Reflexivity and Systemic Risk in Quantitative Finance

Daniel Beunza
Graduate School of Business
Columbia University
db2417@columbia.edu

and

David Stark
Department of Sociology
Columbia University
dcs36@columbia.edu

For presentation at the Summer School on
Self-Reinforcing Processes in Organizations, Networks, and Markets
Freie Universität Berlin July 13-17, 2009

Note to participants: We had hoped to do one last round of editing of this paper before sending it to you. It is very close to a final version. Our apologies for a couple of places that are still rough.

Reflexivity and Systemic Risk in Quantitative Finance

Daniel Beunza and David Stark
Columbia University

Abstract Our study identifies a new form of systemic risk in model-centered financial markets. The literature in behavioral finance has explained modern crises as the outcome of blind imitation, overconfidence, or an unreflective use of models. Our account, by contrast, points to the unintended consequence of reflexiveness. Arbitrageurs, we found, not only use models to take positions but also to check their own views against those of their rivals. This form of reflexive modeling, however, creates a cognitive interdependence: a trader's position becomes his or her rivals' cautionary check. When a sufficient number of arbitrageurs erroneously overlook a crucial aspect of a trade, models project a misleading sense of confidence to the entire arbitrage community, leading to so-called "arbitrage disasters," widespread and oversize losses. To address these questions we conducted ethnographic research in the derivatives trading room of a major international investment bank.

INTRODUCTION

On June 12th, 2001, the European Commission stated its opposition to the planned merger between General Electric and Honeywell International. The ruling by the European authorities effectively cancelled the proposed merger between the two Fortune 500 companies which had been announced in October 2000 (Shishkin 2001: A11). As news of the cancellation arrived on Wall Street, the stock price of Honeywell (the target company in the proposed merger) dropped by more than ten percent. Professional arbitrageurs (hedge funds and investment banks), who had expected the merger to succeed, collectively lost a remarkable $2.8 billion on the deal. The loss was large enough to offset their cumulative profits for the entire quarter. Such heavy losses among sophisticated trading firms is all the more signficant because these firms had embraced the use of models to avoid debacles. The failed merger between GE and Honeywell provides fertile ground to understand the possibilities and pitfalls posed by financial models. What role, if any, do models play in financial disasters?

We address this question with an ethnographic study of financial modeling. Our study concentrates on the daily operations of a major international investment bank, pseudonymous "International Securities," located on Wall Street. Its proprietary trading unit lost $6 million dollars in the GE-Honeywell deal. In examining this outcome, our analytic point of departure is the dilemma posed by models: they offer

its users the possibility of extraordinary returns, but also pose the risk of catastrophic losses. Our study, investigates the ways in which traders grapple with the fallibility of their models. Arbitrageurs, we found, not only use models to take positions but also to check their own estimates against those of their rivals. This form of *reflexive modeling*, however, creates an unexpected cognitive interdependence between financial actors: a trader's position becomes the cautionary sign used by his or her rivals in the arbitrage community. As a result, when a sufficiently large number of arbitrageurs overlook the critical factor driving merger failure, the use of models can provide misplaced confidence within the arbitrage community, leading to widespread and oversized losses. The occurrence of such losses has been well documented, and is referred to in the finance literature as "arbitrage disasters" (Officer 2007).

Our study speaks to a growing interest in the interplay between models, systemic risk, and financial crises. Existing studies, mostly in the field of behavioral finance, have explained crises in terms of unreflexive risk-taking, faulty models, or social dynamics among investors,. None, however, have explored how models mediate these social dynamics, how social processes could be a means for reflexive risk-taking, or how such reflexive use of models could be a source of systemic risk.

One major strand of the behavioral literature on financial failtures, for example, points to overconfidence and excessive risk taking. Experimental studies have documented how individuals are more willing to take risks as their performance improves, even if the underlying probability distribution stays constant. However, this research has not taken models into account. The underlying assumption in these studies is that actors construct a sense of risk and opportunity by direct assessment of their past performance. But modern modeling techniques provide tools to estimate the probability distribution that the decision-maker confronts. Traditional accounts of excessive risk-taking need to be updated to reflect the new quantitative nature of securities trading.

By contrast, financial models dominate accounts of financial crises that focus on the "black swans" of extreme events. Building on the Knightian distinction between risk and uncertainty, several authors have argued that crises can occur when the unquestioned use of financial models leads banks to underestimate uncertainty (Taleb 2007, Derman 2004, Bookstaber 2007). The financial models used by investors, the argument goes, are predicated on the assumption that the future will resemble the past. Investors may assume, for instance, that stock returns will follow a Normal distribution. But the capital markets, according to these authors, are social settings and are subject to unpredictable, extreme events. Instead of a normal distribution, stock returns are more accurately described by fat-tailed distributions in which rare events are a distinct possibility. Thus, to the extent that investors do not incorporate exceptions into their models, their trading strategies will ignore the possibility of black swans -- of rare, unexpected, and high-impact events.

An important shortcoming of this argument is that it presents financial actors as hopelessly irreflexive about the limitations of their models. Traders, according to the proponents of black swans, either ignore what every finance academic already knows, or lack the reflexive capacity to act on it. Confronted by this argument, we ask, why should we deny to financial actors the capacity for reflexivity that we prize and praise in our own profession?

The social nature of systemic risk is brought to the fore in behavioral studies of imitation. This mechanism, known in the literature as "herding", takes place when actors decide to disregard their own information and imitate instead the decisions taken by others before them (Scharfstein Stein 1990; Banerjee 1992; Bikhchandani, Hirshleifer and Welch 1992). In the context of the financial markets, herding adds to systemic risk because it can lead investors to artificially bid up the price of an asset, to the point of creating a bubble (Shiller 1984, 1999 and finally producing a crash.

Theories of herding, however, do not take models into account. Herding actors simply abandon their own opinion. The interdependence between market actors posed by these theories, in other words, is extraordinary enough to remove any room for financial models to enter the picture. By contrast, a theory of systemic risk that takes models into account should explain how actors trade off the conclusions stemming from their own model with the social cues arriving from their environment.

Within sociology, the canonical study of systemic risk is offered by Merton's account of a run on a bank. As Merton (1968) observed, banking is a special form of economic activity subject to positive feedback between beliefs and behavior – and, therefore subject to self-fulfilling prophecies. Because a depositor's decision to draw out his or her funds reduces the liquidity available to other depositors, the collective perceptions of a bank's solvency among its depositors end up sealing the fate of the bank.

Self-fulfilling prophecies, however, are still an incomplete guide to modern systemic risk. In the standard account, these prophecies entail an over-abstracted, almost tautological account of how a crisis happens. *If* a sufficiently large number of depositors fear a crisis, the run on the bank will surely happen. But, as Callon (2007) asks, *how* do these beliefs arise in the first place? One answer might be that these beliefs are a shared convention, but how do depositors coordinate their views around the convention? The answer, Callon suggests, points to the material basis of belief formation. In retail banks, a line forming outside a branch can be enough. Theories of self-fulfilling prophecies, in other words, need to account for the ways in which decision-making tools and techniques coordinate the beliefs of market actors. Our interest is to explore how financial models might be one such source of coordination.

In sum, understanding modern systemic risk not only calls for understanding the existence of social dynamics among investors. It also calls for studying the ways in which financial models mediate this interdependence.

In advancing towards that goal, four core sociological ideas orient our inquiry. First, the embeddedness perspective, developed by Granovetter (1985), Baker (1984) and others, provides a useful framework to conceptualize interdependence. The notion that market transactions are structured by the personal ties of social networks can certainly explain some aspects of systemic risk. It has, for instance, direct bearing on phenomena such as Ponzi schemes, where clients are introduced to the scheme through friendship networks.[1] But the notion of embeddedness, developed before the full impact of the quantitative revolution, needs to be reconsidered when personal networks have been replaced by the socio-technical networks that accompany financial models. Embeddedness presupposes the existence of personal acquaintance among social actors; by contrast, current markets are shaped by deliberate anonymity. To put it succinctly, embeddedness is excessively centered on people. What is the counterpart of embeddedness when all that traders see at work is a screen?

In this respect, the application to markets of the analytic tools of Science and Technology Studies (STS) offers useful guidance. To understand anonymous transactions, argue Callon and his colleagues,[2] we must analyze the materiality of calculation (Callon 1998, 2007; Callon and Muniesa 2005). Once a market has been purposefully designed to depersonalize transactions, actors turn to market devices such as financial models to perform calculations. These models frame decisions and quantify alternatives, thereby exerting a mediating role on the value of goods and securities. But the notion that investors act in complete independence of each other, however, is not attentive to the actual processes in which investors make moves in relation to the actions of other investors. To capture systemic risk, researchers need to contemplate how *interdependence* exists among *anonymous* financiers.

In short, grappling with modern crises calls for an understanding of the forms of engagement introduced by financial models. Models has given rise to a new mode of sociability that is disembedded, yet entangled; impersonal, but social. In conceptualizing this redefinition of "the social," we draw on Knorr-Cetina's (2006) notion of scopes, or observational instruments. Knorr-Cetina draws a distinction between markets where personal relations carry the burden of coordination ("network architectures"), and markets where objects are the central coordination device. In the latter, the actions of investors are projected onto a scope, giving rise to a life-form to which investors can react; their reactions, in turn, will be part of the

[1] The premier example of a Ponzi scheme is of course the recent fraud perpetrated by Bernard Madoff. The key question raised by the Madoff case concerns the ways in which the financier managed to elicit enough trust to obtain not simply thousands, but millions of dollars from investors, without any transparent account of the destiny of the funds.

[2] See also Callon 1998, 2007; MacKenzie and Millo 2003; Mackenzie 2006; for reviews see Fligstein and Dauter 2007; Healy and Fourcade 2007; Ferraro, Sutton and Pfeffer 2005.

representation. Investors, in other words, not only react to each other but to the aggregate traces of each other's actions, projected on the scope. Such new rules of association –aggregation, anonymity, and mediation through a shared representations – offer fertile grounds to theorize the ways in which systemic risk can originate in financial models.

