

Survival of the Reckless: How the U.S. Mortgage Market Evolved Toward Disaster

Michael G. Jacobides
London Business School

Sidney G. Winter
The Wharton School

Prepared for the American Economic Association session
“Evolutionary Perspectives on the Financial Crisis”
Atlanta, GA, January 4, 2010

Preliminary draft, incompletely referenced, for discussion only
December 31, 2009

Abstract

From today's standpoint, it is easy enough to see that the U.S. economy has not been well served by the mortgage lending institutions that became dominant over the past several decades. The financial crisis as a whole seems over-determined in retrospect, it was a "perfect storm." The debate rages over the appropriate weightings of the candidate causes. In this paper, we focus relatively narrowly on the functioning of the mortgage market from which the financial crisis emerged, on the related transformations of the financial sector, and on the structures and incentives that developed in it. Our theoretical approach is firmly in the evolutionary tradition, emphasizing learning as the source of competence and feedback as the driver of behavior --- and disputing the view that the market system can adequately compensate, through the "market test," for the incompetence of individual actors, as well as the view that economic behavior can be entirely understood in terms of calculation of consequences in an imagined future.

Survival of the Reckless: How the U.S. Mortgage Market Evolved Toward Disaster

I. Introduction

The evolutionary perspective on the financial crisis has at least one major theme in common with many other perspectives that have been offered since the crisis flowered in 2008. That theme is, “I told you so!” While there is a significant element of truth in many of the most prominent claims that embrace that central theme, all fall short of being fully convincing on the key points. Some claims are, of course, more convincing than others. There were, for example, some emphatic warnings of the existence of a housing bubble, which were firmly reiterated in the face of dismissive opposition. Similarly, some observers were alert to the broader dangers, underscored by the historical precedents, of debt-financed prosperity. In the end, however, it is hard to sustain the claim that anybody was endowed with a true, detailed, Cassandra-like vision of the full dangers of the developing situation. The actual events revealed the true force of some considerations, particularly the extent and implications of the short-term interconnectedness of the global financial system, that no observer had clearly identified as key.

We share the view that the crisis can reasonably be regarded as a “perfect storm,” its overall scale and consequences being the result of the concatenation of a substantial list of vulnerability-enhancing factors. Although there is a long list of such candidate causes, there is not much dispute about what lies at the center of the storm into which they fed. At the center of the perfect storm, we find the U.S. housing bubble and the speculative mentality that sustained it (Shiller 2005; Shiller 2008; Akerlof and Shiller 2009) . Near that center we find the substantial increase in household indebtedness in the US and the UK, as well as the role of China in absorbing the debt created, and also in keeping global inflationary pressures in check through its aggressive export policies.. These macro-economic points are hard to deny, and we certainly do

not deny them here. We acknowledge their centrality, yet seek to complement this by a focus on the institutional and evolutionary forces that played a considerable and less studied role.

In this paper, we explore the extent to which behavioral factors at the individual and (particularly) the firm and industry levels contributed to the debacle. A number of such factors are discernable, in particular, in the operation of the U.S. mortgage market, which is our principal empirical focus. We also consider the changes in the practices of broader swaths of the financial services sector, where key contributing factors have been the changes in the industrial structure and associated changes in the mechanisms of risk taking and risk re-distribution at the societal level. There has been substantial attention to the fact that some of the practices that emerged in that market, particularly toward the middle of the decade, seemingly reflected a reckless disregard for financial risk. There has not been corresponding attention to the changing institutional structure of the sector, involving the emergence of new types of vertical specialists, and new products and services. While securitization and derivatives have received attention from financial economists in terms of their properties, we know little about how these new practices changed the operation of institutions and individuals within the sector.

Our interest lies in understanding how the limitations of individuals and firms, the character of sectoral structures, and the interactions of these fed into an evolutionary process that contributed to the collapse -- fuelled by (and further reinforcing) the macroeconomic imbalances. Our particular objective is to understand how the financial sector started accepting an extraordinary level of risk. Since the risks subsequently materialized in catastrophic fashion, there is no need to engage the question of whether the practices “really” were risky. They were. What is more interesting, and of broader significance, is the question of *why* the very real risks were overlooked, mis-appraised and deemed acceptable – i.e., the question is why the *reckless*

disregard happened. Is it relevant, for example, that the system that malfunctioned was a relatively new one, in the sense that the institutions of the mortgage marketplace, and of the financial service sector more broadly, were different in 2005 than they were in 1995? In what ways did the changes in the division of labor between different institutions, and the emergence of new products and rules linking institutions lead to this collective meltdown? Should we have expected that such changes could have been collectively detrimental, and if so, why? Is there anything beyond blaming individuals that might help us understand what brought about the crisis, and how could we have averted some of these firm- and sector- level policies that exacerbated, or perhaps led to the crisis? If so, what sorts of policy concerns can we add to the gamut already considered? Does the story of the crisis have important lessons for us regarding the conditions under which possible shortfalls in actor competence are likely to be checked and mitigated by the discipline of the market place? And finally, what conclusions can be drawn regarding the choice of new financial institutions and new directions for economic analysis?

The “I told you so” from evolutionary economics is simply this: There was never any good reason to presume competent actors, *in general*, i.e., independent of context and circumstance. Individuals and organizations do sometimes perform with high competence. On other occasions, including many presented in the recent crisis and the run-up thereto, they perform with incompetence that seems laughable from the viewpoint of any slightly more informed, or somewhat more disinterested, observer. The problem of explaining why we sometimes observe the one and sometimes the other, particularly at the organizational level, has been called “the competence puzzle”(Nelson and Winter 2002). Previously achieved insights into that puzzle are reviewed here in the context of understanding what happened in the mortgage market.

Also, drawing on some recent work on how “industry architectures” evolve, our second central claim is that the structure of the financial services sector, the division of labor between industry participants and the set of rules and roles that emerged, were a major part of the problem. The lack of attention of regulators and policy-makers to the changing ways that labor was divided, and to set of capabilities, incentives and trajectories that subsequently emerged, helped assure short-term returns to some of the industry participants -- but ultimately allowed the system to become vulnerable to collapse.

In the following section, we review some basic principles that are long-established in the research tradition of evolutionary economics, taking note of some of the many vivid illustrations of these principles offered by the history of the financial crisis. Next, we briefly consider the nature and role of industry architectures, and how a better understanding of the division of labor and the links between institutions can inform our understanding of the crisis. We then review the specific example of the evolution of the mortgage banking industry in the U.S., noting the specific drivers of change – and how little those drivers reflected the longer-term hazards that ultimately materialized in that sector. We also consider some of the broader changes in the financial services sector, in particular as they pertain to the mortgage risks at the core of this crisis. We then link our theoretical analysis to the situation at hand. We briefly describe why the architecture of the financial services sector changed; how that mattered; and in what way a more intelligent “industry design” might have averted or mitigated the crisis. In the final section, we attempt to bring the evolutionary perspective to bear on the difficult policy problems presented by the widely acknowledged objective, “How can we make sure that this doesn’t happen again?”

II. The “Market Test” and the Question of Competence

Among economists who are actively concerned with understanding the economy as it is, there probably are very few who would seriously argue that rational choice models provide a good scientific account of economic behavior at the individual actor level. Even Milton Friedman, in his famous essay (Friedman 1953) defending the dominant methodological approach in economics, did not go so far as to defend rational choice modeling in that context. Rather, he implicitly conceded its inadequacy, and went on to defend its pragmatic usefulness as an approach to larger questions. Since Friedman wrote, an enormous amount of systematic empirical evidence has accumulated on the point, and it weighs very heavily against the rational choice scheme.¹ The buildup of this contrary evidence has accelerated in recent years, and some significant alternatives to rational choice have been forward. Meanwhile, the textbooks sustain long-standing traditions on the matter – and to a remarkable extent, advanced theoretical work does as well. Sometimes it appears that almost any disregard of the complexities of the real world is legitimate in economics, provided it paves the way for a rational choice model.

This situation leaves it quite unclear what it is that (reality-oriented) economists might actually endorse as a good scientific account of behavior at the actor level – and it is important here to understand “actor” as including “firms,” consistent with the main theoretical tradition in economics. It may well be that there is a stronger consensus supporting the view that economic inquiry doesn’t really require an answer to this question than there is for any particular answer. There appears to be a strong tendency to accept the general thrust of Friedman’s argument, suggesting that possible deficiencies in the rational actor model are of limited relevance for

¹ See for example the classic work on biases and heuristics;(Kahneman, Slovic et al. 1982; Einhorn and Hogarth 1985) Thaler and xxx, 2009; or observations on individual and collective psychology enumerated by (Akerlof and Shiller 2009)..

practical purposes because of the operation of other mechanisms. From an “instrumental” point of view, therefore, a lot of important truth is captured by the rational actor model – in its innumerable specific manifestations, and notwithstanding any demonstrated shortcomings it may have. We actually share this assessment, up to a point. The question is, what point?

Prominent among the alternative mechanisms that Friedman explicitly invoked, and that many economists still seem to believe in, is some version of a “natural selection” argument. This argument says that actors – now meaning *firms* in particular – who stray too far from the dictates of a rational analysis will be punished, and ultimately suppressed, for doing so. They will be punished by the market environment, for failing the “market test.” The concept of a “market test” for firm behavior is widely invoked, in diverse contexts, and seemingly is intuitive for many economists. Many committed advocates of market-based social organization, who would definitely eschew any automatic endorsement of the rationality of particular actors, nevertheless think the market test is real, effective and profoundly important. And again, we agree with that -- up to a point.

Reviewing the development of the financial crisis, it is easy to identify numerous examples of firm behaviors that passed the market test with plenty of room to spare – for a while. In retrospect, these same behaviors are frequently assessed as destructive of important interests, including those of the actor in question and of society at large. These examples give reason to consider more carefully the degree to which the market test provides a trustworthy warrant for the competence of individual actors (or those who have survived for a period of time), and also a warrant for the corollary proposition that society has little to fear from actions that are, or at least appear to be, powerfully “dis-incented” for the perpetrator. Warren Buffet famously warned

against derivatives as “financial weapons of mass destruction;” here the question is whether Richard Fuld (for example) should be acknowledged as a financial suicide bomber.²

How well, and under what circumstances, does the “market test” protect the system from possible deficiencies of rationality at the actor level? To put it differently, if we reject the rational model because it is substantially disconfirmed empirically, where are we look for an alternative, systematic account of actor behavior? This is an important question from the viewpoint of understanding the sources and implications of the crisis. That sort of context reveals specifically the need for an account of *firm* behavior that reaches to the types of decisions actually made (as distinguished from stylized ones designed for experimental investigation). To illuminate these issues, we can turn first to previous work in evolutionary economics, in which the question of the efficacy of the market test is basic and has a long history. In this section, we reprise the basic analysis of the requirements for evolving high levels of competence in a population of actors who are not individually capable of the sort of careful, foresighted optimization envisaged in standard theory.

A. Selection of Competence: Requisites and Mechanisms³

Sources of behavioral continuity. As noted by Penrose long ago, an evolutionary process in economic life must be based on something that plays the role played by genes in biological evolution. Absent some such mechanism that establishes a strong correlation across time in survival-related characteristics, there is no reason to believe that passing the test imposed by the

² As CEO of Lehman Brothers, Fuld resisted the strong urgings of government officials to sell the company rather than risk having it slip into bankruptcy. Then he is said to have held onto his 10 million shares of Lehman stock until the end, losing about \$1 billion. Since he remained a wealthy man, suffering nothing like the significant lifestyle setbacks that the crisis inflicted on others, it is possible to assess his overall performance as a kind of success, -- perhaps a manageable loss suffered on a rationally accepted gamble. We set aside this line of inquiry here, noting that it leads ultimately to an explanation of the crisis as the result of deliberate risk-taking by top Wall Street executives, with no element of misunderstanding or incompetence involved. We do not accept that as a *primary* explanation, but it is not an easy point to prove.

³ The analysis that follows draws on previous work such as (Winter 1964; Winter 1987; Nelson and Winter 2002).

environment of time t has much to do with the quality of performance to be expected at time $t+1$. In modern evolutionary economics, it is generally supposed that the needed behavioral continuity arises from the relative constancy of what we may call “behavioral propensities” – tendencies to respond to situations in characteristic ways.

In evolutionary theorizing, these “propensities” often go by the name “routines,” but similar implications follow from things called rules, skills, habits, systems, practices, procedures, policies or, sometimes, technologies.⁴ All of these have, by their nature, a broader scope than can be manifested in a single action setting. Behavioral inertia attaches much more securely and fundamentally to these propensities than to specific actions. Thus, while behavioral propensities are sometimes reflected very visibly in sequences of highly similar actions across time, they typically underpin more complex patterns that are systematically responsive to the environment - - obviously including the transactional environment of the marketplace, but also the clock and the calendar.

Although both the denotations and the connotations of the terms differ, all share the important feature of involving a significant disjunction between the factors operative in *creation* and *exercise* – e.g. the distinction between the shift to new accounting systems and the repetitive use of those systems. To the extent that intelligent deliberation plays a role in the establishment of the behavioral propensities, it typically does so in the creation, or learning, phase. Deliberation plays at most an attenuated role in exercise, and sometimes plays none at all – full automaticity prevails. It is these propensities that organizations become committed to, through deliberate investments and by a range of non-deliberate and even subconscious mechanisms such

⁴ The term “behavioral propensities” has the advantage that it avoids the connotations specific to some of the listed terms, which distract attention from the important similarities in the phenomena referred to. It also has the virtue of evoking the idea that the relevant regularities are expressed in statistical tendencies rather than in highly deterministic, formal or purely nominal “laws” of behavior.

as “implicit learning” and habit formation. Because exercise is often non-deliberative, possibilities for dysfunctional exercise arise that simply are not within the appropriate scope of the skeptical question, “Why would anybody *do* that?” – because it was not a thinking person (or firm) that did it. The “good news” on the other hand, is that non-deliberative response is typically much faster and does not impose the costs of “reinventing the wheel” when a familiar situation is encountered.

