Firms, Agency, and Evolution

Armin Schulz

Department of Philosophy

University of Kansas

3010 Wescoe Hall

66045 Lawrence, KS

awschulz@ku.edu

785-864-3976

Abstract

A recent (though controversial) trend in economics has been to appeal to evolutionary theory when addressing various open questions in the subject. I here further investigate one particular such appeal to evolutionary biology: the argument that, since markets select (in a standard biological sense) firms as coherent units, firms should be seen to be genuine economic agents. To assess this argument, I present a model of firm/office selection in a competitive market, and show that there are cases where markets can select for firms/offices as collective units – and thus, as agents of their own – but also that there are cases where they do not. In this way, I try to make the evolutionary argument for the agency-based view of the firm more precise.

Keywords: evolutionary economics, evolution, biology, group agency, methodological individualism

Firms, Agency, and Evolution

There has been a recent trend in connecting evolutionary theory to economics: more and more economists are seeing value in appealing to evolutionary biological considerations in thinking about economic problems (Hodgson & Knudsen, 2010; Vromen, 2009; Witt, 2003; Nelson & Winter, 1982).[1] In this paper, I critically analyse one instance of this trend: namely, a specific evolutionary argument for one position in the debate over the agential status of firms – i.e. the debate over whether firms are genuine economic agents (much like individual people) or whether they are something else (Knight, 1921; Coase, 1937; Nelson & Winter, 1982; Schmalensee & Willig, 1989; Hodgson, 1999, chap. 11; see also Witztum, 2012). The evolutionary argument in question concerns the idea that, since markets select – in the same way that natural selection operates in the natural world – firms as coherent units, the latter ought (often) to be seen as agents of their own (see e.g. Hodgson & Knudsen, 2010).

However, as I try to make clearer in what follows, a problem of the existing work on this topic is that it appeals to the wrong kind of evolutionary model. Further, I try to show that correcting this has two key implications: firstly, it makes the above evolutionary argument for the agential nature of firms less unequivocally convincing, but secondly, it makes more precise when there are evolutionary reasons for seeing firms as agents of their own. The latter kind of insight is especially important, as it can aid in the empirical investigation of the agential nature of firms.

The paper is structured as follows. In section I, I present and evaluate a key evolutionary argument for the agential nature of firms. In section II, I present and discuss a new model of firm

[1] Of course, the general idea here is old – see e.g. Veblen (1898) and Schumpeter (1942).

selection. I present the key implications of this model for the evolutionary argument for the agential nature of firms in section III. I conclude in section IV.

I. An Evolutionary Argument in the Debate over the Agential Nature of Firms

When it comes to theories about the nature of the firm – i.e. about what kinds of things firms are – three broad approaches can be distinguished (Hodgson, 1999, 247-249; Fama, 1980; Nelson & Winter, 1982; Foss & Klein, 2008; Hodgson & Knudsen, 2010; for an interesting legal view of this distinction, see Gindis, 2009): the 'contract-based' approach, the 'agent-based' approach, and the 'routine-based' approach.[2] These three approaches are not mutually inconsistent – the natures of different firms in an economy might differ, and even the same firm may have three 'natures' (i.e. it may, in different contexts, be well described using each of the three of the approaches) – but they represent three very different views of what firms are.

According to the contract-based approach, firms are a special kind of 'transactional space' (Coase 1937; Williamson, 1971; Klein et al., 1978; Kreps, 1990; Hart, 2008; Aghion & Holden, 2011). In particular, on this kind of account, it is posited that there are so called 'transaction costs' that come from taking part in a market; these costs include the costs that come from acquiring the relevant information (where which goods and services are bought and sold, at what price, etc.), the costs that come from drawing up the explicit contracts about what is to be bought and sold, and the costs of monitoring and enforcing these contracts. The (posited) existence of these costs is then taken for a reason for removing some economic transactions from the market: firms can be assumed to grow up to the point at which, for the last transaction in question, the costs that come from taking part in the market are matched by the costs that come from

[2] Note that there are also many other questions surrounding the nature and workings of firms (Coase, 1937; Alchian, 1950; Enke, 1951; Nelson & Winter, 1982; Schmalensee & Willig, 1989; Satz & Ferejohn, 1994; Nickerson & Zenger, 2004; Radner, 2006). These, though, are not so important here.

producing the relevant good or service in house (which include higher input costs, the costs of the acquisition of new skills, etc.). What matters most for present purposes about this approach is that, according to it, firms are a certain kind of economic environment within which economic agents act – they make it possible for economic agents to transact more cheaply than what would be possible on the market. The firms themselves, though, are *not* economic agents – they are merely transactional spaces. As Williamson (1979, 239) puts it: firms are just an 'institutional matrix within which transactions are negotiated and executed' (see also Kreps, 1990, 724).

By contrast, on the agent-based approach to the nature of firms, firms are seen as goal-directed collectives (Knight, 1921; Penrose, 1959; Clark, 1997, chap. 9; Hodgson, 1999, chap. 11). In particular, here, firms are seen as collections of economic agents that have accumulated, on the one hand, a certain set of 'quasi-beliefs' that give them the ability to produce and market some particular good or service, and on the other, a set of 'quasi-desires' that give them the motivation to see through this production and marketing (Knight, 1921; Hodgson, 1999, 251-258; Simon, 1957; Pettit, 2003). Note that these firm-'beliefs' and firm-'desires' can differ from those of any of their employees (they are 'emergent properties' of sorts): firms are thought to frequently 'know more' than (or 'believe' differently from) what their employees know (or believe), and might 'want' to do things that none of their employees want to do. As Winter (1988, 170) puts it: 'it is firms, not the people that work for the firms, that know how to make gasoline, automobiles, and computers.' Note also that, while questions can be raised about the extent to which firm 'beliefs' and 'desires' are similar to the beliefs and desires that feature in intentional explanations of human actions, for present purposes, a detailed treatment of this point is not necessary. All that matters here is that, on the agency-based view of the firm, firms are seen as being goal-directed in a way that marks them as agents of their own. This last point is

key for this essay: according to this approach, firms are coherent economic individuals in their own right (List & Pettit, 2006).

Finally, according to the routine-based approach, firms are seen as the embodiment of a set of behavioural routines that determine how a given good is produced (Nelson & Winter, 1982). These routines may include quite concrete behaviours (e.g. the order in which a product is assembled) and more abstract ones (e.g. how the sales price of the firm's product is set). A firm is seen as successful to the extent that it gets to retain or replicate the behavioural routines that make it up. In this way, the routine-based approach encompasses elements of both the contract-based and the agent-based perspective: on the one hand, firms here are seen as having a genuine (routine-based) identity that goes beyond being an 'institutional matrix' within which economic agents can transact, but on the other, firms here are not seen as decision-making, goal-directed systems, but merely as bundles of behavioural routines that operate automatically.[3] For present purposes, though, it is best to see the routine-based approach as a kind of non-agential view. The reason for this comes out most clearly when it is noted what happens when there is conflict as to which behavioural routines a firm ought to adopt. In a case like this, the routine-based view will only track the routines themselves: the firm then becomes again merely a kind of environment for the interaction of the economically meaningful entities – it is just that, in this case, the latter are behavioural routines (Nelson & Winter, 1982, 107-112, see also the model below). As Nelson & Winter (1982, 134) put it: '[I]t is quite inappropriate to conceive of firm behavior in terms of deliberate choice from a broad menu of alternatives that some external observers consider to be "available" opportunities for the organization. The menu is not broad, but narrow

[3] Note also that Nelson & Winter (1982, chap. 4 and 124-128) argue that human agents can be subsumed under a similar model as well: they are embodiments of skills. If that is so, then this view would have the implication that there are no genuine agents at all that it is worth tracking in an economy.

and idiosyncratic; it is built into the firm's routines, and most of the "choosing" is also accomplished automatically by those routines'.

In contrasting these three views about the nature of the firm, it is also interesting to note that Nelson & Winter (1982, 134-136) analogise behavioural routines to genes, which suggests that their theory of the firm is akin to a 'gene's eye' view of biological evolution. Indeed, one could analogise the three different views of the nature of the firm to three different views about the nature of biological entities like colonies of social ants: one could see these colonies as merely assemblies of convenience of its members (the contract-based view), as 'superorganisms' (the agent-based view), or as more or less homogenous gene-pools (the routine-based view). I will return to these three different views of biological systems below.

Summing up: one can distinguish two broad opposing viewpoints about the nature of the firm: fundamentally non-agential views (comprising the contract-based and the routine-based approaches) and fundamentally agential views (comprising the agent-based approach).[4] While there is much more that could be said about each of these views, for present purposes, it is sufficient to note that, in general, economists are divided over which of these approaches to the nature of the firm to favour (Knight, 1921; Williamson, 1971; Nelson & Winter, 1982; Witztum, 2012).

