Journal of Institutional Economics (2010), 6: 4, 415–443
© *The JOIE Foundation 2010* doi:10.1017/S1744137409990397
First published online 16 July 2010

Inherent Vice: Minsky, Markomata, and the tendency of markets to undermine themselves

PHILIP MIROWSKI*
University of Notre Dame, Indiana, USA

Abstract: Most current explanations of the crisis which began in 2007/8 tend to search for scapegoats, in the format of behavioral flaws. Their treatment of 'risk' is an important signpost to where such theories go awry. This paper suggests a structural theory of the crisis, informed by Institutionalist themes. We insist there is an alternative to a neoclassical macroeconomics, in the guise of possible alternative heterodox microfoundations for Minsky's account of economic crises, beyond the Kaleckian markup model. The sketch is based upon elevation of some formal notions of computational complexity to pride of place, and characterization of crises as a collapse of complexity. It is an attempt to portray a market system evolving to a point of 'inherent vice': an endogenous development which by its very nature, cannot be tamed through conventional insurance or risk models.

> ***Inherent vice*** (n., insurance terminology): An exclusion found in most property insurance policies eliminating coverage for loss caused by a quality in property that causes it to damage or destroy itself.

> . . . the economy is a complex system of interacting individuals – and these individuals themselves are complex systems. Neoclassical economics radically oversimplifies both the individuals and the system – and gets a lot of mileage by doing that; I, for one, am not going to banish maximization-and-equilibrium from my toolbox. But the temptation is always to keep on applying these extreme simplifications, even where the evidence clearly shows that they're wrong. What economists have to do is learn to resist that temptation. (Paul Krugman blog, 5 September 2009)

> The dynamics of a capitalist economy which has complex, sophisticated and evolving financial structures leads to development of conditions conducive to incoherence. (Hyman Minsky, 1986: 10)

Nick Paumgarten (2009), in a *New Yorker* version of irony, has compiled a laundry list of more than a hundred proposed causes of the world economic

*Email: pmirowsk@nd.edu

Apologies to Thomas Pynchon. Thanks to Don MacKenzie, and three anonymous referees for comments.

crisis beginning in 2007 that he had run across, starting with rather generic trends 'technological, cultural, demographic, behavioral and legal' to the more palpable 'ratings agencies, mortgage-backed securities, credit-default swaps, Bretton Woods, Basel II, disintermediation, the SEC' to 'mark-to-market accounting, hedge funds, private equity firms' and finally 'an indulgent, undereducated populace ... debt, greed, hubris'. What is striking is that this laundry list fairly closely tracks the range of explanations offered by economists over the course of 2008/9: while there was little consensus, the majority of accounts tended to focus on culprits rather than systemic causes. In this sense, they have been antithetical in orientation to most Institutionalist approaches. The dominant trope by far was to seek to blame the government one way or the other, either for not regulating selected sectors 'enough', or else for regulating in such a way as to bring about the crisis. Strangely, the vast majority of commentators became fixated upon the crisis as some sort of plebiscite about the nature of government. Examples of the former 'insufficient regulation' school would encompass the repeal of Glass Steagall in the US in 1999, the refusal to regulate credit swaps, the SEC revision of the net capital rule, and the enshrinement of 'market discipline' for international banks through the Basel II accords (Bieling and Jager, 2009); examples of the latter 'misconceived regulation' school were accusations that the government had brought about degradation of mortgage quality through the Community Reinvestment Act and political pressure on Fannie Mae (Wallison, 2009), and accusations that the Greenspan Fed kept interest rates artificially low in the new millennium (abetted by Chinese capital inflows and cut-price consumer goods). The second most popular category was the rush to indict various individuals for their moral failings. First out of the gate were the easy denunciations of the rapacious larceny of the Bernie Madoffs and Joe Cassanos of the world; Simon Johnson (2009) became briefly notorious for denouncing Wall Street as a den of 'crony capitalists' and comparing them to Russian oligarchs. Searches for scapegoats intensified, and there followed the more general indictments of broad swathes of the population. News outlets jumped on this bandwagon before it was out of the barn: for instance, ABC News 20/20 on 10 October 2008 explicitly counseled its viewers not to blame the CEOs of the firms collapsing all about them, and instead acknowledge that the average debtor must learn to face the music and accept the consequences of their improvident gambling. Some economists were a tad more subtle, but essentially converged on blaming the victims. Tyler Cowen (2008) for instance, simply asserted that Americans in general stubbornly refused to save for the future, had taken on too much risk due to 'investor hubris and collective delusion', and then run into some bad luck.[1] In a nominally left-leaning version of essentially the same argument, George

1 Just to show that the freedom inherent in attributions of 'risk aversion' can explain anything whatsoever in orthodox economics, what appears in retrospect a rather untimely article in the New Classical vein argued that consumers and investors in the last decade were *too* risk averse because of their

Akerlof and Robert Shiller (2009) took the name of Keynes in vain in order to argue that new-fangled behavioral economics suggested too many people had succumbed to a bout of 'animal spirits' and had transgressed the orthodox canons of rational behavior.[2] And then in the midst of an otherwise perceptive history of the rise of credit default swaps, collateralized debt obligations and sundry related derivatives, Gillian Tett (2009) ended up suggesting that 'JP Morgan's own history showed innovation need not lead to crises' (p. 246), and therefore the rot that had apparently pervaded the financial sector was merely the consequence of some greedy, improvident or ignorant players. Economist Robert Frank (2009) boils the indictment down to a concise thousand words: the flaw in free markets is the 'human' agent – not the theories of the economists, and certainly not something inherent in the dynamics of capitalism.

The stereotypic narrative line of 'When Good People Do Bad Things' may never run the risk of alienating the audience, and therefore proliferate endless anodyne 'critiques' of 'making bad choices'; yet it does serve the purpose of diverting attention from serious analysis. As the pink-slipped analyst of mortgage-backed securities complained to the CBS Evening News (23 January 2009), 'I don't think I brought the economy to its knees by myself.'[3] Taken one by one, no one, not the analysts, the rating agencies, the quants, the CFOs, the traders, or the watchdogs will shoulder personal responsibility. Conversely, when faceless archetypical generic individuals are said to shoulder the blame, then it is an effortless glide to suggesting no one really bears the onus. (By late 2009, this does seem to be the general direction that conventional wisdom is trending at the very pinnacles of economic policy.) Far from being confined to cable news and talk radio, this abdication of serious responsibility can be found at the very heart of modern macroeconomics, from the 'representative agent' of the dynamic stochastic general equilibrium models to the so-called New Keynesians to the supposed New Neoclassical Synthesis (Leijonhufvud, 2009; Colander *et al.*, 2009; Posner, 2009).

The widespread conviction that the crisis was due either to personal failures or regulatory 'snafus' has led to creeping paralysis of understanding, and bodes ill for the growing conviction that the danger has passed. As Barry Eichengreen (2008) wrote in the heat of the contraction, 'the deeper question is how this belief

unfortunate continuing fixation on the Great Depression of the 1930s (Cogley and Sargent, 2008). This, of course, is simply the reverse of the Minskyan financial fragility hypothesis.

2 Implicitly, this paper also argues that Akerlof and Shiller also took the name of Hyman Minsky in vain when they assert, 'Our line of thinking in this book parallels that of Minsky' (2009: 177, fn 2). Minsky's insights are fundamentally opposed to their neoclassical approach, we shall insist.

3 This sort of defense on the part of the legion of erstwhile traders at Lehman and elsewhere threatens to become another minor genre of exculpation: see, for instance, the quotes in Story and Thomas (2009). 'How do you blame us? A lot of what we did from an origination standpoint was based on investors appetite!' If the purpose of the market is merely to give voice to human desires, how can it possibly go awry?

in the efficacy of self-regulation was allowed to develop. Failing to inquire into deeper forces may lead to regulatory reforms that address symptoms rather than fundamental causes.' The perennial scapegoats Debt and Greed do not explain anything by themselves, because they have always been with us. The quest to uncover deeper forces will have to begin with renunciation of the obsession with culprits and turn more concertedly towards longer-term historical structures, and perhaps even more importantly, the role of economic theory in the framing of those structures.

This paper explores some preliminary conceptual premises belonging to alternative economic traditions[4] from which it would be necessary to depart in order to explore a rather less popular or prevalent explanation of the crisis which began in 2007, the one which Paumgarten (2009) cites one of his informants as calling 'a rise in complexity' of the entire market system.[5] This explanation would appear both be more abstract than the more conventional searches for scapegoats or misguided regulators, and more general than a narrative confined solely to the financial sector of the economy. It would acknowledge that, 'Mathematicians and physicists ... gravitated to Wall Street and began devising ways to measure, price, and package risk. It was a kind of decentralized Manhattan Project.' The resulting 'gadget' (as Los Alamos denizens called the Bomb) did explode, but not in a controlled fashion out in the New Mexico desert. We need to learn to regard this program to ratchet up the level of complexity in certain markets as itself simply one manifestation of a much larger historical trend. It was a case of what is called in the insurance industry 'inherent vice'.