Competence is learning-based. It is an evident fact that actors are sometimes, and in some ways, highly competent within the restricted spheres of their actual activity. Evolutionary analysis directs attention to the sources and limits of competence, and locates these primarily in the way that the behavioral propensities of the actor have been shaped by past learning and experience. Differences in competence, whether between actors or within actors over longer periods of time, reflect differences in history, modified by considerations of current context and social control. At the simplest level, there are typically major differences between improvised performance, in response to problems not encountered previously, and responses to familiar problems that have been well practiced and well tested by the environment. For problems of the latter kind, the relatively automatic responses arising from established behavioral propensities are often quite effective; this is how the world’s work mostly gets done. Even when they are not effective, actors typically believe that they are, which is a point of independent importance. Improvisation, by contrast, naturally tends to have an erratic track record,

The central role of learning and experience is directly rendered invisible by any theoretical scheme that ascribes a sharply defined decision problem to a representative actor, and describes that problem in a timeless setting in which nothing about the history of the actor, the problem, or the relationship between the two, has any chance to intrude into the story – and then

presumes that the actor to be fully capable of getting the right answer on the first try. The more complex the problem, and the greater the role of time pressure on the quest for a solution, the more doubtful such a presumption becomes.

To acknowledge that competence is based in learning and experience is to recognize that innovation, which presents new problems to actors, can be dangerous. Indeed, evolution itself is dangerous because it can, by its own workings, change a system the actors are familiar with into one that makes their past learning obsolete. We pursue the implications of these observations later in the paper.

Evolution takes time. In the long tradition of casual invocation of economic “natural selection” arguments in support of the standard assumptions of “profit maximization,” this simple point is very commonly ignored. One might think, from such accounts, that any momentary departure from a rigorous adherence to profit maximization would quickly prove fatal. “But “profit” is properly understood as a flow concept, and it is clear that a brief interval of deficient performance, relative to the prevailing competitive standard (if not the “maximization” standard of theory), is unlikely to have conclusive effect on the fate of an organization. Thus, the logic offered in these casual arguments essentially fails to survive a units check.⁵ In fact, large companies often seem to persist in a significant way in spite of displaying noticeably deficient performance for quite extended periods – the example of General Motors comes to mind.

⁵ The same problem arises in the not-so-casual, and very well known, paper on irrational consumers by Gary Becker (Becker 1962). Becker offers a quasi-evolutionary argument in support of the conclusions of standard consumer theory. The driving assumption is that even irrational consumers will not buy what they cannot afford – but this is explored in the setting of a conventional single-period consumer model, or alternatively, one that is entirely static. In that context, consumers cannot buy bigger homes they can afford because they cannot take out bigger mortgage loans than they can afford to repay. Perhaps they cannot do that forever, but in the real world context they can do it for a while – and in the run-up to the crisis they did so, on a massive scale.

It is the cumulative effect of competitive shortfalls that ultimately renders the verdict and declares the firm to have failed the market test. The immediate implication is that the market test necessarily falls short of being a substitute for pervasive competence in one key respect: It does not protect against the *appearance* of incompetence; neither does it directly imply anything about how long an appearance might last. “Anything can happen for a while” captures the time-consuming nature of the market test, and the duration of an asset bubble can be quite a while.

Beyond that, the consequences of differential performance at the system level depend on how it is registered organizationally and institutionally. Selection can operate at the firm level, by driving the worst performers to bankruptcy, or by enhancing the relative growth of the best performers. It can also operate directly at the level of behavioral propensities, as in the diffusion through imitation of the practices of the best performers, or the tendency of stress conditions to lessen the grip of inert commitments to established ways of doing things.

Selection via bankruptcy. The idea that there is a long-run environmental constraint on firm behavior, is most straightforwardly illustrated in the case of bankruptcy. Because bankruptcy, and fears thereof, played such a prominent role in the crisis, we elaborate on it here in a way that is disproportionate to the role of bankruptcy in more normal times.

It is clear that the ability of a firm to pay its bills when due is a fundamental consideration that affects firms survival through the institution of bankruptcy law, as well as in other ways. Disastrous behavioral propensities are sometimes checked effectively as the carriers of “inefficient genes” (including propensities to accept poor or excessive investment risks) are weeded out through such inefficient firms becoming bankrupt. Yet this selection mechanism operates, as we shall see, with considerable indirectness. To play the constructive evolutionary role suggested by the “market test” concept, it should reflect the direct financial consequences

generated over time by the repetitive exercise of the firm's routines. In fact, it is often detached from those consequences (as will now be explained.) Thus, whatever the implications for a firm's survival chances arising from its inability to pay the bills in some particular "short run" situation, the status of such a development as a verdict on the firm's established ways of doing things is much less clear.

The detachment is attributable in the first instance to the fact that the cash derived from past successes does not accumulate, in some special account, as a "rainy day fund" to buffer the firm from temporary episodes of difficulty. Rather (in a world of limited liability), it gets paid out, irreversibly, to shareholders as dividends and to other stakeholders in other forms— such as executive bonuses. This means that firms typically operate with a degree of financial vulnerability that is partly attributable, not to the quality of the productive performance, but to the payout policy.

To put it mildly, the underfunding of the hypothetical "rainy day fund" is not a puzzling phenomenon, given the operative incentives. Instead, those incentives tend to produce underfunding much more broadly – underfunding of the reserves that would enable permit an insurer to fulfill its contractual obligations, underfunding of pension benefits, deficient bank capital, and so forth. Strong incentives for stakeholders to "take the money and run" are a feature of the institutional system. They are obviously grounded in self-interest, but not infrequently are supported also by legitimate (though still private) concerns about what will happen to the money if you *don't* take it and run.

Thus, cautious firm behaviors that are reasonable and prudent from an "objective" viewpoint – meaning, any viewpoint not directly contaminated by myopic greed – are not necessarily viable from an evolutionary viewpoint. Firms displaying sound business practices

that generate fluctuating returns may succumb to short-term adversity, given the hazards and difficulties of maintaining an adequate financial buffer. What is worse, the adversity itself may be the consequence of competitive pressure from growing firms that are not burdened by propensities toward prudent behavior; evolution can favor the reckless, at least in the short run. The second source of the detachment referred to is that the state of a firm's balance sheet is ultimately a reflection of expectations about the future. Here, the word "ultimately" acknowledges a contrast with what is true for most firms, most of the time. Within the intrinsic limitations of financial accounting, and subject to the qualifications about payout mentioned previously, a firm's balance sheet is ordinarily a relatively passive, backward-looking measuring instrument. It provides a valuable, though imperfect, indication of the cumulative effect of the firm's past operations. In conditions of exceptional threat or opportunity, however, balance sheet transactions (or the absence of anticipated ones, such as loan roll-overs) may change the picture dramatically over a short period of time.

The emergence and management of the financial crisis in 2008 illustrates the interplay of cumulative experience and payout practices, registered in firm balance sheets, versus changing expectations about the future -- the two different facets of the general "detachment" possibility. Consider, for example, the somewhat contentious question of why government intervention allowed Bear Stearns to avoid outright bankruptcy, but did not do that in the case of Lehman Brothers. According to accounts of the frantic quest for a deal to save Lehman, a major factor was the poor quality of Lehman assets (Cohan 2009; Stewart 2009). It seems that, even against a background of severely depressed expectations, the "hole" that new equity needed to fill was revealed by close examination of the books to be unexpectedly large, particularly because of the weakness in commercial real estate loans. This posed a problem for the potential buyer of the

company, and also forced the Fed to confront the question of the legal limits to its discretion in accepting dubious collateral. Thus, Lehman's "balance sheet" (at the detailed level) did provide a cumulative record of fatally flawed operations in the past, operations that ultimately "caught up with it." That record defeated a major public and private effort to fend off the rising systemic risk by converging on new expectations and saving the firm.

The "detachment" point here is that the operation of these mechanisms on the firm's balance sheet, with their implications for the firm's survival and growth, cannot be regarded simply as a "verdict of experience" on the firm's operations up to that time. In principle, expectations and/or payout practices could operate to overthrow the verdict of experience, in either direction – producing failure where experience indicated success, and success where it indicated failure. At the practical level, of course, experience to date is typically a significant influence on expectations, as the Lehman example illustrates.

The third reason for the detachment has to do with the typical bluntness of the selection mechanism considered as feedback on the performance of a whole organization. A firm, especially a large and diversified firm, has a mind-boggling array of diverse practices, while financial results and bankruptcy-based selection processes operates on the single amalgamated entity. The case of AIG is instructive. A generally well-run (if under-reserved) insurance firm allowed one of its branches to run an illusorily profitable business out of London, issuing credit default swaps. But for the government intervention motivated by concerns for the stability of the system, the CDS business would have led the firm to bankruptcy – in the event, it wound up in partial public ownership. Similarly, Lehman had a set of healthy areas of business unrelated to the risky real estate positions that brought down the entire organization. Arguably, this sort of bundling of strength and weakness in a single firm should not be a problem, provided it is

possible to separate the healthy and unhealthy parts of a business, and ultimately to distinguish the practices that contributed to failure from those that did not.

While efforts to accomplish such separation are often a feature of bankruptcy proceedings (including Lehman's), there are major obstacles to full success. First, it is often very difficult to identify the practices that were, in isolation, dangerous; and, as we will elaborate below, *the very same practices that are healthy in one environment can become toxic in another* – a corollary of the point that behavioral propensities persist, and control much more than individual actions. When the environment changes, there may be delayed or inadequate feedback to provide warning of the dangers. And second, even when it is known what is valuable and what is not, the legal tangles of large financial institutions are making it harder to separate healthier bits from more problematic areas of the firm. Rather than allowing a bankrupt firm to re-organize, weeding out the ineffective parts of it and allowing the healthy parts to survive, it may be necessary to let the entire organization perish.

In sum, if bankruptcy is conceived as the principal “enforcer” of economic selection processes, there are important qualifications to its suitability for the role. We do not mean to suggest that these qualifications seriously detract from the basic idea of the market test, or the significance of bankruptcy as an enforcer. Rather, these qualifications suggest some of the reasons for doubting the faith that the market test will support the most competent actors, and *quickly* suppress incompetent ones. (There are many others.) The qualifications are potentially relevant to relatively detailed questions of institutional design relating to how, precisely, the test should work.

Selection via differential growth. We now consider briefly the role of the other major mechanism of firm-level selection, the differential growth of firms. First, we should consider the

general point that the relative prevalence of particular practices is the result of the relative growth of the firms that carry such practices, or the entry and growth of such firms. Firms that exhibit high profitability are likely to be rewarded by the capital markets, thus obtaining additional resources with which to expand the scope of application of their practices – and this remains true when the profitability is attributable to risky activities, as long as the risks have not materialized. Note here that ownership and governance structures play a key role in this regard, as they mediate between the ability to finance growth and the policy decision to pursue it. *In theory*, these structures could endow provide the firm with the ability to resist the temptation of expanding risky but profitable activities, by introducing a dose of reason or a note of caution into the proceedings.. In practice, the temptation proves hard to resist. In the absence of any *particular feedback* that would arouse concern in the management, there is little reason to expect risky practices to be abandoned – because of the strong tendency for actors “satisfice” with respect to their established propensities (Winter 1971). They might be abandoned by firms receiving negative feedback, those whose realized results make them unable to grow, unwilling to continue engaging in such practices, or both. These restraints, however, are unlikely to be effective in a timely fashion when the story unfolds in a context where the feedback is delayed -- such as in financial services, where risk is typically revealed at a late stage. And it is all the less likely when a number of firms are pursuing similar risky programs, and in the process generating temporary validation for each other’s behavior by bidding up the prices of risky assets.

Selection via imitation. Practices, or behavioral propensities generally, can become more prevalent in the system simply because more firms adopt them. According to a very large empirical literature on the diffusion of innovations, such adoptions become more probable when the practice is seen as contributing to realized success for actors who currently employ it. When

the practice involves the quest for profit from risky investments, the same logic of delayed feedback operates for imitation processes as for growth.

The siren song of temptation is magnified for relatively prudent firms by a decision mechanism featured in some evolutionary models: Recognizing that its returns are inferior to those of its peers, a firm searches for better policies, frequently ending by adopting the policies of the peers. By this mechanism, pressures for growth and profit thus convince other relatively prudent firms to abandon their practices and join in the pursuit of what appears to be “easy money”. Of course, the same mechanisms also operate constructively in the economy all of the time, generally acting to diffuse practices that are in fact superior. But in the buildup to the crisis, it was toxic practices that were promoted at the expense of prudent ones.

Prudent corporate practices can be a career hazard for executives. Consider Jamie Dimon of JP Morgan Chase, now widely considered to be one of the most competent managers in the sector. If the crisis had not started in 2007, and the meltdown of 2008 had been put off by a year or so, it is quite conceivable that Dimon and other similarly disposed executives would have been replaced under the pressure of “the market” (felt and transmitted by an imitation-driven board), for growth and additional returns.

B. The Selection Environment Is Endogenous

The statement, “Everything depends on everything else,” is sometimes offered as a (very) simple bottom-line summary of the message of general equilibrium theory. One of the dependent “things” that depends on everything else is the viability of the practices of an individual firm. That is, there is no sharply defined criterion for the viability of an individual firm that is independent of “everything else” – in particular, of the competitive standard

determined by the aggregate behavior of competitors, and of the terms that buyers and suppliers offer in the market place.

As an analytical exercise, one can construct a simple version of the viability question in the context of an assumed (proximate) equilibrium in the system as a whole. In that restricted context, the criterion for the viability of the individual firm is relatively clear-cut in principle. The market environment provides some signals, including price signals, and the firm responds to those signals with the specific actions dictated by its routines (propensities). If those actions fail to produce at least a break-even financial performance, overall, then the firm is non-viable. Assuming, plausibly, that there are causal connections from profitability to growth/decline and survival/failure, the firm will fade and ultimately disappear -- thus moving the system as a whole an additional small distance toward the full equilibrium.

It should be noted that the viability of the firm in the posited equilibrium would not provide a general warrant for the competitive consequences of its behavioral propensities under other conditions. Further, as discussed above, evolution takes time. Behaviors that will “ultimately” be suppressed by the competitive environment can make a difference in the meantime. There are other theoretical qualifications that could be noted for what is, overall, a generally reasonable argument that is modestly informative regarding the implications of the market test when the market as a whole is near equilibrium. The question of the practical relevance of this analytical exercise remains open, however -- considerably more open than that of its theoretical validity -- since the posited all-but-one-firm, general near-equilibrium would be a hard thing to locate in the real world. It probably it has *some* practical relevance nevertheless.