Given this, it is tempting to appeal to considerations from outside of (what is generally taken to be) the sphere of economics to help push this debate forwards. One prominent such idea is to consider the situation from the point of evolutionary biology: after all, humans are biological organisms, so it is reasonable to suppose that at least some of what we know about biological

[4] Note that this question is related to the plausibility of 'methodological individualism': should we make room in the social sciences for genuine higher level (group) agents, or should we only permit agency at the level of the individual person (Watkins, 1952)? However, since even the formulation of the latter thesis is controversial (Hodgson, 2007), I prefer to phrase the issues here without reference to the latter term.

entities in general also applies to them (Hodgson & Knudsen, 2010; Hodgson, 1999, Witt, 1999; see also Vromen, 2001). Specifically, it is here interesting to ask whether it is possible to use evolutionary theory, as it is standardly understood in biology (Futuyma, 2013; Godfrey-Smith, 2009; Okasha, 2006), to defend the agent-based view of the firm. One way in which such a defence might go is by noting that there may be reason to see markets as selecting firms as coherent units (in the same way that natural selection selects organisms in the biological world). If this is right, then this may be taken to underwrite the idea that firms are economic agents of their own: for what better reason is there to think that firms are independent economic agents than the fact that markets treat them as such?[5]

Before looking at this argument in more detail, though, it is important to ward off a possible confusion. There are many ways to introduce evolutionary biology into economics, and not all of them require seeing firms as economic agents of their own. Indeed, in some contexts, a different view of the firm may make it easier to provide evolutionary economic models of the particular issues that are stake in those contexts, and in some other contexts, the agential nature of the firm might not matter much, so that any view of the firm may be appropriate to use there (Nelson & Winter, 1982, 51-53). Indeed, some of the most well known evolutionary economic investigations concerning economic growth, economic dynamics, and economic innovation

[5] For example, Hodgson and Knudsen (2010, 170-171, footnote 9) write: 'Organizations as here defined have the capacity for goal-directed behavior, irrespective of whether goals are actually declared. In this sense, an organization has the capacity to be a "collective actor" (Knight 1992, 3).' They then note that '[t]he importance of […] firm-specific capabilities and learning effects mean that the firm often has the necessary cohesion to qualify as an interactor' (Hodgson & Knudsen, 2010, 173) – where an interactor is (following Hull, 1988) defined as an entity that 'interacts as a cohesive whole with its environment in such a way that this interaction causes replication to be differential' (Hodgson & Knudsen, 2010, 165, see also 167). Turning this around, this seems to imply that if firms qualify as interactors, then that is (probably) due to the fact that they have the kind of cohesion that gives them the capacity to be a collective actor. However, the textual situation here is not entirely clear, in that Hodgson & Knudsen (2010) also seem to want to embrace Nelson and Winter's routine-based picture of the firm (see e.g. Hodgson & Knudsen, 2010, 173-179). Fortunately, assessing this is not so important here. The goal of the present paper is to discuss an interesting evolutionary biological argument for the agency-based view of the firm; whether this argument is also the one suggested by Hodgson & Knudsen (2010) can be left open here.

adopt the routine-based view of the firm (Nelson & Winter, 1982; Day & Chen, 1993; Hodgson and Knudsen, 2010; see also Boyd & Richerson, 2005).

However, this kind of pluralism about appropriate views of the firm in evolutionary economic models is entirely consistent with the claim that evolutionary theory provides a specific kind of reason to think that firms are often economic agents of their own. Just because the agent-based view of the firm need not always be adopted, there may still be a (specific) reason to do so – though there may also sometimes be other reasons that speak against it. For present purposes, all that matters is whether and when it is indeed true that standard evolutionary theoretic considerations provide *a reason* for seeing firms as agents of their own.[6]

With this in mind, the evolutionary economic argument that is relevant here can be reconstructed as follows.

1. In an evolutionary process, there are many potential units of selection – targets of natural selection.

2. In some contexts, being a unit of selection is sufficient for being a genuine 'evolutionary individual'.

3. Competitive markets work just like natural selection.

4. In competitive markets, firms are units of selection in the way set out in 2.

5. Hence, in competitive markets, firms are genuine evolutionary individuals (from 1-4).

[6] At times, some evolutionary economists also suggest that their aim is more heuristic in outlook (Hodgson & Knudsen, 2010; Witt, 1999; Vromen, 2001). In particular, they note that they are merely interested in providing an evolutionarily grounded 'how possibly' story about how it is *possible* for firms to be agents of their own (for more on 'how possibly' stories, see Brandon, 1990). However, since it was already known that it was possible (in some sense at least) for firms to be economic agents of their own (Richardson, 1972; Bacherach, 1999; Foss & Klein, 2008; see also Hodgson, 1999, chap. 11), this does not appear to be a particularly interesting argument to make (Schulz, 2013). Hence, I will not consider this further.

6. Hence, in competitive markets, firms are plausibly seen as genuine economic individuals (from 5).

Consider the premises and conclusions of this argument in turn.

The first premiss of this argument notes that units of selection – i.e. the entities directly targeted by natural selection – can, in principle, be found in many places of the (biological) hierarchy. This claim is the core assumption of 'multi level selection theory': the idea that natural selection can act on many different levels and target groups as well as organisms or genes (Sober & Wilson, 1998; Okasha, 2006). Now, it is important to realise that this theory, while not universally accepted (West et al., 2011), has become quite firmly entrenched in the evolutionary biological community. However, it is also important to realise that it is generally thought that there are in fact two very different kinds of multi level selection: MLS1 and MLS2 (Damuth & Heisler, 1988; Brandon, 1990; Sober & Wilson, 1998; Okasha, 2006; Godfrey-Smith, 2009).

MLS1 concerns cases where the evolution of particular evolutionary entities – such as individual organisms, individual genes, or traits – depends both on how the fitness of the entity relates to that of other entities *within* the group it is a part of and on the fitness relationships *among* such groups in the overall population (Okasha, 2006; Sober & Wilson, 1998). By contrast, MLS2 concerns cases where natural selection operates straightforwardly at a relatively high level of the biological hierarchy: different groups of some evolutionary entity differ in some *group level traits*, which moreover confer fitness benefits *to the group*, and which are heritable *at the group level* (Sterelny & Griffiths, 1999; Okasha, 2006). It is important to realise how cases of MLS1 differ from cases of MLS2: in the latter, the fates of the members of the higher level entities are, in and of themselves, irrelevant and not further considered in the analysis (unless

they affect relevant features of the higher level entities), whereas in the former, the evolution of these members is precisely what is at issue (Okasha, 2006; Sober & Wilson, 1998; Brandon, 1990).[7]

This distinction between MLS1 and MLS2 is crucial for the second premiss of the above argument: the claim that, in some contexts, being a unit of selection is sufficient for being a *genuine evolutionary individual*. Here, 'genuine evolutionary individuals' are the kinds of things that are the basic subjects of an evolutionary process: it is *they* that differ with respect to some heritable traits that have consequences for *their* survival and reproduction. The reason for the qualification 'in some contexts' lies precisely in the distinction between MLS1 and MLS2 – in particular, the key point to note here is that MLS1-type scenarios are *not* generally seen to constitute cases where a group (or any other higher-level entity) is an evolutionary individual of its own, despite the fact that they might constitute cases where groups are units of selection. The reason for this is that, with MLS1, the focus of the analysis always remains on the individual *members* of the group: as noted earlier, these are cases where there is a group component to *individual* fitness. As such, there is no reason to always see the relevant groups as being evolutionary entities of their own; for the latter to be true, the groups need to evolve in an MLS2 sense (Maynard Smith, 1987; Okasha, 2006; Godfrey-Smith, 2009).[8] Of course, marking this distinction then raises the question of what it takes for something – such as a group of organisms – to start evolving in an MLS2 sense. This is a difficult question to answer (Buss, 1987; Maynard

[7] It is important to note that MLS2 does not require that the fitness of a group is completely unconnected to any of the features of its parts. The point is just that group fitness is measured in terms of the expected reproductive success of groups – not of the individuals making up the groups. See also Okasha (2006).

[8] There is an alternative route to the same conclusion. In particular, one can follow Godfrey-Smith (2009, 39-40), and, *in all cases*, define evolutionary biological individuals as units of selection. This leads to the same conclusion as the one in the text, since, on this picture, cases of MLS1 do not (in general) constitute cases where groups really are units of selection – e.g. they are often to be seen merely as instances of frequency- or (social) environment-dependent selection of individual organisms. Either way, the point is that one needs to distinguish group-level influences on evolution in the MLS1 sense from those in the MLS2 sense. See also Okasha (2006).

Smith & Szethmary, 1995; Michod, 1999; Okasha, 2006, chap. 8; Godfrey-Smith, 2009, 121-128). Fortunately, for present purposes, it can be left open: all that matters here is that showing that some collective is a genuine evolutionary individual requires showing that it evolves in an MLS2 sense, not an MLS1 one.