1. Two alternative microfoundations for Minsky's crisis theory

Without gainsaying any of the various 'explanations' of the crisis which identify imbalances or other mismatching of financial and 'real' variables, I shall suggest that attempting to force each of these imbalances upon the Procrustean Bed of 'disequilibrium phenomena' has not served economists well in coming to understand the systemic nature of our recent experience. At best, it has suggested an account whereby various unexplained 'shocks', ranging anywhere from something as small as a mortgage default to something as large as the rapid industrialization of China, somehow temporarily deranged the ability of

4 Although this is not the place to do the history, I view the inspirations of this alternative approach as dating back to early Institutionalist thinkers, ranging from William Thomas Thornton to John R. Commons to Wesley Clair Mitchell.

5 Other glimpses of the complexity approach can be gleaned from Buchanan (2008), Marsili (2008), Rosser (2008), Bouchard (2009), Arora *et al.* (2009). It can be traced back to an insight expressed at the Santa Fe Institute back in the early 1990s: 'We are myopic ... not because of the costs of computation, but because we are individually too clever; we tend to transform the world in which we are adapting into a yet more chaotic world where we fare less well' (Cowan *et al.*, 1994: 84). In my other work, I contrast this conception of 'complexity' with that found in Hayek, where the term is used as an excuse for why economists will never be able to fully comprehend and explain economic success or failure.

world financial markets to achieve equilibrium. Convenient stories of individual scapegoats and maladapted regulatory schemes are then inserted to motivate the otherwise unexplained sand in the works of the otherwise efficient markets. The net consequence is that almost nothing about the financial crash and contraction is deemed to have been the natural outcome of the operation of those very same markets over the last three or more decades.

The complexity approach would instead explore the alternative hypothesis that as the internal dynamic of market innovation became more complex over the last few decades, the system as a whole evolved to an ever-more fragile structure, until it reached the point that it could be globally vulnerable to the breakdown of some particular market in a particular geographic location. Hence, while it is indisputable that the utter collapse of the set of financial instruments which would supposedly shield investors from the default of the underlying mortgages in their portfolio of derivatives was the trigger of the current crisis (Tett, 2009), it was by no means the necessary and sufficient cause. That cause, instead, would be traced to various indices of the attainment of an untenable level of complexity in the operation of key markets.

The godfather of such an approach, as has often been acknowledged, is Hyman Minsky (1977, 1986). First Nouriel Roubini (2007) and then John Cassidy (2008) have called the experience of living through the crisis a 'Minsky Moment'; whether the economics profession has experienced that Moment yet is disputable (Erturk and Ozgur, 2009). Minsky had repeatedly argued that the enjoyment of long periods of relatively smooth macroeconomic expansion would inevitably lead to endogenous changes in the economy which would push it to the brink of collapse: in blunt terms, the capitalist system was inherently unstable. Further, as evidenced by the quote heading this paper, he pointed to the evolution of increasing complexity in markets as the root cause of the instability. Nonetheless, the problem we shall identify with Minsky's approach is that it worked primarily in terms of a single index of complexity, namely the extent of leverage prevalent in the entire macroeconomy,[6] due to what we shall argue were inadequate microfoundations in his own theory of markets. Given that he was almost exclusively concerned with financial crises, this narrowing of analytic focus was perhaps understandable. Minsky trained his incisive skepticism upon fundamental presumptions that markets can properly assess risk based on complex, backward-looking models; that markets can adequately hedge and shift risk to those best able to bear it; and that market forces will discipline wishful thinking and willful delusion in human decision making. In a world of evolving financial innovation, each of those presumptions has proved to be woefully incorrect. The models used by the actors (and the

6 More precisely, in chapter 9 of Minsky (1986) his index was defined as the relative proportions of business current obligations paid out of current income, balance sheet manipulations and portfolio alterations in a given time period. He never managed to translate this into actual measurements, however.

economists) were almost inevitably based on data generated during an unusually stable era in which losses were small, and required that the structure of the financial system would remain constant. However, as Minsky (1986) argued, an era of relative stability will inevitably encourage behavior that reconfigures the financial structure (he proposed the terms *hedge*, *speculative*, and *Ponzi* finance to describe the proposed stages of transition from sustainable to unsustainable leverage). This evolution, in turn, rendered the risk hedging models increasingly useless, even as they were employed on an ever-expanding scale to justify interest rate spreads that implied virtually no defaults would ever occur. Furthermore, the risk models used by investors could not account for growing linkage among debtors, not to mention gauge the escalating vulnerability to bankruptcy at the firm level, escalating the systemic risk that insolvency in some locations would generate an avalanche of defaults.

Clearly, a Minsky-inspired story of debt finance run wild resonates quite well with many alternative versions of current accounts of our economic predicament (Tett, 2009; Leijonhufvud, 2009).[7] The tendency of both intellectuals and the public to grow complacent and forget the lessons learned in the last Great Depression about unsustainable levels of leverage plainly played a major role in fostering the crisis. But complacency and a truncated time horizon cannot constitute the entire story of systemic fragility. The complexity approach should not be reduced to the simplistic maxim 'Neither a borrower nor a lender be'; for that would clearly be a grievous mistake in comprehension of the economic history of capitalism.[8] The incessant striving to get around budget constraints and finesse borrowing restrictions has been one of the major sources of diversity and innovation in the annals of market economies; indeed, the system of markets has grown up largely through channeling those very predispositions into more

7 But that does not prevent the orthodox profession from continuing to disparage him. See, for instance, Paul Krugman's blog post: 'So I'm actually reading Hyman Minsky's magnum opus, here in Seoul . . . And I have to say that the Platonic ideal of Minsky is a lot better than the reality. There's a deep insight in there; both the concept of financial fragility and his insight, way ahead of anyone else, that as the memory of the Depression faded the system was in fact becoming more fragile. But that insight takes up part of Chapter 9. The rest is a long slog through turgid writing, Kaleckian income distribution theory (which I don't think has anything to do with the fundamental point), and more.' http://krugman.blogs.nytimes.com/2009/05/19/actually-existing-minsky/ (last visited 22 September 2009). Then there is the much more conspiratorial thesis that the people at the pinnacles of finance in fact were fully aware of Minsky: 'the Greenspan–Rubin–Paulson milieu seems to have been post-Minskian. They understood Minsky's theory of bubbles and blow-outs, but believed they could use it strategically for blowing bubbles [and] bursting them' (Gowan, 2009).

8 Recent claims that religious prohibitions upon usury, such as those found in Islamic societies, have insulated those societies against the financial crisis would be the final terminus of such a belief. The problems with such claims are (a) that they tend to be empirically false, in the sense that financial institutions located in those societies have also found ways to engage in modern levered finance without nominally violating religious prohibitions; and (b) to the extent they are participants in the global economy, they have not escaped the fallout from the crisis. Both points are illustrated by recent problems with Islamic bonds in Dubai.

productive activities and pursuits. It would thus seem plausible that the system of markets can sustain the unremitting pressure to manufacture ever-greater degrees of leverage – at least under most circumstances – until something else happens to render the system vulnerable to a liquidity crisis. From this perspective, what the Minskian approach lacks is a deeper grounding in the theory of the evolution of markets, which can help us understand how and why a particular threshold of leverage tumbles over into Ponzi finance at a particular nexus, with consequences that concatenate throughout the system, such that the damage can no longer be isolated or confined to a particular sector. Put in a more suggestive (at least to physicists) terminology, we would like to understand what renders certain aspects of the economy 'scale-free', while others remain 'scale-dependent' (Barabasi, 2009). To approximate that understanding, the Minskian categories would need to be supplemented with a microeconomic theory of market evolution towards greater generic levels of complexity.

The first steps towards a non-orthodox theory of the crisis would compare and contrast two deep fundamental presumptions concerning microeconomics: the *first*, the very heart of orthodox economic theory, is the insistence that all markets are fundamentally alike, in that they each individually grind out the *same* equilibrium prices and quantities for the 'same' commodities.[9] The theory of Walrasian general equilibrium is the clearest expression of this conviction, because it situates all markets on the same conceptual and functional footing. But many econophysics models also share this presumption (Rickles, 2009; Rosser, 2008). If all markets operate in essentially the same manner, then if it transpires that the same commodity is being sold for two or more different prices, then this is an immediate indicator that there is something wrong or 'imperfect' about one of the component parts of the market. Nevertheless, the system is presumed to be extremely fault-tolerant, since once any price discrepancy is detected, dedicated arbitrageurs will swoop down and make a profit, while trading away the discrepancy. Flaws in the system are thus said to be contained and localized.

This brings us to the *second* fundamental presumption of orthodox microeconomics, which is a corollary of the first. Arbitrage pricing theory is one of the ten or so commandments of modern finance theory, and also is enshrined at the heart of neoclassical economic theory. That theory is constructed upon the bedrock assumption that arbitrage opportunities do not exist in 'equilibrium', since they cannot persist.[10] Other components of the theory include the postulation of an equivalent martingale measure for every portfolio, and

9 The stability of commodity identities is thus central to this premise, and, as such, is a deeper ontological presumption, but something we must pass by in this venue. See, however, Mirowski (1986).