The quasi-evolutionary argument just sketched actually tells us very little about might or might not happen in the evolution of a system of diverse firms that are still in the process of

enacting, collectively, a changing environment that is testing all of them at the same time. The feedback received by an individual firm at a particular time, including the signals of profit or loss, is a consequence of the whole ecology of behavioral propensities at that time. Aggregated across the firms, this feedback will largely determine the different ecology of the subsequent time period. Iterated over time, this is the process of economic evolution, and it is intrinsically very complex. While strong theoretical assumptions can isolate analytically tractable sub-cases – the proximate equilibrium story above being one very simple example -- such an intellectual move does not actually make the fact of complexity go away, though it might tend to obscure it.

For the purpose of understanding the financial crisis, there are three ways in which we can specialize the general point about the endogeneity of the environment, each of which deserves attention. The first concerns the local feedback to an actor, the way the signals received by a focal actor depend on the behavior of everybody else. The second involves “bubble psychology,” the way the collectively-generated asset price trends and the shared perception of the economic situation affected a wide range of behaviors. The third is at a still larger scale, involving the mutual dependence between the “industry architecture,” – i.e., the structure shaping the interactions among firms , and the actual content of those interactions. The third is of distinctive importance, and we largely defer its discussion to the following section.

Local feedback. On the local feedback aspect, we start from an aphorism that captures much of the evolutionary perspective: *Feedback, not foresight, drives economic behavior.* The validity of that overly concise summary of the evolutionary view is powerfully underscored by the crisis of 2008 and the events on the long path that led up to it. For many different actors and at many different levels, limited foresight and misleading feedback played a key causal role in the developments. In some parts of the picture, the complexities are such that the inability of

actors to see through them hardly seems surprising. Even assessments aided by hindsight remain controversial, and are inconclusive in that sense at least, although the propounders often exude high subjective confidence. In other areas, however, the basic puzzle that persists relates to how “obvious” considerations could possibly have been so widely missed or so seriously misappraised – even by presumptively sophisticated actors who seemingly had strong incentives to do otherwise, and who were highly favored in terms of access to relevant facts and theoretical logics.

Consider, for example, the roles of the mortgage broker and the mortgage originators with whom the broker deals. It seems obvious enough that the incentives of a mortgage broker, who retains no stake in the transaction once the deal is closed and the commission collected, present an “agency problem” with respect to the quality of the loan. If that is obvious, it must be equally obvious that controls are needed to offset the broker’s incentive to generate low-quality loan business at the expense, ultimately, of the mortgage investor whose capital is at risk. Indeed, back at the institutional origins of the mortgage broker role, the creation of routinized control mechanisms was a prominent part of the story, which doubtless would not have gone forward otherwise. Then, does the chain of obvious implications continue? How obvious is it that the routinized control mechanisms should be reviewed from time to time, with respect to their effectiveness, alignment with changing circumstances, and even their fundamental logic? In the end, the actual practice of the broker role in some markets evolved into an exaggerated manifestation of the “obvious” underlying incentive problems – so exaggerated that, in retrospect, the formal efforts at control seem perfunctory at best and fraudulently deceptive at worst. Not only were flawed practices not suppressed, they were actively encouraged – in particular by the feedback from the originators, who generally pushed the brokers to bring even

more. The originators in turn were pushed by the feedback from their customers up the chain. The locus of ultimate responsibility for the consequences remains obscure, shrouded in the fog covering the hierarchical summits of large financial institutions. But it seems very likely that, at those levels, an important factor contributing to the defeat of the obvious case for caution was the misleading feedback from operations that were temporarily very profitable, though also very risky.

These are considerations of first-order importance, as numerous examples and studies attest. They are not illuminated in simple analyses that ignore the endogeneity issue -- the problem of the evolutionary, systemic determination of the competitive standard that, in the long run, individual firms are (more or less) compelled to meet.

The endogenous environment of the bubble. While the general principle that the environment is endogenous was exemplified in the run-up to the crisis in a number of ways, the overwhelmingly important way was the housing bubble. It was manifestly a joint product of the jointly-determined behaviors of a very large number of individual actors, particularly of the ability of innumerable pairs of buyers and sellers to reach agreement on transaction prices for housing and related assets --- prices that, over time, both traced the historical course of the bubble and sustained the expectations that drove it.

Discussion of the dynamics of the bubble, and of its direct consequences, has emphasized its psychological aspects. Bubble psychology, and more generally the psychology of periods of economic expansion, is seen as a manifestation of the tendency of human societies to undergo shared surges in confidence or optimism. For example, Akerlof and Shiller say, "Confidence is not just the emotional state of an individual. It is a view of other people's confidence. It is also a view of the world -- a popular model of current events, a public

understanding of the mechanism of economic change as informed by the news media and by popular discussions..... Just as diseases spread through contagion, so does confidence, or lack of confidence. Indeed, confidence, or the lack thereof, may be as contagious as any disease.” (Akerlof and Shiller 2009): 55-56). The authors note that the role of confidence as a cause of, economic events is either denied, or at a minimum serious underemphasized, by standard economic theory.

The negative results that ultimately follow from an excess of confidence are, on this account, substantially the result of biased foresight. A too-rosy view of the future leads to actions that will, ultimately prove to be mistaken – though not for those who “take the money and run” in a timely fashion. Apart from the important qualification regarding the presence of the bias, this view of economic decisions departs little from the standard one: decisions are predominantly future-oriented and foresight-based. While feedback is accorded a significant causal role, it is primarily feedback at the macroeconomic level. For example, investments induced by optimistic expectations lead in the short term to increased incomes, which partially justify the optimism and feed it further.⁶

There is another large class of feedback mechanisms that operate at the micro level and which, in line with much evolutionary theorizing, do not presume that actors are concerned involved with imaging and re-imaging the unknown future as they take actions. Rather, they are largely acting out their established behavioral propensities. To explore how these micro mechanisms might have affected the course of the housing bubble, it is helpful to isolate them by means of a thought experiment.

Imagine an economy much like the U.S. economy during the bubble. In Scenario A, the overall course of housing prices follows the actual track of the recent bubble, up to its peak. In

⁶ See (Akerlof and Shiller 2009) on the “confidence multiplier” (pp. 14-17).

Scenario B, housing prices are relatively stable, increasing overall at historically typical rates, but with much micro-level diversity. We consider the behavioral differences between these two scenarios, and we set aside temporarily the problem of closing the causal loop back to the determination of the housing price trajectories.

In this hypothetical economy, nobody ever changes his or her mind about the future level of housing prices – or, indeed, attends to such a large-scale, imponderable, foresight-demanding question at all. They simply behave according to their usual propensities, responding in accustomed ways to prevailing prices and other indications of what is currently happening in the economy (*not* to conceptions of what *might* happen in the future). Further, this assumption rules at the organizational and institutional levels as well as at the individual level. Individuals and firms are highly heterogeneous in their propensities, and in particular in their responses to risky investment opportunities, their willingness to incur debt for various purposes, and their general levels of care in economic decisions.

Scenario A is a set-up for, in the words of our title, the survival of the reckless. Many individuals and firms are rewarded in A for risky behaviors that would be severely punished in Scenario B. Because the price track of the bubble postpones and obscures the downside risks, the risk-prone are rewarded relative to their more cautious neighbors or rivals, who may take note of the example and treat it as advice. The reckless tend to become more active than the others in the housing-related markets, both because they are empowered by greater wealth and because they are encouraged by their success. The mortgage brokers whose loan-marketing practices would backfire promptly in Scenario B, prosper in Scenario A. In Scenario A also, there are diverse pressures on firms whose accustomed ways of doing things have a more prudent tilt. Prudent asset managers may lose clients to rivals who seem to offer superior

returns. As noted previously, prudent corporations and their executives face shareholder pressures to generate returns “competitive” with those reported in firms favored by the bubble. A wide variety of unwary, sloppy and incompetent behaviors are essentially subsidized and rendered inconsequential in A, at least for actors who are positioned appropriately relative to housing assets. The effects of such lapses are, in that scenario, swamped by the apparent wealth increases signaled by the bubble. In Scenario B, such lapses would be consequential. And so on; it is not hard to elaborate the account of such mechanisms, or to model them.

Consistent with the assumptions of this thought experiment, the effects described can occur in the absence of any revision of opinions about the future prices of housing and related assets. They invoke a view of action as governed by established patterns, with attempted foresight playing a role that is peripheral to the determination of behavior, and perhaps largely rhetorical. (Mortgage brokers *say* that house prices will rise; it’s part of their routine to do so.) This view obviously derives from a “psychology” that departs from the standard, future-oriented psychology of optimization – if that can be called a psychology. But it departs from the standard view not in the direction of emphasizing the importance of emotional considerations like optimism, but in the direction of emphasizing what might be termed “habit” – but habit in a broad sense that encompasses all the various sources of behavioral propensities, including skills, routines and systems created deliberately and at substantial cost. Some of the effects also depend, as evolutionary mechanisms generally do, on the existence of variety. If there were no differences in behavioral propensities where risk is concerned, then there would be nothing for the selective aspect of evolutionary feedback to leverage into systemic effects.

It is admittedly quite doubtful that the mechanisms just described suffice to close the causal loop and generate the assumed price track of Scenario A, once something happens to get it

started. As we acknowledged at the beginning, it is hard to resist the claim that “bubble psychology” played an important role, given that there were abundant indications of that in diverse sorts of commentary that were not directly connected to action. We consider, however, that it is as important to make room in the account of economic causation for the role played by established behavioral propensities, which are directly connected to action by their very nature, as it is to make room for the emotional factors coloring images of the future. We take it for granted, also, that a good deal of room will remain for the deliberative and calculative aspects of human economic behavior, which standard theory has tried to capture in the optimization paradigm. As John Dewey emphasized long ago (Dewey 1922), all three elements – the emotional, the habitual, and the deliberative -- are involved, all the time, in all human behavior.⁷ The relative importance of the three varies according to context, and the big problems are to understand how the weights shift according to context, and how the omnipresent three elements are integrated in behavioral expression. Here, however, we have been concerned with the modest objective of illuminating the role of the habitual element – the behavioral propensities – in the causation of the crisis.

The elements of evolutionary analysis reviewed in this section are not independent, but interact very significantly in producing the dynamics of the system. Each, alone, offers a reason for doubting that the lessons of experience and evolutionary processes could combine to produce a system dominated by actors who are highly competent in general, i.e., without regard to context. The combination of them is super-additive (as detailed arguments can easily confirm), which means that they add up to one very big reason for the same sort of doubt. One implication

⁷ Dewey’s actual names for the triumvirate were habit, impulse (emotion) and intelligence (deliberation). Michael Cohen has persuasively argued that today’s social scientists would do well to balance their accounts of human conduct by keeping Dewey’s analysis in view (Cohen 2006; Cohen 2007).

is that, if the belief in general competence is grounded, a convincing grounding for it has yet to be adduced, much less supported by evidence. A second implication is that, in addition to the psychological, and forward-looking mechanisms that have properly gained considerable attention of late, some organizational propensities also need to be understood. In principle, they could by themselves be capable of driving us into (or at the very least not averting) the sort of crisis that was experienced. A third conclusion is that there is likely no prospect for really understanding the dynamics of the system – in a way that might, for example, have foreseen the risks that materialized in the financial crisis – without a substantial inquiry into the specific behaviors involved. But, before we move to juxtapose our theoretical concepts with the history of the crisis, we need one more building block, drawn from some recent research that bridges institutional and evolutionary economics. This conceptual building block emphasizes the structure of the economic environment, and suggests that the selection process itself (and, as a consequence, the selected practices), depends on the industrial structures that emerge. This is of particular note, as (a) such “architectural” elements now tend to evade both academic and policy analysis; and (b) they could possibly become tools for the guidance and exercise of policy, bringing us back to the key question of the practical implications of the evolutionary view we propose.

III. Industry Architectures, Business Models and Evolution towards Disaster

The preceding section referred to organizations, and the extent to which market processes ensures a degree of selection for actor competence. This conjures up an image of relatively similar competitors engaged in relatively similar tasks – collectively known either as “a market” or “an industry”. Yet, upon closer inspection, it turns out that the very nature of markets and industries is itself an important piece of the puzzle. The division of labor in a sector can take

alternative forms (recently dubbed “industry architectures”), our analysis need to take into account the implications of these alternatives and the associated business models. Both play an important role to play in the evolutionary dynamics of a sector.

A. From Textbook “Industries” to Industry Architectures.

For a long time, organizational, managerial and economic analysis has taken industries as given (note that we use the term “industry” here- but one might interchangeably think of “markets” and “sectors”). Whether we look at industries through the lens of industry classification schemes (SIC/NAICS, input-output table structures) or with qualitative measures, we have already implicitly accepted a definition of the general concept of an industry, which, Marshall (Marshall 1920) defined as a set of more or less homogeneous players. This might indeed have characterized some industries in some periods in the past, when stability was predominant and the way money was made did not change much (so that we could focus instead on population dynamics, size, and profit of homogeneous players). It survives as an indispensable element of theoretical and empirical accounts of situations where the homogeneity assumption is crucial to the valid application of the analytical tools employed – and such situations are sought out by scholars precisely because they do permit such application. However, this view is poorly suited to today’s sector dynamics, where the very nature and definition of a sector can change rapidly, as new ways of making money (including new business models) emerge and evolve. A recent contribution (Jacobides, Knudsen et al. 2006) suggests that an “industry architecture is an abstract description of the economic agents within an economic system (in terms of economic behavior and the capabilities that support the feasible range of behaviors) and the relationships among those agents in terms of a minimal set of rules governing their arrangement, interconnections, and interdependence (the rules governing exchange among

economic agents)... Architectures affect industry participants in ways that may be either anticipated and designed in, or unanticipated (2006: 1203).”

The term “industry architecture” (IAs) has an inherent tension. While it includes the word “industry”, the actual intent of the IA research program is to revisit the received idea of an “industry” and emphasize the potential *different ways in which one can organize labor within a sector*. The term also implies that industry architectures are *stable* at any one point in time. This is true in the sense that they define the gamut of options available to participants, but they also *evolve*, endogenously, in response to pressure from current or potential industry participants.