Another key theoretical pillar is the work on the performativity of Black-Scholes (MacKenzie and Millo 2003; MacKenzie 2006; Millo and Mackenzie 2008). In our reading, an economic model is performative when its use improves its predictive ability. As MacKenzie and Millo (2003) described, performativity took hold as investors began to use Black-Scholes in reverse to translate option prices into "implied volatility" (the volatility estimated by other actors in the market.) This move led investors to alter the basis of their interactions, changing how they talked about options. But this use of the model, we add, also leads to interdependence among the investors. Once traders can rely on their anonymous competitors as a source for extra information, a novel mechanism of social influence has arisen. .

In the following pages we abstract from the complexities of modern arbitrage to present the ways in which economic models can lead to financial crises. The first element of our study will be to use our ethnographic observations to reconstruct how the arbitrageurs at International Securities dealt with a particular merger. That is, from numerous hours of observations across many trades, as one aspect of our study, we offer a very detailed analysis of a particular arbitrage opportunity. Our presence in the trading room meant that we could analyze a given merger from the first moment that the merger arbitrage desk learned about its announcement. We will analyze how the traders "set up the trade," starting with studying the PowerPoint presentations and videos from the merger announcement, leading to the use of Excel spread sheets and proprietary databases. With these modeling tools, the traders build a picture of the merger, anticipating the future by drawing analogies to the past.

Second, we use our ethnographic observations to show a counterintuitive aspect of financial modeling. Taking a position is only the first step in the process of using financial models. Our research explores the next step, in which traders cast a sceptical eye on their own estimates. To do so, they exploit the fact that other traders have also taken positions on this trade through a second form of modeling known as "backing out." Knowledge of the implied probability can trigger search processes along new dimensions in previously unexamined territories. We refer to this mechanism as *reflexive modeling*. Gaps, disparities, differences, mismatches can produce positive friction that stimulates re-search. The lack of them gives traders greater confidence that their views are correct.

The final step of our analysis examines the systemic dangers posed by this reflexive process. We consider a phenomenon known as "arbitrage disasters" in the finance literature (Officer 2007). If a sufficiently large number of arbitrageurs simultaneously fail to see a merger obstacle ahead, the use of implied probability will provide traders with false reassurance, leading them to expand their positions

and suffer widespread, potentially catastrophic losses. The reflexive use of models, in other words, creates systemic risk.

RESEARCH METHODS

Research site. The data reported below focuses on our observations from the merger arbitrage desk at International Securities. The firm is a top-ten global bank in equity underwriting (Hoffman 2006) with an active proprietary trading unit. Our observations center on its equity derivatives trading room, located in Lower Manhattan. Proprietary trading units of this kind function as internal hedge funds within an investment bank. We conducted detailed observations at three of the bank's trading desks, sitting in the tight space between traders, following trades as they unfolded and sharing lunch and jokes with the traders. We complemented this direct observation with in-depth interviews of the traders at each desk. In the final year of our investigation we were more formally integrated into the trading room, provided with a desk, a computer and a telephone.

Our study focuses on arbitrage, an ideal site to examine the role of models in systemic risk. Arbitrage seeks to exploit financial mispricing across markets by relying on models to determine similarity across securities that trade at different prices (Beunza and Stark 2004; Derman 2004). While this use of models has proved successful, leading to legendary returns and a seven-fold increase in the number of arbitrage firms in the past two decades (Lo 2008), arbitrage has also been associated with several recent financial crises. These include the market crash of 1987, the crisis of Long Term Capital in 1998 and the hedge fund "mini-crash" of August 2007 (see, respectively, Dunbar 2000, MacKenzie and Millo 2003; Lowenstein 2000, Jorion 2004, MacKenzie 2006; Khandani and Lo 2007).

Our study narrows the scope of inquiry to a particular arbitrage strategy, merger arbitrage. Unlike other arbitrage styles, merger arbitrage allows us to identify failure. That is, it allows us to separate the perceptions of financial actors on Wall Street from the actual events that unfold outside it. Merger arbitrage boils down to informed speculation about the likely completion of corporate mergers. Our traders, in other words are not simply monitoring the positions of others in order to anticipate "where the crowd is moving." Rather, they do so to derive the expectations of other traders about the likelihood of an event –the merger—that will, in the end, happen or not happen. And that event, the merger, is by and large independent of the collective wagers of the arbitrage community. Thus, the specific form of specularity (see also Dupuy 1989) involved in merger arbitrage differs from Keynes (1936) view of financial markets as beauty pageants in that arbitrageurs can collectively be wrong. This makes the strategy ideal to understand financial crises.

We explore the role of models in merger arbitrage with an ethnographic research design. Ethnography is useful to understand the day-to-day practices of calculation, for it places the researcher in the same uncertain position as his or her subjects,

thereby avoiding the danger of retrospectively underestimating uncertainty (Orlikowsky 1992; Barley 1986; Agar 1986; Spradley 1979). Partly for that reason, ethnography has been a method of choice in the social studies of finance literature (Abolafia 2001, Knorr-Cetina and Bruegger 2002, Zaloom 2003, Beunza and Stark 2004).

Period of observation. Our study is based on two central events during our three-year engagement with the bank. This engagement extended to more than sixty visits between December 1999 to March 2003. Specifically we examine an arbitrage disaster (the failed GE-Honeywell merger) in June of 2001 from the standpoint of a different moment in time: from a morning of trading on March 27th of 2003. Our ethnographic explanation of disaster, in other words, is not based on the morning in which the event took place. This decoupling, however, is an advantage. We did not arrive at the merger desk with the objective of studying arbitrage disasters, but to understand quantitative finance. By providing a symmetrical treatment of success and failure in merger arbitrage, our study avoids the trappings of the sociology of error (Bloor 1976), in which "the social" is only seen as dysfunctional. Thus, whereas most studies of failure isolate the negative aspects of models and social interaction, our study explains disasters in the same way that it explains extraordinary financial returns.

ANATOMY OF A MERGER ARBITRAGE TRADE

Distributed cognition

Trading at the merger arbitrage desk at International Securities is a highly quantitative, technologically sophisticated endeavor. We arrived at the desk at 9:00 am of March 27th, 2003, before the US markets opened. We found the arbitrageurs quietly working at their computers. Oswald, the junior analyst among the three, was absorbed in a succession of PowerPoint slides displayed on his screen, isolated from the others by a pair of headphones. Max and Anthony, senior and junior traders respectively, were entering data from a sheet of paper into Excel spreadsheets. Max and Anthony were transposing the details of the collar into their respective Excel spreadsheets, working in parallel to prevent clerical mistakes. As they typed, their conversation turned to data about other ongoing trades. "What's your price for Whitman?" asked one of them. "I've got bad data on it."

This appearance of business as usual was somewhat surprising, for an important merger had just been announced. Career Education Corporation, a private provider of vocational training based in Illinois, had stated its intention to acquire Whitman Education Group, a Miami-based competitor. The news had landed on the Bloomberg terminals of the traders at 5:58 pm of the previous day, with the market already closed. The arbitrageurs confronted the news the following morning, minutes before our visit. Given the announcement, a visitor might have expected to

see the traders engaged in energetic buying, selling or talking on the phone. Instead, all three traders appeared to be immersed in their typing.

The impression of inactivity proved to be misleading. Far from ignoring the merger announcement, the traders were reacting to it in their characteristic way, preparing the trade. The first step in this process was the elaboration of a memorandum. The memo summarized the key details of the Whitman-Career combination. Oswald compiled the memo after listening to the presentation that the merging companies put out for analysts; hence his headphones. The output of his work was a document stating the legal details of the merger: the cash and stock that Career would pay for Whitman, the expected closing date, etc.

Preparing the trade entailed a further step. Having finished the spreadsheet, the traders proceeded by linking the document to yet another Excel spreadsheet, known as the "Trading Summary." This second spreadsheet functioned as a brief of all the trades in which the desk was involved. On the morning of May 27th the traders were active in 31 deals, so the involvement in Career-Whitman meant the addition of a 32nd row to the document. Like the instrumentation panel of an aircraft, the Trading Summary made all financial action readily visible at a glance.

These early observations underscore the importance of quantitative infrastructure in modern finance. A merger trade requires the assembly of electronic scaffolding to supplement the arbitrageurs' mental processes: a PowerPoint presentation, followed by a Word memorandum, followed by an Excel spreadsheet, all of it condensed into a single live cell on a Trading Summary. In short, cognition is distributed at the merger arbitrage desk. Like the pilots and ship crew studied by Hutchins, arbitrageurs can reduce their cognitive overload – the extent of their bounded rationality by turning to the machines and objects around them. Arbitrageurs are aware and understand this process, and refer to it as "setting up" the trade. As we shall see, however, even as arbitrageurs enroll machines in their calculations, taking up a position is far from mechanical, for it involves a good deal of judgment and interpretation.

Judgment

Amidst the hubbub of the data entry, the arbitrageurs attempted sized up the nature of the newly announced merger. Categories, analogies and other references to the past allowed them to engage in pattern recognition that would result in taking a position. At 9:40 am, for instance, Max and Oswald engaged in a dialogue on Whitman and Career. "Do they have regulatory approval?" asked Max, without taking his eyes off the screen. "They do," Oswald replied, looking at his spreadsheet. "Do they have accreditation?" Max inquired. "What schools are these, anyways?" he added emphatically, his eyes squinting at his screen. "Technical, for adults" Oswald responded. "They teach you things such as how to be dentist assistant," he added.

The conversation, a seemingly casual exchange about Whitman, was an effective first step in estimating merger probability. This probability is the figure that arbitrageurs care most about. The basic principle of arbitrage is to exploit situations where two different regimes of value coexist in ambiguity (Beunza and Stark 2004), and merger arbitrage is no exception. In the case of mergers, the ambiguity arises from the fact that a company is being bought. The acquiring firm typically buys the target company at a price well above its market capitalization, leading to two possible valuations: if the merger is completed, the price of the company will rise up to its merger value; if it is not, the price will drop back to the level before the merger announcement. Arbitrageurs exploit the ambiguity as to which of the two will apply by speculating on the probability of merger completion. To the arbitrageurs, therefore, profiting from mergers boils down to successfully estimating a probability.

In their exchange, Max and Oswald established a set of facts that subsequently proved relevant. For instance, they established that the merged company, if completed, would belong to the "for-profit post-secondary education sector." The usefulness of this categorization became clear at 9:45 am, as Max turned to examine a chart of Whitman's sales. "Is it true that there's a summer drop-off in this business?" he asked Oswald, faced with what appeared to be weak summer sales. This mattered, because a common source of merger failure is negative results at one of the merging companies. "It's the summer recess," Oswald replied. The weakness in sales, in other words, was the school holidays -- a normal part of the education industry. Because the companies belonged to the education industry, the cyclical drop-offs in sales were not a relevant merger risk. Categorizing Career and Whitman, we concluded, helped arbitrageurs interpret information that could have material implications for merger completion.

