Flexibility, Profitability and Survival in an (Objective) Model of Knightian Uncertainty*

John Sutton

London School of Economics

Abstract: This paper explores the relationship between a firm's investments in 'capability building', and its prospects for survival. The analysis is carried out using a model of Knightian uncertainty which differs from the standard 'subjective probability' approaches in a number of fundamental respects.

* Walras-Bowley Lecture, Econometric Society, 2004. This lecture was subsequently presented at the UCL Conference in honour of Ken Binmore; I am grateful for comments received on that occasion. My thanks are also due to Ken Binmore, Adam Brandenburger, Andrea Prat and Michele Piccione for their very helpful comments on a preliminary draft.

1. Introduction

When a firm invests in acquiring the know-how to produce some new product, it may turn out at some future date that the know-how it has acquired in this way will be of value in producing further products, whose nature was not foreseeable when the initial investment was made.[1] These more distant, indirect benefits of investing in know-how play a central role in the modern literature on firms' capabilities, as discussed in the extensive literature on 'Schumpeterian' or 'Evolutionary' models pioneered by Nelson and Winter (1982).[2] One aspect of a firm's capability is its 'flexibility': if the firm invests in know-how in a way that makes it 'more flexible' (in the sense that it can respond more quickly or more effectively to changes in future market conditions), then how does this affect (a) its profitability, and (b) its prospects for survival? While a positive link between flexibility and survival is empirically well supported,[3] the connection with profitability is complex. The main substantive contribution of the present paper lies in unraveling the implications of such 'investments in capability' for the firm's profitability, and its survival, respectively.

Profitability and survival are usually taken to be positively associated: poor current profitability may threaten current survival, while strong current profitability makes

---

[1] In some instances, a firm's decision to proceed with an R&D project may rest on an appeal to benefits of this kind. The Indian machine-tool maker described in Sutton (2001a) emerged in the late 1990s as the country's leading producer of the most popular computer-controlled machine tools. A 'next generation' of more sophisticated tools already accounted for a large share of the market in advanced industrial economies. It was clear that investing in the design and production of these new generation tools could not be profitable per se – local demand would remain modest over the life cycle of these machines. The company's decision to undertake a 'loss-making' investment in developing these more advanced machines lay in its belief that, ten years hence, yet newer designs would have evolved, and if it did not develop the capability to produce current leading-edge machines, it would not be able to produce designs a decade hence which could compete against imported machines of whatever new but unforeseeable types the Indian market would at that stage be demanding. To ask the Indian company to specify the future scenario more precisely than this, or to attach probabilities to different scenarios, would be futile – yet the company's view of its survival prospects hinged on this judgment.

[2] This paper is one of three companion papers that offer a simple analytical representation of the 'capability' concept. For an overview, see the first paper in the series (Sutton 2006).

[3] See in particular Schott and Bernard (2005), who link the survival of U.S. manufacturing firms to their propensity to switch from one product group to another.

withdrawal from a market less likely[4]. This conventional view of the association rests implicitly on two notions: that the firm's current profitability is correlated with its likely future profitability, and/or that firms face some kind of capital market imperfection, so that their access to retained earnings protects them from the consequences of profit fluctuations. In what follows, we focus on a different aspect of the relationship: whenever some change occurs in market conditions that require a costly adjustment, each firm's future prospects will depend inter alia on how costly it will be for it to adjust to the change, and so on its underlying capability. This suggests a rather different take on the relationship between profitability and survival: the first implication of the present analysis in that firms that engage in 'capability building' enhance their prospects for survival, but – depending upon the values of underlying parameters – such firms may score either relatively well, or relatively badly, in terms of their (ex-post) lifetime profitability.

A second implication of the present analysis relates to the question of whether profit maximizing firms generate 'too little' or 'too much' investment in capability. A familiar argument turns on the fact that a firm cannot fully capture the gains from such investments, in part because the know-how in which it invests may be non-patentable, and may be held by individual employees who can quit and join other firms, thus generating 'spillovers' for rivals. This generates a bias towards under-investment in capability. In what follows, however, we identify a second, offsetting, bias that leads towards over-investment. This arises because a firm's investment in capability reduces its likelihood of exiting the market; but the exit of any one firm generates benefits for rival firms (or potential entrants) who may capture the departing firm's market position.