One way to grasp the industry architecture concept is to observe the international differences in the way labor is organized and coordinated in the same sector. Take construction in different European countries, for instance, where, as (Winch 2000) observes,

“[there exists] extensive variation in the configuration of [building sectors’ value chains]. Construction business systems have evolved over very long periods, and display well-rooted rigidities, with the balance between the actors in the system hard fought and hard won... [For example,] the French *architecte* has a much more constrained and limited role in the construction process than the British architect; the German *Architekt* has a state-derived role in obtaining building permits which the British counterpart does not, and so on. In the case of some actors such as the German *Prufstatiker*, the British quantity surveyor, and the French *bureau de contrôle*, there is simply no close comparator in other systems.”

Industry architectures, then, encapsulate the ways in which labor is organized and divided in a sector. But in practice this confronts us with some definitional challenges. The first is the definition of the sector itself. For instance, one might want to map “telecommunications” overall as a sector, looking at the different types of participants and how they interconnect; or instead focus on the subsector “mobile telecommunications”. The definition of sector boundaries here is arbitrary, and depends on the policy or strategy objectives at hand; there is no “right” way of delineating a sector *ex ante*. Likewise, the question of which peripheral actors need to be

included or left out is also a matter of judgment.⁸ So, within “mobile telecommunications”, we should arguably include those actors providing mobile connectivity for the growing Mobile Internet Devices segment. Whether we do or not really depends on the importance of that segment, and on whether we are engaging in a policy debate, conducting a strategic mapping exercise for the near term or focusing on the interdependence of different segments. So, again, the question of “just how far one goes” in including various different types of participants does not have a definitive, analytically derived answer, but depends on the specific policy or strategy objectives of our exercise. The reason is that there are architectures attached to particular segments or business ecosystems (e.g., the “mainstream voice mobile telephony” architecture, vs the “data device” architecture), all of which are partly overlapping (in terms of demand and substitution of services). Each ecosystem has an architecture, and we also observe that the ecosystems change, and usually increasingly overlap as time passes. This suggests that one sector may have more than one ecosystem, each associated with one (and sometimes more than one) architecture. This makes the empirical application of the lens of industry architectures as much an art as a science. However, since this accurately reflects the complexity of the structures we observe in the real world, we think the extra difficulty is a price worth paying.

Definitional issues aside, industry architectures provide discrete templates that are stable but evolving. As argued in (Jacobides, Knudsen et al. 2006), “Once an industry architecture emerges and stabilizes, it is difficult to stray from it, for reasons relating to inter-operability (who else is willing to participate in a new architecture, or is capable in so doing); regulation (which reinforces some ways of dividing labor and excludes others); and information (what the customers have learnt to expect). Industry architectures are [partly] shaped by legal and

⁸ One thing that is clear, however, is that the homogeneity of products or similarity of firms is not part of the guiding concept.

regulatory authority, and this explains why in different jurisdictions (different states or countries) we observe different ways of organizing labor. Also, industry participants who stand to benefit from a given architecture usually fight the introduction of new alternatives through legislative or regulatory means (e.g. (Shell 2003)).”⁹

Consider the financial services sector in general, and the mortgage banking area in particular.¹⁰ As Figure 1 illustrates (and as we will discuss in section IV), before the 1970s, the sector was organized in a rather straightforward way, with major participants being S&L associations, commercial and community banks, and some other specialized intermediaries that originated, warehoused, kept and serviced the loan till it was repaid – or until the debtor defaulted. But, over time, securitization broke the link between origination and loan ownership; new participants such as mortgage brokers, mortgage pool insurers and subservicers emerged. With them came new ways to profit from the same underlying mortgage – that is, new business models. In an apparently unchanged sector, new arrangements emerged in terms of the roles (who does what); and the rules that connected the players.¹¹ In other words, a new *architecture* emerged, which in this case complemented the extant architecture without fully supplanting it.

⁹ Interestingly, we have seen a number of countries moving to a more homogenized set of industry architectures—which is what, arguably, underpins the recent trend of “disintegration of production and integration of trade” (Feenstra, 1998)

¹⁰ As noted above, we first need to address the question of how broad the definition of a sector is. We can map a complicated picture with the entire financial intermediation sector, which will consist of a number of partly overlapping architectures; and in this case, we will probably need to make judgment calls on where to stop. (For instance, we may want to include remittances, including Western Union and the Post Office; car finance companies, subsidiaries of automobile manufacturers, which provide credit; and credit insurance firms, that provide substitutes to paying in cash in the context of trade). Next, we need to focus on specific sub-parts of the broadly defined sector, to be able to gain some concrete traction. To do so, we can focus, say, on mortgage banking alone. Within mortgage banking, there will be more than one architecture (ie means of originating and holding and servicing loans), which need to be evaluated one-by-one. Here, we will have the “traditional integrated” form that will compete with the “unbundled / securitized” modes of producing loans, which will grow over time. While it usually is the case that with time, more complexity emerges in these architectures, it also is the case that one “type” of architecture becomes more dominant than another.

¹¹ Much to our surprise, both regulators and industry participants neither had a good understanding of the overall structures in the late 1990’s, nor a sense of the sector evolution. Firms, executives and regulators might know full well their own “part of the puzzle” but did not have much of an understanding of the broader context.

Industry architectures are important because they affect the nature of individual and collective competence; they determine the activities and choices (the propensities we noted in the previous section) that pass “the market test”- because architectures *define and prescribe* what “the market” looks like and shape what “the market” values.¹²

Finally, industry architectures also evolve over time. As illustrated by (Ferraro and Gurses 2009) in the context of the film industry, firms and individual agents within a sector help change the way rules and roles are established (also see (Pisano and Teece 2007)). Changes in financial services, including those recently chronicled by journalists or practitioners (see, e.g., Tett, 2009; Sorkin, 2009; Taleb & Triana, 2009), attest to endogenous pressures for adapting rules of the game – such as the exemption of particular swathes of the sector from formal regulation (investment banks and hedge funds in particular).

B. Industry Architecture and the Evolutionary Dynamics of Firms and Sectors

How does industry architecture shape dynamics at the level of both sectors and firms? The first point is quite straightforward: the division of labor, the creation of different “parts” of a sector and the establishment of boundaries shapes the process of capability development. It drives what (Dosi 1982) calls the “technological trajectories” – which, as research in innovation economics has shown, are empirically important features. Each way of organizing and dividing labor shapes the process through which organizations, and individuals within them, develop their skills. So an industry architecture creates a set of “trajectories” that defines what we can expect from the system. To return to our example of construction, in the UK, the compartmentalization of work between architects, main contractors, quantity surveyors and subcontractors, ratified by

¹² This point has much in common with sociological analyses of markets, and of the construction of meaning in markets; see White, 1981; Abolafia, 1996; and MacKenzie and Millo, 2003, on the market for derivatives and the auto-poetic mechanisms of market creation. The latter also discuss another element which we do not explicitly consider here- the extent to which financial theories *created* a reality, allowing traders to find a legitimized way to increase their influence and profits. On this topic, also see Taleb & Triana, 2009.

regulation, reinforced by tradition and upheld as the dominant “architecture” by which labor was divided and organized, led to a number of shortcomings in the process of building (see Latham, 1979, for a report commissioned by the UK government). The problem was two-fold. First, specialization defined the *parts* of the sector that each firm was “good at”. Architects were evaluated on the basis of their designs, quantity surveyors on their ability to keep track of materials costs and engineering companies on their ability to develop a brief successfully. The division of labor created the yardsticks with which skills were evaluated, and defined what being “competent” meant – even though the individual competence level, for any part of the production process, varied widely. Second, both cognitive limitations (e.g., “what does it mean to be an architect”, or “what makes for a good engineer”), and incentives (in terms of being deemed effective in a particular part of the value-added process and capturing related returns) led the system to focus a little too much on the individual parts of the problem, leaving overall efficiency of the sector behind. Sure enough, the Latham report and other observers noted that a number of individually competent firms with individually competent participants did exist; yet, when these were put together, the entire process was problematic at best. In particular, nothing ensured any emphasis either on timeliness or on “design for buildability”. As a result, the old industry architecture was gradually abandoned, with the sharply defined division of labor and clear hand-off between the roles declining and new, broader firms encompassing more parts of the value adding process emerging. (See (Cacciatori and Jacobides 2005) for details).

To put this last set of observations in more experiential terms, anyone involved in organizations (from firms to government agencies to academic institutions) knows tales aplenty of how excessive specialization leads otherwise intelligent individuals to lose sight of the collective goal. The subordinate goal displaces the ‘big picture’ because it is cognitively simpler

to identify with a limited, partial task. Similarly, everyone has encountered parochial views and cronyism, where local incentives and feedback have led people to attend to their own, narrow part of the environment, to the blatant neglect of collective efficiencies. Incentives and cognition interact leading to substantial, often inescapable limitations. The answer to these problems is, quite often, the exception management imposed by hierarchical intervention (called “leadership” in the field of management). In the absence of effective management, the inescapable consequences of the division of labor proliferate unchecked. In industries the phenomenon is more insidious. Since divisionalization and the definition of industry architectures is quite haphazard, pushed hither and thither by lobbying or ideology (including the aversion to tackling the problem head-on), the same pathologies can end up being quite severe.¹³

The points made above extend the evolutionary analysis on competence in two important ways. First, they suggest that the division of labor shapes the nature of individual capabilities and hence collective outcomes. Second, they suggest that competence is always *local*; that is, it consists of performance as judged in the context of a set of tasks that characterizes a particular sub-sector. Third, since industry architectures tend to be fairly stable at any one point in time, enshrined as they are in regulation and tradition, there is no reason to expect that better

¹³ The problem of divisionalization and re-aggregation has been considered in a number of fields. Simon (1962), for instance, considered how different ways of structuring complex problems made them more or less effective. Firms have confronted a much more applied variety of this problem for long; they need to create separate divisions, which allow work to happen in simple ways (with clear identification of competence of the parts of the organization), yet the emergence of these simple ways of dividing and organizing labor gives rise to a number of problems of either excessive or ineffective departmentalization. Both for reasons of cognitive boundedness and for reasons of following simple incentives imposed by divisionalization, frictions emerge in any pattern of organizing labor in a large firm. And precisely for this reason, re-organizations are frequently undertaken, in order to redress the pathologies associated with the division of labor. Finally, it has recently been argued that management, and hierarchical intervention, exists precisely to redress these issues, whether through re-organizations or exception management (Jacobides, 2006; 2007). Interestingly, in organizations, senior managers have more authority to change the structure than regulators have, or feel they have, to change the structure of a sector. Yet the problem is quite symmetric: Each way of organizing and dividing labor, including incentives, leads to all the issues noted above. And regulators should, yet mostly do not, view themselves as “architects” who engineer and re-engineer the sectors they are responsible for, or those who impose “exception management” when called for.

architectures will emerge simply because the existing ones do not support the achievement of system-wide or superordinate goals.¹⁴

C. Industry Architectures, Business Models, and Feedback

Having considered the role of industry architectures, let us now focus on a particular component of them: the operation of a business model, and the nature of feedback. The term “business model” is still in flux; here, we take it to mean the means through which a firm can monetize a particular source of advantage. As the preceding discussion suggests, we believe that at any one point in time, there are a limited number of business models that may operate in a sector, all of them constrained by, and built on, the prevailing industry architecture(s). Of course, as any sector evolves, new ways of monetizing an advantage emerge. The long and often arduous process of new-market emergence (which requires very substantial investments in infrastructure, as described in (Baldwin and Clark 2003; Jacobides 2005) or North, 1990), are cases in point. New markets change the architecture, or lead to new ways of organizing altogether. And with new architectures come new criteria for engaging in each part of the production process; and, as a consequence, a new collective outcome – though not always a better one, at least in the short run.

In financial services, much of the innovation over the last two decades did revolve around new business models and changing, expanding industry architectures – so much so that the way firms monetize loans today bears little resemblance to what happened twenty years ago, *even where the underlying products (e.g. loans) might not have changed*. Securitization, for instance (for standard 30-year fixed-rate mortgages), not only led to a new set of markets, but also to different ways in which particular types of (new) industry participants monetized the benefits

¹⁴ Also, as we will stress later, even if some such super-ordinate goal might help force an architectural change, the question is whether that goal would be a *collective goal*, or a particular *private goal*, or a *set of goals*, that might suit one or a set of players.

from a loan. And, along with the particular division of labor came a new set of rules, as well as a new selection mechanism and a new definition of what it meant to be “competent”. Loan originators, in this example, established a means to benefit from selling a loan to a warehousing bank, obtaining closing fees upfront, and receiving a (somewhat risk-adjusted) payment. More important, the nature of these arrangements had been shaped in particular conditions of demand and a particular macro-economic context. And the viability of the arrangements would only be put into question if specific feedback makes a change necessary. Even if participants knew that the loans would not ultimately perform, they would only be driven by the actual feedback at hand (which tended to be revenue and profit growth through origination). So an originator would care about finding ways of producing more loans that would sell; not more loans that would ultimately perform.¹⁵

New business models, then, which have been pervasive in financial services over the last two decades, have changed both the way labor is divided and the criteria used in each part of the value chain; and these will not evolve unless and until a specific feedback mechanism arises to cause such change. However, new architectures are not necessarily better architectures – and unless we embed feedback that alerts us to faults or mistakes in the system, we run the risk of spreading inefficient practices and changing a relatively healthy architecture to a less healthy one. Moreover, without massive restructuring and re-regulation, it might be difficult or impossible to put Humpty back together again – particularly after a major crisis such as the one seen in financial services . It also appears that a vast number of new business models emerged in financial services without much supervision. The view that innovation should not be hampered

¹⁵ Recent research by Haselmann, Pistor, & Vig (2009) confirms the fact that securitized loans do tend to be of lower quality in terms of a number of attributes; and Jacobides (2005) found that mortgage broker originated loan did default substantially more than loans originated and held in portfolio. Such findings are aligned with plenty of anecdotal evidence on the different criteria instituted by the division of labor in the 2000’s.

blinded industry participants and regulators to the fact that “new” does not always mean “more effective” or “more competent”. So the financial services sector changed rapidly, as did the nature of the feedback driving firms’ performance.