The third premiss of the evolutionary economic argument concerns the claim that (competitive) markets are selective processes just like natural selection (Hodgson & Knudsen, 2010, chap. 7). In particular, the idea here is that, firstly, market participants can be assumed to differ in various features – such as the extent to which they are profitable, easily replicable, or maximise utilities. Secondly, these features can (at least sometimes) be assumed to have consequences for the probability with which these participants persist and reproduce in the market: participants lacking these features are more likely to be forced to leave the market, those having them are more likely to stay and create other entities like them (Boyd & Richerson, 2005; Alexander, 2007). Finally, these features can (at least sometimes) be assumed to be heritable. In this way, the conditions for natural selection to occur – i.e. the presence of a population of entities that differ in heritable traits with fitness consequences (Godfrey-Smith, 2009; Okasha, 2006) – can be said to be satisfied by competitive markets.

Now, there is no doubt that this premiss has to be seen to be controversial (Godfrey-Smith, 2009, chap. 8; Frank, 2003; Witt, 1999; Rosenberg, 1994; Ghiselin, 1987). In particular (to concentrate on the entities at centre stage in the present context), while it is true that, in at least somewhat competitive and open markets, more profitable firms really do seem to be more likely to persist and give rise to relatively similar daughter firms, not all markets are even reasonably competitive, and the extent to which unprofitable firms are being 'weeded out' by the market can differ greatly across markets (Alchian, 1950; Enke, 1951; Satz & Ferejohn, 1994; Demsetz,

1997; Baum & McKelvey, 1999; Witt, 2003; Frank, 2003). On top of this, the extent to which daughter firms share features with their parental firms need not always be high (Klepper, 2009; Frank, 2003; Witt, 2003; Hirshleifer, 1977). While responses to this charge have been formulated (Nelson & Winter, 1982; Boyd & Richerson, 2005; Witt, 2003; Hodgson & Knudsen, 2010), it seems clear that this premiss could do with some further theoretical or empirical support – e.g. in the form of further models that suggest exactly when we should expect markets to function as selective agents.

The fourth assumption of the above evolutionary economic argument concerns the claim that the way markets select firms is of an 'appropriate' kind. Here, the distinction between MLS1 and MLS2 again becomes relevant: the claim is not just that markets select firms in *some* sense – but that they do so in the MLS2 sense. In other words, this premiss has it that firms are treated by the evolutionary pressures inherent in a competitive market as a coherent entity that survives and reproduces as a whole, not merely as something that affects how well other (individual) economic agents are doing.

Once again, the truth of this premiss is by no means obvious. While it is quite plausible that, to the extent that it is accepted that markets are selective processes at all, they do select *firms*, whether and when they select them *in an MLS2 sense* has not yet been conclusively shown. In particular, the discussion of this point by two key defenders of the above evolutionary argument – Hodgson & Knudsen (2010) – does not help here, as it focuses on MLS1-type models. (The same goes for the discussion by Nelson & Winter, 1982, as they use the routine-based conception of the firm and only consider cases of selection at the level of the behavioural routine.) Hence, there is still more work that needs to be done in order to justify premiss 4.

Given premises 1-4, the evolutionary economic argument tries to derive the conclusion that firms should be seen as evolutionary individuals of their own. However, this inference, too, needs to be handled with a lot of care. In particular, firms can only be considered evolutionary individuals *in an interesting sense* if they can in fact *persist* in the market. If the conditions in the market are thus that firms go bankrupt almost immediately after they have been founded, they cannot be seen to be evolutionary individuals in a theoretically meaningful sense: for something to be an evolutionary individual that it is worth treating as such, it actually needs to be able to *evolve*. If the conditions are too inhospitable to make this possible, talking of evolutionary individuals is not useful or informative. Of course, purely formally, the same causal factors are acting here as in cases of genuine evolution. Still, the fact that the relevant individuals do not persist at all is reason to see them as failing to be theoretically meaningful evolutionary individuals (Godfrey-Smith, 2009).

Finally, what about the inference to conclusion 6? This inference is key, for, as yet, we have at most established that firms are genuine *evolutionary* individuals: they survive and reproduce like other evolutionary individuals do. While this is interesting in and of itself, with a view to the debate over the agential nature of firms, the key issue is whether and when firms can be genuine *economic* individuals: goal-directed systems that make autonomous (at least in some sense) decisions. So what reason is there to think that firms, *assuming that they are genuine evolutionary individuals*, are also genuine economic individuals?

One answer to this question centres on the fact that, like many scientific disputes, the issues here are contrastive (Sober, 2008; Royall, 1997). In particular, we are not considering the agent-based theory of the firm in a theoretical vacuum, but are trying to compare it to the non-agential theories. This matters, in that it is more likely that firms are genuine evolutionary individuals if

the former theory is true than if one of the latter theories is true. If firms are seen as mere transactional spaces or environments for behavioural routines to replicate in, there is less reason to think that they have the necessary internal 'connectedness' to be able to be maintained as genuine evolutionary individuals than if they are seen as agents of their own.[9] Note that this is not to say that the fact (assuming it is one) that firms are selected in an MLS2 sense *entails* an agency-based view of the view (i.e. that it would be inconsistent to see firms as anything other than agents of their own). The point is just that the fact (to the extent that it is a fact) that firms are evolutionary individuals *fits better* to the hypothesis that firms are genuine economic individuals than to the hypothesis that they are mere transactional spaces or environments for behavioural routines to replicate in. In turn, this makes clear that there is indeed *a reason* – though not necessarily an overwhelmingly strong one – to infer 6 from 5 in the above argument: if, in a given context, it turns out that it is reasonable to see firms as genuine evolutionary individuals, then we should be more inclined towards seeing them as genuine economic individuals of their own than as something else. For this reason, the inference to conclusion 6, if read in a modest (evidential) way, can be seen to be tenable.

All in all, therefore: while not completely implausible, the evolutionary economic argument set out above requires further support. In particular, the truth of premises 3 and 4, as well as the inference from premises 1-4 to conclusion 5, deserve to be further investigated.

II. Firms, Employees, and Multi Level Market Selection: A Model

In this section, I outline a model of differential firm survival and reproduction in a competitive market. Before considering this model in more detail, though, it is worthwhile to pause and ask

[9] Note that much the same worry has been raised about the possibility for ecosystems to evolve (Sterelny & Griffiths, 1999, chap. 8; Sterelny, 2001; Godfrey-Smith, 2009).

what, in general, the goal in presenting such a model could be. After all, we are interested in what firms are like in the real world – and not in some toy model world.

In response to this question, it can be noted that what such a model can do this is to make more precise exactly *when* we should expect firms to have the features of genuine agents, and when not. As it turns out, this is far from trivial: even in the best case – where it is assumed that firms have all that is inherently necessary to be genuine evolutionary individuals – external conditions may be too inhospitable for firms to be stable collective entities. In other words, in line with the moderate evidential account defended earlier, the model here is not meant to provide the final answer to the question of whether firms are to be seen as genuine agents; its real value lies in making more precise *the conditions under which* we should expect firms to be agents of their own – it provides a framework within which the agency of real firms can be situated.[10]

1. The Structure of the Model

The model's structure is seen most easily if it is broken down into three stages. Firstly, there is the initial, setting up stage. Here, it is assumed that there is a competitive and open market for a certain good, that there is sufficient demand for the good to support a large number of firms, that there are a number of firms that are looking to access that market, and that there are a number of potential employees looking for work in this industry. Formally: there is an x by y grid representing the market, with each spot on the grid being able to hold one office; initially, $m \varepsilon N$

[10] In this way, the model sits between a purely instrumentalist account of economic modeling and a fully realist one: the claim is not that the model 'proves' that firms sometimes are economic agents, but nor is the idea just to show that it is sometimes possible to model firms as economic agents. The claim is that the fact that it is sometimes possible to model firms as economic agents is *evidence* for them sometimes being economic agents – and that we can say in more detail when, exactly, this is the case.

offices are created, each of which is assumed to belong to a different firm type, and randomly distributed over the grid (for more on the difference between offices and firms, see below).

Each of these firms / offices then hires $n \in \boldsymbol{N}$ employees. In turn, potential employees are of two types: they can be hard working employees that go beyond what is minimally required for the job – the 'expectation-exceeders' in what follows – or they can just provide this minimum level of effort – the 'clock-punchers' in what follows. Firms initially hire expectation-exceeders and clock-punchers in proportion to their overall number in the applicant population (p_E and $1-p_E$, respectively) – e.g. because they want to set up shop quickly and thus do a somewhat superficial search only. (I return below to the reasons for making this assumption about the hiring process.)

The second stage of the model concerns what happens within the firms, once they are set up. Here, the key thing is that employees can change their type as time progresses: they *learn* during their tenure in a given firm's office. More specifically, there is a chance that expectation-exceeders turn into clock-punchers.[11] This chance depends on two factors: the number of clock-punchers in their vicinity, and the length of time during which employees interact before the next phase of the model begins. Formally, for every expectation-exceeder, there is probability $p_{td}=IP+E_1d/n$ per learning round that she will turn into a clock-puncher, where $E_1 \in \boldsymbol{R}$ and in the interval [0, 1], IP=0.01 as a fixed parameter, and d is the number of clock-punchers in the relevant office. It is furthermore assumed that there are r learning rounds.

