`

PREDATOR-PREY: AN ALTERNATIVE MODEL OF STOCK MARKET BUBBLES AND THE BUSINESS CYCLE *(non-technical version)*

Eduard Gracia
(+44) 7879 430 370
egracia@deloitte.co.uk

ABSTRACT

For the last quarter of a century or so, the Real Business Cycle (or "RBC") model has stood as the dominant interpretation of the business cycle in mainstream economics, although throughout this long period a significant number of relevant empirical objections have been raised, both regarding its fit to "real" aggregate data and its consistency of its predictions regarding the expected behaviour of the financial markets throughout the process. All this suggests there is room to explore an alternative model. This paper proposes to base such a model on a predator-prey mechanism, along the lines of the classical Lotka-Volterra model for ecosystem dynamics, in which agency costs play the role of the predatory activity. The result is a system where "producers" (i.e., those in direct, hands-on contact with production) gradually gain control of the productive process when the times are good, and use this control to enhance their own well-being at the expense of the investors (thus increasing agency costs) until this depredation drives the system into a crisis that leads the investors to tighten their bureaucratic controls on the producers, which eventually brings the system back to the growth path. In this context, the introduction of the Efficient Markets Hypothesis in the model of course results in the resulting cycle being fully discounted from the future *expected* path by the rational investors but, as the market rate of return is assumed to be subject to a random walk perturbation, the growth rate that will most approximate the one in a time series of empirical observations (i.e., the one minimising the tracking error) is the *median* path, not the *mean* (whereas the market discounts the future impact of the cycle from the mean, i.e., the "expected" path, not the median) – hence, we should expect the cycle to appear on the observed time series even though market efficiency precludes it from the future expected path. Consistently with this, the resulting model also predicts that observed stock market valuations would present instances of bubbles and crashes more or less synchronised with the Business Cycle, without this implying any irrational behaviour on the part of the investors. Based on a preliminary analysis of the features of this model against well-known stylised facts, this paper suggests that it may be able to perform better against empirical validation than the Real Business Cycle model.

This paper constitutes a non-technical version of another one, available at http://ssrn.com/abstract=549741.

JEL Classification: E32

PREDATOR-PREY: AN ALTERNATIVE MODEL OF STOCK MARKET BUBBLES AND THE BUSINESS CYCLE *(non-technical version)**

Eduard Gracia
egracia@deloitte.co.uk

> *"When in fact a community has overcome many and serious dangers and has reached unquestioned power and lordship, new factors come into play. Prosperity takes its seat in that community and life turns towards luxury. Men become ambitious in their rivalries to achieve magistracies and other distinctions. As this takes place, the aspirations to magistracies, or the protests of those who see themselves rejected, the pride and luxury, will give rise to decadence."*
>
> Polybius (*circa* 140 B.C.) Histories[1]

FINANCIAL BUBBLES & RATIONAL EXPECTATIONS

In one of Stanislav Lem's most memorable tales, cosmonaut Ijon Tichy, his recurrent hero, is sent as a secret agent to a distant planet populated by a community of marooned robots led by a fiendish computer with an irredeemable hate against mankind. Despite being disguised as one of the robots, after a short while Tichy is discovered and captured. The central computer then offers him a bargain: he will stay alive if he switches his allegiance to the machine society, remains in the planet as one of the robots and promises to blow the whistle on any other human spy he manages to discover. Tichy of course accepts the deal but, being a man with an inquisitive mind, he soon starts to wonder how many of the "robots" on the planet would also be men in disguise –only to find out in the end that every single one of them is actually a human being dressing and behaving like the robot he believes everyone else to be.

Written in the late 1950s in communist Poland, Lem's parable was probably as loud a denunciation of the fundamental lie behind the Soviet system as one could possibly have made from that side of the Iron Curtain without ending up in prison. Yet, as it so often happens with art, the tale admits many interpretations applicable to many different contexts. One could, for example, interpret it as a special case of the so-called "beauty contest" mechanism that John Maynard Keynes (1936) put forward to explain the apparent irrationality of financial bubbles:

* This paper is a non-technical version of another one, available at http://ssrn.com/abstract=549741.

> *...professional investment may be likened to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. It is not a case of choosing those which, to the best of one's judgement, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practise the fourth, fifth and higher degrees.*

From a financial markets viewpoint, the implication is of course that everyone could be buying a stock that each one of the investors privately thinks is overvalued, but which nevertheless keeps soaring because the purchase decision is not made on the basis on one's individual assessment of the asset value itself but of everyone else's assessment: in short, this interpretation says that, in a financial "bubble", an asset price would go up because everyone is buying it, and everyone would buy it because it is going up, regardless of its fundamentals. Yet, intuitive as this interpretation may feel, Lem's allegory highlights an implicit assumption in Keynes' metaphor: even in a beauty contest, rational players will only behave irrationally in order to follow the crowd if they expect the "crowd" itself to behave irrationally. The results of the beauty contest would thus not be impacted by the players' attempts to anticipate everyone else's choice if they all regarded themselves as relatively "average" or, in a more financial context, if each one regarded himself as a rational investor surrounded by equally rational investors. In other words, it is not enough that each one of the players try to guess the general opinion by giving more weight to publicly-available information than to his/her own private sources: for the valuation to be biased, either this publicly-available information or the other players' reactions to it need to be biased as well. Thus, in Ijon Tichy's adventure, the core problem is not that everyone finds it safer to behave like everyone else, but that most of the publicly-available data have been purposefully manipulated in order to exclude a key bit of information: that everybody else is also a human being and not a robot.

There is abundant empirical evidence of the existence of financial bubbles from far back as the "Dutch Tulip Mania" in the 1630s to the "dot.com" bubble in the late 1990's. In each one of these instances there were many contemporary voices that first claimed the asset prices to be grossly overvalued and that later, after the crash, intoned a collective "I told you so". Yet this type of experience, however repeated, does not necessarily prove that the bubble was fuelled by irrational expectations on the part of the investors. In fact, numerous empirical studies suggest that this may not be the case. For example, Lo (1991), using a "modified R/S test" specifically devised to distinguish short-term from long-term memory effects, finds no evidence of the existence of long-term memory in US stock market returns –as one would expect to find if those returns were subject to "bubbles" caused by people naively using the latest observed returns to develop their expectations and thus "following the crowd" by demanding an already overvalued stock. This result has been confirmed by subsequent studies based on alternative specifications: see, for example, Donaldson & Kamstra (2000) for testing against Monte Carlo simulations, Engsted & Tanggaard (2004) for application of covariance analysis, or Day

(2004) for a review of the historical evidence about three of the most famous early bubbles (the “Dutch Tulip Mania”, the “Mississipi Company Bubble” and the “South Sea Bubble”) suggesting that, in fact, markets in each one of those instances behaved rationally given the circumstances. While this evidence is not uncontested, at the very least it suggests that financial bubbles cannot be easily interpreted as just another form of collective madness.

There also seems to be a growing consensus in the most recent literature about why rational arbitrageurs would not burst the bubble when it appears, even if they realise it is indeed a bubble. There is a saying among traders: “the market can stay crazy longer than you can stay solvent” –in other words, if credit is limited and there are enough people “acting crazy” in the market, a rational individual will do well to follow the crowd, at least until the time comes when he/she expects the bubble to burst. As we have seen above, this reasoning is consistent with Keynes’ old “beauty contest” argument. In more recent times, a number of papers, following the pioneering work of Blanchard & Watson (1982), have helped to confirm that this is indeed a rational possibility –see for example Abreu & Brunnenmeier (2001) or Allen, Morris & Shin (2003). Indeed, if a trader had bought into the US stock market by the end of 1996, when Alan Greenspan was already denouncing its “irrational exuberance” and Campbell & Schiller had already provided their testimony in the same sense before the Board of Governors of the US Federal Reserve (as described in Campbell & Schiller, 2001), and had stayed until, say, mid-1999, would have quit the game with a handsome profit.

