the behavioral distortions directly, ideally without sacrificing the useful advances made in economics. The neoclassical framework in economics assumes that all agents

maximize a well defined objective function (utility function), aggregate outcomes are described by an equilibrium, and expectations are rational.

The **"Optimal Expectations"** framework (American Economic Review, 2005), jointly developed with Jonathan Parker, proposes a departure from the rational expectations paradigm that all agents' beliefs always coincide with the true objective probability distribution. Optimal expectations provide a structural model of belief distortions that can explain many experimental findings documented in the economics and psychology literature. Optimal belief biases in our model are driven by the desire to look forward with optimism, and are limited by the cost of impaired decision making. Unlike much of the previous literature, beliefs are derived endogenously. This not only imposes discipline but also avoids the outcome that agents with biased beliefs can be exploited without limits by rational agents. Our framework provides a parsimonious explanation for previously ill-understood phenomena such as why (i) agents facing an investment problem are overoptimistic and exhibit risk-loving behavior for lottery-type (skewed) payoffs; (ii) agents bet against each other in an asset pricing setting; and (iii) observed consumption profiles are concavely downward sloping even though agents' expected profiles are not.

Jonathan Parker and I extended this line of research in at least two ways. In **"Optimal Beliefs, Asset Prices and the Preference for Skewed Returns"** (American Economic Review P&P, 2007) also joint with Christian Gollier, we study the asset pricing implications of our optimal expectations framework in general equilibrium. Our theoretical model explains why households' portfolios are tilted towards individual stocks with positive skewness, which consequently yield lower returns. In **"An Economic Model of the Planning Fallacy''** (also with Filippos Papakonstantinou) we address the phenomenon that people systematically underestimate the time required to complete a task and the optimality of intermediate deadlines. We show that self-imposed deadlines are less stringent than deadlines externally imposed by an objective observer. We test our theory using existing evidence on expectations and behavior in different experimental settings.

Another way of departing from the neoclassical framework is to assume a different objective function for the economic agents. Prospect theory does exactly this. My note **"Learning to Re-optimize Consumption at New Income Levels: A Rationale for Prospect Theory"** (JEEA, 2004) provides a rationale for three elements of prospect theory.

One can also deviate from the canonical utility function by making it state dependent, e.g. by introducing preferences which depend on past consumption habits. If agents' habit preferences are additive, then a positive (negative) wealth shock leads to a transitory decrease (increase) in agents' relative risk aversion. Such preferences have been proposed for models with a single representative agent. These models are generally viewed as a "rational" explanation of the equity premium puzzle and the predictability puzzle of stock returns. My article with Stefan Nagel **"Do Wealth Fluctuations Generate Time-varying Risk Aversion? Micro-Evidence on Individuals' Asset Allocation"** (American Economic Review, forthcoming) investigates whether there is micro-level evidence in support of this proposed (negative) relationship between wealth shocks and relative risk aversion. To this end, we analyze two decades of panel data on household asset allocation from the PSID surveys. Using a variety of specifications, we find that the share of financial assets that households invest in risky assets is unaffected by shocks to their wealth. Our findings are important because they show – contrary to common belief – that the assumed time-varying risk aversion of the representative agent is not simply a reflection of the same preferences at the microlevel. Hence, a different microfoundation is needed.

Recently, I completed a first draft of the project **"Leadership, Coordination and Mission-Driven Management"** with Patrick Bolton and Laura Veldkamp. The core idea is that leaders (i) help a group of agents coordinate their actions and (ii) adapt to changing circumstances. A good leader coordinates his followers by making a precise mission statement, which credibly communicates the optimal course of action. In practice, leaders learn about the optimal action over time. Learning creates a time-consistency problem because the leader has an incentive to commit to a mission to achieve coordination and then adjust it when new information arrives. Overconfidence is a valuable attribute in such a setting, since it helps the leader stick with his prior belief. Even with a costly commitment technology available, overconfident leaders still facilitate better coordination and teamwork. The drawback of overconfidence is that it inhibits learning from the followers' actions. Hence, overconfidence is costly when followers have sufficiently valuable signals.