The reason for this way of modelling the situation is that, on the one hand, it can be assumed that being a clock-puncher is the utility maximising strategy for an employee. In the first place, it is plausible that being an expectation-exceeder takes time and effort on the part of the employee

[11] Alternatively, in terms of the routine-based approach, there is a chance for the behavioural routines supported by the expectation-exceeders to fail to be maintained *within* a given firm.

– i.e. it comes with utility costs. Furthermore, it is not clear that there is any obvious upside to being an expectation-exceeder: it can be assumed that firms do not treat expectation-exceeders and clock-punchers differently, as distinguishing these two types may be very difficult for them (for example, both types of employees might work the same number of hours per week – just with different levels of 'focus'). Also, it is assumed that there are ample equivalent employment opportunities, so that an individual employee does not care about losing her job.

On the other hand, it is also reasonable to assume that people often will not choose the utility maximising strategy from the get-go: they are only boundedly rational, and need to figure out what this strategy is (Radner, 2000; Gigerenzer & Selten, 2001; Witt, 2003). Therefore, it is plausible that the probability of becoming a clock-puncher depends on the number of clock-punchers in the firm (i.e. d/n) and the number of interactions among employees (i.e. r). In particular, the more clock-punchers there are – and the longer a given expectation-exceeder gets to interact with them – the more likely she is to pick up on the minimal effort level required for the job (Boyd & Richerson, 2005; Sterelny, 2012).

The third and final stage of the model concerns the interactions among firms on the market. These interactions are determined by the *profitability* of a firm. In turn, a firm will be more profitable the more expectation-exceeders there are among its employees: the employees will work more efficiently, creatively, and cooperatively, and thus create a better product. Specifically, firm profitability influences inter-firm relations in two ways.

The first concerns office survival. Unprofitable offices are less likely to be able to pay their running costs, cover tax bills, or unexpected expenditures of one kind or another, and thus are more likely to go bankrupt. (As will be made clearer below, less profitable offices are also more

likely to be bought out by more profitable offices from other firms.) Formally: for every office, there is a probability $p_b = E_2 d/n$ that the office goes bankrupt, with $E_2 \varepsilon \boldsymbol{R}$ and in the interval [0, 1].

The second effect of profitability on the interactions between firms concerns the firm's chances of opening up new offices (assuming they have not gone bankrupt). More profitable firms have the cash flow to lay out the down payment for a new lease of new premises, start a new hiring process, etc..

To understand both of these effects better, one needs to understand a key distinction made in the model: that among firms and offices. Firms here are a *type* of organisation, the instances of which are its *offices*. In other words, a 'firm' here is merely a higher-order grouping of a set of offices: the actual businesses 'on the ground' are all different offices (some of which might belong to the same firm). In the model, it is furthermore assumed that offices have a lot of independence. In particular, every office gets to make its own decisions about opening up a new office of the same firm type, and every office is on its own as far as its continued existence in the market is concerned. The idea, then, is that the managers of the firm as a whole take a very hands-off approach to managing the offices that make it up: the only instructions that they impart to the managers of a new office is that they (the new managers) ensure the office stays in business for as long as possible, and that they open up a new office when that is feasible.

It should be clear that this way of managing offices is not the most common way in which firms operate. Indeed, firms typically seem to make decisions in a more centralised manner. However, this does not mean that the above assumption of a highly decentralised approach is never plausible. In fact, decentralised decision making has been found (a) to be quite efficient, and (b) to describe various real firms reasonably well – for example, some franchising setups

could be seen to fit into this model, as well as some large production conglomerates.[12] Hence, this assumption, while restrictive, does not mean that the model has no realistic applications whatsoever.[13] Note also that I here do not consider the question of *why* this particular organisational structure was chosen by the firms' owners – my concern is with establishing what happens *given* this particular structure (as noted earlier, this is not an inquiry into a form of a 'major transition' – it assumes that firms / offices already are individuals in sense required for an MLS2 approach to be applicable at all). However, as will become clearer below, much can be derived from this alone.

Two more points are important to note about how the present model treats the process of opening up a new office. Firstly, opening up a new office can be done either by starting one up from scratch, or by buying out less successful offices of other firms and turning them into offices of one's own firm (an office cannot buy out an office of its own firm, as it is not in competition with offices of its own firm in the same way as it is in competition with offices of other firms).

Secondly, the nature of the hiring process is different here from that of the initial offices. In particular, offices opening up a new office are no longer assumed to sample the applicant pool randomly, but to *replicate the composition of expectation-exceeders and clock-punchers in their own office*. The main reason for this assumption is that it is plausible that offices with more expectation-exceeders will not only be more profitable, but also be better at hiring new employees. The expectation-exceeders will work better and harder together not just in producing the relevant good or service, but also in screening applicants, interviewing them, and making the selection. Hence, they are more likely to hire 'high-quality' applicants. In turn, this makes it

[12] For more on decentralised organisational structures, see Radner (1993), Carley (1996), Witt (2003), and Gindis (2009).

[13] This thus mitigates the concerns raised in Hirschleifer (1977) and Frank (2003): the present model presents a case where it is at least *possible* to make sense of the idea of office reproduction, rather than just firm growth. For a good recent overview of some of the work on firm spinoffs, see also Klepper (2009). See also below.

plausible that newly opened up offices will tend to look like their 'parental' offices in terms of the composition of their employee types.

Formally, the process of opening up new offices of the same firm type is therefore modelled as follows. After determining which offices go bankrupt, the surviving offices assess whether they are in a position to open up another office of the same firm type. This happens with probability $p_{nf}=E_3(1-d/n)$, with $E_3 \varepsilon \boldsymbol{R}$ and in the interval [0, 1]. If the office thus determines that it can open a new office of the same firm type, it searches randomly either through the entire grid or through their neighbourhood (a parameter that can be changed, but which does not seem to alter the conclusions of the model much), and stops when it has either found an empty spot or one occupied by a office of a different firm type that has more clock-punchers among its employees. When a suitable spot for the new daughter office has thus been found, the new office will hire employees in such a way that they match the composition of clock-punchers and expectation-exceeders in the parental office.

The model overall then consists of repeated cycles of stages 2 and 3 above: the employees within the various offices interact with each other and (potentially) adjust their strategies, after which the offices' profitability is assessed and given the (probabilistic) power to influence (a) whether a given office gets to survive and (b) whether it can open up a new office of the same firm type. The cycles will end if a given number of offices exist (i.e. the market is saturated), or the number of existing offices falls below its initial state *m*. In the former case, no new offices can be opened (though old ones can still go bankrupt as before), and in the latter case, a number *v* of new firms / offices are created as in stage 1, so as to bring the number of existing offices back to *m*. Then, the rotation of stages 2 and 3 begins anew.

2. *The Results of the Model*

The model's structure is sufficiently complex so that an exhaustive analysis of its behaviour is difficult. However, it is possible to get a fairly good description of the major patterns of these results by simulating runs of the model under a wide range of different parameter values.[14]

To do this, I have done the following. Firstly, I have held fixed the following parameters: the size of offices (i.e. n) was 10 employees; the initial number of firms (i.e. m) was 10; 50% of employees in the global pool of hireable employees were expectation-exceeders (i.e. $p_E = 0.5$), and offices could only open up new offices in their immediate neighbourhood. The maximum number of firms a market was assumed to be able to contain was set at 400. Each run of the model lasted 1000 inter-office interactions. Given this, I then considered a number of different scenarios. In particular, for three settings of r (1, 4, and 7), I have successively set E_2 at 0.1, 0.5, and 0.9, and then, in each of these cases, let E_1 and E_3 vary from 0.1 to 0.9 (in 0.1 increments); finally, I have switched E_2 and E_3 and repeated the exercise (i.e. set E_3 at 0.1, 0.5, and 0.9, and then, in each of these cases, let E_1 and E_2 vary from 0.1 to 0.9 in 0.1 increments). This yielded three main patterns, depending on how the different parameters were set.[15]

Firstly, if the rate of firm creation is sufficiently faster than that of firm-internal learning, a firm will eventually establish itself and take over the entire market.[16] Call this the *positive scenario*. A typical plot of 10 variations on this theme looks like this (the y-axis represents the number of offices in the economy, and the x-axis the number of inter-office interactions):

[14] This was done using NetLogo; the code is available at http://tinyurl.com/ofurolb.
[15] These results have also been confirmed with a midlevel analysis that simultaneously varied r between 1, 4, and 7, E_1 between 0.01 to 0.41 in 0.1 increments, and E_2 and E_3 between 0.4 and 0.8 in 0.1 increments. The details of the results are available at http://tinyurl.com/pp4ec62.
[16] This last fact should not be overemphasised – it is a straightforward consequence of the fact that this is a selection-based model with a monotonically increasing fitness function. Note also that there can be long periods where several firms coexist in the market.