Arbitrageurs complement categorizations with analogies to past mergers. At 9:50 am, the conversation shifted to a discussion of another company in the for-profit education sector. "This guy Edison," Max explained, "a few years ago wanted to manage the primary school system. But then went down in flames." The entrepreneur mentioned by Max was Christopher Whittle, founder of Edison Schools. Edison began operations in 1995 with the promise to bring private-sector discipline to the bureaucratized education industry. But the company saw its stock price plummet in 2002 amidst accusations of corruption (Denison 2002). A scandal of the type that Edison experienced would immediately ruin the merger at Career and Whitman, so the probability of a scandal had to be factored in.

In this discussion, the traders exploited the power of analogies to anticipate possible merger obstacles. Like categories, analogies allow them to glean the future from the past. "We look for patterns," Max explains, "precedent, similar deals, either hostile or friendly, degree of product overlap, and earnings variability. We look at all the ways to slice the factors that weigh into the merger." In the case of Career and Whitman, the analogy associated two merging secondary-education firms with a for-profit *primary* education company, Edison Schools, previously marked by

corruption. The analogy associated the merging firms with another one from outside their industry. But the association between Edison Schools and the merging companies prompted a new concern: it led the arbitrageurs to focus on the honesty of the management teams at Career and Whitman. The use of partly overlapping categories and analogies underscores that arbitrageurs do not just passively fit mergers into boxes.

In stressing the importance of flexible association, Max points to his success in anticipating the failure of a past deal by drawing non-obvious parallels with other deals. Carl recalls a merger between two junkyards that had incompatible databases. In the low-tech world of junkyards, one might not anticipate information technology to be a key factor in derailing a merger. But, Max added, "if the point of a junkyard is to find that door for the 1996 Volvo," says Max, "you can imagine how important databases are. We had another deal with similar proprietary databases in a different industry [that] reminded me of that junkyard deal." The arbitrageurs correctly predicted the failure of the merger between the junkyards and closed their positions early enough to avert any losses. As Max concludes, "drawing parallels and linkages and saying 'this reminds me of that' is at the heart of what we do."

The traders, however, do not just rely on their own memory to draw those associations. At 9:55 am Max called up a black-and-white window on his screen. The screen displayed a set of old fashioned, 1980s-style Microsoft DOS characters. Pressing a combination of commands keys, Max obtained information on Edison to look for patterns that were similar to the Whitman-Career deal. The screen corresponded to a proprietary database that Max has meticulously assembled over the years, with information about all past mergers in which the desk has been involved, classified along numerous dimensions. This gives "thumbnail" information about each company that merged. "You think you would remember," Max says about it, "but you don't. Memory is very deceiving." Like the other arbitrage artefacts presented above, the database contributes to distributed cognition at the trading desk. Specifically, by providing a costless system of storage and retrieval of past information, the database helped arbitrageurs mobilize past deals to make sense of current ones.

After two hours of establishing associations, the arbitrageurs were beginning to develop an overall impression of the Whitman-Career merger. Max explained,

> There may be many issues with this company, but I can invest right away by knowing that they're a $5 million and a $2 million companies. This means it's not one company acquiring another that it's the same size, which right away means that there are not financing issues involved. If there were, it would be a whole different game.

In other words, even though the industry – for-profit education – was tainted by a past scandal, the traders were still encouraged by the lack of other obstacles.

At 10:15 am, the market opened on Whitman Education with a price of $13.95. The arbitrageurs' spreadsheets showed the spread to be a generous ten percent, signaling to the traders a potential opportunity. "I'd like to have a presence in the deal," said Max almost immediately. "Let's bid $13.60 for 10,000" he added. Following the instruction, Anthony lifted the headset from his phone turret and called the block trader to place an order. Thus, barely two hours after starting to work on the deal, the merger traders at International Securities took a position in the Whitman-Career merger.

Why take a position within minutes of the opening? Arbitrage, we observed, is a game of speed. The longer arbitrageurs take to adopt a position, the more time their competitors have to seize the opportunity before they do. As in Occam's razor, arbitrageurs take into account as many factors as they need to take a position, but not more. Taking a position, then, involves a successive winnowing of the possible contingencies involved in the merger as the arbitrageurs think through the deal. The traders search through a form of mental decision tree in which each specific merger is considered in relation to similar deals that they encountered in the past. Max explains, "it's almost like you've been in this road before and [the past incidences] direct you." The advantage of this system, which Max describes as a "process-driven arbitrage," is that numerous issues need not be taken into account. Arbitrage, in other words, is fast, light and deploys resources in a strategic manner.

The arbitrageurs, in other words, are not simply performing a routine task of recognition – classifying mergers into pre-existing categories – but a far more active task of *re-cognition*. That is, changing, expanding and going beyond the existing categorical structure to ascertain the key merger obstacles in a given deal. It is for these reasons that the arbitrageurs have not attempted to automate the process of sizing up a merger. According to Max, "it's impossible to turn [the process] into a purely quantitative exercise. There's judgment."

Representing the collective rival

Our analysis so far has established two related observations about merger arbitrage. First, it is clear that the arbitrageurs deploy quantitative tools and models. Second, as they do so, they also understand that these models are fallible and that they need to exercise their judgment. This points to the dilemma that the traders encounter: models offer the possibility of extraordinary returns, but also pose the danger of oversized losses.

As we shall see, the traders solve this dilemma by simultaneously exploiting the categories and procedures that guided them to an initial position, and distancing

themselves from these. This, however, is easier said than done. Mental awareness of the limits of one's view does not automatically provide a check against these limits. Traders gain some cognitive distance from their categories,, we found out, by exploiting the fact that other arbitrageurs have also taken positions on this trade. It is to the second moment of a distributed cognition – across a socio-technical network outside the trading room – that we turn.

At 10:30 am, the conversation between Max, Oswald and Anthony shifted from Career and Whitman to another ongoing merger. Five months before the focal morning, Hong Kong and Shanghai Bank (HSBC) had announced its intention to acquire Household International, an American bank specialized in subprime mortgages. The traders at the merger desk had been "playing" this deal. At 10:40 am Max typed a command in his Bloomberg terminal, producing a large black and blue graph on his screen. The chart, reproduced in Figure 1 below, displays the evolution of the "spread" between HSBC and Household. The spread is defined as the difference in the prices of the merging companies, adjusted for the terms of the merger. In this case the spread corresponded to the difference in the prices of HSBC and Household over the five-month period in which the merger unfolded, weighted by the stock conversion ratio agreed by the merging partners: 0.535 shares in HSBC for each share in Household International.

[Figure 1 and Figure 2 about here.]

Visualizing merger likelihood. The graph, known as the "spread plot," plays a key role in the conversation among the traders. Movements in the spread signal changes in the likelihood of merger completion. If a merger is completed and the two merging firms become a single entity, the difference in their stock prices --the spread-- will disappear. Thus, when the spread narrows before merger completion date, it is because arbitrageurs collectively anticipate assign a greater likelihood of merger completion. Conversely, if the merger is canceled and the equivalence between the two firms disappears, the spread will widen closer to its level before the merger announcement. When the spread widens before completion date, it is because arbitrageurs expect a lower likelihood of completion.

Using the spread plot involves semiotic sophistication. In this complex system of signs, for the trader at a particular desk, the spread plot is an indirect sign of the likelihood of the merger, achieved by signaling the aggregate of his or her rivals' assessment of that likelihood. For the very reason that they are deeply proprietary, the trader does not have access to the proprietary databases through which *particular* other rivals constructed their own independent probability estimates. And, indeed, to have such access would result in cognitive overload: how could one gain cognitive distance from one's own models if one had to engage in the time-consuming task of comparing them with those of dozens of other traders? The spread plot reduces that cognitive complexity by representing the *aggregate* of the expectations of other traders.

The arbitrage trader, however, is not interested in the spread plot as a sign of what others are doing in the market but as a sign of an event that will or will not happen in the world – the merger. The promising aspect of *this* sign is that it is quasi-independent of his or her own estimates of that probability. That is, the arbitrage trader is not a technical trader who, like the fashionista who monitors others to anticipate the hottest clubs, seeks to profit by anticipating market trends. Instead, arbitrageurs use the movements of their rivals *as a check* on their own independent opinion. Just as office employees with a limited view of the sky can anticipate the likelihood of rain by checking through the window for pedestrians carrying umbrellas, arbitrageurs check for merger obstacles by monitoring the aggregate actions of their rivals.

The spread plot of HSBC and Household International illustrates how changes in the spread signal merger obstacles (see chart on Figure 2). The chart shows two clear spikes along a descending line. These correspond to instances in which market participants lost confidence in the merger. The first, on November 22nd 2002, was inspired by funding concerns: was HSBC simply buying Household to get funding? In other words, was HSBC a sound company? This surge in the spread subsided after a general market rally. The second spike took place on March 20th 2003, following news that Household International was shredding documents. This reminded arbitrageurs of similar shredding at Enron years before. The spread then fell again after the company received its approval from the financial authorities, and once HSBC reassured investors. The two spikes illustrate how plotting the spread brings out the crisis points in a merger. Even more interestingly, the obstacles underscored by the spread plot may not have been taken into account by the arbitrageurs in their initial decision process; that is, these concerns might have been unexplored as a abandoned branch as the traders advanced in their tree-like decision pattern. Checking the spread plot, then, is a way to avoid cognitive lock-in by the decision tree.

Translating prices into probabilities. A second key concept mobilized by the arbitrageurs is the "implied probability" of a merger. By implied, the arbitrageurs refer to the probability of completion that rival arbitrageurs assign to the merger. Quantifying this probability entails manipulating the basic regularity governing arbitrage, the Law of One Price, in a process known as "backing out."