At the core of these models, however, there is always the explicit assumption that some form of irrationality or information asymmetry has caused and keeps fuelling the bubble for, as shown by Tirole (1982), in an efficient market where all players behave rationally, no financial bubbles would ever be possible even if traders were risk-neutral and they held different pieces of information. This brings us to the core question of what causes and then keeps feeding the bubble in the background. In recent years, the hypothesis of behavioural finance, according to which most investors do not behave fully rationally (or, technically speaking, behave according to “bounded rationality” rules), has gained a lot of following –see for example Schiller (2001) or De Grauwe & Grimaldi (2004). Thus, according to this school of thought, there would be nothing wrong in the information the market valuations are based on, but on the way investors process it to set their expectations. Yet the postulate of “bounded rationality” is just a soft way to assume irrational expectations –which contradict the abundant evidence supporting the overall rationality and efficiency of the markets. And, despite the now-abundant empirical evidence of relatively short-term departures from the efficiency markets hypothesis (largely attributable to transaction costs and other forms of friction), from a higher-level, more long-term viewpoint the Efficient Markets Hypothesis (EMH) is remarkably well supported by the facts… So there is room for exploring other potential explanations of the financial bubbles phenomenon that would be compatible with the EMH.

The purpose of this paper is to offer an alternative interpretation, a bit more along the lines of the allegorical situation described in Lem’s allegory. Specifically, the aim of this paper is to put forward a relatively simple model of how more-or-less cyclical financial bubbles could be observed in an efficient stock market where asymmetric information (resulting in agency costs) between the producers or entrepreneurs that control the means of production, on the one hand, and the lenders that finance part of the enterprise assets with their credit, on the other, plays a role similar to the fantastic delusion described in Ijon Tichy's adventure –only, in this case the "lie" is not fostered by a single, central authority but by the independent actions of numerous individuals. We will thus develop a model to explore the possibility that bubbles be caused and fed by the information asymmetries inherent to the principal-agent relationship between creditors and producers –in other words, that they ultimately constitute a form of agency cost.

More specifically, the key components of the model we will put forward here are the following:

1. Information asymmetry among the economic players generates an incentive for non-cooperative strategies. Specifically, we assume that the producers' control on the means of production of the companies allow them to extract rents from them at the expense of the sector of investors that do not enjoy such control (the creditors), who in turn impose limits on this "predatory activity" by subjecting the corporate leadership through creditor controls,[2] resulting in a dynamic process analogue to the classical Lotka-Volterra Predator-Prey model.[3]

2. If the market instant rate of return demanded from an asset is assumed to follow a normally-distributed Wiener process (i.e., the simplest form of random walk), and therefore the asset value growth follows a geometric Brownian motion with drift, then the result we should expect to obtain if we calculated its average growth rate on a past valuation time series is the distribution's median, not mean, path[4] (for an example of an application of this classical theorem in the context of portfolio management theory, see Roll, 1992). Yet, under rational expectations, the expected future path is of course the mean, not the median. Thus, under these conditions it would be possible to find a distribution such that its median path displayed observable cyclical patterns despite their being fully discounted out of the mean (i.e., "expected") path, thus remaining consistent with market efficiency.

Note that, by assuming that the cause of the phenomenon resides in the information asymmetries between producers and creditors in the "real" economy, we are not only saying implicitly that the market valuation at any given point in time is actually "right" on the basis of the information at that time (which is why the expected path displays no cycle), but also that we should expect it to be linked to observable fluctuations in the levels of "real" economic activity –in other words, we are linking it to the business cycle. In this sense, this interpretation of financial bubbles differs significantly from others that also interpreted financial bubbles as an agency cost –most notably Allen & Gorton (1993) where fund managers knowingly take an overvalued asset because they share with the investors they represent in positive profits but not losses, and therefore are more willing to run the risk of riding a bubble that would be in the best interest of the investors they represent. Indeed, by associating financial bubbles to changes in the "real" economy, this interpretation aligns itself to the well-known fact that, in spite of Paul Samuelson's famous quip that "the market has anticipated 9 of the last 5 recessions", financial bubbles are often associated to observed swings in the business cycle… which in turn suggests that we take a closer look at the merits of the model that has dominated business cycle theory for the last quarter of a century, and to which this interpretation would represent an alter: the Real Business Cycle (or "RBC") model.

WHY ANOTHER MODEL OF THE BUSINESS CYCLE?

Interest in the business cycles tends to be cyclical itself. When the economy is booming, there are always voices raised to claim that the business cycle is a ghost of the past, a result of system imperfections that the most recent developments (be it Keynesian-style government intervention, as in 1950's and 60's, or the process-optimizing capabilities of the "New Economy", as in the 90's) will turn as obsolete as Richard Malthus' old demographical theories; conversely, when the downturn comes (and the point is that in the end it *always* comes), the public find it easier to accept that there may be a cyclical pattern of some sorts behind it and that, as Peter Navarro (2004) put it at the Sloan Management Review earlier this year, "while the predictability of the business cycle remains very much a debate among both academics and managers, it seems quite beyond debate that the line between corporate success and failure is often defined by the decisions that are made around key turning points and movements in that cycle."

To be sure, the ever-changing feelings triggered by our ephemeral economic fortunes do not prove anything by themselves. The fact is, nevertheless, that empirical evidence of the existence and relevance of business cycles has steadily accumulated at least since the late 19th Century, and has been confirmed by numerous contemporary studies (see for example Zarnowitz, 1992, or Diebold & Rudebusch, 1999), although, at the same time, the evidence also strongly suggests that the cycle is, in the words of Dore (1993), "recurrent but nonperiodic" –i.e., although a cyclical pattern can be observed *on average* in the time series, this past information cannot be used to predict the timing of future swings with any acceptable degree of reliability. From this viewpoint, the business cycle looks strikingly similar to the stock-exchange mean-reversion patterns whose existence has been highlighted in a number of empirical works at least since those by Poterba & Summers (1988) and Fama & French (1988); furthermore, the stock market role as a lead indicator of the business cycle also suggests that there may well be a causal link, and not just a spurious structural similitude, between this mean-reversion and the cycle.

As mentioned in the preceding section, it is probably fair to state that, for the last quarter of a century or so, the Real Business Cycle (or "RBC") model has stood as the dominant interpretation of the business cycle in mainstream economics (see for example Long & Plosser, 1983, or King, Plosser & Rebelo, 1988, *a* & *b*). Its central idea is that, in a system where the production function presents decreasing returns of scale respective to the stock of capital, the long-term optimal level of capital investment can be modified by an unexpected technology shock, but then, as the move to the new "optimal" level requires a technical process of accumulation of "physical" or "real" means of production (hence its name), it cannot take place instantly. The resulting model is compatible with rational expectations and market efficiency and, in addition, a number of papers have shown, primarily on the basis of calibration/simulation techniques (thus following the steps of Kydland & Prescott, 1982), that the patterns of key variable co-movement the RBC model predicts are generally compatible with observed aggregate data. Yet there are also significant empirical objections to this model. In particular:

1. A number of studies have produced results that do not support the model's basic hypothesis that technology shocks are the primary factor driving the business cycle. The lack of empirical support for this assumption was already pointed out by Summers (1986), and the observation has subsequently been reinforced by a number of more recent empirical papers such as Galí (1996), according to whom "the pattern of economic fluctuations attributed to technology shocks seems to be largely unrelated to major postwar cyclical episodes", or Shea (1998), who concludes that "favorable technology shocks do not raise measured TFP at any horizon. Taken at face value, this suggests that observed procyclical variation in TFP is (…) not at all due to procyclical technology."

2. Similarly, recent papers aimed at verifying the presence of long-term memory in aggregate GDP series along the lines predicted by the RBC model (notably Haubrich & Lo, 2001) have also yielded negative results[5] –in fact, Haubrich & Lo suggest in their paper that the greater power of the statistical tool they apply (the so-called "Modified R/S" test) may explain why their results contradict earlier work that purported to find long-range dependence.