[Figure 1: A Set of 10 Typical Runs in the Positive Scenario]

The second scenario to consider here comes about when expectation-exceeders are turned into clock-punchers relatively easily – i.e. if the rate of firm-internal learning is sufficiently faster than the rate of firm creation. Then, genuine market selection of different firms does not take hold: firms / offices constantly get created and go out of existence, with none of them having the ability to persist for an extended period of time. Call this *negative scenario*. It is important to note that this scenario is in fact quite common: in many cases of the model, there is constant firm / office ending and firm / office creation, without any firm / office getting to persist long in the market at all (the total number of firms that is being created here quickly moves into the thousands). A typical plot of 10 variations on this theme looks like this (x- and y-axes are as above):

[Figure 2: A Set of 10 Typical Runs in the Negative Scenario]

Finally, if the rate at which new offices can be opened is sufficiently balanced by the rate at which individuals learn, a cyclical behaviour ensues: one or a few firms temporarily take over the market, collapse, and a new set of firms takes over the market. Call this scenario the *neutral scenario*. A typical plot of 10 variations on this theme looks like this (x- and y-axes are as above):

[Figure 3: A Set of 10 Typical Runs in the Neutral Scenario]

Now, it turns out that there are three parameters that are particularly important for the occurrence of a positive scenario (as opposed to a negative or neutral one):

(a) the number of interactions among employees (i.e. parameter *r*);
(b) the elasticity of offices going bankrupt on the number of expectation-exceeders (i.e. parameter E_2);
(c) the elasticity of offices opening up daughter offices on the number of expectation-exceeders (i.e. parameter E_3).

What is furthermore interesting to note is that, among these parameters, E_3 matters much more than E_2. For a positive scenario to occur, E_3 needs to be quite high; if it is too low, then no matter how high E_2 is, a firm does not establish itself in the market. Higher values of E_2 matter only to the extent that, if E_3 is above the threshold value, they make the occurrence of positive scenarios more likely. In other words: differential office replication is much more important that differential office survival. This is in fact quite in line with what should be expected from general evolutionary theoretic considerations: survival matters only in so far as it affects reproductive success, whereas the latter affects the evolutionary process directly (Sober, 2000). Another interesting point about this model (though this may be more due to its specific details, rather than being a general finding) is that the occurrence of a positive scenario does not depend much on the elasticity of learning of expectation-exceeders on the number of clock-punchers around (i.e. parameter E_1). All that really matters is the length of the learning period (i.e. *r*) – the higher the value of *r* is, the higher the threshold value of E_3 is that allows the occurrence of

positive scenarios. At the extremes, firms either have a hard time establishing themselves at all or will nearly always do so.[17]

In a bit more detail, the relationships between r, E_1, E_2, and E_3 in the creation of positive scenarios can be illustrated with the following graphs.

[Figure 4: Distribution of Positive Scenarios for r=1]

[Figure 5: Distribution of Positive Scenarios for r=4]

These graphs bring out four main conclusions:

(1) For a given value of r, the proportion of the E_1 / E_3 space (if non-zero) that allows for the evolution of stable firms increases as E_2 increases.

(2) If E_3 is less than a given threshold level, positive scenarios do not occur, no matter what the value of E_2 is.

(3) The threshold value for E_3 to lead to the occurrence of positive scenarios increases as r increases.

(4) There were no positive scenarios for r=7.

[17] The number of expectation-exceeders there are in the initial pool of applicants (i.e. p_E) also matters, but it seems primarily to concern the length of time it takes before a firm establishes itself in the market.

III. The Evolutionary Argument for the Agent-Based Theory of the Firm

What does the above model imply concerning the evolutionary argument for the agent-based theory of the firm? To answer this question, I first connect the model to the premises of this argument, and then analyse its solutions in relation to the latter's conclusions.

1. The Model and the Premises of the Evolutionary Economic Argument

The first thing to note is that the above model, to some extent at least, can help shore up assumptions 3 and 4 of the evolutionary economic argument for the agent-based view of the firm. In particular, it presents a case where a market really can be seen to provide the background to a selective mechanism as found in traditional biological settings. In the model, there are entities – offices – that differ in features – profitability – that alter the probability of their surviving and having offspring – i.e. not going bankrupt and opening up new offices. Furthermore, the model's assumptions are not so implausible as to make it unlikely that any actual market could reasonably be said to correspond to them: in fact, as noted earlier, there seem to be several real world situations that are well modelled in this way (Radner, 1993; Carley, 1996; Klepper, 2009).

Importantly, it is also the case that the model can be interpreted in terms of genuine MLS2: the entities – offices – that survive and reproduce in the model can be seen to be collectives (made up of individual employees), but fully coherent evolutionary entities nonetheless.[18] This becomes even clearer when it is noted that the fate of the individual employees need not be seen to be of direct consequence to the model at all: the total number of expectation-exceeders in the market can be taken to matter only to the extent that this influences the probability with which

[18] Here it is also important to recall (see also note 7) that MLS2 does not require group fitness to be completely independent of any property of the lower level (here, it is dependent on the number of expectation-exceeders in an office).

offices survive and reproduce – individual employees need not be presumed to reproduce in this model at all.

Note also that this conclusion is not altered by the fact that, if one focuses on the evolution of the two *strategies* (expectation-exceeder / clock-puncher) – i.e. if one switches to considering the issue from the perspective of the routine-based view of the firm – the model appears to represent an MLS1-type scenario (within an office, the expectation-exceeder strategy does worse than the clock-puncher one, but assuming expectation-exceeder strategies cluster in some offices, more new expectation-exceeder strategies will be created than clock-puncher ones *overall*). This change of perspective towards the evolution of strategies does not change the conclusion reached here, as it remains true that there is no *need* to see strategies as being the only replicators in the model (a point that is further reinforced by the fact that this is an agent-based model, not a strategy-based one; see Alexander, 2007, for more on this distinction). Put differently: all that I need for the argument here is that it is granted that the model can reasonably be interpreted as concerning an MLS2-scenario; whether there are also other interpretations is not so relevant. (It is interesting to note that the situation here is analogous to many straightforwardly evolutionary biological cases: in the latter context, it is also often possible to keep track of evolutionary changes in terms of changes in gene frequencies – however, this need not always be taken to be the best or only way of understanding the evolutionary mechanisms at play. See also Sober & Wilson, 1998.)

However, having said all this, it also needs to be noted that the model has several peculiar features. In particular, the kind of independence with which offices are assumed to operate is certainly not found in every market. Equally, the fact that all employees can find new jobs easily is far from universally plausible, as goes for the fact that all firms produce highly substitutable

goods. These peculiarities of the model are important to note, for at least some of them are in fact *necessary* for the modelled market to operate like a selective mechanism. In particular, unless offices have the kind of independence set out above, it will not even be *possible* for them to be genuine evolutionary individuals (Godfrey-Smith, 2009).

In this vein, the complexity of the status of the evolutionary economic argument already becomes clear: while the assumptions needed for it to get off the ground are not completely unreasonable, they cannot be seen to be innocuous or very easily satisfied.[19] Putting all of this together, this therefore means that the above model provides some – highly qualified – support to assumptions 3 and 4 of the evolutionary economic argument.

2. The Model and the Conclusions of the Evolutionary Economic Argument

When it comes to conclusion 5 there are two things that need to be noted. Firstly, the model shows that the evolutionary economic argument presented earlier can indeed sometimes be cogently spelled out. In particular, it shows that offices really can, at least sometimes, be seen to be genuine evolutionary individuals – and that is so even though they are composed of individual agents with their own and to some extent office-opposing agendas.[20]

However, secondly, it also needs to be noted that the opposite is possible as well: a negative scenario exists, where offices *lack* the kind of internal coherence to be genuine evolutionary individuals. In particular, it is possible that the agendas of employees constantly pull the offices apart, causing their destruction and the formation of a new set of offices or firms. As noted above, there is little theoretical sense in speaking of offices as evolutionary individuals of their

[19] In this way, the present model partly contradicts and partly supports claims made in Enke (1951), Hirschleifer (1977), and Frank (2003): there are cases where talking of firm reproduction does seem plausible. See also Klepper (2009).

[20] It is worthwhile noting that the model displays, in the first instance, *office* selection. *Firm* selection is merely a by-product of office selection. This, though, does not alter anything of importance for the present discussion.

own here – genuine *evolution* does not take place at all. This second conclusion is especially noteworthy, as the model was to some extent designed to make it as easy as possible for firms to evolve as collective agents of their own. The fact that this does not necessarily happen even in this case is thus an important point to note: even where it is assumed that firms can, in principle, evolve in an MLS2 sense, it is not guaranteed that they are able to evolve at all.