In analysing issues this kind, it is reasonable to model the part of the firm's payoff associated with the product it plans to produce in terms of a conventional 'probabilistic' analysis. It is, however, less plausible to imagine the firm having some model relative to which it might form a well-founded expectation of the additional benefits it may derive from 'future generation' products whose nature is not yet known. A more plausible

[4] The model that follows relates to the special case of a firm operating in a single market, so that the exit of the firm, and its withdrawal from some market (its closing down some line of business) are synonymous.

representation of the firm's decision-making in this setting is one that incorporates 'true' or 'Knightian' uncertainty.

A key claim in the 'Schumpeterian' literature is that there is a fundamental difference between the environment faced by a firm within a 'probabilistic' world, as against a world of 'Knightian uncertainty' in which the evolution of the system is driven by a sequence of unique, unrepeatable circumstances that cannot be represented in a probabilistic way. This claim is difficult to formulate, and address, within the now standard 'subjective probability' tradition that runs from Savage (1954) to Schmeidler (1989), since the question it poses relates to the nature of the 'true, underlying model' faced by agents, a question that is set aside in the standard approach.[5] What the question amounts to, in formal terms, is this: what difference does it make to the (range of) choices that a rational agent might make, if the different scenarios he or she may face can not validly be modeled as arising with some (known or unknown) probabilities? An appropriate way of addressing this issue is to introduce a setup in which the objective environment faced by the agent is fully specified, and in which we can examine the effect of relaxing the key assumption of measurability on which any probabilistic analysis must rest. This is the approach taken in what follows: firms are expected profit maximizers, but the environment they face may be one in which expectations cannot be taken across certain sets of outcomes. (The state space contains some non-measurable subset). The result of this is that the set of actions that may be

[5] The conventional 'subjective probability' approach began with Savage (1954) and Anscombe and Aumann (1963), and its most widely accepted exemplar lies in the work of David Schmeidler and his co-authors. See in particular, Schmeidler (1989), Machina and Schmeidler (1992, 1995), Gilboa and Schmeidler (1993, 1995), Dow and Werlang (1994). In the original Savage approach, the choices of agents are taken as a given, and these choices are rationalized by attributing to agents a set of preferences and (subjective) probabilities over states of the world under which these choices are consistent with rationality. Some strands of the modern literature depart from the Savage approach of taking choices as 'givens' and inducing the probabilities; for example, one strand begins by equipping different agents with different prior beliefs as to the likelihood of various states of the world ('non-common priors'; on Savage's views on 'common priors', see the remarks of Aumann (1987)). Different agents then make different choices, because they hold different priors. The effect of using a non-common priors approach, in the present setup, is noted on page 19 below. Other strands in the post-Savage literature include the approach of Bewley (1986, 1987), the 'incomplete preferences' approach of Aumann (1962), and the 'lexicographic preferences' of Blume, Brandenburger and Dekel (1991).

taken by a rational agent will in general be widened, relative to that obtaining in the probabilistic setup. It will not in general be the case that an appeal to rationality can serve to confine a set of agents with identical preferences to some unique best action.

Specifically, we introduce two devices ('machines') that generate outcomes, one of which is probabilistic, and the other of which is non-probabilistic. We compare the case of an agent facing a machine of the first type, as against the second. If the machine is used only once, then there is no operational difference between the two types of machine, in the sense that the set of actions that may be chosen by a rational agent is the same in both cases. Once the machine is used more than once, however, the use of the probabilistic machine has the effect of constraining the range of actions that a rational agent may validly take, relative to the non-probabilistic setting ('reasonable actions' in what follows).

By introducing a representation of Knightian uncertainty in which a rational agent faces a fully specified objective model of the environment, the present paper reverts to what Savage labelled the 'frequentist' or 'British' view, and which he set out to replace with his 'subjective personality' approach. In reverting to this viewpoint, I make no general claim for its merits beyond noting that it offers an appropriate vehicle for addressing the 'Schumpeterian' claim noted above.[6] A full description of the several ways in which the

---

[6] A more general way of motivating this exploration of an 'objective' representation of uncertainty is to note that modern economic analysis rests on the notion of analyzing the actions of rational agents faced with a model of the world that is fully specified up to some probabilistic description (Sutton, 2006). A failure to 'close' the model, by leaving some aspects of the objective description of the agents' environment unspecified, has in the past been seen as a defect. Most famously, the use of extraneous 'expectations' or 'expectations formation mechanisms' (such as static, extrapolative or adaptive expectations) became unfashionable in macroeconomics from the 1970s onwards. It was agreed that agents' expectations are in practice formed by reference to their experience of their environment, and so the appropriate modelling technique is to provide a full (objectively accurate) description of the environment, and to require that the agents hold 'rational' expectations, relative to this 'true model' given their (complete or incomplete) information regarding models parameters.
Relative to this standard paradigm of the 'rational agent plus fully specified model', the conventional literature on uncertainty, at least in the strict Savage interpretation, remains an exception; here, we work in terms of subjective probability and may choose to leave open the objectively correct description of the market. While the usefulness, and fruitfulness, of this approach is not in question, it does nonetheless seem reasonable to explore another formulation, in which we seek the same kind of 'closure' in respect of a non-probabilistic setting of Knightian

present approach differs from the standard (subjective probability) approach to Knightian uncertainty of Schmeidler et al. will be found in Section 5.