3. Last but not least, a number of authors argue that, even at a stylised fact level, the RBC does not match the observable facts, and in particular that business cycle is closely linked to market disequilibrium phenomena (most conspicuously involuntary unemployment and undesired accumulation of inventories) that the RBC model, to the extent it postulates a frictionless market, simply cannot explain, as the fluctuations it portrays are caused by a gradual adjustment to a desired stock of capital *along an equilibrium path*. Although additional explanatory hypothesis for these disequilibria have been proposed (based, for example, on the assumption of indivisibility of labour, as in Hansen, 1985, or Rogerson, 1988), they are neither supported by hard facts nor really consistent with the "frictionless" spirit of the RBC model itself (for a more extensive exposition of this critical view on the RBC model see for example Dore, 1993). Some have stated this objection rather eloquently: for instance, Michael Mussa, former chief economist at the IMF and now at the Institute for International Economics, described it as "the theory according to which the 1930s should be known not as the Great Depression but the Great Vacation.[6]"

The evidence does not seem to be conclusive either way. On the ones side, the results of studies like Galí (1996) suggest that the cycle may be due to sticky prices, adaptative expectations and other autocorrelation phenomena that could be interpreted as incompatible with full market efficiency (and consequently strengthen the case for a Neo-Keynesian revival). On the other, works like Haubrich & Lo (2001) indicate that such inter-temporal autocorrelation phenomena may not exist, and would thus even lend support to the view that perhaps the business cycle is a statistical mirage –although this interpretation would, in turn, contradict the long-accumulated evidence suggesting that the business cycle, whatever its cause, is a real phenomenon. There is hence a case for exploring alternative models that may present a better fit to the observed facts; in the context of this debate, the purpose of this paper is to introduce an alternative model that may contribute to reconcile these disparate sources of evidence.

There is in economics a long tradition, going back at least to the glorious days of Alfred Marshall, of hiding mathematical developments in footnotes and appendices so that they do not scare away the non-specialist reader. I would like to think I am following so illustrious a precedent when I state that it is not the purpose of this paper to develop an explicit, analytical model but, instead, to explain the underlying reasoning and its logical implications from a more discursive, intuitive, accessible perspective. For those interested in the "hard" analytical proof of some of the statements made below, however, a version of this paper containing their formal development is available at the Social Sciences Research Network, and can thus be found at http://ssrn.com/abstract=549741. Commenting on Marshall's works, John Maynard Keynes wrote that the true economist would do well to study the footnotes and appendices while having no more than a cursory look at the main text; once again striving to humbly follow the Masters' steps, I would similarly like to encourage those readers with an analytical background whose appetite may have been opened by the following pages to have a look at this more formalised version of the paper.

PREDATORS AND PREYS

Originally put forward independently by Alfred Lotka and Vito Volterra in the 1920's, the so-called Lotka-Volterra predator-prey model is not difficult to understand from an intuitively viewpoint. As Volterra first proposed it to explain the observed dynamics of certain fish catches in the Adriatic, it probably makes sense to use a marine ecosystem to illustrate it with an example. Thus, imagine an aquatic community with only two types of fish: sharks and sardines. The sardines eat plankton, whereas the sharks feed on the sardines. For simplicity, we assume the plankton supply to be unlimited; thus, in the absence of sharks, the sardine population would grow exponentially. In the presence of sharks, however, the sardine population will grow at a slower rate the higher the total population of sharks eating them. On the other hand, the population of sharks will grow faster the larger the proportion of sardines over sharks and, in the absence of sardines, will gradually starve away.[7] Under these conditions, there should logically be a value for the shark and sardine populations such that they stay in equilibrium; unless both populations start with precisely the numbers that would be required to stay in equilibrium, however, their interaction will result in a cycle. Indeed, as the sardine population grows so will that of sharks and, as the number of sharks grows, their predatory activities will slow down the growth rate of the number of sardines until bending it down to nil and then to negative. But, at the same time, as the sardine population declines, the sharks themselves begins to starve and reduce their own numbers until, eventually, their numbers become small enough for the sardine population to recover its positive growth rate, thus resuming the cycle. Over time, therefore, the population of sharks and sardines will evolve as in Figure 1:

Figure 1: Predator-Prey Dynamics over Time

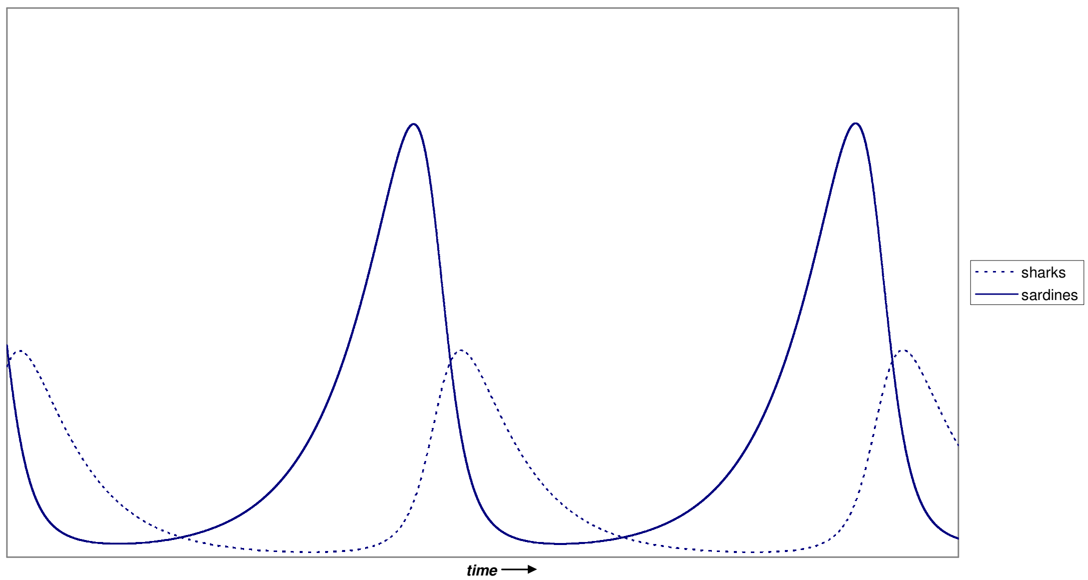

For this particular example I have obviously chosen quite an extreme (and therefore also unrealistic) set of parameters in order to make the features of the model visually clearer. As we see, the predator-prey interaction not only results in a self-perpetuating cyclical motion, but also in a pattern where the prey population experiences relatively gradual growth periods punctuated by periodic, relatively spectacular “crashes,” whilst the predator population mirrors this pattern by having sudden increases coinciding with these crisis, followed by relatively gradual periods of decline. If these “crisis” look like the periodic burst of a sudden epidemic fever (where the “sardines” would now be the host population, be it cattle or human beings, and the “sharks” would be the disease germs), it is by no means a coincidence; indeed, since its inception, the Lotka-Volterra system (as well as its more sophisticated descendants) has proven very successful as a model of biological population patterns where one species predates on another –as it is the case, for example, of epidemic diseases.

One obvious objection could be raised here: as described so far, this is a deterministic model –so why are we not able to predict with perfect accuracy, for example, when the next cattle disease will burst? The simple answer is of course that, in the real world, there is no such thing as “other things being equal,” convenient as this qualification may be for modelling purposes; external, unexpected factors intervene by introducing changes in the population of predators and preys that may therefore radically change the timing of the next system crash.