So, what are we to conclude concerning the validity of conclusion 5? Given the above, it seems clear that, once again, only a highly qualified assessment is plausible: sometimes, evolutionary theory does support the agent-based view of the firm, but at other times, it does not. Moreover – and, as noted above, this is the key finding of the model – we now have some idea of *when* the agent-based view of the firm is evolutionarily favoured. In particular, this happens when office *creation* – not just office *survival* – is sufficiently strongly related to the presence of high-effort employees to counterbalance the learning of these expectation-exceeders within offices. That is, the agent-based theory of the firm is evolutionarily favoured when the destructive pressures on a firm (exerted by the individually beneficial actions of its employees) are balanced by sufficiently strong market forces that reward highly profitable firms with increased presence in the market.

In this way, two important lessons can be learned. Firstly, even if it is assumed that firms have all the necessary *intrinsic* features to be genuine evolutionary – and thus (or so it is assumed here) economic – individuals of their own (such as the ability and motivation to replicate independently), they might fail to be genuine evolutionary individuals due to the fact that the *extrinsic* conditions are such that the persistence of these individuals is impossible. Secondly, the latter case is more likely to occur when the selective forces on the office level are not strong enough to balance the individual-level forces that lead to office destruction. This is important, for

especially this last kind of knowledge has so far been missing in this context (Schulz, 2013): while it has been known already that firms may *sometimes* be collective agents, what has not been clear so far is exactly *when* they are.

Precisely this is what the present model can show: the more real firms instantiate the conditions of the positive scenario of the present model, the more plausible it is to see them as genuine agents of their own. Examples of this may include firms like coffee shops or organic grocery stores: these firms often seem to have employees that remain committed for a long time (perhaps for ideological and moral reasons) and fierce pressures to expand in the market (people like their coffee shops and organic grocery stores small, near them, and not overly busy). By contrast, companies with large employee turnover and less immediate market expansion pressures – perhaps such as discount stores or fast food restaurants – are less well seen as agents of their own. Of course, further confirmation of this would require detailed empirical studies to see to what extent these firms and markets really do match the conditions of the positive and negative scenarios of the above model. The important point to emphasise here is just that the present model should make these empirical investigations easier, as it can suggest what we ought to look for.[21]

IV. Conclusion

I have critically assessed and extended a specific evolutionary economic argument – based on the market selection of firms as collective entities – in support of the agent-based theory of the firm. In particular, I firstly hope to have shown that the existing work on this argument needs to be

[21] Of course, the model can also be used – perhaps in conjunction with other models – to investigate various other questions, e.g. concerning economic growth or innovation (see e.g. Nelson & Winter, 1982). As noted earlier, this is not the aim in the present context, but of course, there is nothing speaking against using the model for these other purposes as well.

clarified and corrected, as it is based on the wrong account of multi-level selection. Secondly, I have presented a model that does precisely this – in particular, it can be seen to be based on genuine MLS2. In this way, thirdly, I hope to have made clearer exactly when – and when not – there are evolutionary reasons to see firms as agents of their own.

Bibliography

Aghion, Philippe, and Holden, Richard (2011). 'Incomplete Contracts and the Theory of the Firm: What Have We Learned over the Past 25 Years?'. *Journal of Economic Perspectives* 25: 181-197.

Alchian, Armen (1950). 'Uncertainty, Evolution, and Economic Theory'. *The Journal of Political Economy* 58: 211-221.

Alexander, Jason (2007). *The Structural Evolution of Morality*. Cambridge: Cambridge University Press.

Arrow, Kenneth, and Hahn, Frank (1971). *General Competitive Analysis*. San Francisco: Holden-Day.

Bacharach, Michael (1999). 'Interactive Team Reasoning: A Contribution to the Theory of Co-operation'. *Research in Economics* 53: 117-147.

Baum, Joel, and McKelvey, Bill (eds.) (1999). *Variations in Organizational Science*. London: Sage.

Boyd, Robert, and Richerson, Peter (2005). *The Origin and Evolution of Cultures*. Oxford: Oxford University Press.

Brandon, Robert (1990). *Adaptation and Environment*. Cambridge: Cambridge University Press.

Buss, Leonard (1987). *The Evolution of Individuality*. Princeton: Princeton University Press.

Carley, Kathleen (1996). 'A Comparison of Artificial and Human Organizations'. *Journal of Economic Behavior and Organization* 31: 175-191.

Clark, Andy (1997). *Being There*. Cambridge, MA: MIT Press.

Coase, Ronald (1937). 'The Nature of the Firm'. *Economica* 4: 386-405.

Damuth, John, and Heisler, Ida L. (1988). 'Alternative Formulations of Multi-Level Selection'. *Biology and Philosophy* 3: 407-430.

Day, Richard, & Chen, Ping (1993). *Nonlinear Dynamics and Evolutionary Economics*. Oxford: Oxford University Press.

Demsetz, Harold (1997). 'The Firm in Economic Theory: A Quiet Revolution'. *The American Economic Review* 87: 426-29.

Enke, Stephen (1951). 'On Maximizing Profits: A Distinction between Chamberlin and Robinson'. *The American Economic Review* 41: 566-578.

Fama, Eugene (1980). 'Agency Problems and the Theory of the Firm'. *The Journal of Political Economy* 88: 288-307.

Foss, Nicolai & Klein, Peter (2008). 'The Theory of the Firm and Its Critics: a Stocktaking and an Assessment'. In E. Brousseau and J.-M. Glachant (eds.). *Handbook of New Institutional Economics.* Cambridge: Cambridge University Press.

Frank, Joshua (2003). 'Natural Selection, Rational Economic Behavior, and Alternative Outcomes of the Evolutionary Process'. *Journal of Socio-Economics* 32: 601-622.

Futuyma, Douglas (2013). *Evolution*. Third Edition. Sunderland: Sinauer Associates.

Ghiselin, Michael (1987). 'Bioeconomics and the Metaphysics of Selection'. *Journal of Social and Biological Structures* 10: 361-369.

Gigerenzer, Gerd, and Selten, Reinhard (eds.) (2001). *Bounded Rationality: The Adaptive Toolbox*. Cambridge, MA: MIT Press.

Gindis, David (2009). 'From Fictions and Aggregates to Real Entities in the Theory of the Firm'. *Journal of Institutional Economics* (2009), 5: 25-46.

Godfrey-Smith, Peter (2009). *Darwinian Populations and Natural Selection*. Oxford: Oxford University Press.

Gould, Steven Jay (1996). *Life's Grandeur: The Spread of Excellence from Plato to Darwin*. London: Jonathan Cape.

Groenewegen, John, and Vromen, Jack (1999). *Institutions and the Evolution of Capitalism: Implications of Evolutionary Economics*. Cheltenham: Edward Elgar.

Hart, Oliver (2008). 'Reference Points and the Theory of the Firm'. *Economica* 75: 404–411.

Hirshleifer, Jack (1977). 'Economics from a Biological Viewpoint'. *Journal of Law and Economics* 20: 1-52.

Hodgson, Geoffrey (1999). *Evolution and Institutions*. Cheltenham: Edward Elgar.

Hodgson, Geoffrey (2007). 'Meanings of Methodological Individualism'. *Journal of Economic Methodology* 14: 211-226.

Hodgson, Geoffrey, and Knudsen, Thorbjorn (2010). *Darwin's Conjecture*. Chicago: University of Chicago Press.

Hull, David (1988). *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago: University of Chicago Press

Klein, Benjamin, Crawford, Robert, and Alchian, Armen (1978). 'Vertical Integration, Appropriable Rents, and the Competitive Contracting Process.' *Journal of Law and Economics* 21: 297–326.

Klepper, Steven (2009). 'Spinoffs: A Review and Synthesis'. *European Management Review* 6: 159-171.

Knight, Frank (1921). *Risk, Uncertainty, and Profit*. Boston: Houghton Mifflin.

Kreps, David (1990). *A Course in Microeconomic Theory*. Princeton: Princeton University Press.

List, Christian, and Pettit, Philip (2006). 'Group Agency and Supervenience'. *Southern Journal of Philosophy* 44: 85-105.

Maynard Smith, John (1987). 'How to Model Evolution'. In: J. Dupre (ed.). *The Latest on the Best: Essays on Evolution and Optimality*. Cambridge, MA: MIT Press, pp. 119–131.

Maynard Smith, John, and Szethmary, Eoros (1995). *The Major Transitions in Evolution*. Oxford: Oxford University Press.

Michod, R. (1999). *Darwinian Dynamics: Evolutionary Transitions in Fitness and Individuality*. Princeton: Princeton University Press.

Nickerson, Jack, and Zenger, Todd (2004). 'A Knowledge-Based Theory of the Firm: The Problem-Solving Perspective'. *Organization Science* 15: 617-632.

Nelson, Richard, and Winter, Sidney (1982). *An Evolutionary Theory of Economic Change*. Cambridge, MA: Belknap Press.

Okasha, Samir (2006). *Evolution and the Levels of Selection*. Oxford: Oxford University Press.