Financial services presents some interesting examples of new feedback mechanisms resulting from new industry architectures. First, financial services, by its very nature, exhibits a *delayed feedback* process: for many products (loans in particular), some of the benefits might be booked in advance, whereas the “cost” in terms of risk is only revealed later. Of course, if financial institutions do not materially change their practices, or if the architecture (*or* the macro-economic environment) does not change materially, then there is no reason to expect that any major distortion will take place. But if the architecture and the business model change, and / or if the types of debtors change, then it may well be that what appears to be a profitable new practice will turn out to be unsustainable. Securitization, and other financial engineering tools that break up the loan, exacerbate the problem, especially as they bring up (they front-load) some of the benefits the loan is expected to bring. And we often won’t know whether the new structure was worth it for a while- or whether, when the risk catches up with the return, it ends up outweighing it.¹⁶

In other sectors, innovations face easier, more incontrovertible, direct and quicker tests. An unworkable concept for an aircraft simply will not fly, right there and then. But financial services has no such direct means of dismissing obviously silly ideas, either at prototype stage or during actual operation. An instrument carrying unknown risk will only reveal itself to be dangerous later; and so will a new way of originating a loan whose qualities won’t be tested

¹⁶ Of course, the danger here is that the risk never quite “catches up” with the previously overly generous return, as those who take the returns are long gone, and the risk gets re-distributed or socialized. The risk, as has partly happened now, may be visited upon us all in the form of social support for the failing of the financial services sector, well after any return has accrued to the firm that took less-than-stellar decisions, or to that firm’s executives.

before there is a cyclical hardship. More subtly, some financial services innovations may make sense in one context (e.g., the use of brokers and warehousing specialists for the production of conservative, conforming loans for Fannie Mae and Freddie Mac) but not in another (sub-prime mortgages where the same arrangements might become problematic). Another problem might be pursuing a change that seems acceptable based on local feedback (“produce more loans for given loan types”) that ultimately produces very bad apples indeed, once allowed to operate in the periphery of prime loans, or when the conditions in terms (e.g., house prices), change. So we might (as we did) see many innovations whose true costs in terms of risk were not visible for some time. And in that sense, the financial services sector can be more “dangerous” than many others, not only because of the vital role of the sector in terms of the operation of development of the economy, but also because of its particular fragility and evolutionary risks. New business models, with the attendant new ways of measuring performance, can result in some ultimately misleading indicators that may persist for a while. And they may help *shape* both the sector and the wider economy, as this new type of feedback drives decisions about which activities “make sense” and which should be undertaken. This is particularly dangerous in the context of *a new and innovative way of organizing the sector*, such as the one we observed in financial services through new innovations such as securitization. Changing industry architectures can create an entirely untested model that holds the seeds of collective failure, which is why the financial services crisis was such a perfect storm – and an edifying and sobering illustration of some key business and economic principles.

IV. How the Firms and Institutions Evolved

This paper builds on a number of research projects undertaken by both authors, independently and jointly. The data reported in this section draw on an in-depth study of the

mortgage banking sector, which was undertaken between 1996 and 2000; as well as follow-on research in the mortgage banking sector and more recent work on the evolution of the financial services as a whole undertaken from 2007 onwards.

Our account of the changes in the architecture of the sector, and the evolutionary dynamics that they triggered, starts with the transformation of mortgage banking in the US. We then briefly consider the growth of securitization, the subsequent development of CLO / CDO (Collateralized Loan / Debt Obligation) instruments, the expansion of hedge funds that traded in them and the emergence of CDS (Credit Default Swaps), and we briefly review the changing role of industry participants such as the ratings agencies and regulators. We conclude our account with the implications of this analysis.

A . *The Transformation of Mortgage Banking.*¹⁷

Mortgage banking was transformed from the 1970s onwards, and especially in the last decade, by the institution of a series of new markets along the previously integrated value chain. The first such market broke up the traditional mortgage banking value chain, separating the creation and servicing of a mortgage loan (done through a mortgage bank) from holding the capital claim to that loan – initially done by investors through the newly-created secondary market for mortgage loans, an innovation that enabled the growth of mortgage banks (i.e., non-depository financial institutions.) Although mortgage banks existed prior to securitization, their role was different, their scope limited, and their market share very small. They were used to generate and service mortgage loans that were held on the books of insurance companies or the federal government (which wanted to subsidize particular needy groups). The big change in the

¹⁷ This section draws heavily on (Jacobides 2005), which provides more detailed evidence on the mortgage banking sector, drawing on archival sources and over 180 interviews.

housing finance industry happened during the late 1960s and early 1970s. Because of shortages in housing finance, especially in the quickly developing Sunbelt (Florida, Texas, and California), the government sought additional means of providing housing finance. Limits in the federal budget made it attractive to rely on capital markets, as opposed to government funds, to support the availability of housing finance (Lederman, 1985; Tuchman, 1986).

Thus, by 1970, the United States had created three government-sponsored enterprises – Fannie Mae, Freddie Mac and Ginnie Mae – to facilitate the provision of mortgage finance. Mortgage bankers issued and serviced loans according to particular criteria on behalf of these enterprises. In 1972, Ginnie Mae created the “mortgage-backed security”: it created a pool of loans, sold shares in these pools, and insured the pool against default (so that, from an ultimate capital provider’s perspective, a default would be treated as a prepayment, given the insurance). Henceforth, as securitizers the three government agencies would specify in broad terms what loan bundles they wanted; they would then buy such loans from mortgage bankers; and then they would package the loans they had bought into new bundles and sell these to the capital markets. Thus, an alternative to the prevailing structure of integrated S&Ls and banks came about; this new mode relied on capital markets for funds, on securitizers for access to funds and relationships with lenders, and on mortgage bankers for the origination of loans, their servicing, and the payment to securitizers (Baldwin & Esty, 1993; Fabozzi & Modigliani, 1992). The government, having set up the government-sponsored enterprises, soon reduced its involvement, providing them with public company charters. (Fannie Mae was listed on the New York Stock Exchange between 1970 and 1972; Freddie Mac was listed in 1989). Although the initial loans the three agencies securitized were government-sponsored loans of the Federal Housing Administration and the Veterans Administration, from 1972 onwards the agencies focused on

loans made by mortgage banks, which were *not* part of any government initiative (Fannie Mae, 1992).

Then Fannie Mae and Freddie Mac began to sell bundled loans to investors in the capital markets, starting in the 1970s. This created tradeable securities –the first “MBS” (Mortgage-Backed Securities), sold to investors who did not have the ability to generate mortgage loans, yet had an appetite for the returns offered for carrying mortgage-related risk, usually for 15- or 30-year mortgages with standard characteristics. Once the basic securities were “understood”, a template for evaluating and assessing them emerged; and capital market participants gradually learnt how to use these securities effectively. With greater understanding came greater use; but also more competition. As such, creativity led to the creation of ever more complicated and less secure loans. This expanded the market, and provided to those who would package and sell them with a healthy source of revenues. New feedback cycles emerged for a variety of industry participants. Fannie Mae and Freddie Mac were subject to increasing competition from banks that creating similar securitized loans, and the market quickly expanded. The fact that the pool available for lending could be expanded without direct budget costs was a major appeal of this line of policy. As for capital markets, this was also warmly welcomed. Securitization was welcomed by institutions that were interested in particular asset classes to which they had no access; an investment fund, could, in a securitized world, benefit, e.g., from direct access to low-risk housing 30-year loans in Oklahoma, even with no origination and servicing skills or presence in the US. The ability to match sources of capital (on the one hand) to uses (on the other) led to a substantial new market. Regulators favoured these new structures in a number of direct and indirect ways. For instance, accordingly to Basel regulation, securitized debt was considered half as risky (in terms of the bank’s ability to offload it) as non-securitized debt.

Thus, companies that securitized and kept the very same loans on their books were able to hold lower reserves- another important reason for the market to grow.

This break-up of mortgage lending through securitization was further hastened both by yet more markets on the mortgage banks side, and on the capital market side. On the mortgage bank side, firms started relying increasingly on mortgage brokers (and specialized sub-servicers), and providing the brokers with the ability to underwrite loans on a given set of specs, thus outsourcing parts of the mortgage banking value chain. Mortgage banks therefore became increasingly reliant on outside providers for origination.¹⁸ Perhaps more important, underwriting came to be done to the specifications set by the mortgage securitizers, and from the late 1990s onwards mortgage origination would increasingly rely on automated procedures, largely controlled by the securitizers (the automatic underwriting systems of the two GSE's, in some combination of mortgage banks' own systems, being the dominant such tools in the sector). This arrangement facilitated market transactions and created a seemingly "objective" system, wherein each actor in the value chain could focus on its own area – and largely ignore the issue of loan quality. The agency problems thus created were not serious initially. Checks and balances did exist, and worked for a long time, until the competition to generate new loans (demanded by the capital market) drove down underwriting requirements.¹⁹ We will soon examine why this

¹⁸ It is interesting to note how this segment came about – and how the architecture of this sector changed, to suit the industry participants involved. As Jacobides (2005) notes, "The new mortgage broker segment (and the corresponding intermediate market) emerged partly as a result of the recession of 1979–81, when banks laid off loan origination staff but maintained some flexible arrangements with former employees. When the business cycle improved, in 1982, these arrangements became more permanent (Garrett, 1989). A newly independent segment of brokers, generally former loan officers who would identify a loan customer and prepare the "loan lead," started formally transacting with banks. As this segment grew, a number of mortgage banks began reducing their in-house origination staff and using the market to procure the loans, restricting themselves to the wholesale side of the business. By 1997, the share of this market-based arrangement between mortgage bankers and the new species of brokers had reached 63 percent of the total volume of loans produced (LaMalfa, 1990)."

¹⁹ In yet another episode of un-bundling, a market for servicing rights (MSR) was established, largely as a result of the need of the Resolution Trust Corporation's need to sell the right to service all the mortgages it had gathered on its books, as the result of the Savings & Loan crisis, which led numerous S&L's into bankruptcy. The anatomy of

happened. Figure 1 illustrates the successive changes in mortgage banking, which set in motion a broader transformation; and that transformation, for reasons that we will see shortly, became a debilitating weakness for the very dis-integrated structure that emerged.

B. From Changes in Mortgage Banking to Financial Services Transformation

On the capital market side, vertical breakup and the emergence of new markets also progressed as new markets gave rise to new types of revenues for those involved in securitizing revenues and claims on future receipts- from credit card receivables to student loans to parking ticket future collections, to pub sales in the UK, to concert revenues. The other important change was the gradual progression of such securities from the safest, guaranteed mortgage-backed securities to riskier “sub-prime” securities, issued initially by private investors and later by Fannie Mae and Freddie Mac, that were offered to compete in these new segments. As the market grew, so did the “private-label securitization” (i.e, securitization by major banks), which grew rapidly until 2007, when concerns over its quality led the market to collapse.

In addition, substantial sophistication emerged in the 1980s and especially in the 1990s onwards, with investment banks buying MBS and, on the back of these securities, creating new products. These initially consisted of “tranches”: the mortgage pool was artificially structured so that potential investors could selectively buy more or less risky tranches of the mortgage pool, exposing themselves to more or less risk depending on the riskiness of the “tranche” purchased. The investment banks that created these tranches and made such increasingly complex (“synthetic”) securities managed to make a nice margin every time they did so – ultimately because the demand for such securities was very strong.²⁰

the S&L disaster is yet another source to strengthen our argument, but the institutional detail is too complex to include in this paper.

²⁰ A number of underlying factors increased the demand for these structured products, and also for CDO's, as we will elaborate further. International differences in how banks are regulated were a contributing factor. European

These fees and the demand for more complicated structures motivated more and more creativity in allowing the previously discrete entity called a “loan” to be financed by multiple investors through a host of co-specialized entities. A financial game emerged that involved taking bundles of loan or debt obligations and re-packaging them in any imaginable way, so that the risks and returns would rest with not one institution or investor, but with a great number of them – a bewildering variety and number of them.

These developments brought to the fore the role of rating agencies, who were the *de facto* arbiters of risk (and thus, indirectly, of access to different pockets of capital). Rating agencies had grown out of (and are still to this day partly owned by) publishing houses, and were originally mere compilers of information. For a fee (initially modest), they would check available information on potential borrowers or issuers of debentures and report it back to the investing community. From these humble beginnings, the “ratings” industry was born, with three non-regulated rating firms (Moody’s, Standard & Poor’s and Fitch) assessing the potential risk of anything, from mortgage loans to sovereign risk.²¹ Until the 1970s, rating agencies were compensated by the investor, and would publish confidential, modestly priced reports for those interesting in buying particular debentures. The proliferation of photocopying machines in the 1970s, and the associated illegal copying of their reports put a substantial dent in the rating

banks had largely adopted Basel II regulations, which placed greater emphasis on the riskiness of the assets; so banks could be substantially leveraged, inasmuch as the assets were deemed to be safe. American banks, in contrast, faced an environment that discouraged leverage, yet did not discriminate as much on the quality of the underlying assets. As such, an evident gain from (regulatory) trade emerged, with the European banks wanting to over-sample on the safer parts of a synthetic product, leveraging themselves significantly; and the US banks did not mind picking up the riskier parts of the loan structures, as they were constrained in terms of how much of them they could carry, regardless of quality. This regulatory difference meant that both parties perceive themselves to be better off by creating, and trading into, such structures. In addition to this, other macro-economic and international forces were at play, such as China’s growing appetite for US-denominated debt of various forms, not only at the level of sovereign debt, but also structured products. Thus, both regulatory choices and international macro-economic conditions facilitated the growth of this market.

²¹ There were more ratings agencies overall, but these three (Fitch having acquired IBCA and Duff & Phelps last decade) have been granted quasi-regulatory status in terms of the way their ratings affect either capital requirements, or sources of invested capital by regulated entities.

agencies' revenues, so they successfully changed their pricing model (their business model, in essence), shifting from investor-paid to issuer-paid fees. With the increasing complication of the financial products that emerged, the ratings agencies' purview broadened, and became instrumental in the development of securitization. An important point is that the banking regulations (Basel I and in particular Basel II) ascribed a very important role to ratings agencies, as the ratings were used to determine how much capital a firm would need to hold; the lower the risk ascribed, the less the provisioning. More important yet, regulations and national guidelines in terms of the investment of pension fund assets made the three agencies recognized by the NRCRA the arbiters of how much money could be invested in a particular set of loans: Pension funds have generally limits on how much of the money can be invested in more or less risky ventures, and it is the three ratings' agencies assessments that thus becomes the decider of whether an investment opportunity or a loan can be used by a fund manager. As Partnoy (1999) observes, the US government gave the three CRA's "regulatory license" – and one which de facto expanded outside the US, too. So ratings agencies were vested, perhaps unwittingly, with quasi-regulatory authority, and that authority extended to the rating of securitized debentures. Thus, given permissive policies in the rating agencies, new markets could easily take off, with the indirect blessing of pensions and banks' regulators. And originators of loans had now some very good reasons to "baptize" some of their loans as AAA; as this gave them access to broader pools of capital. Hence the interest of tranching CDO's, so as to address different parts of the market, as we will see below.