10 'In my view, to specify an environment in which arbitrage opportunities go unexploited . . . is simply bad microeconomics' (Durlauf, 2005: F237). This was one of the rallying cries of rational expectation macroeconomics: no $50 dollar bills left on the sidewalk. On its role in modern microeconomics, see Mirowski (2006). We must here bypass the weaknesses of the 'behavioral' literature, such as Shleifer (2000) and DeLong *et al.* (1990), that introduces arbitrary frictions into the arbitrage process.

hence a derived valuation attached to any package of contingent claims. In option pricing theory, for instance, the price of an option is rendered determinate by the presumption that it can be replicated precisely as a continuously adjusted portfolio that will have the same payoff as the option in all 'states of the world'. It is therefore the basis for the doctrine that markets are the means whereby traders bundle and repackage 'risk' so that it can be conveyed to those most capable of bearing it, thereby diminishing global levels of risk through diversification. The dual theses of the fundamental uniformity of markets and the validity of arbitrage pricing theory are indispensable building blocks of modern orthodox economic theory. *Without arbitrage pricing theory, there would be no neoclassical theory of financial 'risk'*. Curiously, arbitrage pricing also tends to be yoked to a chiliastic vision of the end of history, with evolution negated through the arrival at the summit of perfection of financial organization.[11] And yet, the conviction that risk can always be effectively tamed and priced through rational trading strategies was one of the misbegotten tenets that Minsky identified as fostering financial instability.

Now consider the alternative heterodox proposition, that various markets are generally *not* all alike, but operate according to radically different principles which can be specified in purely formal terms (Mirowski, 2002, 2007). This is a proposition which has been entertained in certain precincts of the contemporary economics profession, such as experimental economics, or the so-called 'market microstructure' theory in finance; but the profound implications are rarely pursued. In these various corners of economics, the structural differences between posted price markets, dealer mediated markets, sealed bid auctions, and continuous double actions are intensively studied and acknowledged. Yet, under those conditions, it is known that if the 'same' commodity were to be offered for sale under the auspices of two different market formats, then it is highly unlikely that the price and quantity consequences will always and everywhere be identical between them. Imbalances of supply and demand therefore need not lead in general to identical prices and quantities sold, in this view.

Negation of the *first* fundamental proposition of orthodoxy leads inexorably to revision of the *second* fundamental principle as well.[12] Patently, some traders will approach the diversity of market prices as an opportunity for further

11 Just so the reader does not think I am engaging in bootless caricature of orthodox economics, here is one example, *not from the often scorned Chicago School*: 'market institutions appear to have converged to relatively stable forms, surviving relatively unchanged through thousands of years of economic evolution . . . I argue that evolutionary processes lead to economic institutions that tend to simplify the computational problems facing economic agents' (Rust, 1997: 27). It is enough to make a historian throw up his hands in despair.

12 It is at this point precisely that we diverge from the modern literature of the 'social studies of finance'. As one representative writes, 'the social studies of finance does not set out to debunk conventional financial economics. Orthodox finance theory, for example, gives an account of markets that is in many respects perfectly successful' (Buenza *et al*., 2006: 740). One wonders if the authors might have written that so confidently four years later.

interventions. These traders may engage in what they consider to be arbitrage operations because they see an opportunity to link the disparate markets, buying the same or similar commodities low and selling high, but this is not the rectification of an 'imperfection' as much as it is the creation of a third, more 'complex' market out of the previously disjoint components. Indeed, the arbitrageur may trade making use of a different market format than either of the previous two: this happens when a 'futures' market based upon the mechanisms of a version of the double-sided auction is used to link various spot dealer-mediated markets in a designated commodity, or when an over-the-counter 'derivatives' market is superimposed upon an existing auction market.[13] In any event, although arbitrageurs may exist and operate, their activities do not render the underlying component market forms more 'alike', nor do they therefore make the opportunities for arbitrage permanently 'disappear'; this is one point where the orthodoxy goes off the rails. Heterodox theory entertains the possibility that this sort of bridging activity may not necessarily render the entire system more stable, for evolutionary reasons explored below. In this framework, attempts at arbitrage are symptomatic of one large class of market innovation, and not the rectification of some temporary flaw in an otherwise efficient economic system.

These revisions bring us to the heart of the matter, the treatment of risk. The irreducible diversity of market formats and the reframing of the meaning of arbitrage have a corrosive effect on arbitrage pricing theory, that is the orthodox presumption of an equivalent martingale measure for every portfolio, and a derived valuation attached to any conceivable package of contingent claims. Most assertions of the possibilities of 'hedging' are thus built upon shaky foundations. Whenever some entrepreneurial firm packages together a set of contingencies and asserts they are 'the same' as some other existing commodity, *they are in fact creating something new*, which cannot behave in ways identical to the prior situation in every respect. This cannot be properly understood as merely 'cognitive uncertainty about risk'. Patently, when you couple a posted price mechanism to a double auction, the hybrid result cannot appeal solely to *either* component market as dictating the unique 'correct' price for the derivative commodity. According to our first precept, there is by definition no unity of correct price that can be restored through arbitrage between two suitably distinct markets. But, further, there is no reason that the stochastic profile in the new market need resemble the prior stochastic profiles of the previous components. More generally, under the rubric of packaging and reselling risk, *new connections are*

13 This was exactly the case with credit default swaps over the last decade. They were traded over-the-counter to evade any collateral requirements, but were purportedly a repackaging of insurance and mortgages sold through more conventional fixed price and auction markets (Tett, 2009). Orthodox theory told the inventors of CDS that the format of the market didn't matter – only the notion that they were completing a process of 'arbitrage' (MacKenzie, 2009b: 67). Hence orthodox arbitrage theory lies at the heart of the conviction that the shadow banking system could purchase 'protection' as a substitute for bank shareholder capital on the balance sheet (Johnson, 2009).

being forged between previously decoupled market formats, or, in other words, the complexity of the entire system of markets has been enhanced. Equivalent notions in physics of the evolution of complexity have been formalized in the literature on networks of automata.[14] Things which are formally asserted to be identical in attempts to repackage 'risk' constitute in fact the debut of systemic novelty, both in the definition of the new composite financial instrument and in the institutional rules under which the new composite is traded.

This scenario is not the only means by which the augmentation of complexity levels can occur in an economy, but it is the one most relevant to the Minsky-style narrative of exacerbated fragility of the financial sector. As one insider put it, 'the so-called Too Big To Fail policy is intimately intertwined with derivatives regulation policy ... complex derivatives render these behemoth institutions Too Difficult To Resolve' (R. Johnson, 2009). Approaching Minsky fragility as a subset of the larger evolution of market complexity endows that school of macroeconomic thought with a different, and arguably better-suited, suite of microfoundations than some awkward appeal to optimization. We shall discuss some pertinent formal measures of this complexity in the next section.

It should be noted that Minsky explicitly renounced Walrasian price theory as bearing any relevance to his theory (1986: chapter 7), but felt that he could replace it with a generic markup theory of price, which he attributed to Michal Kalecki and others. While one can appreciate his motives in opting for that tradition, I would argue it did not allow him to escape the basic flaw of orthodox theories of financial risk, namely the presumption that all markets operated in the same manner, and, thus, to inadvertently sanction the orthodox arbitrage pricing theory of risk. In effect, Minsky wrote as if the entire economy was one big posted price market, a solecism nearly as debilitating as presuming one big Walrasian auctioneer. The key to an evolutionary approach to complexity is the appeal to the diversity of markets to explain the appearance of macroeconomic instability that cannot be reduced to a calculation of risk that can itself bear a determinate price.

Why could this distinction be so crucially central to the complexity account of the crisis? Briefly, the orthodox theory of homogeneous identical markets essentially posits a *symmetry* that renders the model of market operation much more tractable. Because it posits that market operations are all essentially alike and arbitrage-free in the absence of temporary second-order disturbances, it then follows that the price of any contingent claim, no matter how exotic and composed of far-flung heterogeneous constituents, can be computed and realized independent of any posited level of 'risk'. The entire market could therefore be replicated by a finite portfolio of instruments such as bonds and derivatives. In a sense, the orthodox account of arbitrage pricing bears the

14 Crutchfield (1994), Gornerup and Crutchfield (2006), Newman (2003), Percus *et al.* (2006), Barabasi (2009).

built-in presumption that the economy is inherently and continuously scale-free: any arbitrary subset (suitably apportioned) can reproduce the behavior of the whole: you can encompass the entire market in a single portfolio. But, if that is true, then there is no way that a crisis can start as a local and confined outbreak of defaults, only then to subsequently propagate throughout the system. Thus, there are only two dichotomous states of the entire macroeconomy in orthodox theory: incomplete, or perfect arbitrage-free equilibrium.

Furthermore, since orthodox theory rarely takes into account any considerations of formal Turing computability, this price purportedly can be realized upon any suitably programmed computer whatsoever. This symmetry principle in essence collapses any and all complexity measures to a single uniform featureless landscape. In this sense (and responding to Krugman's meditation), *there can be no relevant indices of complexity in the orthodox arbitrage pricing theory* of modern finance. Either the market is deemed 'complete', or else it is 'imperfect'. Dichotomy pre-empts complexity. States of grace poised somewhere between these two extremes are literally inconceivable. This perhaps explains why Minsky's attempt at a graduated sequence of increasingly fragile states of leverage has never held much if any intellectual attraction for orthodox finance economists or macroeconomists. And, patently, that explains why a neoclassical theory of market complexity and macroeconomic fragility has been an oxymoron.