Penrose, Edith (1959). *The Theory of the Growth of the Firm*. Oxford: Blackwell.

Pettit, Philip (2003). 'Groups with Minds of Their Own'. In F. Schmitt (ed.). *Socializing Metaphysics: the Nature of Social Reality*. Lanham: Rowman & Littlefield, pp. 167-193.

Radner, Roy (1993). 'The Organization of Decentralized Information Processing'. *Econometrica* 61: 1109-1146.

Radner, Roy (2000). 'Costly and Bounded Rationality in Individual and Team Decision-Making'. *Industrial and Corporate Change* 9: 623-658.

Radner, Roy (2006). 'Neo-Schumpeterian and Other Theories of the Firm: A Comment and Personal Retrospective'. *Industrial and Corporate Change* 15: 373-380.

Richardson, G. B. (1972). 'The Organisation of Industry'. *The Economic Journal* 82: 883-896.

Rosenberg, Alexander (1994). 'Does Evolutionary Theory Give Comfort or Inspiration to Economics?'. In P. Miroski (ed.). *Natural Images in Economic Thought*. Cambridge: Cambridge University Press.

Royall, Richard (1997). *Statistical Evidence – a Likelihood Paradigm*. Boca Raton, FL: Chapman and Hall.

Satz, Debra, and Ferejohn, John (1994). 'Rational Choice and Social Theory'. *The Journal of Philosophy* 91: 71-87.

Schmalensee, Richard, and Willig, Robert (eds.) (1989). *Handbook of Industrial Organization*. Amsterdam: North-Holland.

Schulz, Armin (2013). 'Beyond the Hype: The Value of Evolutionary Theorizing in Economics'. *Philosophy of the Social Sciences* 43: 46-72.

Schumpeter, Joseph (1942). *Capitalism, Socialism, and Democracy*. New York: Harper.

Simon, Herbert (1957). *Models of Bounded Rationality*. Cambridge, MA: MIT Press.

Sober, Elliott (2000). *Philosophy of Biology*. 2nd Edition. Boulder: Westview Press.

Sober, Elliott (2008). *Evidence and Evolution*. Cambridge: Cambridge University Press.

Sober, Elliott and David Sloan Wilson (1998). *Unto Others: The Evolution and Psychology of Unselfish Behavior*. Cambridge, MA: Harvard University Press.

Sterelny, Kim (2001). 'Darwin's Tangled Bank'. In *The Evolution of Agency and Other Essays*. Cambridge: Cambridge University Press, pp. 152-178.

Sterelny, Kim (2012). *The Evolved Apprentice*. Cambridge, MA: MIT Press.

Sterelny, Kim, and Griffiths, Paul (1999). *Sex and Death*. Chicago: University of Chicago Press.

Veblen, Thorstein (1898). 'Why is Economics Not an Evolutionary Science?' *Quarterly Journal of Economics* 12: 373–397.

Vromen, Jack (2001). 'The Human Agent in Evolutionary Economics'. In J. Laurent and J. Nightingale (eds.). *Darwinism and Evolutionary Economics*. Cheltenham: Edward Elgar, pp. 184-208.

Vromen, Jack (2009). 'Advancing Evolutionary Explanations in Economics: The Limited Usefulness of Tinbergen's Four-Question Classification'. In H. Kincaid (ed.). *The Oxford Handbook of Philosophy of Economics*. Oxford: Oxford University Press, pp. 337-368.

Watkins, John (1952). 'Ideal Types and Historical Explanation'. *British Journal for the Philosophy of Science* 3: 22-43.

West, Stuart; Mouden, Claire El, and Gardner, Andy (2011). 'Sixteen Common Misconceptions about the Evolution of Cooperation in Humans'. *Evolution and Human Behavior* 32: 231-262.

Williamson, Oliver (1971). 'The Vertical Integration of Production: Market Failure Considerations'. *American Economic Review,* 61: 112-123.

Williamson, Oliver (1979). 'Transaction Cost Economics: The Governance of Contractual Relations'. *Journal of Law and Economics* 22: 233-261.

Winter, Sidney (1988). 'On Coase, Competence, and the Corporation'. *Journal of Law, Economics, and Organization* 4: 163-180.

Witt, Ulrich (1999). 'Bioeconomics as Economics from a Darwinian Perspective'. *Journal of Bioeconomics* 1: 19-34.

Witt, Ulrich (2003). *The Evolving Economy*. Cheltenham: Edward Elgar.

Witztum, Amos (2012). 'The Firm, Property Rights, and Methodological Individualism: Some Lessons from J. S. Mill'. *Journal of Economic Methodology* 19: 339-355.

Figures

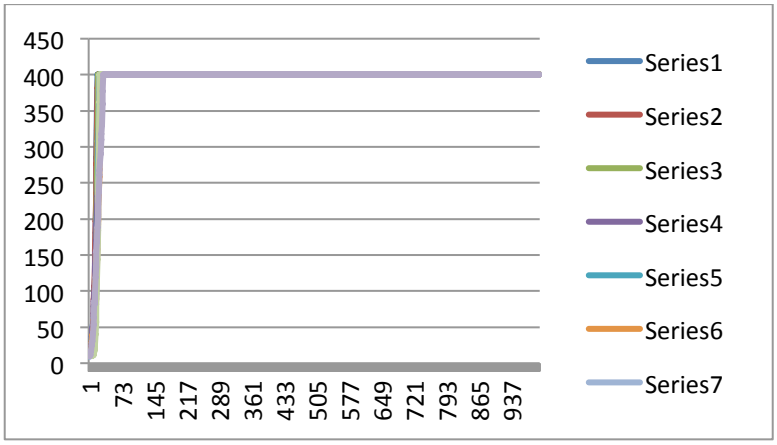

[Figure 1: A Set of 10 Typical Runs in the Positive Scenario]

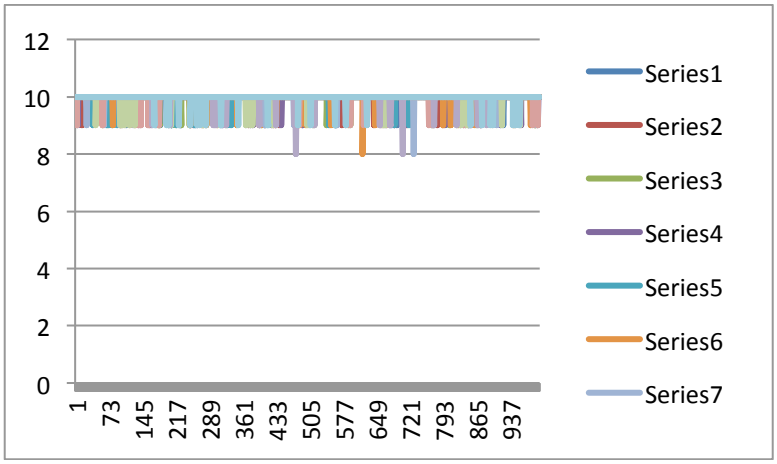

[Figure 2: A Set of 10 Typical Runs in the Negative Scenario]

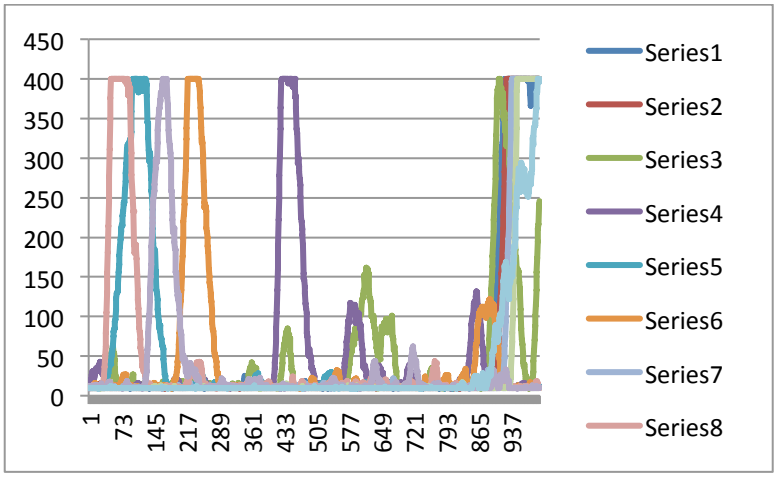

[Figure 3: A Set of 10 Typical Runs in the Neutral Scenario]