As securitized structures took hold, and capital market actors learned how to describe and understand them, ever more complicated structures emerged, with new firms wanting to find ways to make money. Rating agencies would provide their assessment of how risky each new

financial instrument, and each tranche of each instrument was, and get a revenue by doing it. The only thing that possibly stood in the way of very permissive assessments by the rating agencies was a presumptive concern with “potential reputation issues” that might arise down the road. As far as the immediate feedback was concerned, it was clear that in competing for new business, agencies would be expected a degree of “constructive behavior” that would be welcomed by the firms footing the ratings bill.²²

Another key development was the creation of the CDO which established a sub-sector that would ultimately rival the “traditional” banking sector in scope. CDOs were bundles of loans, put together through a special vehicle (known as a “Special Purpose Vehicle” or “Structured Investment Vehicle” (SPV/SIV). Such SIVs were structured either by an investment bank such as Bear Sterns or Lehman Brothers, or by hedge funds.²³ The idea was similar to that of the original securitization, on which it fed. In simple (so-called cash) CDOs, the issuer collects a pool of loans or other debts (such as bonds), or securities based on debt, and re-tranches, giving pieces of the CDO at different “prices”, offering low returns for the least risky part of the loan (the “tripe-A” segment) and higher returns on the risky bits. Then money is made in two ways: first, through arrangement fees (for putting the deal together- including potential allocations of part of these CDOs), and second by trading particular parts of the debt structure that may be in demand. It could be that the riskiest (but most profitable) bit, if cleverly arranged and sold, would not require much (if any) capital to be put in by the arranger – leaving some pure upside. Of course, this provided the incentives for substantial growth in the CDO market – and for the

²² With asset prices having only risen in property over the last twenty years, rating agencies were oblivious of the risk of asset prices decrease – and when such an issue was raised, it was, as we found out from the recent congressional hearings in October 2008, swiftly dismissed. Rating agencies would thus continue certifying all such securitizable debentures (which quickly grew to include credit card receivables, student or car loans, and anything from port authority receivables to future concert revenues) since financial institutions and investors would buy them, in the expectation of returns superior to what they would get from more traditional placements.

²³

underlying loans, inasmuch as investor demand could easily be found. And the rating agencies were happy to tag along, using the scant historical information in a boom period for the underlying asset as sufficient grounds to grant “triple-A” status to what they thought were the best CDOs. Meanwhile, documentation demands became increasingly light, as the entities that created these synthetic structures lobbied and cajoled raters.

Given the ability to tranche part of the sub-prime mortgages and call them “triple-A”, and the dearth of quality “prime” mortgages, the market for securitized sub-prime mortgages quickly took off. In terms of demand, an environment of low interest rates and low yields, as well as the desire for growth from firms in various parts of the financial services sector, all meant that there was a fight to get higher rates of return. The “triple-A” CDO tranches offered substantially higher returns than “triple-A” national debt traded at. Since both were “triple-A” rated, and the former would yield more return, it was close to impossible to resist the temptation to invest in them, from a banking perspective. Given the feedback mechanisms in place, it was no surprise that the demand for them skyrocketed; and no surprise either to see creative bankers taking their PowerPoint presentations offering loans to companies that neither needed them nor could afford them, in the knowledge that the banks could structure any debt the companies could be induced to float, and thereby create (at least the illusion of) profit. Whether those structuring the debts knew the whole enterprise was unsustainable is neither clear nor something we can really establish (after all, dissonance would distort the view of the present or the past). What is crucial to note is that the positive *feedback* initially produced by this new set of business models, the visible success of the new ways of making money, was driving the behavior of the sector as a whole. (And, interestingly, much of the feedback system, and of the rules of operation of participants such as the ratings agencies, has yet to change.)

A related development was the institution of a market for Credit Default Swaps – another instrument that changed the way money was made. CDSs are, in and of themselves, a useful tool, whereby one company would assume the responsibility for the possible credit default of another company- so that if, for instance, Lehman goes bankrupt but a Lehman creditor holds a CDS on Lehman’s debt (issued, e.g. by AIG), then the issuer of the CDS (AIG) pays the full amount (the “underlying”) to the creditor when Lehman defaults. In the abstract, a CDS is a form of insurance that a company can offer, taking a payment upfront (or in scheduled intervals), keeping the risk on its books and losing the gamble if the CDS goes bad. Yet it is statutorily different from insurance in that it offers a fixed payment for a fixed event (as opposed to an “incident-based-loss”) and as such was essentially unregulated, not being deemed to be insurance proper. And, while CDSs were alleged to spread the risk more evenly, what they actually did was re-shuffle risk between different institutions in ways that were very hard to decipher, making everyone very suspicious of their counterparties and their risk exposure if things were to go bad (see Tett, 2009). But it added to the environment yet another hard-to-resist tool, profitable upfront for the sellers, yet potentially explosive when things turned sour -- another innovation whose short-term feedback was positive, and whose costs were hardly visible (and surely hardly relevant, for the executives structuring them) in the distant mist of the future.²⁴ The simple (and novel) economics of the business made CDOs an easy way to make money by structuring or

²⁴ Combining these two innovations, another new way of making money that emerged was synthetic CDOs. These combined a CDO and a CDS: the arranger got investors to fund a CDO that also had a CDS attached to it. That is, the CDO entity (technically, a “Special Purpose” or “Structured Investment Vehicle”) would gather the funds of investors and get the loans or debts or securitized debts as collateral, but would also sell a CDS, i.e. some credit default protection to another company. This way, investors would make a bigger return – though they would be potentially hit if the CDS bundled in the CDO went bad, as the CDS would have first claim on their collateral. This exposed the investors in these CDOs to very substantial risks to their principals; yet because the risks of “insurance” through CDSs were limited in the good times, investors would largely forget about them, and focus instead on higher returns, oblivious of the risks that are now on the rise. Related developments were opaque structures such as CDOs squared (CDOs with CDOs as opposed to loans on their balance sheet), or even CDOs cube (CDOs with CDOs of CDOs which did have some actual loans or debt...)

investing, without needing to put in any capital. CDOs were an unregulated and highly profitable business that exceeded \$57 trillion in total face value by 2007.

Given the huge demand from CDO issuers, the development of this set of instruments and of the hedge funds that administered them created, down the line, a lot of pressure to issue more loans that could be securitized, tranced, sliced and diced in these new, creative ways. CDO producers needed their “raw material” to structure these profitable products, and issuers were pressed to produce more – and issuers quite naturally became increasingly lax in their underwriting standards, extending credit to the good and the not- so- good, often becoming pure predators. After all, part of the new loans they marketed would be packaged into “triple-A” tranches, so they could be sold along, up the vertical system to the ultimate investors. Risk was creeping into the system – and, given the emerging set of rules and structures, this was not particularly shocking.²⁵

The final structural driver for the crisis was the changing roles, the new role of the investment banks, and the rapid growth of hedge funds. This involved yet another instance of vertical dis-integration, which led to the diffusion of new practices concerning how money could be made, and new conceptions as to who had a claim on the money. First came the increasingly lax supervision of investment banks, instrumental for the development of new products, which was sealed through the 2004 exemption of investment banks from regulatory

²⁵ The incentives of all other parties in this new, securitized architecture, were such that they would also support further growth and expansion of risky securities. Whether that was done in a cynical, cold-blooded way or whether it was driven by a genuine belief that value was added since “innovating firms” were responding to their local feedback, remains unclear and cannot be established. We believe that given the emerging architecture, it was not odd, in Kane’s words (2009: 104), that “such... financial engineers’... claimed that by their joint intervention they could accomplish the quasi-magical task of turning extremely risky mortgage loans to under-resourced households into riskless securities. These would-be financial alchemists (accountants, appraisers, investment banks, derivatives dealers, credit raters, statistical model-builders, credit insurers, and financial servicers) cooperated in overstating collateral values and understating institutional leverage and other risks.” We believe that such value-laden, perhaps cynical views do not necessarily represent the reality as seen from the “local” vantage point of, say, an ambitious appraiser that considers herself to be responding to a “market need”; but the fact remains that this ecosystem helped hasten the road towards disaster.

oversight with regards to the way they used capital. This led to further increase in leverage in investment banks, which reputedly went up to 35, and at times 45 times asset values. Like any other leveraged operation, this introduced a very substantial amount of risk- much greater than in the past. This was particularly important as investment banks' revenues shifted from advisory (in terms of debt issuance, M&A advisory, etc), to proprietary trading and investing, becoming more like hedge funds, but with the additional benefit of a very high, and cheap leverage.

More consequentially yet, the last decade saw the huge growth of hedge funds, often born in the bowels of investment banks. Senior or talented traders decided that, rather than making a living taking on risky but generally profitable positions with the investment banks' money, they could strike out on their own and appeal to experienced investors, focusing on "alternative" investment approaches (such as shorting stocks, using synthetic products, complicated algorithmic trading, strategies focusing on volatility etc). So, partly born from the vertical disintegration of investment banks, the hedge fund industry emerged, with a different, new set of rules on how to make money – often charging very high fees upfront (2%), locking funds in and also asking for substantial success fees (20%), while avoiding much in the way of and "underperformance penalty". The newly booming sector of hedge funds instituted new "rules of interaction" with clients, thus raising incentives for risky moves. Banks and technology providers were quick to support the growing segment of hedge funds with co-specialized services, changing the landscape in investment banking itself, and adding more reasons for growth of derivatives and synthetic products, in an unregulated environment.

Equally important, the establishment of the hedge fund payment schemes from investors, combined with the fact that these were small, entrepreneurial entities free to redesign how success was rewarded, changed the way successful executives would be paid. Aggressive

payment for performance (immediate performance, and not broadly risk-adjusted) gained acceptance in hedge funds. Based on what textbook economics would argue in terms of “optimal incentives”, hedge fund managers who generated immediate profits (not generally corrected for risk) would receive substantial proportions of these profits. This new set of values concerning appropriate remuneration then diffused to investment banks, which decided they had to catch up on the way they would compensate their executives, and particularly on the proportion of revenues that would be paid out as bonuses. Afraid that they would otherwise hemorrhage their best bankers (and keenest risk-takers), they underwrote, sometimes unwillingly, a culture of risk taking and very high returns.

The world had changed, and nobody was minding the risks of the system as a whole – let alone understanding or mapping the increasingly bewildering variety of risks that were leading to short-term growth, but were doomed to come to grief. In the words of the Pew Report on financial reform (Pew 2009), published in December 2009 and reflecting the views of a remarkable set of contributors,

“The crisis revealed both gaps in regulation and unanticipated interconnections among different types of financial institutions and markets. Yet no one was charged with understanding these interconnections, looking for gaps, detecting early signs of systemic threats and acting to mitigate them. During the years preceding the crisis, no regulator was tasked with monitoring and understanding the overall health of institutions and markets and the connections between them across the entire breadth of the financial system. Nor was any regulator charged with taking the lead in responding to any early signs of systemic risks. So, for example, several years ago there were widely recognized signs of unusual credit expansion and increases in leverage associated with an unprecedented rise in housing prices. These developments signaled the beginning of a bubble with the potential to destabilize the entire system. No action by any government agency was taken to address this.”

In other words, the excessive compartmentalization that led to the crisis, also plagued regulators: Architectural problems were evident not only in the private precincts of the sector,

but also in the way it was regulated. Here one needs to consider the principal pressures that regulators feel they have to respond to. As Kane (2009: 101-2) notes,

“For government supervisors, incentive conflicts trace not only to short horizons, but also to clientele influence and to pressure to support the expansion of homeownership for low-income households. As credit spreads increased in 2007 and 2008, these incentive conflicts led authorities to temporize by adopting policies that risked allowing the depth and duration of the crisis to increase”

In our terms, policy-makers, too, were responding to their own selection pressures, shaped by the regulatory architecture and their own, bounded view of reality.

Figure 2 provides a partial schematic representation of the forces that led to the crisis, summarizing parts of the discussion above.

C. Assessing the Transformation of the Financial Services Sector

The discussion above provides a concrete sense of what we mean by changing the architecture of a sector, and also of the changes in business models that operate within an architecture. Some of these architectural changes had the express blessing (or were initiated by) the US government, such as the initial push towards securitization. Others were the natural outgrowth of interaction between different firms, motivated by finding ways of profiting in the near, and sometimes the longer term. Other structures emerged without the direct support of the government (such as the markets for , CDSs, CDOs and, derivatives) but their development was affected by the lack of intervention or regulation. (Of course, as recent journalistic accounts discuss in full colour, there was an effort of regulators to step in, but senior policy makers, and very strong and effective lobbying managed to forestall or change regulation; see Tett, 2009, for details). Whatever the process, the result is clear, and demonstrates lack of effective intervention, or sometimes unwitting change. In the case of ratings agencies, the government made them

arbiters of fund allocation (for example, by enshrining their opinions into regulation of where pension funds can invest), or capital structure (through the Basel agreements, themselves a weak self-regulation supported by the US).

This brief historical analysis shows how different types of organizations promoted different business models: new ways in which profit can be booked and then disbursed. One very important change was the ability to “front-load” part of the presumed value of a loan by paying it upfront as fees – whether this relates to securitization, to structuring CDOs or to CDSs where the fee is accepted upfront and most of the value happens by repeated trades of different types of risk.