The alternative to orthodoxy proposed herein starts market innovation which initially appears to the participants as a straightforward repair of minor market imperfection – in finance, usually under the banner of arbitrage; elsewhere as an exercise in risk reduction – but inadvertently introduces excess complexity into the system, eventually rendering it vulnerable to system-wide breakdown. Curiously, each participant believes that they are simply further protecting their own portfolios through more elaborate mechanisms of risk sharing, when in fact they are introducing the potential for a type of stochastic behavior that resists all forms of insurance: technically, what is called in the insurance profession 'inherent vice'. The class of models that attempt to capture the phenomenon of inherent vice fall under the rubric of theories based upon evolutionary complexity, rather than appeals to misguided or irrational behaviors cited at the outset. There is nothing special about contemporary human frailty, which is why the latter models are unavailing in explaining the timing and specific outlines of the current crisis. We need to isolate what is peculiar to the current crisis to effectively understand it.

There have been a few quasi-orthodox attempts to capture the intuitions behind the evolutionary complexity accounts of the crisis, but most of them are still rooted in the older Walrasian framework. One is the attempt by Buz Brock and his colleagues Cars Hommes and Florian Wagener (2009) to argue that more hedging instruments may actually destabilize markets. They suggest this by means of the insertion of 'non-fully rational agents' with mechanical reinforcement learning dynamics into a standard rational expectations macro model. Since it is no surprise that insertion of arbitrary 'irrationalities' create problems in standard macro models, and, furthermore, that their setup entirely ignores the phenomenon of innovation in *markets* per se, their models utterly fail in capturing the phenomenon of inherent vice. Somewhat more pertinent attempts have been proffered by the econophysicist Matteo Marsili (2008, 2009). He adapts a model from statistical mechanics in order to demonstrate a situation where financial engineering becomes a part of the problem, rather than a presumed part of the solution, as in orthodox finance theory. He explicitly conceptualizes a crisis as a phase transition in the entire system, due to individual 'rational' contributions. Nevertheless, a residual commitment to the orthodox model of Arrow securities leads Marsili to posit an 'increase in complexity' to be measured by a simple index of the total number of idealized identical 'assets' divided by the number of the states of the world: when this version of complexity approaches unity in the model, then the economy approaches a crisis without reaching full neoclassical equilibrium. The notion of identifying a financial crash with the realization of full Walrasian equilibrium is rhetorically cute, yet practically unavailing. Again the constraints of orthodox models tend to hamper the attempt to model serious breakdown: in the case of Marsili, all financial markets and all financial instruments are still treated as operating essentially alike, differing only in the preordained states of the world covered (an artifact of recourse to statistical mechanics in the first place – all molecules must be treated as interchangeable in order to apply statistics), so nothing really *novel* appears in this strange simulacrum of 'financial innovation'. Furthermore, there is no serious evolution in the model, and no computational considerations, largely due the familiar stress on individuals and their utility functions. Hence, in this instance, there is no well-formed connection to natural science connotations of 'complexity' in any formal sense.[15]

The recourse to physical models need not inevitably misrepresent legitimate notions of complexity in the economy, however. There is a burgeoning research tradition in computer science which has explored the profound commonalities between computational intractability and models of statistical mechanics (Percus *et al.*, 2006). It has been known for decades in computer science that the 'satisfiability problem' [or k-SAT problem] for Boolean functions written in conjunctive normal form where the number of literals [k] is greater than or equal to 3 is NP-complete; that is, a common form of logical function is computationally intractable, in the sense that worst-case versions would not

15 The failure of the model in what must be stressed is an author who has so clearly seen the nature of the flaw in orthodox financial theory raises the issue of whether an appropriate modeling strategy can be expected to emanate from econophysics, when indeed econophysicists have played such a significant role in precipitating the current crisis in the first place (Lohr, 2009; Jones, 2009; Cho, 2009). Unfortunately, we cannot address that question in this venue; but see Rosser (2008), Rickles (2008).

be solvable in polynomial time. More intriguingly, recent research has shown that the probability that a random k-SAT formula is satisfiable depends upon a function of a simple threshold variable: generally, probability of solution falls off drastically as one approaches the threshold, while computational 'costs' rise exponentially. Further, dramatic progress was then made when it was noticed that this mathematical structure was isomorphic to simple models of magnetism and phase transitions in statistical physics.[16] This isomorphism has subsequently provided the basis for adaptation of a number of models from statistical physics to help clarify problems of computational intractability (sometimes misleadingly called 'complexity'): under generic conditions, many problems undergo 'phase transitions' from regions where computation is easy in cost terms to where it is NP-hard or NP-complete. This crossover tradition has not yet been applied within economics, if only because the bulk of the orthodox profession has been unsympathetic to the importance of computational complexity or intractability as a serious source of economic instability.[17]

The theory of the economy as a diverse ecology of markets differentiated according to their formal computational structures, or as I have called it elsewhere (2007), the theory of markomata, is therefore much better suited to express the Minsky scale-dependent insight that increased complexity of finance can lead to increased fragility of macroeconomic stability. What does 'market failure' look like in this tradition? It posits a level of complexity attained by an interlinked network of markets whereby, past a certain threshold, no amount of normal calculation can stipulate what some subset of derivative goods in composite hybrid markets are actually worth: markets begin to fail in their primary function of attribution of price. This in turn compromises all manner of existing hedge finance, pushing it into speculative and even Ponzi territory. What starts out as a local outbreak of defaults upon contingent claims defies all attempts at containment, and the consequent inability to value financial instruments begins to spread. Derivatives become decoupled from their purported components; vast quantities of value evaporate with distressing alacrity; fear spreads that there is no perceptible floor under certain prices; and some markets have to be shut down temporarily because they have become untenable. No amount of leverage can be justified through existing market signals. Price collapses seemingly initially limited to one sector (say, house mortgages) are discovered to have ripple and contagion effects in sectors

16 For those who recognize such things, the k-SAT problem is isomorphic to the Hamiltonian for the simple Ising model, and in a more complex case, the model for a spin glass. While spin glass models have been proposed in economics (Durlauf and Brock have written many papers in this vein), they have *not* made use of this isomorphism to computational complexity, which is central to our current concerns.

17 For examples of open hostility, see Rust (1997), Durlauf (2005), or indeed, many of the participants in the symposium issue of JEBo devoted to Mirowski (2007). For general historical discussion of the phenomenon, see Mirowski (2002, 2007).

seemingly far removed (e.g., insurance premiums, interbank lending rates) through channels which are obscure to the participants. What had previously perdured as a scale-dependent process now threatens to become scale-free: in physics terminology, the system has undergone a phase transition. It now subsists in a state which, by its very nature, can *never* be subject to insurance: the technical state of inherent vice. No one can trust the validity of anyone else's balance sheets, and many financial instruments become hot potatoes in the absence of consensus mechanisms for validation. The terminus to the downward spiral comes with: (a) some agent, usually some organ of the state in its guise as the lender of last resort, simply dictating what values should be; combined with (b) profound changes in the existing rules of the game, which often means widespread revision of the formal structures of markets (Anderson, 2007). Once the crisis grows acute, markets cannot heal themselves on their own; some outside social authority or external entity must push the reset button.

In Section 2, we informally sketch the model which attempts to describe a crisis of increased complexity in market innovation due to Minskyesque inventions of more complex financial instruments, based upon computational notions of complexity. An account of certain specific historical events from 2008/9 in light of this alternative theoretical account will be postponed until a subsequent paper,[18] as will the full elaboration of the mathematical principles behind the recourse to formal attempts to characterize the turning points in a revised Minskian model.

2. The crisis as a collapse of market complexity

When discussing the role of the financial sector in the macroeconomy, it would seem less of a stretch to suggest that the individual market formats should themselves be treated as sets of abstract algorithms, that is as *automata*.[19] After all, much of the riotous effulgent 'innovation' in the run-up to the current crisis was initiated by econophysicists and mathematicians bringing their models to banks and hedge funds, there to be enshrined in automated pricing schemes and programmed as automated trading mechanisms. Until recently, it was the subject of boasts by these 'quants':

> Options submit to the well-known Black–Scholes formula for pricing ... it is possible to go beyond this analytic formalism using the algorithms of computational intelligence. In [one case], neural networks are used to push past the constraints of the Black–Scholes equation and price these complicated assets appropriately ... Fuzzy logic and evolutionary processes are used in [another case] to calculate trading rules for markets... One approach to asset allocation

18 Although a good start can be stolen on the conventional wisdom by consulting MacKenzie (2009a, 2009b) and Muolo and Padilla (2008).

19 This section is a revised version of a more elaborate description in Mirowski (2007).

> is to control the value-at-risk of a given portfolio. As a control problem this is well suited to ADP and other computational intelligence techniques. (Seiffertt and Wunsch, 2008: 29)

Much of the contemporary financial system had come to rely upon explicit algorithms in order to carry out any trade whatsoever: banks had their 'value-at-risk' packages to propitiate the regulators and their shareholders, the ratings agencies had their own programs to validate their dicta, and the only reason mortgages could be packaged into derivatives was because they were said to be backed up by credit scoring algorithms. The output of one set of algorithms became the necessary input for another class of algorithms. As one science studies scholar described it:

> The models used for pricing CDOs run fairly quickly on modern computers, but calculations of hedging ratios and risk-management parameters are more demanding. Even using grids of several hundred interconnected computers the risk calculations can take several hours. But then a physical constraint becomes relevant: heat. Many banks want crucial computations to be performed in or close to main offices, trading floors, and risk managers, because even fibre-optic connections are still too slow ... (MacKenzie, 2009a:17)

The many faces of financial innovation had mostly become tied to the computer and the notion that information itself could be securely packaged, processed, and sold alongside the novel commodities that had been innovated.