Backing out merger probabilities builds on the notion of risk-neutral pricing. The core idea behind this practice is to extract useful information from mispricings in markets where are arbitrageurs are present (Cox and Rubenstein 1979, Harrison and Kreps 1979). As the Law of One price argues, the presence of arbitrageurs will remove unjustifiable differences in prices across markets. Thus, in the absence of transportation costs, the price of gold in London would not be able to systematically differ from that of gold in New York without inviting the activity of arbitrageurs. Conversely, the difference in prices between New York and London can be interpreted as the cost of transportation. Merger arbitrageurs apply this idea to a different setting, corporate mergers. When a merger is certain and arbitrageurs are

present, the value of a merger target and acquirer should be the same. Thus, the difference in prices between the two –the spread-- can be read as a measure of the uncertainty that arbitrageurs assign to the merger.

Backing out the degree of uncertainty amounts to inferring a probability distribution on the basis of the payoffs. Consider the toss of a coin. If a gambler is not told the probabilities of heads or tails but is the payoffs involved in the bet and is told that they are *fair* --that is, they are those that would make both players even after multiple rounds-- the gambler will be able to infer those probabilities. For example, if the payoffs are $1 (heads) and minus $1 (tails), a gambler can infer that the probabilities are 50-50. As Max says, "I would not need to describe the game to you at all, just the payoffs. And backing out from the payoffs you would be able to infer the probability."

Backing out probabilities, however, is a delicate move. In accomplishing this translation, arbitrageurs make two key assumptions: first, they assume that movements in the spread are dominated by merger considerations. That is, if the spread changes for some a reason unrelated to the merger, the interpretation of the move as a change in merger likelihood would be erroneous. Second, the translation assumes that markets equilibrate rapidly (in the coin example, that the prices are "fair"). Unless rival arbitrageurs have seen prices, compared them to their own information and acted upon it, the spread does not convey the private knowledge of these rivals. As we shall see, arbitrageurs are mindful of these two conditions, and come back to them repeatedly whenever prices do not behave in an understandable manner.

This scheme gains additional complexity in the case of mergers that have "collar" provisions. A merger collar is a legal clause that alters the ratio at which the merging firms will exchange their shares if the price of the acquirer changes substantially. As a result, the deal's value fluctuates based on the price of the acquirer's stock. This structure is beneficial for the acquirer because it limits the overpayment risk, but it makes it more difficult to calculate the implied probability on the basis of prices (see Appendix for details).

Paradoxically, the difficulty posed by the collars benefitted Max by creating opportunities that only he was able to see. As he explains,

> I was one of the first people to use pretty elementary probability theory combined with not-so-elementary log normal distribution and expected value to establish the right approach to collar deals. Initially, there were aberrations in the pricing. Over time, they disappeared.

Max, in other words, is acutely aware of the power of financial models. His pioneering efforts in this direction won him a reputation around Wall Street, leading

to invitations to lecture in universities and routine appearance in the financial press as an authority on ongoing deals.

From personal networks to financial models. The tools described so far mark a sharp break with the traditional clubbiness of Wall Street. Whereas the typical strategies of investors entailed accessing information ahead of their competitors (Abolafia 2001), merger arbitrageurs base their advantage on financial models. Max emphasized this difference in terms of the price movements of the merging companies. "Look at this jump," he said, in reference to the price of one company on the day its merger was announced. He added,

> This is the value that the [mutual] fund managers and the guys on the street are after. Once the jump has taken place that, it's a matter of pennies. The value investors don't have the fine-tuned tools to position themselves in this spread, to determine if it's too wide or too narrow for them. We do.

Thus, in other words, the arbitrageurs eschew the fat margins that can be found by correctly anticipating the merger announcement, and only come into the trade once the deal is officially announced. This deliberate strategy is a telling sign of the calculative orientation of the arbitrageurs, for while the merger *announcement* can only be anticipated by pursuing rumors from a traders' network, merger *completion* can be anticipated with the modeling tools used by the arbitrageurs, that is, the spread plot and implied probability. Indeed, Max and his desk of arbitrageurs see themselves as part of a different breed of Wall Street professionals than those that filled the trading rooms of banks before the quantitative revolution. As Max says, "I was a math major as an undergraduate. I would not have found a place on Wall Street before [the arrival of] models in the 1980s."

In line with this long-running shift from insider knowledge to models, the traders do not see information, but interpretation, as the source of their advantage. When asked about the reason for the disparity between their own assessment of merger probability and the merger spread, Max argued that it stemmed from a differential interpretation of the data. Max said,

> The reason why the spread is large is that other traders have their own proprietary models for it. And they can all be right. At this point, it's all about the future, and we don't know the future. So their assumptions on volatility, for example, could be different than ours. Or their assumptions about timing.

The opportunity that Max saw, then, was not the result of privileged information. As Max said, "right now, the data is all on the Internet, even the SEC filings." That is, most of the information is widely available and therefore does not confer any advantage. To him, it resulted from his desks' distinct interpretation of publicly available data.

Our account thus far presents the arbitrage desk as a smart, smooth and highly knowledgeable team. Like a well-trained SWAT-team, they come across as authoritative and almost self-sufficient in their range of tools, experience and skill. But, as we shall see, the merger arbitrageurs depart from a Callonian world of seamless calculation for they are aware that formatting places limits on cognition. There is very little that is smooth about arbitrage, expertise is constantly called into question by the singularity of ongoing events, and the collective knowledge of the merger desk is not enough to carry them through. The arbitrageurs address their shortcomings by reaching out to the rest of the market and opening up to alternative interpretations.

Gaining distance

"Are we missing something?" By 12:00 pm, the spread between Whitman and Career remained at the same wide margin it displayed two hours before, ten percent. Early on, a ten percent spread signaled an opportunity. But its persistence posed a puzzle for the traders, for it could now be interpreted very differently. It could mean, first, that other professional arbitrageurs were not "playing" the deal for some genuine reason: "are we missing something?" Max asked. Second, it could also mean that there were incorrect assumptions in the traders' model. "After all, the spread is only wide if my assumptions are right," Max said. These assumptions were "based on the 20-day moving average of the prices before closing date, but if the deal closes on a different date the price would be different." Finally, the wide spread could also mean the reverse of a threat: a better-than-expected opportunity. "Can it be," Max asked, "that the deal has gone under the radar screen of other traders?" The persistently wide spread, in other words, could be signaling missing information, incorrect modeling, or a profit opportunity. Establishing which of these applied was crucial to the traders. The spread, in other words, was a wake-up call that prompted arbitrageurs to think twice.

The conundrum faced by the traders is symptomatic of the disruptive role of the spread plot. Arbitrageurs, the chart reminded them, should not blindly trust their probability estimates, because they hinge on a representation of the merger (implicit in their databases) that may be incorrect. Given this, the spread plot provides traders with a much-needed device for doubt: by displaying their degree of deviation from the consensus, the spread plot provides arbitrageurs with timely red flags.

Responding to dissonance. Max and his colleagues responded to the discordant spread by plunging into a search for possible merger obstacles that they might not have anticipated. They first turned to databases: at 12.10 pm, one of them typed the names "Whitman" and "Career" on an online proprietary database. Like a Google keyword search, the database presented them with several hits ranked by relevance. Skimming through the sources of each result, the trader was reassured to see familiar newspapers. The search, in other words, did not show anything they did not know in advance.

The database search is an instance of the way in which arbitrageurs respond to the discrepancy induced by the spread plot. Having observed the dissonance between their own probability estimates and the implied probability, the traders went back to search for missing information. In doing this, the database helped even though the traders hardly knew what they were looking for: by including news from local media that the national media might have overlooked, it provided leads for issues that need to be dug deeper.

The traders' approach contrasts with early neo-institutionalist views of markets. In the classic account, the availability of social clues leads actors to economize on their search costs by imitating others (Meyer and Rowan 1976; DiMaggio and Powell 1983). In contrast, knowledge of the spread stimulated the arbitrageurs to search *more*. The discrepancy illustrates an important point about arbitrage. The material tools allow traders to come up with more sophisticated answers than traditional investors precisely by inducing skepticism about the tools. (We expand on this below). Arbitrageurs, in this sense, are persistent but skeptical users of calculative devices.

A cultural artifact on Max's desk sums up this professional *habitus*. Taped to Max's Bloomberg screen, a cartoon drawing showed Snoopy in full pilot gear, goggles, helmet, and scarf flapping in the wind. Sitting atop his doghouse, Snoopy extends his arms to hold an imaginary plane yoke. The fictional dog is seemingly piloting a plane. On a basic level, the cartoon illustrates the sophisticated, self-deprecating humor of the trader – if you think I am a Master of the Universe steering the world financial markets, think again. I am a plain dog in trader gear. On a different level, it illustrates the nature of Max's job: like flying, trading requires maneuvering through uncertainty. But the artifact goes far beyond that, in that it can also be seen as Max's reminder that the opportunities he sees on screen depend on a constructed magnitude, implied probability. And if this probability is not applicable to the case at hand, the opportunity is fictitious.

This interpretation of the cartoon is consistent with the ways in which Max speaks about implied probability. To him,

> It's a reality check. It's a number that's out there and it challenges everyday when you come in to have 85 percent confidence in this deal, whatever that is. You could have a little sign saying, are you challenging yourself in every day on every deal.

Bob, the manager of the trading room, confirmed this interpretation by remarking on Max's ability to sustain self-doubt. "Max," said the manager, "calculates the most sophisticated Bayesian formulas to get at a probability number, and then he'll say, 'it's all guessing'."

Confronting uncertainty with recourse to the network. Following the inconclusive search on Whitman, the arbitrageurs got on the telephone. At 12:20 pm, Anthony lifted the headset of his phone turret and called the floor broker who handled orders for Whitman at the exchange. "John says buy this WIX [for Whitman], no one's really hedging it," he said to Max as he finished the conversation. No other arbitrageur, the floor broker implied, was active in the Whitman trade. From this, Max concluded that the merger had passed "under the radar screen" of other arbitrageurs. He reacted by increasing the desks' exposure to the merger. "Let's work another ten [thousand], but pick your spots" he said to Anthony, asking the junior trader to purchase additional shares in Whitman, but to do so carefully to avoid inflating the stock price.

Why did the arbitrageurs call up their contacts? Until 12:00 pm, the traders had interpreted the spread as the implied probability of the merger. The persistent discrepancy between the wide spread and the traders' estimates, however, created a dissonance that led them to question their own interpretation. Having re-checked the database, they decided to inquire about the identities of the shareholders, partially lifting the veil of anonymity that protects securities trading. In doing so, the arbitrageurs were seeking to clarify whether backing out made sense in this context: that is, was the spread reflecting the information in the hands of rival arbitrageurs? The traders concluded it was not.