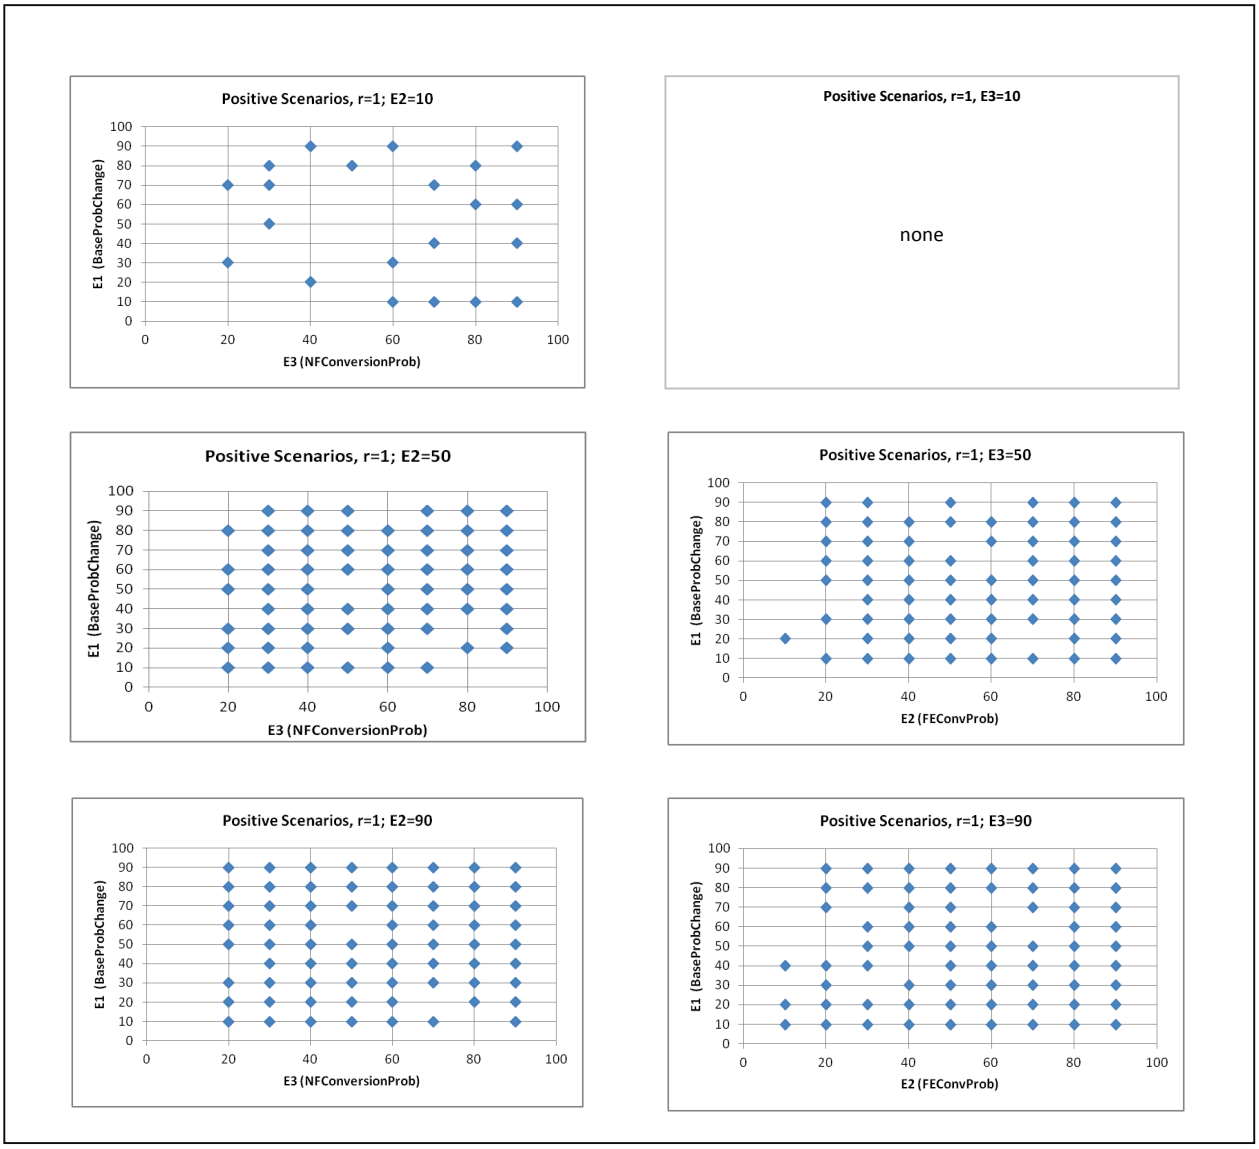

[Figure 4: Distribution of Positive Scenarios for r=1]

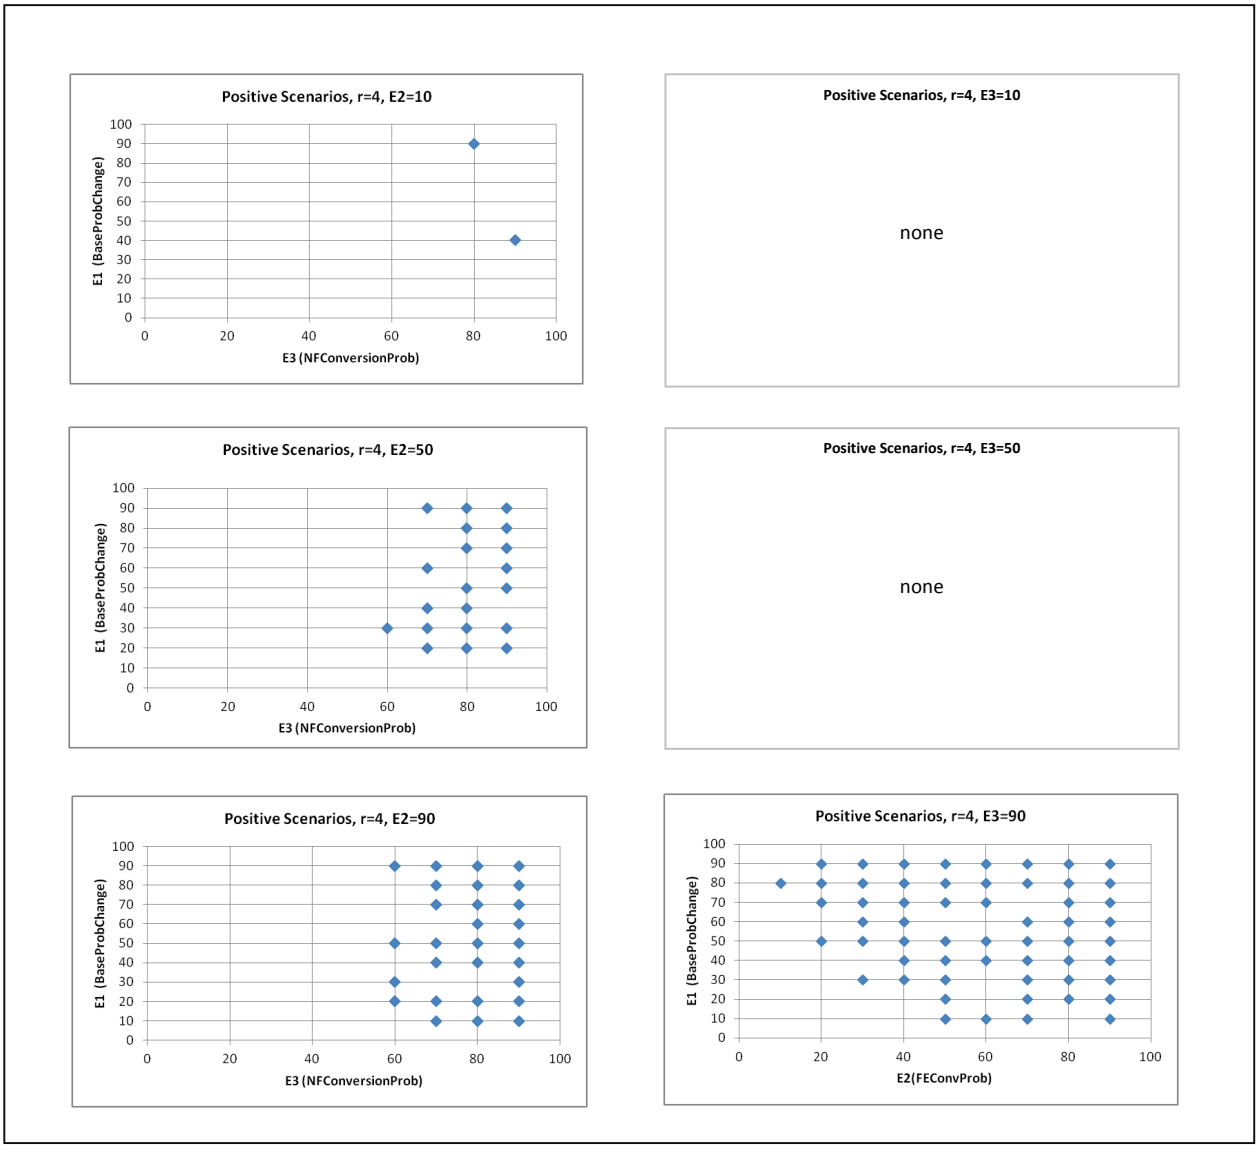

[Figure 5: Distribution of Positive Scenarios for r=4]