The reaction of most economists to this ebullience has generally taken one of two formats: either to insist that the justifications of the algorithms found everywhere in finance be conceptually grounded in prior orthodox economics in order to be taken seriously, or else to dismiss them as no better than expensive high-tech Chartism. In this paper, we take a third path: that the sorts of computationalism found recently in the rarified precincts of finance are in fact especially clear cases of a more general phenomenon, namely the phenomenal manifestation of the vast majority of markets in the modern world as abstract automata.

Markomata defined

In this heterodox microeconomics, the most rudimentary description of a market begins with the notion of a finite automaton. In computer science, a finite automaton Ш is defined over an alphabet $\alpha = \{\alpha_1, \ldots, \alpha_m\}$ with states $\theta = \{\theta_1, \ldots, \theta_n\}$ is given by a function T– called a transition function which maps each pair (θ_i, α_j) into a state θ_k; a subset of states $\Theta = \{\theta_k\}$ called final accepting states causes Ш to halt. A finite automaton can be thought of as an extremely limited computing device with no external memory capacity but a single 'working tape', which it can read only once. After reading a symbol on the 'tape', it either accepts or rejects it, depending upon the state that the device is in; it then enters the next state prescribed by the transition function. If the transition function

T maps an existing state of Ш into more than one state, then it is called a non-deterministic finite automaton. Automata are very simplified abstractions of the basic components of a computer, but they are also the simplest formal description of abstract language recognition.

Markets never just do a single thing, but in fact simultaneously perform many different functions, often bundled together in the operation of a single institution.[20] Suppose, for purposes of illustration, we set out to formalize one function of one simple market as an automaton. In one (arbitrary) initial economic example, the order execution function of a very rudimentary market, such as the posted- or fixed-price market, will be modeled as a non-deterministic finite automaton. A single unit of the commodity is offered at a single price, where the alphabet concerned is the rational numbers: order execution either matches that number as bid by the purchaser, or is rejected. At this stage, it is important to keep in mind that it is merely the order execution function that is captured by this particular NDF, and not the entire range of functions potentially performed by any real-world instantiation of the posted-price market in question. Any real-world market is a complex catenation of quasi-modular algorithms, further bedeviled by human meddling.

Data dissemination, order routing, clearing, record-keeping, and all the rest might themselves be described as being composed of automata of various degrees of computational capacity; any real-world market is formally characterized by the composition of these component automata. Pertinent to our current concerns, this begins to reveal the true combinatorial explosion of forms inherent in the theory of markomata. Evolution deals in profusion, and not sleek Platonic forms. Even restricting ourselves (for purposes of brevity and simplicity) solely to tasks of order matching and execution, the possibilities present in any real-life situation begin to outstrip our capacity to subject them to formal abstraction. Can buyers themselves bid, or only respond to the sellers' ask? Are there multiple buyers/sellers, and can they initiate/respond in real time? Can they react to one another, as well as to the opposing side of the market? Can they communicate through channels other than the order execution algorithm?

First pass at economic complexity: Chomsky hierarchy

The explosion of potential functions is partially mitigated by subjecting markomata to the computational and complexity hierarchies propounded within automata theory. The heart of this theory is the treatment of complexity; but even computer scientists use the term loosely, so we need to be very precise when importing it into economics. The first, and most important, computational

20 This is discussed in detail in Mirowski (2007); because of our current interest in macroeconomics, we must pass them by here. However, Minsky suggested something similar: 'prices cannot be treated as though their only function is to allocate resources and distribute income. Prices must also be related to the need for cash flows to validate capital assets, financial structures, and the business style of the economy' (1986: 142).

Table 1. Markomata hierarchy of order execution

Automaton type	Recognizes lang.	Memory	Markomata
Finite III	Regular	None	Posted-price
Pushdown	Context-free	Pushdown stack	Sealed bid
Linear bounded	Context sensitive	Finite tape	Double auction
Turing Machine	Recursively enumerable	Infinite tape	None

hierarchy is known in computer science as the 'Chomsky hierarchy' (Davis *et al.*, 1994: 327–329). It conventionally relates the complexity of the language recognized to the memory capacity of the class of automata deployed. It is summarized for the order execution function in Table 1 below.

One implication of the Chomsky hierarchy is that some problems, which are unsolvable at the lower levels of computational capacity, can be shown to be solvable at the higher levels. Furthermore, there exist some problems that cannot be solved even at the most powerful level of the hierarchy; some strings are Turing non-computable on the Turing Machine. However, the hierarchy is inclusive, in the sense that the more powerful automaton can perform all the calculations of the automaton lower down in the hierarchy, because it can *simulate* the operation of machines of lesser computational capacity. This leads to the important analytical notion of 'markomata simulation'.

The idea of one markomata simulating the operation of another is quite familiar to market practitioners, even though it has been absent up until now in orthodox economic theory. For instance, the futures market for red no. 6 wheat 'simulates' the spot market for red no. 6 wheat, in the sense that it can perform the same operations, augmented by other related operations, in the course of 'tracking' the wheat market. Likewise, the dealer-organized wholesale market 'simulates' the posted-price markets of the retailer, while superimposing other functions. In an abstract computational sense, the futures market 'encapsulates' the model of the spot market within its own algorithms. Markets for derivatives also simulate the markets for their component assets. This would be the case even if the derivatives markets were operated as a double auction, whereas the spot markets were operated as a sealed-bid clearinghouse auction. The plague of Collateralized Debt Obligations [CDOs] were inventions created and sold by banks over the counter, which were supposed to be based upon loans sold to homeowners through dealer-mediated markets. The implicit claim (now proven false) was that the market for CDOs could adequately simulate the operation of the market for mortgages (R. Johnson, 2009). This claim was rendered recursive when banks packaged and sold CDOs composed of prior CDOs, or the notorious 'CDO Squared'. The theory of computation informs us that certain specific market forms can simulate other market forms *as long as* they are composed of markomata of greater or equal computational capacity. That condition was flagrantly violated during the last decade of the sale of CDOs. Nevertheless, in general, the reason that the global markomata hierarchy does not collapse down to a single flat uniformity – say, all markets operating like a single massive computer, as in the Walrasian tradition – is that more computationally complex markets situated higher up in the Chomsky hierarchy perform other functions over and above those performed by the markets that they simulate: for instance, futures markets may seek to facilitate arbitrage of price discrepancies as well as track the spot markets in their purview.

Table 1 suggests that some forms of automata may be mapped into different formats of order execution. While the posted-price format possesses no memory capacity and therefore qualifies as a finite automaton, a sealed bid auction requires the comparison of a submitted bid to an ordered array of previously entered bids stored in a memory, and therefore qualifies as one of a number of *k*-headed pushdown automata.[21] Sealed bid order execution requires an ordering of submitted bids, which can be captured by a first-in first-out memory stack; hence the 'pushdown'. The standard double auction requires even more prodigious memory capacity, given that sequences of bids and asks stored in different identifiable memory locations must be retrieved and compared, and therefore should exhibit the computational capacity of (at least) a linear bounded automaton. Table 1 also suggests that no isolated individual markomata format has the power of a Turing Machine – in other words, the entire economy could not be run as a single auctioneer-mediated sealed bid auction (Walras' original vision), or indeed as any other monolithic market format. However, the entire market system, conceived as a network of individual automata passing price and quantity messages back and forth, may indeed display such capacity from the macro perspective.[22]

Because the same physical commodity or financial instrument can be and often is sold through different markomata, sometimes even within the same spatiotemporal coordinates, different markomata display different price and quantity profiles, it follows that there can be no such lemma as the 'law of one price' in computational economics. (This will be central to the subsequent complexity account of the crisis.) If there might be a universal terminus toward which all automata tend, it is toward their internally defined 'halting conditions'. But even here, one can easily overstate the predictable mechanical character of market automata. It is a theorem of computational theory that:

> There is no algorithm that, given a program in the language $L(\alpha)$ and an input to that program, can determine whether or not the given program will eventually halt on the given input. (Davis *et al.*, 1994: 68)

21 The proof is sketched in Mirowski (2002: 571).

22 What is required under this model description would be an existence proof for a universal Turing machine, that is for an abstract machine using the same language as the markomata, but able to simulate any and all markets in the entire network of linked markomata. For an informal sketch of the structure of such proofs, see Arora and Barak (2009: 20–21).

The undecidability of the halting problem bears direct relevance for the ambitions of an evolutionary computational economics. The impossibility theorems of computational theory do not somehow prevent the construction of specific markomata for attainment of specific targeted functions (since this is the practitioner's notion of the 'predictability' of the market); they merely prohibit the economist from making any ironclad predictions about the inevitable outcomes of the price system *as a whole*. As individual markomata become increasingly networked, their computational powers become increasingly augmented in complexity (in the Chomsky definition), and transcendental guarantees that a particular market format will continue to operate as it has done in the past are repeatedly falsified. In markomata economics, the very notion of 'market failure' thus assumes an entirely different meaning from its usual economic connotations. When a markomata fails, in this conception, it appears unable to halt. Prices appear to have no floor (or ceiling, in the case of hyperinflation); the communication/coordination functions of the market break down. Hence there exists in the real world the phenomenon of 'circuit-breakers', which call a halt to market operations – something that makes eminent good sense in a computational economics (even as they are disparaged in neoclassical finance theory). Bank holidays, stock exchange closures, and bankruptcy itself are examples of such resets. Earlier generations of market engineers had appreciated the need for a manual override when there were 'bugs' in the system.