The traders, however, were emphatically not mimicking their rivals. Theirs was not a case of classic isomorphism or herding. Instead, they were attempting to disentangle overall market movements from the actions of the players who, in their view, were the only ones who really counted: their rivals, namely, other professional arbitrageurs. On learning that no other real player was hedging the stock, they concluded that the spread could *not* be interpreted as a measure of implied probability. Thus, reflexivity at the merger arbitrage desks cuts both ways: whereas an hour or so earlier, the spread plot had stimulated Max and his team to raise doubts about their database, here their phone conversation stimulated doubts about the meaning of the spread plot, the device for doubt itself.

In light of the above, consider now why Max told Anthony "pick your spots." The expression reminded Anthony to cover his tracks as he increased the desks' position on Whitman, with the aim of avoiding an increase in its stock price. The traders' efforts suggest that Max and colleagues felt they were being observed by other arbitrageurs through the lens of the spread. Just as Max and his team engaged in a calculated game of guessing, so were rival arbitrageurs at other firms. Preserving an opportunity that had gone "under the radar screen" of rival traders required avoiding warning competitors.

Reflexivity as a socio-technical process. The developments described above suggest that the traders' caution unfolds as the confrontation between two related magnitudes. A trader's ability to mobilize prices for greater precaution hinges on

the encounter between the probability of merger, estimated at the desk, and implied probability, obtained from the rest of the market. This comparison provides an invaluable advantage: it signals to traders the extent of their deviation from the market, warns against missing information, motivates additional search, prompts them to activate their business contacts, and provides the necessary confidence to expand their positions.

This distinct combination of internal and external estimates points to a novel use of economic models, which we refer to as reflexive modeling. The expression denotes the process whereby dispersed market actors employ economic models to confront their own estimates. This confrontation pits a trader's estimates against those of his or her rivals, thereby introducing dissonance in his or her calculations. This dissonance is attained through the construction of implied probability. This variable is a representation of an economic object that does not have a price,[5] is otherwise not observable, and is co-produced by the positioning of actors who use it to confront their interpretations and re-evaluate their positions. Collectively produced, the spread plot is a device for dissonance. Reflexive modeling, then, denotes a heightened awareness on the part of the arbitrageurs about the limits of their own models in representing the economy.

Arbitrage disasters

The mechanisms of reflexiveness that lead to correction can also pose perils for arbitrageurs. Our observations suggest that when sufficiently few arbitrageurs anticipate the failure of a merger, reflexive modeling can inflict wide-ranging and oversize losses, a phenomenon known as "arbitrage disasters" (Officer 2007). These disasters are induced by reflexive modeling. A trader's misplaced confidence augments that of other traders, leading to a situation where no single arbitrage desk does the necessary research on a possible merger obstacle.[6]

Perhaps the best example of a disaster is the aborted combination of General Electric and Honeywell International in 2001. The cancellation of this merger

[5] The peculiar nature of implied probability is best understood in comparison to market prices. Hayek (1945) famously described markets as a self-organizing system of coordination in which prices were the key means and mode of communication. Because an increase in the quantity demanded of a good typically leads to an upward movement in its price (and vice-versa), market actors coordinate their plans by simply buying and selling, without any other explicit communication. Implied probability, however, differs from Hayek's prices on three counts. First, unlike coordination for the purpose of resource allocation, arbitrageurs do not seek to coordinate with other arbitrageurs but to deviate from them in pursuit of speculation. Second, the problem of the arbitrageurs is not only dispersed information but, most acutely, heterogeneous, disparate interpretations. Third, there is no market for mergers: that is, there is no market in which the event, a given merger, has a price.

[6] The losses from arbitrage disasters are dangerously large because they are incurred in instances of high levels of confidence and, therefore, high exposure.

surprised the arbitrageurs with large positions, exerting collective losses of $2.8 billion on the various arbitrage funds active on the deal. The losses were large enough to offset the cumulative quarterly profits of most of these funds (Sidel 2001). As it turns out, the traders were misled by precedent: in the past, the anti-trust authorities of the United States and Europe had always coordinated, and the American authorities had already signed on the deal. Never before, in other words, had a merger authorized in Washington been blocked in Brussels (Bary 2001: 43). This time, however, the precedent was broken by Mario Monti, the European commissioner, on the grounds that the merger would give the combined entity an ability to engage in anti-competitive "bundling."

[Table 1 and Figure 3 about here.]

Our examination of this episode points to reflexive modeling as the cause of the disaster. We make our claim on the basis of three observations. First, our conversations with the head of the trading room at International Securities –which was active in this deal-- show that the trade led to a disaster (as opposed to simply a minor loss) because the desk enlarged its exposure on it. This exposure, in turn, was the result of high confidence in the deal. According to Bob,

> Max traded it … everyone's database lacked a field, and the field was "European regulatory denial." … I encouraged him to increase his size … you have confidence, all of your fields are fine… so instead of four million, I said six million.

In other words, the desk lost six million. The loss was not a problem for the fund, but it illustrates the size of the potential losses. Each arbitrage disaster, Bob explained, is "sharp and traumatic." Disasters, we conclude, are the result of excessive confidence.

Second, the traders' confidence was not a psychological disposition but a calculated response to a narrowing spread. As the narrow GE-Honeywell spread on Figure 4 shows, in the months before May 2001 arbitrageurs assigned a very large implied probability to the merger.

[Figure 4 about here.]

Third, and most importantly, the arbitrage community by and large missed the eventual cause for the cancellation of the merger, European regulatory opposition. This can be seen from a comparison between the merger spread plot and the media response to the news of concerns at the European Commission. Figure 4 shows the instances of weekly articles in the business press (*The Wall Street Journal*, *The Financial Times, The Economist*, etc.) that included in their text the words "Honeywell" and "Monti," the European anti-trust commissioner. The spike in the number of articles on February 27th 2001 indicates the media had genuine concern about European opposition, even anticipating that its resistance would take the form

of concern over bundling. "Brussels," the *Financial Times* wrote on February 28th, "tells GE to re-do merger homework" (Hargreaves and Hill 2001: 10). "For Mr. Monti," *The Economist* added, "the danger is that GE-Honeywell will be able to use the muscle it gets from a range of related products to squeeze rivals out of the market" (*The Economist*, 2001:43). And yet, even as the press voiced the concern, the arbitrageurs did not seem to be picking up on it. As Figure 5 shows, even as the press articles about Monti accumulated, the spread between the merging companies barely inched.

In sum, our anatomy of the failed GE-Honeywell merger points to reflexive modeling as the cause for the widespread arbitrage losses. Specifically, our analysis suggests that the losses stemmed from an underestimation of its key derailing factor, and that such underestimation was created by the traders' focus on implied probability. Because most arbitrage funds overlooked the risk of European opposition, the spread plot did not react to news of Mario Monti's opposition, suggesting that it was not an issue. This gave a a false sense of confidence to arbitrageurs, leading them to enlarge their positions.

How serious are arbitrage disasters? The consequences of the GE-Honeywell fallout were substantial. The total dollar loss across all merger arbitrageurs in this deal was $2.8 billion, according to Officer (2007); see Table 1. Industry-wide, the merger arbitrage community has experienced several arbitrage disasters in the past three decades. Furthermore, these disasters have not abated over time, and their size appears to be growing larger.

Reflexive modeling and its accidents. Arbitrage disasters point to the dark side of reflexive modeling. Disasters start when active arbitrage funds overlook a potential cause of merger failure. Or as Max puts it, "when there is a first impression and people don't have a basis for handicapping it properly." This initial oversight is then compounded by the fact that each individual arbitrage fund erroneously takes the other funds' lack of visible concern (the absence of spikes in implied probability) as reassurance that the merger will be completed. This added confidence leads each fund to increase its position, compounding its losses when the merger is canceled. Reflexive modeling, in other words, amplifies the limits of modeling when all arbitrage funds think alike.[7] Whereas the *dissonance* induced by implied probability typically improves the arbitrageurs' estimates, reflexive modeling can also lead to a disastrous form of *resonance*.

[7] Khandani and Lo (2007) explain the crisis of August 2007 in the similarity in strategy across hedge funds.

THE SYSTEMIC RISKS OF REFLEXIVITY

The findings of our study of merger arbitrage address two fundamental questions about the nature of markets and financial crises in the era of quantitative finance. How are markets social when the dominant tools of work are impersonal? And, how do markets collapse when careful reflexivity is the norm?

Our notion of reflexive modeling offers a new way to conceptualize sociability and interdependence in markets shaped by models. Traditional sociological accounts prove to be limited. The embeddedness perspective provides a useful point of departure, but our observations at the merger desk make it clear that those personal ties are absent in arbitrage: the traders were not cognizant of the identity of their trading counterpart.[8]

In that sense, our arbitrageurs can be said to realize the grand vision behind the strawberry market analyzed by Callon (1998). As with French strawberry traders, the introduction of anonymity and calculating tools disentangled arbitrageurs from the rest of their profession, refocusing their attention on their own tools. Mathematical sophistication, in other words, liberated our arbitrageurs from the obligations of constant reciprocity and the penury of the networking dinner. Nonetheless, our traders were too sophisticated to place all their trust on their models.

The skepticism of the arbitrageur, therefore, requires creativity, for the classic solution to the problem of uncertainty – isomorphism – is unavailable to them. The potential candidates for imitation are nowhere to be seen, thanks to the disentanglement that the models set in motion. Instead, arbitrageurs turn back to models for reflexivity. Using models in reverse, the traders find out what their rivals are collectively doing. And as they react to this knowledge, their models introduce a degree of interdependence that is absent in Callon (1998). In short, economic models give back with one hand the interdependence that they took away with the other. These two, however, are not the same. Arbitrageurs are not embedded in personal ties, but they are not completely disentangled either. Decoding the enormous challenges posed by modern finance calls for an understanding of this elusive interdependence.

In advancing along that intellectual path, we draw on Knorr-Cetina's notion of scopes. Arbitrageurs, our analysis suggests, engage with the market through their models, equations, and visualizations. Scopic engagement is different from embeddedness in that it is anonymous, therefore impersonal. It is also more subject to ambiguity in the signs, as a spread of ten percent can mean different things at

[8] Indeed, this is not an accident of technology but a historical shift in the professional background, values and dispositions of Wall Street traders. As Ivan Boesky walked into jail in the late 1980s, merger arbitrageurs shifted their strategy from insider trading to mathematical calculation of merger probability.

different points in time. This ambiguity, and the skill required to handle it, makes scopic engagement more technical, dispassionate, intellectual, detached, than traditional trading. It also makes it less community-oriented, and more organizationally bound. But despite these differences, scopic engagement shares with embeddedness a basic form of interdependence.