We human beings are also biological entities, and have an uncanny ability to exploit each other. Save for a very low number of criminal exceptions, of course, human depredation on other humans does not usually take the brutal form of eating each other –yet this by no means makes human predatory activity less relevant as a mechanism of resource redistribution. In the words of the famous historian William McNeill (1982), "disease germs are the most important microparasites humans have to deal with. Our only significant macroparasites are other men who, by specializing in violence, are able to secure a living without themselves producing the food and other commodities they consume." For the purposes of this paper, we will thus define as "depredation" or "exploitation" a relationship between two economic agents such that one of them takes resources away from another other without providing any goods on exchange as part of a mutually-agreed transaction. Basic, highly intuitive examples of this could be the burglar that steals from his neighbour or the security guard that runs away with the money; but other, somewhat more sophisticated examples are probably much more important from an economic viewpoint. Taxes, for example, constitute a predatory activity where the prey are those who pay them (and have to do so whether they like it or not) and the predators are those who benefit from the government's redistribution activities. Similarly, credit defaults also constitute a form of depredation of those who took the money and did not return it against those who lent it and never got it back; and equally predatory is the case of the employee who shirks his responsibilities, or the director that misleads his shareholders. The point is that, although the expected costs of depredation can of course be factored in their expectations by the economic players, and although the payments are often made voluntarily to avoid worse consequences (as in the man who pays taxes to avoid going to prison, or the shopkeeper who bribes the local gangsters to avoid having her shop vandalised), this type of economic transaction cannot be assimilated to a market exchange because, in the case of depredation, the overall result is suboptimal respective to what would have been achieved in an ideal situation where information were perfect and the players were willing to cooperate for an agreed price (technically speaking, is "Pareto-inefficient[8]"). Armies, police forces, disciplinary rules of any kind, even the humble locks that we all put at the front doors of our homes are but a few among many examples of resources that could be devoted to more productive tasks if information were perfect and people were not opportunistic.

The logic underlying this idea is easy to see using what is probably the most famous device in non-cooperative game theory: the Prisoner's Dilemma. The classical exposition of this famous intellectual game is just like the plot of an old cops' movie. Two accomplices in crime go to jail and are interrogated by the police in separate cells. They know that, if no one confesses the crime, they will both go free because the police has no other evidence against them, whereas if they both accuse each other they will both stay in jail for, say, five years. But if one of them accuses the other and is not accused by the other one at the same time, then the accuser will go free whilst his "buddy" goes to prison for, say, ten years. Once isolated, and if the game is only going to be played once, the rational decision for each one is of course to betray the other because, regardless of whether the other criminal has betrayed him back or not, the outcome in either case is equal or better if he betrays –in other words, the rational decision is not to cooperate, even though the outcome of both being rational is to end up worse off than if they had both cooperated. In a world populated by rational agents, the result of this game will be Pareto-inefficient for, even if the prisoners could have agreed an exchange beforehand (in which the price for each one's loyalty would be that of the other), the agreement would become void simply because there would be no way to enforce compliance –in other words, because information is not perfect and, at the time of playing, no one can really know how the other is going to behave. It should now be easy to see that this conclusion can extend to the different forms of depredatory game we described above: the tax game, the default game, the shirking employee game, etc.

Depredation mechanisms along these lines are actually quite common in economic theory, as well as in the wider realm of social sciences. Free-rider models, agency theory and many market failure theories, for instance, all fall within this group. Agency theory, in particular, deals with a number of depredation instances where what enables the predator to exploit the victim is not the threat of raw violence (as it is the case of government taxes) but the possession of privileged information by the "agent" that therefore enables him/her to extract resources from the "principal" (what we call an "information asymmetry"). Thus, for example, the employee that shirks his responsibilities is able to do so because the principal (in this case, the shareholders of the company that pays his salary) has only an imperfect knowledge of his real productivity. To control this form of cheating, the principal can, for example, impose bureaucratic and disciplinary controls, and can also increase the employee's salary above market level so that the cost for the employee of being dismissed because of having caught shirking becomes higher –i.e., in order to provides an additional incentive to stay honest.[9] From the viewpoint of the community as a whole, however, every one of these approaches leads to a form of inefficiency: bureaucratic controls not only represent extra costs, but also tend to slow down the production process as a whole, whilst increasing the salaries of those with a job above the market clearing wage level obviously generates unemployment, which will be larger the higher the difference between the wages actually paid and those that would equal labour demand and supply (see for example Shapiro & Stiglitz, 1984, or Phelps, 1994). Similarly, credit opens the door to a straightforward trick where an entrepreneur could simply set up a limited liability company, negotiate a credit, spend it carelessly and then declare bankruptcy and restart the cycle. Against this possibility, the creditors not only need to charge a premium on the basic interest rate to compensate for this risk, but also establish a system of (inevitably costly) bureaucratic controls to limit the ability of the debt issuers to cheat, or to increase the personal cost for them of doing so; but these bureaucratic controls also represent a cost of non-cooperation, and the risk premium above the basic interest rate means that there will also be legitimate investment opportunities that will not be pursued because the market cost of the funds to finance them is too high, particularly when the economy is in a downturn and thus all credit ratings tend to be lower (see for instance Lowe & Rohling, 1993).

RATIONAL PLAYERS

So, if the concept of predatory behaviour is not really alien to standard economics, one would wonder, why are predator-prey dynamics not being considered as a potential model of the business cycle? The short answer is, of course, that it already has been considered: this is precisely what Richard Goodwin did in his famous paper, first presented at the First World Congress of the Econometric Society in 1965 under the beautifully simple title of "A Growth Cycle," and subsequently published as Goodwin (1967). Today, nearly forty years thereafter, Goodwin's model is still held in very high regard; Dore (1993), for instance, considers that "in the history so far of business cycles theory, the Goodwin model must be seen as a remarkable achievement". The model manages to generate an endogenous, self-sustaining cycle primarily on the basis of assumptions that up to the mid-seventies were generally accepted as part of standard macroeconomics. The result is a Lotka-Volterra cycle like the one we saw for the sharks and sardines example, where it is assumed that all wages are consumed and all profits reinvested, and where wages grow at a faster rate the lower the market level of unemployment (the so-called Phillips curve), so that, in effect, wage growth plays the role of shark population growth, and production growth induced by profit reinvestment that of sardine population growth. From the viewpoint of modern standard economics, however, Goodwin's model is objectionable on three main grounds. First, it relies on the assumption that all wages are consumed and all profits are reinvested, whilst providing no theoretical or empirical reason why this should be so. Second, it also relies on the Phillips curve, a statistical relationship between the rate of unemployment and that of salary growth that became empirically discredited during the stagflation years in the 1970's, when salary growth accelerated in parallel with unemployment rates. And third, it does not postulate rational behaviour, for investors seem to be willing to reinvest the totality of their profits regardless of the expected rate of return.

The overall objection of non-rationality is by far the most serious, as it represents a challenge not only for this specific model but also for any other portraying the cycle as a self-sustaining or otherwise predictable oscillation along the expected path. Indeed, as this argument goes, if investors were rational and they expected rates of return to experience an abnormal raise tomorrow, they would logically invest more until the abundance of supply cancelled this extraordinary profit opportunity by lowering the return –and vice versa, if they expected a fall they would divest (or sell short) until rates fell again. Thus, if people are rational and the capital markets they operate in are efficient (i.e., prices are always able to adjust instantly to clear supply against demand, so that they always end up reflecting the players' expectations at any point in time), they would arbitrage against expected future cyclical swings until ruling them out from the expected path. The strength of this argument relies mainly on the fact that there is very strong empirical evidence supporting the postulate that not only economic players develop their expectations rationally, but also that capital markets are generally efficient,[10] to the point that, already in 1978, Jensen felt entitled to state that "there is no other proposition in economics which has more solid empirical evidence supporting it than the Efficient Markets Hypothesis." Thus, even if non-financial markets are slower to adjust, one would expect the business cycle to be discounted out of the financial market valuations, instead of being reflected under the form of a mean-reversion pattern, as Poterba & Summers (1988) found; and if, as the "behavioural finance" school suggests, investors behaviour is non-rationally biased in a number of critical aspects, one must then wonder why would they nevertheless be rational in other areas, or why have these irrational players not yet been driven to extinction by smarter competition. Following this reasoning, the RBC model, despite its many shortcomings, logically trumps Goodwin's simply because it does not contradict the postulate of people's rationality.