Far from existing as a purely abstract theoretical possibility, a stark instance of markomata failure materialized after this article was written and accepted in this journal. On 6 May 2010, American equity and futures markets suffered a sharp collapse of many prices (some to mere pennies in price), with an equally sharp recovery soon thereafter. A *New York Times* article dubbed this a 'flash crash' and quoted one participant: 'There was no pricing mechanism', Mr. Clancy said. 'There was nothing. No one knew what anything was worth. You didn't know where to buy a stock or sell a stock' (Schwartz, 2010). The utter breakdown of the price assignment function in a matter of minutes further undermined confidence in the operation of financial markets. More disturbing, more than a month after the incident there was still no consensus explanation of the causes of this very real market failure (US Securities, 2010). Conventional economic theory stood flummoxed by what participants experienced as a halting problem.

In computer science there is no 'halting problem' specifically for finite automata; here we posit that the entire global network of markomata could evolve to the point that it can itself possess the power of a Turing Machine, and hence the global networked automata become subject to the halting theorem. This might be regarded as a further exemplification of the emergent properties of the market system at different levels of analysis. Market forms start out isolated and operating at very low levels of complexity: market failures (in the current sense) are relatively unknown; innovation turns them into ever-more elaborate markomata. In the absence of severe macroeconomic contractions, the pace of complexification accelerates. It is characteristic of the dynamic of economic development that human participants, from their individual parochial vantage points, attribute any observed 'failure' to a particular localized market, whereas from an analytical point of view, it should more appropriately be attributed instead to the entire network architecture.

Nevertheless, the hierarchy of computational capacities of markomata does appear to suggest that there might just exist an arrow of time inherent in economic history: that is, mutation of markomata eventually induces the appearance of markets of greater computational complexity than their predecessors, as well as more elaborate hierarchies of 'small worlds' networks. Here computational evolutionary economics would concern itself with the evolution of diverse algorithmic market formats within an environment of irreducibly diverse human beings. To a first approximation, the individual 'species' of markomata would be defined according to the algorithmic specifications; 'genera' would be classified according to taxonomies such as the Chomsky hierarchy of computational capacities. 'Selection' occurs through humans promoting the differential use and reproduction of specific markomata in distinct spatiotemporal locations. Just as in biology, selection pressures are exerted at various different levels: at the level of the individual component algorithm (say, alternative protocols of order routing), at the level of the individual markomata (say, the conventional hypermarket posted-price markomata), and at the group level (say, the network of markomata related by their computational ability to simulate each other's operations). Some markomata become established in certain limited areas (such as double auctions in finance) because they are perceived to bundle an array of functions deemed particularly suited to the community constituted to promote certain objectives (futures and options, arbitrage, rapid market clearing), whereas others become more evenly distributed across the landscape (such as posted price) due to their relative robustness to a wide array of situations. 'Mutation' is present when humans try to 'bend the rules' or otherwise circumvent prior market structures, as in the example that prefaced this section. A small amount of mutation is beneficial for the system, since it is a major source of novelty; however, pervasive mutation in the sense of lawlessness or corruption is ultimately harmful to the continued existence of the deme, if not the entire ecosystem. An empirical evolutionary economics would therefore confront the possibility that, as a general trend, markets as a whole have attained a higher degree of computational complexity throughout time, even though most individual markomata might still operate at a relatively rudimentary level. Average markomata complexity might therefore rise much more slowly than a complexity characterization of the macro aggregate.

Finally, we come to the fundamental narrative about how Chomsky complexity can be integrated into a theory of macroeconomic crisis. Commodities markets of low complexity become increasingly integrated over an expansionary phase, and financial entrepreneurs think they uncover opportunities for arbitrage and the injection of greater degrees of leverage by inventing hybrid market forms and novel financial instruments: this is a process starting long ago with bills of exchange, bonds, shares of joint stock companies and futures markets, but then graduating to derivatives, derivatives of derivatives, trading of indices, derivatives of indices, and so forth. The extension of arbitrage is intimately associated with increased magnitudes of leverage on balance sheets, mainly because arbitrage positions generally incur losses for some period of time before they become profitable (Buenza *et al.*, 2006). Hence, Minsky's stress on the progressive deterioration of balance sheets is intimately tied up with the innovation of financial market structures of enhanced Chomsky complexity; indeed, this is one reason that they initially appear so alluring and lucrative to the entrepreneurs who bring them into existence. Perhaps this might be further encouraged by lax regulation, but this could hardly be the root cause of the increased systemic fragility (Gowan, 2009). Because there can be no law of one price wherever one superimposes markomata of differing Chomsky complexities, market innovation superimposes something novel on top of existing markomata, generally in the form of structures of higher computational complexity in order to encompass and integrate operation of the diverse component markets. Hence, the overall dynamic of market arbitrage is not inherently stabilizing, but may introduce greater price instability. Because augmented complexity is often justified as permitting arbitrage and risk reduction, the growing pyramid of markomata is accompanied by a parallel pyramid of leverage, because arbitrage activity is always accompanied by some form of bridge finance. Since attention is rarely drawn to the structural changes in markomata, the rise in market complexity is almost always perceived in retrospect as manifesting itself as a symptomatic progression through the Minsky stages of hedge, speculative and Ponzi finance. Few have the time or inclination (or theoretical capacity) to ask: why don't the second-order and third-order markets result in a more stable structure of valuation?

Some risks are diverted to other markets, to be sure, but the pyramid of increased complexity introduces a looming danger which cannot be insured away: the computational incapacity of the markomata linkages to themselves adequately mirror what is going on in their component markets. Furthermore (as happened with CDOs in the recent bubble), if financial markets of lower Chomsky capacity in trading the derivative have been superimposed upon markets of greater complexity that retail the component assets, then it is a theorem of automata theory that the architecture of networked automata will exhibit computational failures that become more frequent and more severe. The prices of derivative instruments necessarily become decoupled from prices of the underlying assets. The eventual solution to these problems is not *more* financial innovation, or even the issuance of further debt, but rather precipitation of a crisis of complexity, which appears as delinkage of some markets from the network, and wholesale closure of others, the bankruptcy or disappearance of some providers of financial instruments, all leading to a pronounced retreat from high-complexity markomata in the realms of finance. Some of the nouveaux exotic instruments may disappear, and, in other cases, the rules governing markomata are tightened, producing more reliable computation.

The buildup to a crisis of markomata complexity is a root cause of macroeconomic instability in a literal economics of complexity, but it would not constitute the sole cause. Indeed, computational problems would rarely become manifest initially as the full-blown appearance of formal 'halting problem', but rather as local increases in computational intractability in specific markets, which then reach a tipping point, precipitating a widespread crisis of computational complexity and liquidity. Hence, the theory of markomata suggests there are *at least* two types of complexity involved in the evolution of the long cycle in economic development.

Second reference of complexity: intractability

One obstacle that the computer science literature presents for economists is that sometimes the term 'complexity' is used to reference the Chomsky hierarchy of automata, as discussed above; but other times it is used to refer to the theory of computational intractability, which we already briefly encountered in the previous section.[23] In this latter sense, 'complexity' refers to the cost and difficulty of solving a designated class of problems upon a *specific automaton*, namely the abstract Turing Machine. In the former, it refers to computational capacities of different classes of automata. Because the economic theory of markomata must necessarily distinguish between simpler automata and the Turing Machine, it is of the utmost importance to insist upon keeping the terminology separate. To forestall confusion, henceforth the former phenomenon will be called computational complexity, and the latter, computational intractability. The first is more structural, whereas the latter verges more naturally upon economic notions of resource usage in the course of normal market operations. The distinction between complexity and intractability, while the subject of two distinct literatures in mathematics and computer science, actually overlap conceptually to some degree due to the fact that the Chomsky hierarchy is sometimes portrayed as distinguishing automata by their memory capacities, whereas distinctions between solutions in polynomial and non-polynomial time (or space) are linked to the amount of resources available to the Turing Machine. However, it will behoove us to remember that Chomsky complexity is an attribute of the 'machine' itself (and hence the

23 For one recent example: Arora and Barak (2009: xx) present the two distinct issues as 'computability vs. complexity', which can only flummox novices. Of course, those authors have no interest in automata with powers less than a Turing Machine.

markomata), whereas intractability is an attribute of the class of problems fed into the 'machine', that is the tasks that are expected to be carried out.[24]

Of course, different markomata endowed with different algorithmic mechanisms will encounter different levels of intractability when confronted with any given market task. The question relevant to a Minsky-style macroeconomics is whether any theoretical generalizations can be made about the types and stringency of intractabilities encountered. Due to space limitations, we can here merely suggest two possibilities. The first has to do with the specification of whole classes of economic problems that exhibit severe intractability in general; and the second has to do with the appearance of intractability as a symptom of the onset of crisis.