Scopic engagement as cognitive interdependence

In its simplest form, merger arbitrage rests on the contraposition of two screens. The goal of the traders is to assess the likely success of corporate marriages, but they cannot observe these companies directly. As a result, they turn to two representations of the deal and exploit the dissonance between them. Arbitrage disasters happen when the two representations march in misleading lockstep.

The first of the two representations, an Excel spreadsheet, summarizes how the traders think about the merger. This so-called Trading Summary builds on a web of associations, including categories and analogies, that lead up to the key issue facing the deal. Every arbitrage desk has one of these spreadsheets, but they cannot see each other's.
The second representation, the spread plot, is shared by all arbitrageurs. The spread plot captures how competitors think about the merger. The spread plot conveys the difference in the prices between the merging companies.

Why does reflexive modeling work? It works, we argue, because the friction between the two screens clue arbitrageurs that they might be missing a material obstacle to the merger. The premise of the system is that whenever arbitrageurs miss a relevant merger obstacle, some rival arbitrageur with a different view will prompt a dissonant chord (a spike in the spread plot), leading to additional search and correction. Thus, the dual screen system lets arbitrageurs "test" their own associations against those of their rivals. Arbitrage, then, is a more iterative, knowledge intensive approach than single-screen trading strategies such as momentum trading or value investment.

The two models/two screens system, thus, provides arbitrageurs with dimensions that are not contained in any single screen. Arbitrageurs compare the two representations and exploit the difference for reflexive purposes. Reflexive modeling, in other words, can be seen as a form of stereoscopic vision, as in the two-eyed form of vision that characterizes human sight. Just as the human brain gains a third dimension –a sense of depth, indeed distance – by the comparison of two flat images (left, right eye), arbitrageurs gain a sense of opportunity and risk.

As the practice of using a model to gain cognitive distance from another model, reflexive modeling is a cognitive process. But it is not taking place in the heads of the traders, as if cognition could be turned back onto itself. Just as the cognitive process of deriving their own probability estimates is socially distributed across the tools and instruments at the arbitrage desk, so reflexive cognition (Stark 2009) is a

socio-technical process of distributed cognition triggered by the spread plot – a device for dissonance that is itself a socio-technically constructed object. Thus, the reflexivity of the traders is not a mental process or a solipsistic practice. The traders we observed were not engaging in some heroic mental feat, splitting and twisting their minds back on themselves like some intellectual variant of a flexible contortionist. Instead, as we saw numerous times in a single morning at a single trading desk, the taken-for-granteds of their models were cognitively disrupted by devices for dissonance.

As we saw, the strength of reflexive modeling is based on the fact that it leverages the cognitive independence among dispersed and anonymous actors. But this same process suggests the possibilities of *cognitive interdependence* among the rival traders in the professional arbitrage community. Just as reflexive modeling can typically be a source of correction, so this same cognitive interdependence among traders can, in rare but dramatic instances, lead to the amplification of error.

Thus, the same mechanism of reflexive modeling that makes arbitrage successful also explains collective failures. When the requisite variety in perspectives is missing from the Excel spreadsheets of the different arbitrageurs, a misrepresentation of the merger will not clash with the spread plot.

Reflexive modeling and systemic risk

Our analysis of arbitrage disasters point to a mechanism of crisis that departs from the dominant views in behavioral economics. Instead of ascribing crises to overconfidence, unpredictability or herding, reflexive modeling offers a mechanism where problems unfold in spite of repeated reassurances, early warnings and an appreciation for independent thinking.

Our findings challenge the behavioral explanation of market crises as overconfidence. Laboratory experiments have shown that repeated betting success leads to increasing rates of risk-taking, even if the underlying probability distribution of the bet remains constant. These results are typically explained in terms of affect or an overly narrow consideration of the data. But this explanation is hard to reconcile with our observations at the trading desk. The arbitrageurs have detailed proprietary databases and model-driven mechanisms of reflexiveness that combat overconfidence. But, as noted, it is precisely this reflexiveness that can induce disaster.

Our analysis of the arbitrage disaster challenges the black swan theory of financial crises. This theory contends that models cause crises by ignoring the possibility of extreme events. By contrast, our observations cast doubt that models are used unreflexively. The traders at the merger desk are keenly aware of the fallibility of their own models, and apply a forward-looking measure to overcome such limitation – implied probability. At the same time, the GE-Honeywell case suggests that the organization of modeling can pose serious dangers. The risks of European

opposition to the merger was not unknown to arbitrageurs, yet they chose not to take it into account. The implication is that disasters are not created by implausible merger obstacles, but by those that arbitrageurs overlook. Put differently, difficulties do not result from the mental rigidity created by the model, but by the cracks in the cognitive structures held by most arbitrageurs.

Reflexive modeling also calls into question the behavioral notion of herding. Models of herding, cascades, and rational imitation explain how a dysfunctional measure of imitation can set into a market, brewing trouble in the future. The merger arbitrageurs that we observed, however, make use of social cues in a far more selective manner. As we noted in the case of the Whitman-Career merger, arbitrageurs use implied probability to widen their search, not to abandon their positions. And, indeed, this is consistent with set of values at the desk that look down on aping competitors. Our notion of reflexive modeling suggests that arbitrage is about developing a view and updating it based on that the rivals. Crucially, such updates might lead to convergence with the market view, but they might also lead to divergence. For that reason, arbitrage disasters are not simply a function of imitation.

The contrast between behavioral accounts of crises and our own points to the benefits of incorporating the tools of trading into the study of quantitative finance. Implicit in the behavioral accounts of systemic risk is a disapproval of investors on moral and intellectual grounds. Ultimately, investors come across as reckless gamblers, mindless lemmings or foolish users of models they do not fully understand. By contrast, our detailed examination of the tools of arbitrage –the models-- removes the need to find character flaws in arbitrageurs to explain crisis.

Arbitrage disasters, we add, point to the ultimate paradox of financial modeling. Disasters befall on traders in the very process of pursuing extraordinary performance. Unless arbitrageurs enlarge their exposure when implied probabilities confirm their estimates, their returns will not exceed the average. But this can lead to oversize losses. Thus, it is *because* arbitrageurs insist in improving their own estimates by turning to models that they risk receiving false confirmation of their views. The dangers of reflexive modeling, in other words, go hand in hand with its benefits.

CONCLUSION

Our discussion so far has important implications for the public-policy debate on the regulation of arbitrage and its various components: derivatives trading, investment banks and the hedge fund industry. Since the late 1980s, banks and other financial institutions have replaced regulatory administrative controls with model-based risk management, leading to a situation of self-regulation by modeling. The persistence of arbitrage disasters points to the limits of this setup, for the arbitrageurs themselves are clearly unable to avoid these disasters by simply trading off less risk

for higher returns. As the GE-Honeywell merger shows, disasters strike precisely when arbitrageurs mistakenly believe they are playing it safe.

If self-regulation is not the answer, what alternatives do our findings suggest? Instead of specific regulatory proposals, our analysis suggests an orienting framework. Our findings should not be read as a denunciation of financial models. Certainly, the models used in backing out implied probabilities are necessary to produce reflexive disasters. But these models are averting other difficulties precisely by allowing the arbitrageurs to be reflexive.

First, recognizing the fallibility of models underscores the importance of diversity of perspectives. Reflexive modeling is based on the assumption that there exist a requisite variety of perspectives among the arbitrage community. Policies that favor participation in arbitrage trading would add to this diversity, while policies that restrict participation will reduce this diversity.

Second, our findings provide a way of thinking about the relative advantages of regulatory disclosure and oversight of arbitrage. Granting a central government agency access to the positions of all arbitrageurs could conceivably let regulators see disasters as the brew up. At the same time, such move is not without its costs. By removing the veil of anonymity, personal ties between the regulator and the arbitrageurs would resurface. Engagement would shift from scopic to embedded. Such embedded regime could mark a return to the 1970s, when the Federal Reserve Bank of New York regularly met with the "money center" –the ten largest New York banks—and occasionally exerted monetary policy through informal suasion.

Third, the regulatory solution to the risks posed by models necessarily involves greater use of models. Given the breadth and complexity of the current financial system, it is inconceivable that regulators could reform the system and prevent future systemic failures without the help of modeling techniques such as stress testing and network analysis. As this reform takes place, further research in economic sociology will be required to understand the interplay between models used for the purpose of profit-seeking, for risk management and for regulatory objectives.

References

Abreu, D. and M. K. Brunnermeier, 2003. "Bubbles and Crashes." Econometrica, 71: 173-204.

Agar, M. H. 1986. *Speaking of Ethnography*. Beverly Hills, CA: Sage.

Baker, W. 1984. "The Social Structure of a National Securities Market." American Journal of Sociology, 89: 775-811

Barley, S. 1986, "Technology as an occasion for structuring: evidence from observations of CT scanners and the social order of radiology departments." Administrative Science Quarterly, 31: 78-108.

Bary, A. 2001. "Deal Me Out." Barron's, 81(26): 43.

Beunza, D. and D. Stark. 2004. "Tools of the Trade: The Socio-Technology of Arbitrage in a Wall Street." Industrial and Corporate Change, 13: 369-400.

Beunza, D. and R. Garud. 2007. "Lemmings, Calculators or Frame-Makers? The Intermediary Role of Securities Analysts." In M. Callon, Y. Millo and F. Muniesa (eds.) Market Devices, special issue of Sociological Review 55(2): 13-39.

Callon, M. 1998. "Introduction: the Embeddedness of Economic Markets in Economics." In M. Callon (ed.), The Laws of the Markets: 1-57. Oxford: Blackwell.

_____ and F. Muniesa. 2005. "Peripheral Vision: Economic Markets as Calculative Collective Devices." Organization Studies, 27:1229-50.

_____ 2007. "What Does It Mean to Say That Economics Is Performative?" in D. MacKenzie, F. Muniesa and L. Siu (eds.), *Do Economists Make Markets? On the Performativity of Economics*: 311-358. Princeton NJ: Princeton University Press.