Reasonable as the argument above may look, the fact is that there is a hole in it. Indeed, in a rational agents' world with an efficient financial market, any cyclical patterns or, for that matter, any forms of market disequilibrium should be ruled out from the future expected path; yet from this does not necessarily follow that systematic cyclical patterns or market disequilibria will also be precluded from the observed path, except in those particular cases where the observed path explicitly tracks the expected path (as it is the case, of course, in deterministic models). In particular, it is a well-known phenomenon that, in accumulative stochastic processes (such as the compounded accumulation of returns on an asset reinvested over time) where the accumulation rate follows a random walk, the path that empirical observations should be expected to track (technically, the path that minimises the tracking error respective to the observations) is the median, i.e., the path that leaves 50% of the distribution on either side, not the mean (that is, not the expected path); and these two paths, in asymmetric distributions, can be quite different.[11]

An example may help us to understand this mechanism from an intuitive viewpoint. One such accumulative game is good old double-or-nothing, so widely popularised by television shows. Starting with an initial investment amount (say, $1), the player multiplies it by two on the basis of a given random event (say, answering a question correctly) and loses the whole capital if it falls on tails. If, for simplicity, we assume the random process is such that probabilities at every round are evenly distributed between the outcomes of "double" or "nothing" (say, like tossing a coin), it is thus easy to see that, regardless of whether the game is played once or a thousand times, the expected value at the end will still be equal to the original investment, and thus the accumulated net return will be zero. It is also straightforward to see that, for the first round, a return of zero will also be the median. However, if the game is to be played more than once (say, 10 times), the distribution of the outcomes changes: there is now less than a 0.098% probability of ending up with $1,024 after an uninterrupted sequence of successful games, and over a 99.9% probability of having lost the whole capital at any of those successive ten games. Hence, the mean value after ten games will still be $1, but the median will be zero. Now, if an external observer, unaware of the internal rules of the game, simply tried to analyse empirically its results, what sort of pattern would she unveil? Intuitively, one would expect the results to be closer to this median than to the mean (for a more rigorous, analytical development of this reasoning see http://ssrn.com/abstract=549741, Appendix 1; for an example in the context of portfolio management theory, see Roll, 1992). Yet this result would not imply that the player's expectations of a 0% return (instead of a –100%) were irrational: simply, the expected value was an average on a strongly asymmetrical probability distribution function.

This is critical to cycle theory because, even if prices do indeed adjust automatically to preclude market disequilibria and extraordinary profit opportunities along the expected (i.e., mean) path as soon as they appear, there is no reason to assume they would at the same time cancel them along the median path and thus, to the extent the observed path tracks the median and not the mean, they would appear on the observations. Of course, every time unexpected disequilibria appear, prices will realign to preclude them going forward ***along the expected path*** –but, to the extent the path observed over time is closer to the median, such disequilibria will still consistently appear. In such a system, what would tell us that, despite the presence of cyclical patterns and persistent disequilibria, the economic players are still placing their bets rationally would be that autocorrelation tests would nevertheless tend to reject the hypothesis of long-term dependency, and that attempts to predict these patterns with enough accuracy to arbitrage against them would ultimately be bound to fail. This is, indeed, what has been observed in the real-life business cycle, which, as we have explained at the beginning, can be statistically detected and measured (as in Zarnowitz, 1992, or Diebold & Rudebusch, 1999), but whose swings cannot be traced to an underlying long-term dependency mechanism, neither in the stock market (Lo, 1991) nor in the GDP series (Haubrich & Lo, 2001). Furthermore, to the extent the median path presents a cycle whereas the mean does not, for many (albeit not necessarily all) functional specification, one should also expect empirical data to display cyclical patterns in their sample variance and skewedness (i.e., degree of asymmetry), simply because the mean path is not cyclical, whereas the median (as well as the other key statistical paths) is, so that, when the median path is close to the mean (i.e., in the expansion part of the cycle) the whole distribution is compressed, with the corresponding reduction of on variance and skewedness –and vice versa during the downturn of the cycle. This phenomenon also seems to be consistent with the observations in most of the literature on financial markets variability –see for example Granger & Poon (2003).

So it should be possible to devise a predator-prey system such that a cyclical pattern appears on the median path but not on the mean; in fact, one could conceive not one but many such models. The following section simply describes one of them.

A PREDATOR-PREY MODEL OF THE FINANCIAL CYCLE

The model described below tries to portray how financial market valuations in an efficient market could follow, along their median path, a predator-prey dynamic process where agency costs played the part of the predatory activity. As in the rest of the paper, the description will here be restricted to a discursive form.

Imagine a market with two kinds of individuals: some people that lend funds to firms without directly participating in the production process (whom we will call "creditors" or "debt holders") and others who, in addition to investing their own resources as equity in the company, have direct control of the productive process (and hence will be referred to as "producers" or "entrepreneurs"). Of course, in the absence of any form of control mechanism by the creditors on the producers (be it the laws against fraud, the threat of a tainted credit record or the requirement to keep and periodically audit the company's accounting records, for example), this information asymmetry would open an easy route for the producers to exploit the creditors simply by raising the money, spending it for their personal purposes or transferring it to their own private accounts and then declaring bankruptcy. Against this danger, therefore, the creditors implement a system of punishments and rewards, i.e., a set of rules against the producers' shirking (e.g., a legal system, a set of internal bureaucratic controls, the request that the producers provide collateral guarantees against their personal assets, the threat that a bad credit history may represent on their future funds raising capabilities, etc.) whose purpose is to, if not completely prevent, at least make it more difficult for the producers to appropriate the resources of the company. To the extent these deterrents are not 100% effective, though, depredation will take place under the form of a "rent" (which we will call a "quasi-rent" to distinguish it from rents resulting from the mere market price of a good) composed of earnings received by the entrepreneurs above and beyond the market clearing price of the services they provide to the firm (i.e., of the productive resources they invest in it). Consistently with the efficiency wage theory, these quasi-rents will be higher the greater the degree of control enjoyed by the producers on the assets of the firm (for, in order to act as an effective deterrent of shirking activities, the quasi-rent must be such that the net present value of the future income it generates be equal to the expected profit the producer would obtain from shirking and running the risk of being caught), i.e., there will be a trade off between external controls imposed by the creditors and quasi-rents earned by the entrepreneurs. We will also assume that the establishment and maintenance of these surveillance and control mechanisms against shirking is not free of costs for the creditors and, therefore, they will only impose them up to the point where their marginal costs equal their marginal benefits (i.e.,. the expected reduction of future default costs).

The appropriation of company resources by the producers takes place as follows. For a given production process structure, the entrepreneur has a certain degree of control that translates into a given level of quasi-rents paid as a percentage of the output. At every point in time, a certain number of opportunities to modify this productive structure will randomly appear. Other things being equal, the producer's decisions will of course be biased in favour of those alternatives that generate a higher level of quasi-rents, but their level of discretionality in this respect is limited by the controls imposed by the creditors. As the purpose of these controls is to offset the probability of default, they will be stronger for those companies whose solvency ratio (i.e., the total asset value divided by external debt, which in the case of this model is equal to the part of the invested capital that is not owned as equity by the entrepreneurs) is lower. Thus, other things being equal, the speed at which the quasi-rents increase over time will be faster the higher the firm's solvency ratio, and vice versa. Note, therefore, that this mechanism is playing the role of the growth of the shark population in the predator-prey model.

As we have already discussed, under market efficiency the future expected (or "mean") path of the firm's value will be non-cyclical, for the market is assumed to be able to discount the future evolution of quasi-rent depredation upfront. Yet if we assume that the asset rate of return is stochastic, and its distribution function is such that the observed path is closer to the median than to the mean (as it is the case in the most widely used asset return models used in finance, starting with the most basic of them, the Wiener diffusion process), then there is nothing to prevent more complex patterns from appearing on this observed path. Along the median path (or simply along any path different from the mean) the agents' expectations consistently fail to be met, not because they were not set rationally and on the basis of all the information previously available, but simply because dice rarely, if ever, produce their "expected" value. Hence, along the median path, asset return expectations will consistently be proven wrong and thus need to be continuously revised. The impact of the market inefficiencies caused by the depredation mechanisms generated by the information asymmetries between producers and creditors described above will thus equally fluctuate with the unexpected circumstances and, as the market prices will at every new realignment of expectations adjust instantly so that the future ***mean*** path remains cycle-free, there is nothing to prevent the predator-prey cycle from appearing on the ***median*** path.