These changes in the institutional structure led, in evolutionary parlance, to a very real difference in the nature and operation of feedback mechanisms. So, whatever the foresight (whether driven by “animal spirits” or intelligent expectations), it was the feedback that changed. There is abundant reason to believe, on evidence from within the story of the crisis and far beyond it, that feedback as opposed to foresight drives behavior. As Chuck Prince, CEO of Citibank, famously said in 2007, “When the music stops, in terms of liquidity, things will be complicated. But as long as the music is playing, you got to get up and dance. We’re still dancing”. This remarkable quote encapsulates the root of the problem: *Even if foresight, of an intelligent sort or otherwise, would pointed in to one direction, it was feedback in terms of booking profits and registering growth that really drove behavior.* As a consequence, the structural changes that we described, fuelled by and then amplifying the macro-economic imbalances that we all know about, led to an increasingly dangerous spiral. As Kane (2009) remarks,

“when commissions and other fees for service are paid upfront, managers and line employees of firms that originate, securitize, rate, or insure loans fear that they are

passing up short-run income whenever they nix a questionable deal. At the same time, accountants, appraisers, and even government supervisors know that they can win business from competing enterprises in the short run by establishing a reputation for not challenging a troubled client's dodgy representations about asset values or not conscientiously assessing its efforts to transfer risks off balance sheet." (2009: 101)

What about the watchdogs? Decision-makers within the policy community could not (or, at least, did not) escape the general rule. Feedback on the monetary policy side was compared to the legitimate and long-standing objectives of economic growth and low inflation. While the risk of a major correction existed as a theoretical possibility, and concerns were raised, actions based on the primary, it was far from the case that the primary feedback called prevailing policy into question. The role of the systemic risks that emanate from the structure of the industry's architecture and the related business models was given little, if any, attention, as new theories dismissed the risks under the heading of "the great moderation". Faith in the ability of "the markets" to get it right, could not be wrong, and a steadfast refusal even to consider that an innovation in financial services might be more lethal than beneficial, contributed fundamentally to a reluctance to take precautionary actions based on foresight..

The structural transformation presented multiple examples of the hazards presented when innovation and sectoral evolution undercut the premises of established practices. Some of these hazards came to destructive fruition. Consider the routines used by mortgage banks to deal with loans generated by their brokers. Such loans always had a greater incidence of default than the ones generated by the banks' own branches; the trade-off was generally understood in the mortgage banking community. Yet as demand for securitized loans grew, these standardized practices -- , which were a little risky but still sensible for 30-year, fixed-rate mortgages, at 85% loan-to-value ratios --, were evoked in very different transactions. They were ill-suited for such eventualities as 120% home equity lines of credit with adjustable rate mortgages. However, in a

period of rising house prices, the hazards of extending an established practice into a new domain were hidden by the low overall incidence of defaults, and the growth cycle. So, whatever the foresight that some actors may have possessed, the sort of feedback that might have provoked reconsideration of the application of this policy was absent.

A parallel development happened in terms of the *changes* in the way business was made. As demand for loans grew, “good” loans became scarce. This was because more investment bankers and hedge fund managers wanted to structure CDOs, which were, in turn, demanded by long term investors enamored by the idea of AAA tranches that magically had better returns than AAA gilts.. When wholesalers saw they could no longer satisfy demand, they came up with innovative change in their practices.. Sensibly, they followed the feedback they got, and proposed loans that later became known as “documentation-light”- virtually the definition of an understatement.²⁶ The emergence of documentation-light loans to NINJA customers (No Income, No Job or Assets), irrational as it may seem today, was locally and temporarily sensible – there was demand for loans, and no clear feedback demonstrating just how risky this could be. Furthermore, ratings agencies, which were becoming increasingly profit-oriented, had no reason to turn down good business, and only used ratings developed under the rather bizarre assumption of consistent house-price rises – a key failure in the aggregate risk-assessment system. As a result, regardless of the forward-looking views of individuals, or institutions within the sector, the feedback reinforced the behavioral propensities and the behavior we now deem reckless. Quite simply, the system moved toward recklessness according to familiar evolutionary logics, , and nobody was attending to the hazards implicit in changing architecture of the sector.

²⁶ Example of 2007 Q1 vintage documentation light loans, included loans with verbal verification (!) of income; with use of Google Maps to assess property value- the height of the shadow of a building used as a measure of stories; and ZIP code as a surrogate of property evaluation. The best example the authors saw was a \$750,000 loan to an illegal immigrant from Mexico, whose strawberry picking income turned out to be \$14,000 a year. But he did look better before it all turned sour.

Observers of changes in underwriting standards, in the division of labor between firms, in the rules that connected different organizations, and in practices linking institutions and their employees, might all have been giving dire warnings. But, thanks at least in part to the under-attention of mainstream economics to such institutional and evolutionary dynamics (and even the simple incentives associated with them), few considered a reaction was needed.

In this sense, the crisis was a perfect storm not only from a macro-economic perspective, but also because of the micro-level institutional and evolutionary dynamics that helped create and amplify the deadly mix with which we are still dealing. In particular, the following points are worth stressing. First, the very nature of the sector allowed the feedback mechanism to stray, and to do so for very considerable periods of time. Also, most of the built-in “checks” on hazardous propensities were not “hard realities” (like those of a physical product with defects that could be inspected), but rather other actors’ perceptions and actions. This made the endogenous element of the selection environment and the complementary endogenous element of bubble psychology particularly hard for actors to acknowledge and resist. More important, the last two decades, and in particular the last few years, saw a remarkable, unchecked change in the industry architecture of the sector. Yet the established regulatory and institutional controls were allowed to lapse into obsolescence as policy makers, in the name of nurturing innovation, abstained from even loosely controlling the new structures, rules, links and products (such as derivatives), or providing the contours for interaction (such as the exemption of investment banks from regulation, and the full exemption of hedge funds from any substantial reporting).

Given that little advance or remedial action was taken as the system changed in dangerous directions,, it is not so surprising that it ultimately collapsed. Yet there are some grave social issues attached to the episode, quite apart from the consequences that continue in the form

of the recession. A very substantial wealth transfer happened before, and leading up to the collapse. Given that this wealth transfer rewarded innovations that proved not only worthless, but downright catastrophic, the usual defense of “the market having assigned them value” appears, as we noted in section II, to be tenuous at best. Further, the episode provides vivid illustrations of the payout-induced financial vulnerability of firms, as great institutions collapsed while their managers successfully “ran” with massive compensation and severance packages. The concern is that governments, and ultimately taxpayers have now supported the sector as a whole, preserving many of its nefarious practices, since the feedback processes of the current architecture have not substantially changed.. The injection of massive direct and indirect support, indispensable as it was in the short run, now risks becoming counter-productive as it increasingly serves as an appeal for a retake of the same market test, whose late-arriving result is now clearly seen..

Plainly put, the last decade helped build the “engine” of the financial system that would amplify macro-economic imbalances, and reinforce the animal spirits of the bubble. Although we accept that that psychological factors played a role in the crisis, the underlying architecture and the basic evolutionary forces at work were just as important. Such an analysis of this firm- and industry- level aspect of the perfect storm comes with a number of corollary implications that are worth noting. First, it suggests that, in addition to the macro-economic and psychological factors, we need to consider organizations and industries in terms of their actual institutional structure, and comprehend the forces shaping their evolution. To do so, we need to be mindful of the *very imperfect* forces shaping the evolution of economic and social entities; and of the architecture of sectors. We may also need to rebalance our research to increase our understanding of organizations and their industrial context, and consider their differences from the more or less advanced theoretical idealizations we often take for granted -- much as we have

marched under the “behavioral” banner toward integrating a better understanding of decision-making and psychology in our economic and financial analysis.

Such a “socialized” view of economics and of the evolution of firms and institutions would cast regulators too in a new role, as the (witting or unwitting) architects of an industrial system. The objective of policy would be to ensure that the prevailing architectures and rules in a sector do not lead to predictable crises; and to ensure that inefficiencies are weeded out quickly, and that the structures we have supported allow for a robust evolution to take place.²⁷

V. Looking Ahead: From Analysis to Policy- and Theory

In the wake of the crisis, vigorous discussions of its causes and implications have erupted very broadly, in the U.S. and in other countries. Under the “implications” heading, there is first of all a drive to determine the nature of constructive financial reform that would protect against a repeat of the debacle. Those discussions go forward in diverse media and with diverse sponsorship, but in particular in the U.S. Congress, where a reform bill was recently passed in

²⁷ Interestingly, many of the innovations we have found effective in the past, in periods of growth, are now providing an extremely serious, and understudied hazard. Mounting evidence from the area of turnarounds suggests that the creation of a very big number of different classes in organization’s capital structures makes any negotiations for restructuring an organizations’ debt a very complicated affair. A decade ago, there would be a handful of participants in a modest firm’s capital structure; now, it is common to have over ten different layers of debt, and each layer is held by a variety of parties (often indirectly, such as in the case of CDO’s). As a result, companies that might get in trouble in a downturn are facing very substantial problems in any effort to re-organize, and society pays a remarkable amount in fees of intermediaries involved in restructurings and turnarounds. The innovations that appeared great in periods of growth (such as the fine chopping up of the capital structure or risk) are now posing an evolutionary threat in periods of decline or down-scoping, as they do not allow for an effective evolution. Policy-makers and economists near total lack of attention on an innovations’ impact in times of decline, which are crucial from an evolutionary perspective, are now exacting a high price. This links to some very real policy problems now, ex post. Initial evidence on securitized loans, for instance, suggests that securitized loans not only are riskier, but also that the more “mechanical” structures of servicing such loans, or limits in restructuring them in cases of default, leads to further problems in downturns such as the one we are experiencing, as the leeway of the servicer is more limited, leading to a reduced recovery rate, as a consequence of a more limited ability to restructure the debt effectively, or negotiate changes to the original conditions – see Vig (2009). A related issue, which we will return to in our concluding section, is the mounting danger of building large institutions which, even if they engage in pervasive silliness, cannot die, because they are too big or too interlinked to fail. Arguably, much of the reaction of policy makers towards banks has been of that ilk, and this runs against what an evolutionary / institutional view would argue. To this topic, we will return shortly.

the House and an alternative is taking shape in the Senate. Another major discussion concerns the implications for the economics discipline, and the possibility that economic theorizing, in particular, has seriously neglected the analysis of the major dysfunctions of contemporary capitalism (Colander, Goldberg et al. 2009).

Our purpose in this final section is to engage with a very small fraction of these two branches of the implications discussion, noting the relationships to our analysis of the causes of the crisis. That analysis was itself focused, as we emphasized, on a limited range of the operative causal mechanisms, ones that seem less prominent in the discussion than they should be. We set out to provide an evolutionary and institutional account of the crisis, one that complements much of what has been considered to date. Our suggestions here obviously reflect that orientation, and we are deliberately setting aside many considerations we know to be important.

A. Implications for Policy

When economists do policy analysis, they naturally make use of the tools that their discipline has to offer. This generally means that the assumption of actor competence is taken as an unquestioned premise, inherited from the textbooks. There are significant exceptions to this, particularly in domains involving consumer choice about matters affecting health and safety. Serious engagement with the validity of the premise sometimes does occur, particularly if the conclusions of the analysis have to be effectively marketed to non-economists. But certainly the bulk of economic policy analyses, not to speak of economists' informally expressed opinions about policy, maintain the premise. It is probably more vigorously maintained with respect to firms than to consumers, and the idea of the "market test" plays a role there.

While the market test is a real influence on firm behavior, our analysis identifies

(reviews) a number of reasons why it cannot provide a general warrant for competence. In approaching policy questions, we should dispense with the view that business practices are acceptable if they have passed the market test. The “market test,” especially in sectors and activities with delayed feedback, can be highly imperfect and may turn out to be deceiving. The problem illustrated by the crisis is that we can (as we did) find this out only after massive wealth redistribution, and destruction, has taken place.

The implication is that the question “Do they really know what they are doing?” should be a perfectly legitimate question about private actors, not one for which an automatic default “yes” is supplied, whether for reasons of presumed rationality, market testing, or whatever. It is particularly legitimate in sectors where the interests of the wider society are fundamentally at stake in the behavior being questioned. It cannot be set aside on the ground that individual actors are “small,” in the system, since the coordinated errors of numerous small actors can produce results as bad, or worse, than the errors of a large actor of comparable size. Plausible mechanisms, rational and irrational, that might produce the coordination, are easy to identify and are well illustrated in the crisis. (Thus the notion that society would be much better off if the largest financial institutions were broken up into small ones is one that deserves skeptical scrutiny.)²⁸

This stance has policy implications in the financial services sector, particularly regarding the scope and nature of regulation. First, it seems clear that leaving significant actors outside of the regulatory fence involves serious hazards, and this is not exclusively a matter of whether the regulation is needed as a direct restraint on undesirable behavior. It is also a matter of economic intelligence in a broad sense, of whether the private actors are confronted with hard questions in

²⁸ Such scrutiny would, however, have to take into account the possibility that, after a crisis hits, the situation with the smaller actors would not present such a big bailout problem.

regulatory reviews, and whether adequate information about the evolution of business practices flows into the regulatory agencies. For the full benefits to be realized, the micro level information has to be gathered and processed in ways that reflect macro-prudential purpose – not just guided by the idea of verifying “safety and soundness” one institution at a time. The process has to be able to detect changing practices and business models, to aggregate across actors – and to analyze the implications.

A similar perspective applies to recent proposals to require financial institutions to prepare “living wills” and periodically get regulatory approval for these plans. The first level merit of the proposal, analogous to the direct restraints of regulation, is to reduce the moral hazard problem with respect to large institutions that get into trouble. As it is (and has been), a troubled institution faces the regulators and the government more broadly with the prospect that “chaos would ensue” if the institution were allowed to fail. The living will is essentially a contingency plan for unwinding the firm when it is failing, thus assuring that chaos does *not* ensue. And therefore, in this logic, the government is freed to abstain from intervention, the bankruptcy threat becomes more real, and large firms are thereby incited to be more risk-wary in the first place.

The second level merit is again, the informational one. Believers in pervasive competence may doubt that the institutions would learn important things about themselves by going through the living will process – but skeptics like us are quite confident that they would. And if the institutions themselves would not become better informed, the regulators certainly would be. The process could provide a very helpful vantage point from which to view the evolution of the industry architecture – especially, again, if it is appropriately aggregated and analyzed. The analysis might ultimately take place at the proposed Financial Services Oversight

Council, which is an institutional reform response directed to the massive inadequacy of regulatory understanding of the evolving system (Pew 2009).

Our analysis also points very strongly to the need to prevent the short-term feedback from running the system out of control, defying even the restraint that accurate foresight could in principle provide. It is not a problem that one can easily prescribe for, but there are policy options available. It can be approached from the angle of executive compensation, raising such questions as whether performance-based pay, in particular, should be effectively escrowed. But it also involves the pressures for short-term performance that boards, and CEO's, feel from analysts -- some of the villains of this crisis, who have bizarrely escaped criticism. Analysts and investors (and, in consequence, boards) have not paid particular attention to the dangers inherent in (apparently) profitable growth, as measured according to accounting practices that everybody understands to be deeply flawed. To the extent that corporate boards promoted, or failed to fend off, unreasonable pressures for short-term performance, there was a partial failure of corporate governance in the etiology of the crisis. We should consider how to strengthen governance and institutionalize more of a long-run orientation, at least in boards of institutions that have a systemic impact.