Cox, J. C., S. Ross and M. Rubinstein. 1979. "Option Theory: A Simplified Approach." Journal of Financial Economics 7(3): 229-263.

Davis, G. and S. Stout. 1992. "Organization theory and the market for corporate control: a dynamic analysis of the characteristics of large takeover targets, 1980-1990." Administrative Science Quarterly, 37: 605-633.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. 1990. "Noise trader risk in financial markets." Journal of Political Economy, 98: 703–738.

Derman, E. 2004. *My Life as a Quant: Reflections on Physics and Finance*. New York: John Wiley and Sons.

DiMaggio, P. J. and W. W. Powell. 1983. "The Iron Cage Revisited: Institutional Isomorphism and Collective Rationality in Organizational Fields." American Sociological Review, 48: 147-60.

_____. 1991. "Introduction." In W. W. Powell and P. J. DiMaggio (eds.), The New Institutionalism in Organizational Analysis: 1-38. Chicago, IL: The

University of Chicago Press.

Dunbar, N. 2000. *Inventing Money: the Story of Long-term Capital Management and the Legends Behind It*. New York: John Wiley & Sons.

Dupuy, J.P. 1989. "Common knowledge, common sense." Theory and Decision, 27: 37-62.

Ferraro, F., J. Pfeffer, and R. I. Sutton. 2005. "Economics Language and Assumptions: How Theories can Become Self-Fulfilling." Academy of Management Review, 30: 8-24.

Fligstein, N. 1990. *The Transformation of Corporate Control*. Cambridge, MA: Harvard University Press.

Fligstein, N. and L. Dauter. 2007. "The Sociology of Markets." Annual Review of Sociology, 33: 105-128.

Fourcade, M. and K. Healy. 2007. "Moral Views of Market Society." Annual Review of Sociology, 33: 285-311.

Garcia-Parpet, M. 2007. "The Social Construction of a Perfect Market: The Strawberry Auction at Fontaines-en-Sologne." In D. MacKenzie, F. Muniesa and L. Siu (eds.), Do Economists Make Markets? On the Performativity of Economics: 20-53. Princeton, NJ: Princeton University Press.

Granovetter, M. S. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." American Journal of Sociology, 19: 481-510.

Hanson, N. 1958. *Patterns of Discovery: An Inquiry into the Conceptual Foundations of Science*. Cambridge: Cambridge University Press.

_____. 1963. *The Concept of the Positron*. Cambridge: Cambridge University Press.

Harrison, M. and D. Kreps. 1979. "Martingales and Multiperiod Securities Markets." Journal of Economic Theory 20: 381-408.

Hayek, F. A. 1945. "The Use of Knowledge in Society." American Economic Review 35: 519-30.

Hoffman, A. 2006. "A good year for some…" Euromoney October: 33-36.

Hutchins E. and T. Klausen. 1996. "Distributed Cognition in an Airline Cockpit." In D. Middleton and T. Engestrom (eds.), Communication and Cognition at Work: 15-34. Cambridge: Cambridge University Press.

_____. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.

_____. 1990. "The Technology of Team Navigation." In J. Galegher, R. Kraut and C. Edigo (eds.) Intellectual Teamwork: 191-220. Hillsdale, MI: Lawrence Erlbaum Associates.

Jorion, P. 2000. "Risk Management Lessons From Long Term Capital Management." European Financial Management, 6: 277-300.

Khandani A. and A. Lo. 2007. "What Happened to the Quants in August 2007? Journal of Investment Management, 5(4): 5-54.

Knight, F. H. (1921) 1971. *Risk, Uncertainty and Profit*. Chicago, IL: University of Chicago Press.

Knorr Cetina, K. and Bruegger, U. 2002. "Global Microstructures: The Virtual Societies of Financial Markets." American Journal of Sociology, 107(4): 905-950.

_____. 2005 "How Are Global Markets Global? The Architecture of a Flow World." In K. Knorr Cetina and A. Preda (eds.), *The Sociology of Financial Markets*: 38-45. Oxford: Oxford University Press.

_____ and A. Preda. 2007. "The Temporalization of Financial Markets: From Network to Flow." Theory, Culture and Society. 24: 116-138.

Krippner, G. 2005. "The Financialization of the American Economy." Socio-Economic Review, 3: 173-208.

Lo, A. 2008. *Hedge Funds An Analytic Perspective*. Princeton, NJ: Princeton University Press.

_____ and A. Khandani. 2007. "What Happened to the Quants in August 2007?" Journal of Investment Management, 5(4): 5-54.

Lowenstein R. 2000. *When Genius Failed: The Rise and Fall of Long-term Capital Management*. Random House: New York.

MacKenzie, D. and Y. Millo. 2003. "Constructing a Market, Performing Theory: The Historical Sociology of a Financial Derivatives Exchange." American Journal of Sociology, 109: 107-145.

_____. 2006. *An Engine, Not a Camera: Not a Camera: How Financial Models Shape Markets*. Cambridge, MA: MIT Press.

_____. 2008. "End-of-the-World-Trade." London Review of Books, May 8.

March, J. G. 1987. "Ambiguity and Accounting: The Elusive Link between Information and Decision Making." Accounting, Organizations, and Society, 12:153-68.

Malkiel, B. 1973. *A Random Walk Down Wall Street*. New York: W. W. Norton & Co.

Officer, M.S. 2007. "Are performance based arbitrage effects detectable? Evidence from merger arbitrage." Journal of Corporate Finance, 13:793-812.

Orlikowski, W. 1992. "The duality of technology: rethinking the concept of technology in organizations." Organization Science, 3: 398-427.

Preda, A. 2006. "Socio-technical agency in coming financial markets: The case of the stock ticker." Social Studies of Science, 36: 753-782

Reverre, S. 2001. *The Complete Arbitrage Deskbook*. New York: McGraw-Hill.

Ricardo, David (1817) 1951. "On the Principles of Political Economy and

Taxation." In P. Sraffa (ed.), The Works and Correspondence of David Ricardo. Vol. 1. Cambridge: Cambridge University Press.

Ross, S. 1976. "The Arbitrage Theory of Capital Asset Pricing." Journal of Economic Theory, 13: 341-360.

Samuelson, P. 1965. "Proof that Properly Anticipated Prices Fluctuate Randomly." Industrial Management Review, 6: 41-49.

Shleifer, A., 2000. *Inefficient Markets: An Introduction to Behavioral Finance*. New York: Oxford University Press.

_____ and R. W. Vishny. 1997. "The Limits of Arbitrage." Journal of Finance, 52: 35-55.

Sidel, R. 2001. "Takeover Traders Are Battered By GE-Honeywell Deal Fallout." Wall Street Journal, June 15: C1.

Spradley, J. 1979. *The Ethnographic Interview*. New York: Harcourt Brace Jovanovitch.

Stark, D. 2009. *The Sense of Dissonance: Accounts of Worth in Economic Life*. Princeton, NJ: Princeton University Press.

Suriowiecki, J. 2004. *The Wisdom of Crowds: Why the Many Are Smarter Than the Few and How Collective Wisdom Shapes Business, Economies, Societies and Nations*. London: Little, Brown.

Taleb, N. 2007. *The Black Swan: The Impact of the Highly Improbable*. New York: Random House.

Zaloom, C. 2003. "Ambiguous Numbers: Trading Technologies and Interpretation in Global Financial Markets." American Ethnologist, 30: 268-272.

Table 1. Arbitrage disasters, 1990-2003.

Acquirer	Target	Cancelation date	Percentage holding by arbitrageurs	Implied total losses, $000s
General Electric Co	Honeywell International Inc	10/2/2001	53	2,798,376
American Home Products Co	Monsanto Co	10/13/1998	45	2,335,367
British Telecommu-nications PLC	MCI Communications	11/10/1997	40	1,908,240
Tellabs Inc	CIENA Co	9/14/1998	34	1,179,412
Investor Group	AMR Co	10/16/1989	36	712,042
Staples Inc Inc	Office Depot	7/2/1997	44	558,804
Investor Group	UAL Co	10/18/1989	29	542,058
Abbott Laboratories	ALZA Co	12/16/1999	46	525,194
Tracinda Corp	Chrysler Co	5/31/1995	42	458,918
Revlon Group	Gillette Co	11/24/1986	25	286,371
Mattel Inc	Hasbro Inc	2/2/1996	228	228,557
McCaw Cellular Communications	LIN Broadcasting	10/10/1989	50	219,937
Amway Co	Avon Products Inc	5/18/1989	29	165,816
Investor Group	Goodyear Tire & Rubber	11/20/1986	25	145,344

This table contains details of the fifteen largest merger arbitrage disasters from 1985 to 2004. All dollar arbitrage losses are in 2004 dollars. Arbitrageurs' percentage holding is the percent of target shares outstanding reported as owned by arbitrageurs at the first quarterly 13F reporting date after the bid announcement date. Implied dollar arbitrage loss is the total arbitrage loss multiplied by arbitrageurs' percentage. Source: Officer (2007).

Figure 1. Charting the implicit probability of merger.