It is well known that whole classes of computations can be shown to have worst-case attributes that cause them to be classified as tractable (of class **P**) or intractable (class **NP-complete**) when run on a Turing machine.[25] Roughly, if the algorithm can run in a time frame which is a polynomial function of the size of the program input, as the program input is increased, then it is considered to have an efficient algorithm, and fall into 'polynomial' class **P**. Conversely, if the time requited balloons exponentially as a function of input size, but the correctness of the output can be verified in polynomial time, then the algorithm is said to fall into the relatively intractable class of 'nondeterministic polynomial' or **NP** problems. The very hardest problems from the viewpoint of computational intractability are dubbed **NP-complete**. An efficient tractable solution to any extant **NP-complete** problem would imply **P** = **NP**, and therefore render every all other **NP** problems tractable in principle.[26] It is still an open question whether there exists a mathematical proof demonstrating **P** = **NP** or the converse.

The classes of fundamentally intractable problems that fall into the category of **NP-complete** are well taxonomized, and include such workhorses of orthodox economics as the travelling salesman problem and the simplex algorithm in linear programming. Many calculations of Nash equilibria are also intractable (Conitzer and Sandholm, 2008). More relevant to the Minsky tradition is a demonstration that attempts to search for arbitrage profits between a highly connected set of markomata, each of which providing its own 'price' for a finite set of commodities, turns out to be isomorphic to the Hamiltonian cycle

24 This distinction will turn out to be crucial from a phenomenological viewpoint. Failures provoked by computational complexity could only be observed from an external Olympian viewpoint regarding the entire network of markomata as a whole, whereas failures of computational intractability are readily observed by participants on the ground at the level of individual markomata.

25 There is some variability in definition between whether the intractability is defined in terms of time elapsed or size of the tape (or space) needed for the computation. We ignore those distinctions here.

26 The current status of the mathematical conjecture that **P** = **NP** is covered in Fortnow (2009) and Arora and Barak (2009).

problem in graph theory, which is also known to be **NP-complete**.[27] This result does not suggest that arbitrage profits are impossible to compute – far from it – but rather that as the size of the markomata graph grows, it becomes progressively more daunting to determine whether arbitrage opportunities exist within the current system. This simply demonstrates that the doctrine of arbitrage pricing theory, that any and all price differentials are rapidly dissipated by arbitrageurs, is implausible in this theory, *pace* Robert Lucas' armchair observation that one finds no $50 bills on the sidewalk. The theory of computational intractability suggests that in worst-case scenarios, the participants would not even be able to calculate many arbitrage opportunities in real time. In those circumstances, it would be implausible to claim that traders 'know' whether or not arbitrage options exist. This, in turn, undermines the very notion of risk management built into the foundations of orthodox macroeconomics.

A nicely worked out example of what these sorts of intractability results imply is a recent paper by (Arora *et al.*, 2009). To make it simple, they imagine a trivial case where assets are defined as being either 'normal' or 'junk', posit derivatives composed of complex combinations of such assets, as well as various practices reminiscent of the current crisis such as 'trenching' the derivative, and then ask: how difficult would it be for a computationally bounded purchaser to detect whether the derivatives were truly randomly pooled, versus those inordinately biased towards 'junk'? The answer they provide is that it 'may be computationally intractable to price derivatives even when buyers know almost all the relevant information... even in very simple models of asset yields'. They accomplish this by demonstrating that the search for price is isomorphic to the intractable planted dense subgraph problem.[28] One could regard this as one small instantiation of what may be a much more endemic problem of intractability in the larger phenomenon of arbitrage.

The observation that search for arbitrage profits is **NP-complete** is a 'static' proposition, one which, however, does not differentiate between periods of expansion and periods of macroeconomic crisis. But the most salient lesson of more than a decade of research on network theory is 'the realization that the structure and evolution of networks is inseparable' (Barabasi, 2009: 413).

27 The proof is sketched in Mirowski (2007: 231). The generic Hamiltonian path problem is covered in Arora and Barak (2009: 53).

28 There are a few technical niceties which distinguish their approach from the current markomata theory. Those authors simply presume there is one 'true' stochastic value for each asset (thus removing all considerations of markets from their purview), whereas we do not. Also, they present a result for average case intractability, whereas the results we cite in this paper are for worst-case intractability. Here we observe the tradeoff between writing down a manageable simplified model that satisfies the mathematician, at the cost of misrepresenting the phenomenon, and sketching an impressionistic model that is guaranteed to convey the most important theoretical economic concepts.

Hence, there is another class of analytical models, mentioned in the previous section, which seek to link models of intractability to changes in the state of the automaton itself, usually relative to some forcing parameter. This is research tradition in computer science that has explored the profound commonalities between computational intractability and models of statistical mechanics (Gomes and Selman, 2005; Percus *et al.*, 2006). The paradigm for this research is the k-SAT problem for Boolean functions written in conjunctive normal form, where the number of literals [*k*] is greater than or equal to 3 is **NP-complete**; recent research has shown that the probability that a random k-SAT formula is satisfiable depends upon a function of a simple threshold variable. Written as a function of the order parameters, generally probability of a solution falls off drastically as one approaches the threshold, while computational 'costs' (usually time) rise exponentially. In thermodynamics, the order parameters are denominated in temperature; in computation, the corresponding order parameter is average connectivity in a random Boolean network. The beauty of this literature is that it often is able to uncover a 'phase space' for problems that differ in a small number of parameters, and demonstrate that there is often a 'phase transition' in that space between problems/regions where computation is relatively tractable, towards a singularity where problems become exponentially intractable. In the computational literature concerning the k-SAT problem, it has been discovered that the threshold value can be written as a function of *k* when $k > 2$. This holds out the promise that for other known classes of **NP-complete** problems, equally general threshold values might be discovered for the transition to intractability.

The application of this class of models to macroeconomics could be promoted via the markomata approach to arbitrage. If the arbitrage problem can be framed as finding and calculating a Hamiltonian cycle on a network of markomata, then one might expect to find phase transitions from low to high computational intractability, which would link the onset of the crisis to the topology of the network of markomata.[29] Hence, one can imagine a pseudo-dynamical system where computations are chugging along quite nicely until small changes in a parameter of the system – say, alterations in the average connectivity of a network of markomata, brought about through financial innovation, or closer to Minsky, the average number of balance sheets altered by one markomata in a network – precipitate a phase transition into a situation where suddenly the same format of calculations become intractable. A similar scenario in markomata

29 Although phase transitions have been documented for Hamiltonian circuit problems, as yet very little work has been done applying statistical mechanics to analysis of the phenomenon. The closest analogue so far has been work on phase transitions in the 'number partitioning problem' (Mertens in Percus *et al.*, 2006: 126 et seq). The topology of the network of the financial system has recently become a question of interest to bank regulators (Altman, 2009).

theory could provide the observable trigger mechanism for precipitation of a Minsky-style crisis.

Imagine a small set of networked markomata jointly solving market algorithms: assigning prices, quantities, and keeping records. If the problems solved are already in the class **P**, then changes in their linkages and inputs do not produce any noticeable alterations in their computational abilities. However, some of those problems, like finding profitable arbitrage, or calculating the risk profile of a portfolio of derivatives, have been demonstrated to be **NP-complete**. In some states of connectivity, these particular problems are still tractable; but as 'financial innovation' causes different classes of markomata to become linked in novel and unprecedented ways, the entire system is driven into a region of phase space where it starts to 'freeze up': performance of the 'same' market calculations now slows down appreciably, and markomata start to chew up increasing amount of resources. Here financial innovation, either because of complexity pathologies covered in the previous section, or else simply because of approach to the intractability phase transition, brings the system nearly to a halt. The theory of intractability thus may permit the theoretical prediction of turning points in a Minsky-style collapse, as a function of the topology of the markomata network.

Up until now, most of these models have been written for simpler cases of dichotomous Boolean variables, and derived from simple spin glasses in physics, while the domain of markomata clearly extend over the rational numbers. Furthermore, spin glass models tend to be defined on simple lattices, whereas the topology of markomata graphs is bound to be much less regular. It needs to be stressed that spin-glass models per se (the perennial favorite of many denizens of the Santa Fe Institute) cannot be directly appropriated to do double duty as economic models of markomata. Nevertheless, the analogy between statistical mechanics and computation has led to some empirical generalizations across different fields, which are starting to be proposed with a view towards a general science of the detection of such phase transitions, deploying physical models of critical phase transitions. One recent survey of this literature (Sheffer *et al*., 2009) suggests the following empirical symptoms of a phase transition: approach to the critical point slows response to small perturbations; increased variance is observed in the time series of system outputs; increased asymmetry of fluctuations and 'flicker' phenomena start to appear; and a 'spatial' characteristic becomes noticeable where numerous coupled units tend to occupy the same states with greater correlation.

While these symptoms can only be suggestive at this stage, they do seem to evoke some rather striking empirical hallmarks of the current crisis: the sharp bunching of bursts of variance of prices in the run-up to the crisis, the tendency of the trends of most markets to 'phase lock' by contrast to their behavior prior to the crisis, and of course the heightened vulnerability to what would have been small insignificant shocks in normal times.

3. Conclusion

This paper has been an attempt to answer those who insist there is no alternative to a neoclassical macroeconomics, and to suggest possible alternative heterodox microfoundations for Minsky's account of economic crises, beyond the Kaleckian markup model. The sketch is based upon elevation of some technical notions of computational complexity to pride of place in the account. It is an attempt to portray a market system evolving to a point of 'inherent vice': an endogenous development which by its very nature, cannot be tamed through conventional insurance or risk models. It is avowedly faithful to the spirit of Minsky, who asserted that, 'The speed at which financial innovations . . . occur and spread is a governor that regulates the pace of movement out of hedge and into speculative finance' (1986: 212). It has provided a sketch of the construction of formal models of the crisis, though much work remains to be done.