Another important policy issue relates to the predominance of the public corporation, which has displaced other forms of organization. Arguably, several of the problems we encountered in the crisis were exacerbated, perhaps even created, by the fact that the stock market places a premium on short-term profits and growth, and that long-term risks were discounted. In investment banking, partnerships gave way to publically listed firms; and with the changes in governance came changes in the behavior of these organizations. This change, profitable for the senior partners who monetized a collective reputation for their private benefit,

hardly helped the stability of the sector.²⁹ The problem is that a publically listed company comes with several governance issues attached, and especially a short-term bias. Whatever the implications of this change in investment banking, the case of the ratings agencies is much more troublesome. They should have incentives that reflect their fiduciary obligations toward the long-term interests of investors, yet they are increasingly run as aggressive profit-maximizing businesses.³⁰ As we re-consider how to organize the financial services sector, it might be the case that in some parts of the value-adding process there should be limits on the types of organizations that are allowed to take on certain roles – perhaps reinstating partnerships, or other non-publically listed organizations, as a consequence.

We cannot pursue in our own discussion all the inviting directions that the broader discussion suggests. In general, we consider that society should not expect a socially reasonable sector architecture to emerge naturally from sector evolution – least of all when the sector is financial services. Such an expectation involves a faith in “the market” that goes well beyond the naïve faith in the “market test.” Policy analysis for the sector should therefore include a careful look at the sustainability and the dynamic effects of a given industry architecture. Policies should aim to promote an architecture that will not lead to disaster. Of course, any such policy line risks “distorting” the operation of the “free market”, but, again, the “free market” is not geared to social welfare, given the structure of the financial services sector. While we acknowledge the risks of being too invasive, we think that the downside risk is too substantial to ignore.

²⁹ It is interesting to note that Goldman Sachs, perhaps the most effective investment bank, sailed through the crisis with limited casualties. Goldman, despite being publically traded, has retained a number of the partnership incentive structures for senior Managing Directors (called, not coincidentally, Partner Managing Directors).

³⁰ In this regard, the triumphant declaration of S&P’s CEO, in 2004, that revenues per employee rose from \$186,000 to \$768,000, was entirely understandable for a publically listed company, yet should have immediately raised an alarm for regulators. The idea that the guarantors of quality in the financial services sector would declare victory when minimizing costs (of those who evaluate and stress-test) and maximizing revenues (of their evaluatees) seems simply bizarre. This is a joint effect of a misguided regulatory architecture, and of the governance form used.

Whatever their scope, the proposed changes should happen fast- as we risk “wasting the crisis” (see Portes, 2009). Without the sense of the “burning platform” that the crisis provided, instigating such deep incisions to our financial and regulatory system appear, alarmingly, increasingly unlikely. Yet they are direly needed. Much of the support offered to the financial system was one-sided: We ensured bank survival, without thinking about a plausible mechanism to ensure that banks would change their propensities. As such, a remarkably proximate repeat of the exact same failures appears to be quite plausible from an evolutionary vantage point: Whereas foresight might have changed, feedback has not. And we will, see a tendency for similar mistakes to be repeated.³¹

Summing up, reform policies should include a focus on organizations and the way these take positions. This can complement recent thinking by economists on how new markets can help redress the issues we face (e.g.(Shiller 2005; Shiller 2008). Our analysis suggests that we should complement our usual analysis of markets with a look at what these markets do to the organizations that are involved. And we need to do so by tracking the specific institutional environment that shapes the firms’ and sectors’ evolutionary trajectory.

B. Directions for Theory

Perhaps the most important theoretical implication of our analysis is that we need to adjust the base-line model of economic action. This of course is a long-familiar refrain among those who are impressed, as we are, with the inadequacy of the currently dominant model. Our

³¹ In the credit markets, the last few weeks have sent us chilling messages that make a repeat of the same mistakes we made in 2007 likely. After a period of collective restraint and drying up of the debt market, bond issuance and even high-yield markets are on the rise. Senior bankers have noted their astonishment in the behaviour of long-term investors who are now being asked to generate returns, and to do so they all rush to the debt market. Much like in 2007, there is a sense that screening the specific risks of particular debentures is not important; being able to place the funds is, and the ratings agencies rating is used as a surrogate of risk. As we did not change the underlying structure of the system, it is not so strange that similar reactions would appear. And, again, although bankers and investors might *privately know* that this is a bit of a folly, their organizational incentives and the feedback mechanisms are such that they will not change their behaviour accordingly, with few exceptions.

more specific suggestion is that economic analysis should expand beyond its traditional emphasis on foresight, and the interesting discussion of what drives it (rational expectations, animal spirits, or a combination thereof?), to give more attention to feedback and the way it shapes individual and collective actions. This is particularly relevant to the behavior of large organizations, which tend to have overwhelming “Cassandra problems”: The ability of an isolated individual or team to foresee the future accurately is quite irrelevant when the short-term feedback is dominating the attention of everybody else in the organization. To better understand how feedback works in a sector, we would need to inquire into its architecture and prevalent business models. From the viewpoint of evolutionary analysis, the interest in these derives largely from the fact is that they shape the nature and operation of feedback mechanisms.

Our analysis also underscores the fact that a stronger grounding of economics in realistic views of human psychology and decision making is something we badly need, and have long needed – but it is not the only thing we need. We acknowledge the substantial progress made in these domains ((Akerlof and Shiller 2009); Thaler & Sunstein, 2009), and in the related areas of behavioral finance. Yet at the same time we want to stress, that these new lines of inquiry do *not* address the regularities that we observe at the level of organizations; and in the ways organizations are embedded in sectors and their architectures. Neither do they address, at the individual level, the distinctive psychological mechanisms that underlie skill – the real manifestation of the form of high competence of which humans are actually capable (Cohen and Bacdayan 1994). While the discipline is gradually moving in a more behaviorally friendly direction, it does so in a limited way, focusing primarily at the individual level (or even the neural level), and ignoring the very substantial issues that emerge when we look at social


aggregates such as organizations. We argue that much more attention should be paid to achieving a realistic view of the dynamics operating above the individual level. .

Provocatively put, before seeking a deeper understanding of the endocrinology of risk, we might want to know a thing or two about how organizations and institutions *actually do evolve*. We might want to see the economics discipline regain some of this ground from journalists, and occasionally other business school academics who are now covering it- yet covering it thinly indeed. It is disconcerting to see that as a response to the inability of mainstream economic theories to accurately depict important social phenomena, there is a rise in enthusiasm for inquiry into the neuro-chemical basis of individual behavior. While this is undoubtedly interesting, it will not fill the gaps we have in understanding organizations and their interaction. More realism is badly needed in diverse areas, but primarily on the important social phenomena we miss as a profession. And while this calls for “messier” research, without the appearance of “hard” scientific method that a set of scan images or mouth swabs may provide, confronting such messiness is the price of entry to serious engagement with the social problems. As Paul Krugman noted in his recent *New York Times* op-ed piece on the state of economics (Krugman 2009), “the economics profession went astray because economists, as a group, mistook beauty, clad in impressive-looking mathematics, for truth...It’s much harder to say where the economics profession goes from here. But what’s almost certain is that economists will have to learn to live with messiness”. In this regard, we hope that this paper will provide a first step, and a stimulus for a more thorough discussion

FIGURE 2


An Illustration: The Dis-integrating Mortgage Banking Sector (from Jacobides, 2005)

Original Structure: Integrated Housing Finance Provision




Integrated Banks and Savings & Loans

First Value Chain Change - Securitization and Secondary Market for Loans (1978=> 1988)



Second Value Chain Change: Creation of Broker Segment (1983=> 1987)



Third Value Chain Change – Creation of Market for Mortgage Servicing Rights (1989=>1993)


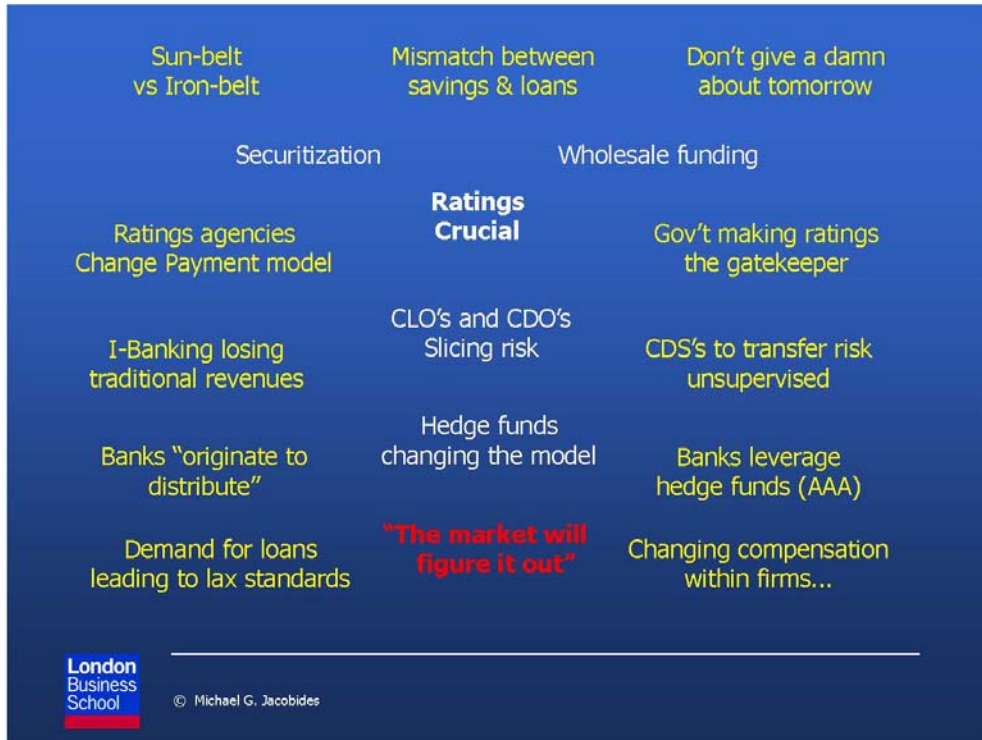


FIGURE 2: A Simple Picture of the Changes in the Sector



References

- Akerlof, G. A. and R. J. Shiller (2009). Animal Spirits: How Human Psychology Drives the Economy and Why It Matters for Global Capitalism. Princeton, Princeton University Press.
- Baldwin, C. and K. Clark (2003). Where do transactions come from? A perspective from engineering design. Boston, Harvard Business School.
- Becker, G. S. (1962). "Irrational behavior and economic theory." Journal of Political Economy **70**: 1-13.
- Cacciatori, E. and M. Jacobides (2005). "The dynamic limits of specialization: Vertical integration reconsidered." Organization Studies **26**: 1851-1883.
- Cohan, W. D. (2009). House of Cards: A Tale of Hubris and Wretched Excess on Wall Street. New York, Doubleday.
- Cohen, M. and P. Bacdayan (1994). "Organizational routines are stored as procedural memory." Organization Science **5**: 554-568.
- Cohen, M. D. (2006). "What's different is routine." Industrial and Corporate Change **15**: 387-390.
- Cohen, M. D. (2007). "Reading Dewey: Reflections on the nature of routine." Organization Studies **28**.
- Colander, D., M. Goldberg, et al. (2009). "The Financial Crisis and the Systemic Failure of the Economics Profession." Critical Review: A Journal of Politics and Society **21**(2): 249 - 267.
- Dewey, J. (1922). Human Nature and Conduct: An Introduction to Social Psychology. New York, Random House.
- Dosi, G. (1982). "Technological paradigms and technological trajectories." Research Policy **11**: 147-162.
- Einhorn, H. and R. Hogarth (1985). "Ambiguity and uncertainty in probabilistic inference." Psychological Review **92**: 433-461.
- Ferraro, F. and K. Gurses (2009). "Building architectural advantage in the U.S. motion picture industry: Lew Wasserman and the Music Corporation of America." European Management Review **6**: 233-249.
- Friedman, M. (1953). The methodology of positive economics. The Methodology of Positive Economics. Chicago, University of Chicago Press.
- Jacobides, M. (2005). "Industry change through vertical disintegration: How and why markets emerged in mortgage banking." Academy of Management Journal **48**: 465-498.
- Jacobides, M. G., T. Knudsen, et al. (2006). "Benefitting from innovation: Value creation, value appropriation and the role of industry architectures." Research Policy **35**: 1200-1221.
- Kahneman, D., P. Slovic, et al., Eds. (1982). Judgment Under Uncertainty: Heuristics and Biases. Cambridge, Cambridge University Press.
- Krugman, P. (2009). How did economists get it so wrong? New York Times. New York.
- Marshall, A. (1920). Principles of Economics, 8th ed. New York, Macmillan.

- Nelson, R. R. and S. G. Winter (2002). "Evolutionary theorizing in economics." Journal of Economic Perspectives **16**: 23-46.
- Pew (2009). Principles on Financial Reform: A Bipartisan Policy Statement. T. F. o. F. Reform, Pew Charitable Trusts.
- Pisano, G. and D. Teece (2007). "How to capture value from innovation: Shaping intellectual property and industry architecture." California Management Review **50**: 278-296.
- Shell, G. R. (2003). Make the Rules or Your Rivals Will. New York, Crown Business.
- Shiller, R. J. (2005). Irrational Exuberance. Princeton, Princeton University Press.
- Shiller, R. J. (2008). The Subprime Solution. Princeton, Princeton University Press.
- Stewart, J. B. (2009). Eight days: The battle to save the American financial system. The New Yorker: 58-81.
- Winch, G. M. (2000). "Institutional reform in British construction: Partnering and private finance." Building Research & Information **28**: 141-155.
- Winter, S. G. (1964). "Economic 'natural selection' and the theory of the firm." Yale Economic Essays **4**(Spring 1964): 225-272.
- Winter, S. G. (1971). "Satisficing, selection and the innovating remnant." Quarterly Journal of Economics **85**(May 1971): 237-261.
- Winter, S. G. (1987). Competition and selection. The New Palgrave: A Dictionary of Economics, v. 1. J. Eatwell, M. Milgate and P. Newman. New York, Stockton Press. **I**: 545-548.