The main novelty of the present framework is that it offers a simple way of bridging the perspective offered by 'equilibrium' and 'evolutionary' models respectively. In standard (Walrasian or Nash) equilibrium models, the preferences, and so the choices, of optimizing agents pin down market outcomes. In a standard evolutionary model agents are not assumed to make optimizing choices; rather their actions are determined by their 'types', and equilibrium outcomes are driven by selection effects that operate differentially across the various types of agent. In practice, both these features of markets are important. In the setup developed here, agents are rational maximizers with identical objectives who face a known environment, but the nature of the environment they face is such that an appeal to rationality does not, in general, lead to a complete ordering of actions. In contrast to the now conventional approach which introduces some additional behavioural postulate ('uncertainty aversion' or 'maximin behaviour') in order to pin down a unique outcome, we appeal only to the rationality of (expected) profit maximizing agents, and we note that there will, under Knightian uncertainty, be some subset of actions, any member of which can be chosen by a rational agent. In this setting, rationality only partially constrains outcomes, while selection effects operating across agents also play a role in pinning down the final outcome. The present approach offers one, very simple, way of combining the roles of optimizing behaviour and selection effects as drivers of market outcomes.

2. Modelling Knightian Uncertainty

Consider a 'machine' which generates a discrete outcome, z, taking the value 0 or 1, as follows. The machine operates in two steps. In step 1, a continuous random variable w is drawn from a uniform distribution on [0,1]. Step 2 involves a set, S, defined on [0,1], the construction of which is described later. If the random variable w lies in the set S, then x=1; otherwise, x=0. In what follows, an agent will be required to make a choice

uncertainty as was achieved for probability (risk) in the 'rational expectations' literature. One virtue of the approach offered here is that it provides a very simple way of doing this.

between two available alternatives. The payoff received by the agent will depend on the action chosen and on the sequence of future outcomes generated by this 'machine'.

In specifying the construction of the set S, we distinguish two cases, and we identify two distinct machines accordingly. The first case represents a standard 'risk' model: Machine I is constructed by choosing as the set S some measurable subset of [0, 1]; the measure of S is denoted p. 'Machine II' involves constructing S as a non-measurable subset of [0,1]. It will be a useful aid to intuition in what follows to recall some standard examples of non-measurable sets:

1. (Vitali): Map the unit interval onto the circumference of a circle, so that 1 coincides with 0. For any pair of points x and y, let x ~y if the length of the arc joining them, |x-y|, is rational. The equivalence relation ~ defines a set of equivalence classes; select one member from each class (this step requires the Axiom of Choice); the set of members selected form a (non-measurable) set, which we denote as S_0. (For a formal statement, and proof, see for example, Billingsley (1995), p.45).[7]

2. (A sequence of outcomes): Consider $S_0 \times S_0 \times \ldots \times S_0$ as a subset of $[1,0] \times [1,0] \times \ldots \times [1,0]$. The set $S_0 \times S_0 \ldots \times S_0$ is non-measurable (Burrill (1972), Sec. 5 – 73, p. 80).

3. (Halmos): Building on example 1 above, Halmos (1950) constructs a non-measurable set M which has the property that its inner (Lebesgue) measure is zero, while its outer (Lebesgue) measure is unity. (The 'ultimate non-measurable set', in the words of Gelbaum and Olmsted (1964)). This feature carries over to subsets of [0, 1] in the following sense: for any Lebesgue measurable set E, the outer and inner Lebesgue measure of $M \cap E$ satisfy.