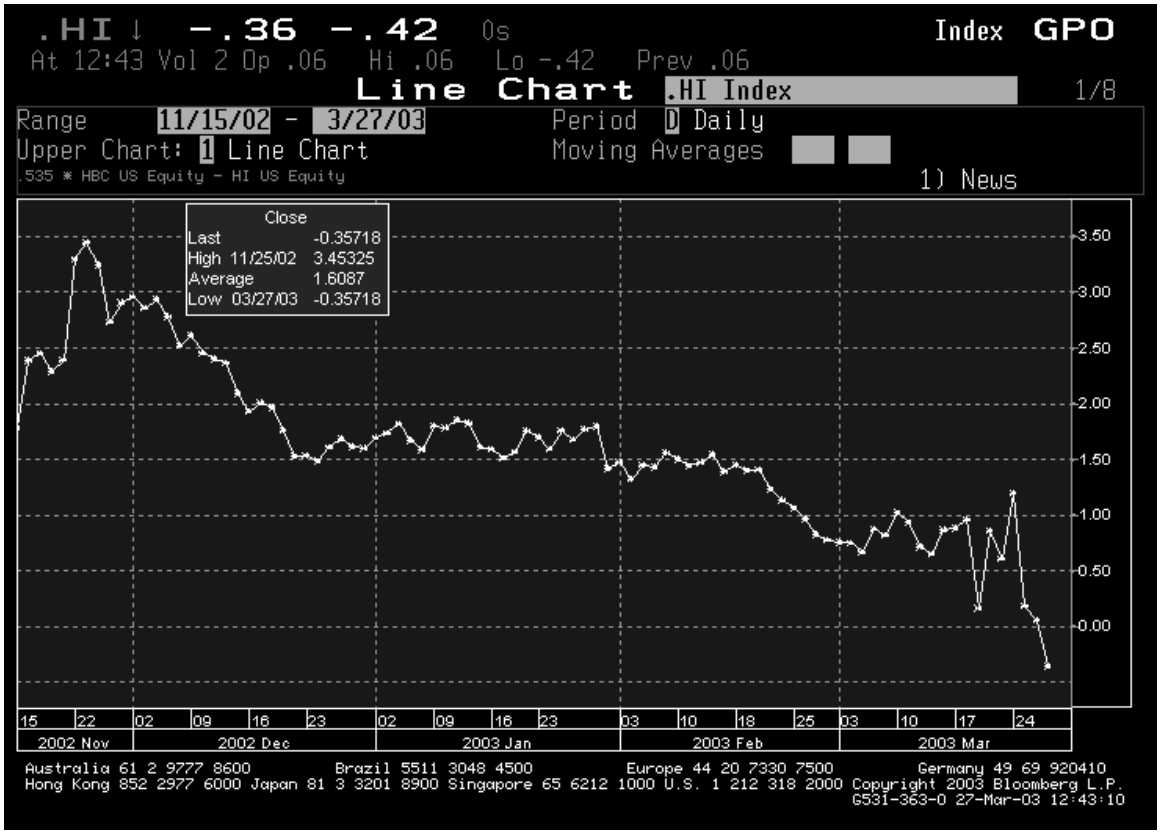

Screen shot of a Bloomberg terminal showing the spread plot of Household International and HSBC Bank, November 2002 to May 2004. Source: International Securities.

Figure 2. Charting the implicit probability of merger.

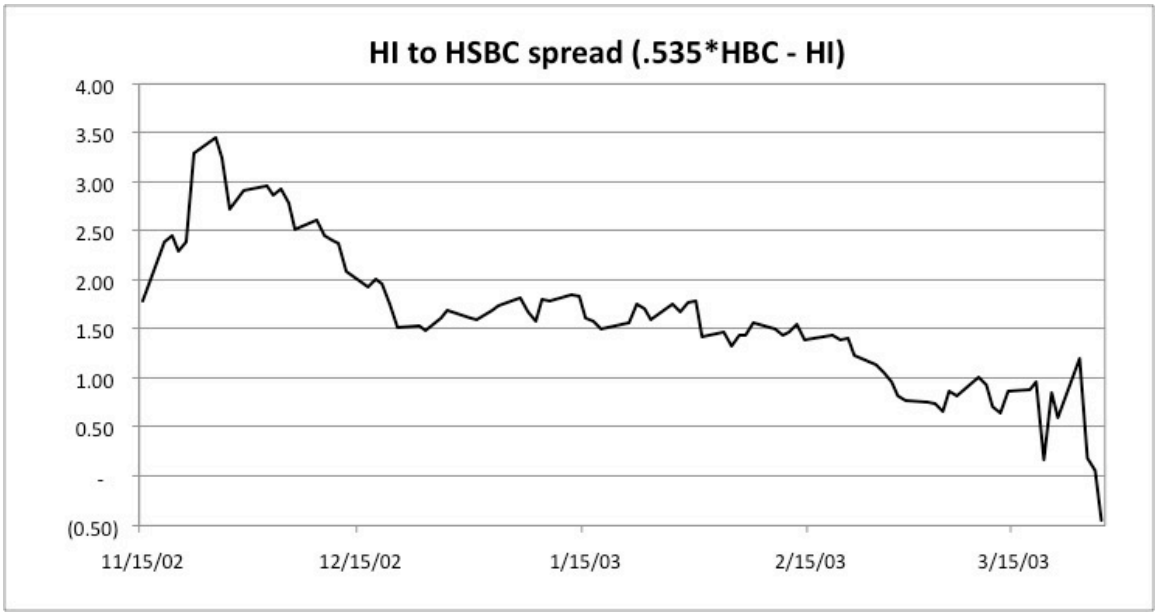

Spread plot of Household International and HSBC Bank. The two spikes in the figure, November 22nd and March 20th, correspond to events that called the merger into question. 2002 to May 2004. Source: Bloomberg.

Figure 3. Merger arbitrage disasters

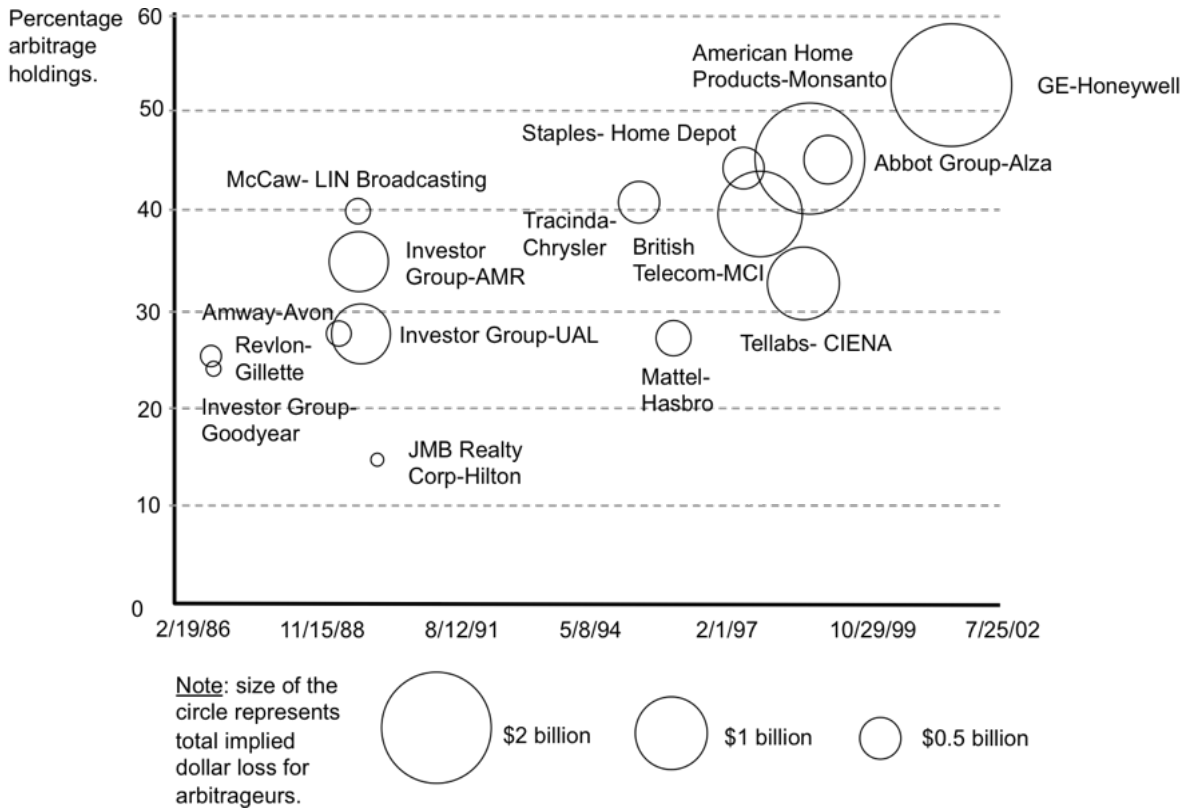

Failed arbitrage deals, with total losses incurred by arbitrageurs (circle size) and relative participation of arbitrageurs in (y-axis). Source: Officer 2007: 27.

Fig. 4: Arbitrageurs overlooked the danger of European opposition.

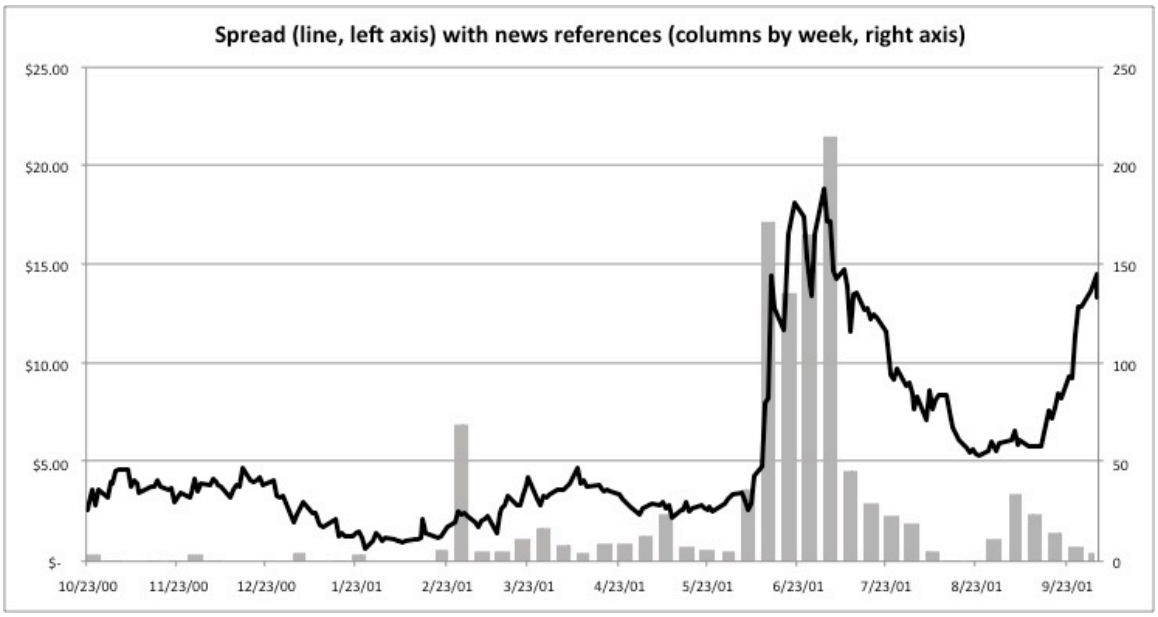

Spread between GE and Honeywell (line) and media concern over EC opposition to the merger. The graph shows that the surge in media concern was not matched by a corresponding increase in the merger spread. Source: Bloomberg and ABI/Inform.