Thus, in this model, when the economy grows at a rate lower than expected (and therefore moves away from the equilibrium path), the solvency ratio of companies falls, thus increasing the risk of default for the creditors, which makes it more cost-effective for them to strengthen the control measures they impose on the firms' producers. This curbs the level of discretionality these enjoy on future changes in the productive process, and therefore also the growth rate of the percentage of the output they subtract under the form of quasi-rents, until the positive effect of this on the finances of the firm overturns the crisis and reignites growth. Then, of course, the process repeats as the solvency ratio improves and the incentive for the creditors to keep tight control mechanisms reduces, which also gradually allows the producers to increase their control on the assets on the firm until the weight of their increased quasi-rents starts to bring the solvency ratio down again, thus restarting the cycle.

Note that the perturbation that explains the difference between mean and median, and is thus responsible for fuelling the cycle, could in principle come from any source, including technology shocks as well as many others. For the analytical development of this model, for example, precisely in order to emphasise its independence from technology considerations, the assumptions were purposefully devised so that the only source of perturbation in it would be the variability of the money market interest rate. Using the interest rate as the source of uncertainty also has the advantage that, as it constitutes a price that is visible to all players in the market, it also provides an example of how the market could induce synchronicity on the cycles experienced by different firms and industries. Nevertheless, there is no necessary assumption in this model that this be the only cycle-inducing source of uncertainty in the real world or that the risk premium would need to be constant; more complex and probably more realistic models could no doubt be developed.

STYLISED IMPLICATIONS

So, what happens when we pull all these ideas together? As mentioned above, a formal analytical development can be found at http://ssrn.com/abstract=549741; for our purposes in this paper, however, it should suffice to say that the result will be a Lotka-Volterra cycle where the "prey" are the investors' assets (where the probability that their credits be defaulted by the firm is reflected in the company's solvency ratio) and the "predators" are the producers or entrepreneurs who gradually increase the size of their quasi-rents over the total returns produced by the firm, as represented graphically in Figure 2:

Figure 2: Financial Cycle Model - Median Path over Time

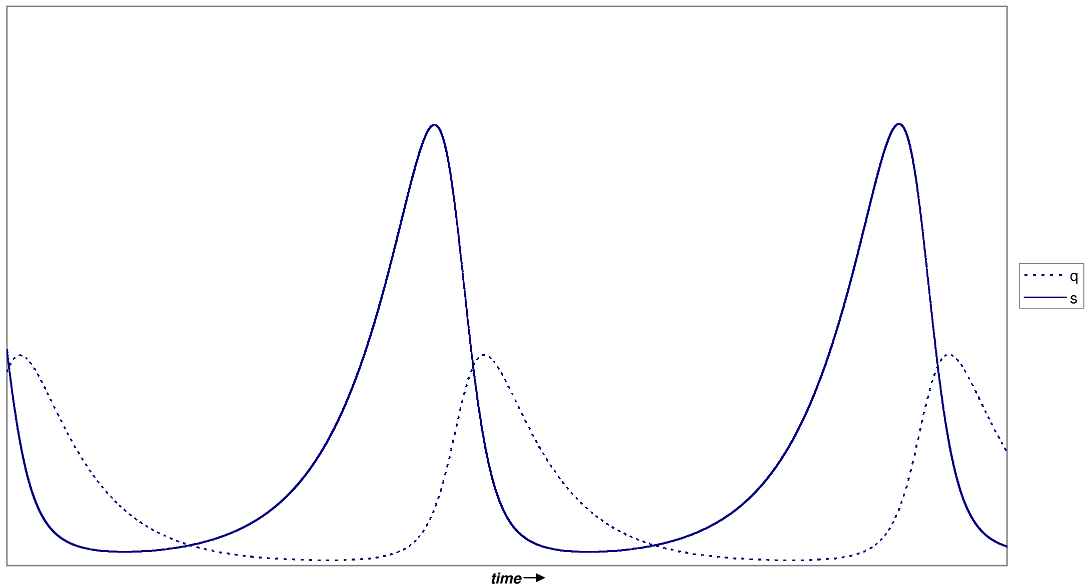

Where '*q'* (i.e., the dotted line) represents the ratio of total producers' quasi-rents divided by investors' returns and '*s'* (the solid line) represents the solvency ratio, i.e., total capital invested divided by total creditor debt, always at market values. To those familiar with the behaviour of financial markets and national economies over the business cycle, the patterns represented in this diagram should look fairly familiar. Indeed, on the one hand, the solvency ratio *s* increases gradually with good times (the "bubble") and then falls down much quicker than it went up (the "crash"). This, incidentally, is consistent with the empirical evidence available (for a recent study see, for example, Koopman, & Lucas, 2003). On the other, the quasi-rent ratio *q* shoots up precisely at the time of the crash, and then only gradually falls down again; although measuring what portion of the income is distributed as quasi-rent as opposed to prices at market-clearing levels is a tricky business, we can use other metrics as a proxy, such as the unemployment rates (because, if the quasi-rents paid to workers[12] above market-clearing level are very high, the market will remain further away from equilibrium, i.e., will present a higher level of unemployment). This, again, is consistent with well-known stylised facts.

Note that this is, to a large extent, a cycle of market expectations. Even if the series is deflated to constant prices of goods and services, the value of the stock of capital (and therefore also of its first differential respective to time, net investment) is impacted by future expectations. It is thus on the basis of these expectations that the solvency ratio is calculated by the creditors, who then use it to decide how much effort should be devoted to contain the producers' ability to gradually change the productive process so that it results in an increase of their quasi-rents. Because of this, the cycle cannot take place on the mean (i.e., expected) path, and thus cannot be forecasted at any point in time, even though the model allows us to predict an average length of wave that could be observed over time. At every given point in time, the current prices are in equilibrium vis-à-vis the most current expectations of future returns (i.e., in respect to the mean path from that instant onwards), but in disequilibrium respective to the median path. What path will be followed afterwards is, of course, unknown; what we know, however, is that an analysis of the series of observations it generates will be closer to the median than to the mean path and, therefore, will be different from that mean, equilibrium path. To the extent the cycle is fuelled by unexpected shocks reflected by the market in its pricing and discount rate structure, many firms that have not suffered the impact directly on their productive processes would experienced it anyway through the impact of the market, and therefore we should also expect the cycles experienced by different companies to synchronise with each other (in other words, we should expect to observe simultaneous market valuation bubbles and crashes of large groups of stocks) even when the relationship between their business activities is in principle small or nonexistent –in other words, even when the co-movement is not justified by the firms' fundamentals.

Furthermore, although one could in principle estimate the future median path more or less the same way one would estimate the mean, this path changes its shape at every next point in time so that it always begins at the same starting position –in other words, one cannot determine with any degree of accuracy how long it is going to take to clean out the "parasitical structures" off the system before growth can be resumed, simply because the most likely path has already been discounted from the current market valuation so that the expected future path remains cycle-free. The cycle in this model thus postulates the existence of an average cycle wavelength to be found in past time series, but does not allow to predict the timing of its future behaviour with any useful degree of accuracy.