References

Akerlof, G. and R. Shiller (2009), *Animal Spirits: How Human Psychology Drives the Economy*, Princeton, NJ: Princeton University Press.

Altman, D. (2009), 'The Network', *New Republic*, 9 October.

Anderson, J. (2007), 'Debt-Market Paralysis Deepens Credit Drought', *New York Times*, 7 October.

Arora, S. and B. Barak (2009), *Computational Complexity: A Modern Approach*, New York: Cambridge University Press.

Arora, S., B. Barak, M. Brunnermeier, and R. Ge (2009), 'Computational Complexity and Information Asymmetry in Financial Products', Princeton Center for Computational Intractability, http://www.cs.princeton.edu/~rongge

Barabasi, A. (2009), 'Scale-Free Networks: A Decade and Beyond', *Science*, **325**: 412–413.

Bieling, H. and J. Jager (2009), 'Global Finance and the European Economy: The Struggle over Banking Regulation', in B. van Apeldoorn, J. Koupil, and L. Horn (eds.), *Contradictions and Limits of Neoliberal European Governance*, London: Palgrave.

Bouchard, J.-P. (2009), 'The (Unfortunate) Complexity of the Economy', www. arXiv:0904.0805v1

Brock, W., C. Hommes, and F. Wagener (2009), 'More Hedging Instruments may Destabilize Markets', *Journal of Economic Dynamics and Control*, **33**: 1912–1928.

Buchanan, M. (2008), 'This Economy does not Compute', *New York Times*, 1 October.

Buenza, D., I. Hardie, and D. MacKenzie (2006), 'Price is a Social Thing: Towards a Material Sociology of Arbitrage', *Organization Studies*, **27**: 721–745.

Cassidy, J. (2008), 'The Minsky Moment', *New Yorker*, 4 February.

Cho, A. (2009), 'Econophysics: Still Controversial after all these Tears', *Science*, **325**: 408, 24 July.

Cogley, T. and T. Sargent (2008), 'The Market Price of Risk and the Equity Premium: A Legacy of the Great Depression?', *Journal of Monetary Economics*, **55**: 454–476.

Colander, D. *et al.* (2009), 'The Financial Crisis and the Systematic Failure of Academic Economics', *Critical Review: A Journal of Politics and Society*, **21**(2): 249–267.

Conitzer, V. and T. Sandholm (2008), 'New Complexity Results about Nash Equilibria', *Games and Economic Behavior*, **63**: 621–641.

Cowan G., D. Pines, and D. Meltzer (eds.) (1994), *Complexity: Metaphors, Models and Reality*, Reading, MA: Addison-Wesley.

Cowen, T. (2008), 'Three Trends and a Train Wreck', *New York Times*, 17 October.

Crutchfield, J. (1994), 'The Calculi of Emergence', *Physica D*, **75**: 11–54

Davis, Martin, R. Sigal, and E. Weyuker (1994), *Computability, Complexity and Languages*, San Diego: Morgan Kaufmann.

DeLong, J. B., A. Shleifer, L. Summers, and R. Waldmann (1990), 'Positive Feedback Strategies and Destabilizing Rational Speculation', *Journal of Finance*, **45**: 379–395.

Durlauf, S. (2005), 'Complexity and Empirical Economics', *Economic Journal*, **115**: F225–F243.

Eichengreen, B. (2008), 'Origins and Responses to the Crisis', http://emlab.berkeley.edu/users/webfac/eichengreen/e183_sp07/origins_responses.pdf

Erturk, K. and G. Ozgur (2009), 'What is Minsky All About, Anyway?', *Real-World Economics Review*, **50**: 3–15.

Fortnow, L. (2009), 'The Status of the P vs NP Problem', *Communications of the ACM*, **52**(9):78–86.

Frank, R. (2009), 'Flaw in Free Markets: Humans', *New York Times*, 14 September, Business section.

Gomes, C. and B. Selman (2005), 'Can Get Satisfaction', *Nature*, **435**: 751–752.

Gonerup, O. and J. Crutchfield (2006), 'Hierarchical Self-Organization in the Finitary Process Soup', www.arXiv:nlin/0603001v1

Gowan, P. (2009), 'Crisis in the Heartland', *New Left Review*, **55**: 5–29.

Ito, M. (2004), *Algebraic theory of Automata and Languages*, Singapore: World Scientific.

Johnson, R. (2009), Testimony before US House Committee on Financial Services, 7 October.

Johnson, S. (2009), 'The Quiet Coup', *Atlantic*, May.

Jones, S. (2009), 'The Formula that Felled Wall Street', *Financial Times [London]*, 24 April.

Leijonhufvud, A. (2009), 'Out of the Corridor: Keynes and the Crisis', *Cambridge Journal of Economics*, **33**: 741–757.

Lohr, S. (2009), 'Wall Street's Math Wizards Forgot a Few Variables', *New York Times*, 13 September.

MacKenzie, D. (2009a), *Material Markets*, Oxford: Oxford University Press.

MacKenzie, D. (2009b), 'The Credit Crisis as a Problem in the Sociology of Knowledge', Edinburgh working paper www.sps.ed.ac.uk/_data/assets/pdf_file/0019/.../CrisisNew19.pdf

Marsili, M. (2008), 'Eroding Market Stability by Proliferation of Financial Instruments', www.arXiv and http://videolectures.net/ccss09_zurich/

Marsili, M. (2009), 'Spiraling toward Market Completeness and Financial Instability', www.arXiv.0906.1462v1

Minsky, H. (1977), 'A Theory of Systemic Fragility', in E. Altman and W. Sametz (eds.), *Financial Crises*, New York: Wiley.

Minsky, H. (1986), *Stabilizing an Unstable Economy*, New Haven: Yale University Press.

Mirowski, P. (1986), 'Mathematical Formalism and Economic Explanation', in P. Mirowski (ed.), *The Reconstruction of Economic Theory*, Boston: Kluwer.

Mirowski, P. (2002), *Machine Dreams: Economics Becomes a Cyborg Science*, New York: Cambridge University Press.

Mirowski, P. (2006), 'Twelve Theses on the History of Demand Theory in America', in W. Hands and P. Mirowski (eds.), *Agreement on Demand*, supplement to *History of Political Economy*, **38**: 343–379.

Mirowski, P. (2007), 'Markets Come to Bits: Markomata and the Future of Computational Evolutionary Economics', *Journal of Economic Behavior and Organization*, **63**: 209–242.

Muolo, P. and M. Padilla (2008), *Chain of Blame*, New York: Wiley.

Newman, M. (2003), 'The Structure and Function of Complex Networks', *SIAM Review*, **45**: 167–256.

Paumgarten, N. (2009), 'The Death of Kings', *New Yorker*, May 18: 40.

Percus A., G. Istrate, and C. Moore (eds.) (2006), *Computational Complexity and Statistical Physics*, Oxford: Oxford University Press.

Phelps, S., P. McBurney, and S. Parsons (2009), 'Evolutionary Mechanism Design: A Review', *Autonomous Agent and Multi-agent Systems*, 10.1007/S10458-009-9108-7

Posner, R. (2009), *A Failure of Capitalism*, Cambridge, MA: Harvard University Press.

Rickles, D. (2007), 'Econophysics for Philosophers', *Studies in History and Philosophy of Modern Physics*, **38**: 948–978.

Rickles, D. (2009), 'Econophysics and the Complexity of Financial Markets', in J. Collier and C. Hooker (eds.), *Handbook of the Philosophy of Science*, vol. 10, North Holland: Elsevier.

Rosser, J. B. (2008), 'Econophysics and Economic Complexity', *Handbook of Research on Complexity*, Edward Elgar.

Roubini, N. (2007), 'Are We at the Peak of a Minsky Credit Cycle?', www.rgemonitor.com/blog/roubini/208166 (last accessed 15 November 2009).

Rust, J. (1997), 'Dealing with the Complexity of Economic Calculation', paper delivered to Santa Fe Institute workshop on 'Fundamental Limits to Knowledge in Economics'.

Scheffer, M. *et al.* (2009), 'Early Warning Signals for Critical Transitions', *Nature*, **461**: 53–59.

Schwartz, N. (2010), 'Surge of Computer Selling after Apparent Glitch Sends Stocks Plunging', *New York Times*, 6 May.

Seiffertt, J. and D. Wunsch (2008), 'Intelligence in Markets: Asset Pricing, Mechanism Design', *IEEE Computational Intelligence Magazine*, November: 27–38.

Shleifer, A. (2000), *Inefficient Markets*, New York: Oxford University Press.

Story, L. and L. Thomas (2009), 'Tales from Lehman's Crypt', *New York Times*, 13 September.

Tett, G. (2009), *Fool's Gold*, New York: Free Press.

US Securities and Exchange Commission (2010), 'Preliminary Findings Regarding the Market Events of May 6, 2010', www.sec.gov.

Wallison, P. (2009), 'The True Origins of this Financial Crisis', AEI Policy Paper, February, http://www.aei.org/docLib/04-23894%20OTI%20Wallison-g.pdf