Computer simulations performed on this model indicate that, given a constant basic asset rate of return, the frequency of the cycle can be reduced by manipulating the rate of growth of the credit available (in macroeconomic terms, the money supply) and/or the portion of the returns of the companies that is paid as net cash to investors[13] as opposed to being reinvested (in macroeconomic terms, consumption as opposed to investment). Nonetheless, this reduction in frequency comes together with an increase in the amplitude, i.e., in the difference between cycle highs and lows, until, if the credit burden grows faster than the rate of earnings retention (which is of course the difference between the assets rate of return and the part of it that is distributed as cash), the system collapses at the first crisis never to recover again. It is intuitively easy to see why: the difference between the rate of earnings retention and the debt growth rate represents the rate of accumulation of equity the company would experience if there were no quasi-rents; therefore, a negative value implies that it would never be able to restore its solvency after the first shock and thus, in the long run, it would inevitably end up in bankruptcy. Taken strictly at face value, this would also suggest that policies of credit injection and stimulation of demand would not so much work by avoiding the crisis as by postponing them, as they would essentially push the system into a regime where crisis are less frequent but more serious. The reason becomes readily apparent as we inspect the assumptions this model is based on: in essence, it postulates that ***a growing economy gradually accumulates inefficient internal depredation mechanisms, and the crisis is the way the system cleans itself of these parasitical structures*** –thus, by postponing the crisis one such system keeps feeding these parasites and thus implicitly sets the scenario for a more violent adjustment when the crisis eventually arrives. The association of cycle swings with unexpected events such as technology innovations or political events would thus simply be due to the fact that, had these events been expected, an efficient market would already have discounted them away. One could say the unexpected event does not, strictly speaking, “cause” the crisis –it just triggers it.

FINAL CONSIDERATIONS

As we have already pointed out, Goodwin (1967) had already used a predator-prey system to model the business cycle in an analytical development that, with its reliance on the Phillips curve and its assumption of a labour market with a very rigid clearing mechanism, can safely be classified as closely related to the Keynesian family. However, as we have also seen throughout this paper, a predator-prey model can be specified so that it becomes compatible with market efficiency, as long a we recognise that the cycle will be absent from the expected path but not necessarily from the observed one. The inefficiencies can take any or all of the forms identified by the labour market literature: unemployment above frictional level as in Shapiro & Stiglitz (1984), influence-seeking activities as in Milgrom & Roberts (1988) or simply the cost of strikes and trade union activism –at least in its present, highly simplified form, the model does not need to make any specific assumptions about how the inefficiency will manifest itself. This would thus seem to support a Keynesian view of the cycle, while remaining compatible with rational expectation. Yet in the critical area of public intervention, the model, as we have seen, invites to draw rather gloomier conclusions that traditional Keynesian thought: intervention can indeed be used to reduce the frequency of the cycle, but only at the expense of making the crisis deeper once it arrives: one can delay Judgement Day, but only by accumulating interests on the debt that will then need to be paid.

From this viewpoint, the model we have described in this paper can also be interpreted as providing a rational expectations basis to the so-called Austrian School model of the business cycle. Indeed, the model that Friedrich Hayek, Ludwig von Mises and later Joseph Schumpeter put forward in the early 20th Century was based on the postulate that during growth periods companies tended to overinvest and that, when it eventually became evident that the return on these investments was going to be lower than initially expected, an overproduction crisis would ensue –what Schumpeter, in an expression that would equally fit the model we have presented in this paper, called “winds of creative destruction.” The problem with this theory has traditionally been that, unless there is a source of inefficiency operating somewhere in the background, the Austrian model is in principle no more compatible with market efficiency than Goodwin’s: in a rational expectations world, one would expect the investors to rationally forecast the likely return on their investments and then act accordingly. But empirical evidence suggests that overinvestment or “empire building” can also constitute a form of depredation: for instance, in a recent empirical study, Hennessy & Levy (2002) find “strong evidence in favour of empire building incentives, with the effect being strongest when founder status is used as a proxy for empire references”, which is what one would expect if the reason for overinvestment were a form of depredation related to the implicit moral hazard that is always associated to credit (i.e., when the game is such that, if the investment goes well, the upside is for the equity holder whereas, if it goes badly, the result is a default to the creditor); thus, a manager-owner will thus logically feel more tempted to overinvest than a hired CEO with only limited participation in the firm's profit.

One is tantalised at this point to explore other potential applications, particularly in reference to longer-range cycles related to political phenomena; after all, we have already seen that taxes constitute a form of depredation supported by the government's control of the main instruments of violence in a country, and we have also quoted McNeill's (1982) view that "our only significant macroparasites are other men who, by specializing in violence, are able to secure a living without themselves producing the food and other commodities they consume." History is full of examples where a predator-prey mechanism seemed to operate behind the rise and fall of states and societies. For instance, the parasitic dependency of the Central Asian steppe nomad kingdoms on the Chinese Empire, where the former would grow strong as the Chinese economy developed and the emperors could afford paying them tribute, and then would precipitate the crisis by plundering the Empire like starving wolves when the bad times came, is already quite well understood. Similarly, the interpretation of the fall of the Roman Empire as the consequence of the gradually increased weight of a bureaucratic system that ended up weakening that formidable government structure until it was too feeble to resist the pressure of the barbarians is also well established now (see for example Grant, 1990). It is therefore tempting to conjecture that perhaps both the "classical" overproduction cycle and the "Keynesian" market-rigidity crisis model may simply constitute different manifestations of the same type of underlying depredation mechanism –a process belonging to the same family as the crisis of pre-industrial empires, the periodic reappearance of epidemic infections in human and animal populations… or the recurrent incidence of "irrationally exuberant" bubbles in the financial markets.

****

REFERENCES

Abreu, Dilip & Brunnenmeier, Markus Konrad (2001) "Bubbles and Crashes" *Princeton Working Paper.*

Allen, Franklin & Gorton, Gary (1993) "Churning Bubbles" *Review of Economic Studies, 60.*

Allen, Franklin; Morris, Stephen & Shin, Hyun Song (2003) "Beauty Contests, Bubbles and Iterated Expectations in Asset Markets" *Cowles Foundation Discussion Paper*.

Blanchard, Olivier J. & Watson, Mark W. (1982) "Bubbles, Rational Expectations and Financial Markets" *NBER Working Paper*.

Campbell, John Y. & Shiller, Robert J. (2001) "Valuation Ratios and the Long-Run Stock Market Outlook: An Update" *Cowles Foundation Discussion Paper*.

Day, Christian C. (2004) "Chaos in the Markets - Moral, Legal & Economic Signals in Three Fantastic Bubbles" *SSRN Working Paper*.

De Grauwe, Paul & Grimaldi, Marianna (2004) "Bubbles and Crashes in a Behavioural Finance Model" *CESIFO Working Paper No. 1194*.

Diebold, Francis X. & Rudebusch, Glenn D. (1999) Business Cycles: Duration, Dynamics and Forecasting. Princeton University Press.

Donaldson, Glen R. & Kamstra, Mark (2000) "Estimating and Testing Fundamental Stock Prices: Evidence from Simulated Economies" *SFU Working Paper*.

Dore, Mohammed H. I. (1993) The Macrodynamics of Business Cycles. Blackwell.

Engsted, Tom & Tanggaard, Carsten (2004) "Speculative Bubbles in Stock Prices? Tests Based on the Price-Dividend Ratio" *EFA 2004 Maastricht Meetings Paper No. 1804*.

Ellingsen, Tore (1997) "Efficiency Wages and X-Inefficiencies" *Scandinavian Journal of Economics,* Vol. 99.

Fama, Eugene & French, Kenneth (1988) "Permanent and Temporary Components of Stock Returns" *Journal of Political Economy* 96.

Galí, Jordi (1996) "Technology, Employment and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *NBER Working Paper No. w5721*, August.

Goodwin, Richard M. (1967) *"A Growth Cycle"* in Feinstein, C. H. *ed.* Socialism, Capitalism and Economic Growth: Essays Presented to Maurice Dobb. Cambridge University Press.

Granger, Clive & Poon, Ser-Huang (2003) "Forecasting Volatility in Financial Markets: A Review" *Journal of Economic Literature*, June.

Grant, Michael (1990) *The Fall of the Roman Empire* (reprint, 1996) Weidenfeld & Nicholson.

Hansen, Gary D. (1985) "Indivisible Labor and the Business Cycle" *Journal of Monetary Economics* 16.

Haubrich, Joseph G. & Lo, Andrew W. (2001) "The Sources and Nature of Long-Term Memory in Aggregate Output" *Economic Review 57.*

Hennessy, Christopher A. & Levy, Amnon (2002) "A Unified Model of Distorted Investment: Theory and Evidence" *American Finance Association 2003 Meetings*.

Jensen, Michael C. (1978) "Some Anomalous Evidence Regarding Market Efficiency" *Journal of Financial Economics.*

Keynes, John Maynard (1936) *The General Theory of Employment, Interests and Money* Macmillan.

King, Robert G.; Plosser, Charles I. & Rebelo, Sergio (1988) *(a)* "Production, Growth and Business Cycles: I. The Basic Neoclassical Model" *Journal of Monetary Economics*, March.

King, Robert G., Plosser, Charles I. & Rebelo, Sergio (1988) *(b)* "Production, Growth and Business Cycles: II. New Directions" *Journal of Monetary Economics*, March.

Koopman, Siem J. & Lucas, André (2003) "Business and Default Cycles for Credit Risk" *Tinbergen Institute Discussion Paper*.

Kydland, Finn E. & Prescott, Edward C. (1982) "Time-to-Build and Aggregate Fluctuations" *Econometrica* 50.

LeRoy, Stephen F. (1973) "Risk Aversion and the Martingale Property of Stock Prices" *International Economic Review* 14.

Lo, Andrew W. (1991) "Long-Term Memory in Stock Market Prices" *Econometrica*, September.

Lo, Andrew W. & MacKinlay, Craig A. (1999) *A Non-Random Walk Down Wall Street*. Princeton University Press.

Long, John B. & Plosser, Charles I. (1983) "Real Business Cycles" *Journal of Political Economy*, February.

Lowe, Philip & Rohling, Thomas (1993) "Agency Costs, Balance Sheets and the Business Cycle" *Reserve Bank of Australia Economic Research Discussion Paper 9311*.

Lucas, Robert (1978) "Asset Prices in an Exchange Economy" *Econometrica* 46.

Mazzarino, Santo (1959) *The End of the Ancient World* (translated to English by George Holmes, edition 1966) Faber & Faber.

McNeill, William H. (1982) *The Pursuit of Power*, The University of Chicago Press.

Milgrom, Paul & Roberts, John (1988) "An Economic Approach to Influence Activities in Organizations" *American Journal of Sociology* 96.

Navarro, Peter (2004) "Principles of the Master Cyclist" *Sloan Man. Review*, Winter.

Phelps, Edmund S. (1994) *Structural Slumps: The Modern Equilibrium Theory of Unemployment, Interest, and Assets* Harvard University Press.

Poterba, James M. & Summers, Lawrence H. (1988) "Mean Reversion in Stock Returns: Evidence and Implications" *Journal of Financial Economics*.

Rogerson, Richard (1988) "Indivisible Labor, Lotteries and Equilibrium" *Journal of Monetary Economics 21.*

Roll, Richard (1992) "A Mean/Variance Analysis of Tracking Error" *Journal of Portfolio Management* 18.

Schiller, Robert J. (2001) "Bubbles, Human Judgement and Expert Opinion" *Cowles Foundation Discussion Paper.*

Shapiro, Carl & Stiglitz, Joseph E. (1984) "Equilibrium Unemployment as a Worker Discipline Device" *American Economic Review* 74.

Shea, John (1998) "What Do Technology Shocks Do?" *NBER Macroeconomics Annual*.

Summers, Lawrence H. (1986) "Some Skeptical Observations on Real Business Cycle Theory" *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall.

Tirole, Jan (1982) "On the Possibility of Speculation under Rational Expectations" *Econometrica, 50.*

Zarnowitz, Victor (1992) *Business Cycles: Theory, History, Indicators and Forecasting*. The University of Chicago Press.

NOTES

[1] Cited by Santo Mazzarino (1959).

Notes

[2] By representing this "predatory" activity under the form of corporate agency costs subject to a control mechanism through creditor controls we are implicitly following the lead of authors like Lowe & Rohling (1993).

[3] The Lotka-Volterra system of equations was originally proposed in 1925 by Alfred Lotka and Vito Volterra to model the observed behaviour of fish populations in the Adriatic. It is not too difficult to see intuitively how it works. Imagine a fish community with only two species: sharks and sardines. The sardines eat plankton, whereas the sharks feed on the sardines. For simplicity, we assume the plankton supply to be unlimited; thus, in the absence of sharks, the sardine population would grow exponentially. In the presence of sharks, however, the sardine population will grow at a slower rate the higher the population of sharks eating them. On the other hand, the population of sharks will grow faster the larger the proportion of sardines over sharks and, in the absence of sardines, would gradually starve away. Given a set of structural parameters for these equations, there is logically a value for the shark and sardine populations such that they stay in equilibrium. Unless both populations start with precisely the numbers that would be required to stay in equilibrium, however, as the sardine population grows so will that of sharks. As the number of sharks grows, their predatory activities will slow down the growth rate of the number of sardines until bending it down to nil and then to negative. But, at the same time, as the sardine population declines, the sharks themselves begins to starve and reduce their own numbers until, eventually, their numbers become small enough for the sardine population to recover its positive growth rate, thus restarting the cycle.

Analytically, the system takes the following form:

$$\left.\begin{aligned} \frac{dx_t}{x_t} &= \alpha dt - \beta y_t dt \\ \\ \frac{dy_t}{y_t} &= \gamma x_t dt - \delta dt \end{aligned}\right\}$$

Where, in this basic example, x_t would represent the population of sardines, y_t the population of sharks and $\alpha, \beta, \gamma, \delta > 0$ would be positive parameters.

[4] Remember that the distribution mean is the sum of all possible outcomes multiplied by their probability, whereas the median is the value that leaves 50% of the probability at either side.

[5] In their paper, Haubrich & Lo (2001) explicitly present a version of Long & Plosser (1983) as an example of the type of model specification they intend to test, and whose rejection their empirical results suggest.

[6] Quoted by *The Economist* on 26th September 2002.

[7] Analytically, the model takes the following form:

$$\left.\begin{aligned} \frac{dx_t}{x_t} &= \alpha dt - \beta y_t dt \\ \frac{dy_t}{y_t} &= \gamma x_t dt - \delta dt \end{aligned}\right\}$$

Where, in this basic example, x_t would represent the population of sardines, y_t the population of sharks and $\alpha, \beta, \gamma, \delta > 0$ would be positive parameters.

[8] The concept of Pareto-efficiency was introduced in 1906 by Vilfredo Pareto (1848-1923); according to it, we say the resource allocation in a system is "Pareto-efficient" if it is not possible to improve the welfare of at least one of the participants without reducing that of another one.

[9] This is the core argument of the efficiency wage theory. As an introduction, I personally find that Ellingsen (1997) contains a particularly clear analytical development; more general exposition of efficiency wage theory can be found in standard manuals like Milgrom & Roberts (1992), for example.

[10] It should be noted, in this respect, that even the subsequently uncovered evidence highlighting instances of departure from the random walk does not lead to necessarily the logical rejection of the efficient markets hypothesis, except perhaps in the case of transaction costs and other frictions on the short run, for, as Lo & MacKinlay (1999) rightly point out, "the random walk hypothesis need not be satisfied even if prices do fully reflect all available information", as was already proven by LeRoy (1973) and Lucas (1978). The (rather common) view that market efficiency necessarily implies random walk prices and vice versa is thus a dangerous misperception. Indeed, in their comprehensive collection of previous articles, Lo & MacKinlay (1999) illustrate both the departures from the Efficient Markets Hypothesis observed in the short run (as one would anyway expect, for transaction costs and other forms of friction prevent market prices to adjust strictly instantaneously) and the fact that the hypothesis still holds in the long run (cf. particularly Lo, 1991).

[11] For an example of usage of this principle in modern financial theory see, for example, Roll (1992).

[12] The model as defined does not establish a distinction between the resources invested in a company under the form of labour or of any other input. Thus, in using unemployment as a proxy of the quasi-rent ratio we are implicitly assuming that workers have some degree of control on the well-functioning of the business (even if it is only because they can threaten to interrupt it by going on strike), as the efficiency wage theory postulates, and therefore also get their share of quasi-rents.

[13] Once again, remember that the model admits an interpretation where workers are just investors that contribute an input called labour.