[7] The idea behind the proof is that any translation of S_0 by a rational number r_k produces a congruent set S_k. The proof that S is non-measurable lies in noting that there is a countably infinite number of these (disjoint congruent) sets whose union coincides with [0,1], and if S is measurable, each must have the same measure. But this measure must be zero, whence [0,1] has measure zero; or it must be positive, whence [0,1] has infinite measure; whence a contradiction.

$$\mu^*(M \cap E) = \mu(E) \text{ and } \mu_*(M \cap E) = 0$$

To construct the set S used in Machine II we proceed as follows: we introduce two disjoint Lebesgue measurable sub-sets of [0, 1], labelled E_1 and E_3, the sum of whose measures is at most 1. We define the (Lebesgue-measurable) set $E_2 = (E_1 \cup E_3)^c$. We define S as follows:

$$S = E_1 \cup (M \cap E_2) \tag{1}$$

so that $\mu^*(S) = \mu(E_1) + \mu(E_2)$ and $\mu_*(S) = \mu(E_1)$. In the special case where $\mu(E_1) = \mu(E_3) = 0$, we have 'pure Knightian uncertainty', while in the case where $\mu(E_1) + \mu(E_3) = 1$, Machine II is equivalent to the 'probabilistic' Machine I, with $p = \mu(E_1)$. The outer and inner Lebesgue measures $\mu^*(S)$ and $\mu_*(S)$ provide a counterpart to the non-additive probabilities of Schmeidler (1989): to see this, define the random variables Z_t, taking the value 1 if $w \in E_1 \cup E_2$ and zero otherwise, and W_t, taking the value 1 if $w \in E_1$ and zero otherwise, whence the partial sums satisfy

$$\frac{1}{n}\sum_1^n Z_t \geq \frac{1}{n}\sum_1^n x_t \geq \frac{1}{n}\sum_1^n W_t$$

where the first and last terms converge almost surely to $\mu^*(S)$ and $\mu_*(S)$ respectively. It follows that we can associate the non-additive probabilities $p = \mu_*(S)$ and $q = 1- \mu^*(S)$ with the events 1 and 0 respectively, so that $1 - p - q = \mu(E_2)$ provides a measure of the degree of uncertainty in the sense of Schmeidler.

The following intuition provides the motivation for the definition of 'reasonable actions' on which the present approach rests: we may think of the agent considering various possible scenarios, each represented by a sequence of random draws of w, (not observed by the agent). We may imagine an agent who 'thinks through' each such scenario, and arrives at a (correct) view as to the outcome $(x_1, ..., x_T)$, that will arise under this scenario. But there is an uncountable infinity of such possible scenarios to consider, and in the (pure) uncertainty setting (recalling Example 2 above), there is no 'grouping' or 'ordering' of such scenarios that provides a shortcut to classifying them. Exploring any countable number of scenarios will not lead to any 'settling down' in the fraction of occurrences of some particular sequence $(x_1, ..., x_T)$, and so continued introspection by the agent will not lead to any settled view as to the relative likelihood of this sequence. It is plausible that no agent can examine more than a countable infinity of scenarios, and so different agents, each having examined different sub-sets of possible scenarios, may choose different actions. In particular, there is no basis for imposing on each agent a requirement to choose actions as if $(x_1,x_2) = (1,0)$ is no more or less likely than $(x_1,x_2) = (0,1)$: a continuing exploration of scenarios leading to (1,0), or to (0,1), can produce no settled view as to their relative likelihood. This is the key difference between Machine II and the 'probabilistic' Machine I, under which these events are equally likely. [8, 9]

This remark suggests an obvious question: what would happen if, instead of using the same Machine I repeatedly, we used a different version of Machine I each time? Could we not thereby relax the cross-restrictions on the probabilities of different sequences of outcomes? The answer to this turns on the question: where do these 'different' versions of Machine I come from? Two cases arise. In the first case, the different versions of

---

[8] This interpretation may suggest an analogy with the recent literature on complex environments, which begins from the question: what if two contracting parties face an environment which is 'too complex' to admit of the writing of a complete contract? (See, for example, Al-Nagar, Anderlini and Felli (2006)).

[9] It is worth commenting on the role played in the present setup by the 'stationarity of structure' introduced by way of repeated draws from Machine II. The idea here is to pin down a key distinction: if there are observable differences in the agents' environment over successive periods, then the agent might condition a choice of action on such differences. In any empirical setting, there will always be such differences across successive periods. We abstract from any such conditioning here, in order to focus attention on the key idea: in seeking to capture the idea, prevalent in the older literature on Knightian uncertainty, of a sequence of outcomes driven by unique, unrepeatable events, we set aside such aspects of this uniqueness and unrepeatability as might be used to condition a choice of action by a rational agent; and we embed what remains of this uniqueness and unrepeatability in the realizations of w for successive periods. The point is that stationarity in the structure or environment does not here imply any 'nice' properties in that stationarity in the structure or environment does not here imply any 'nice' properties in regard to the sequence of outcomes that will be generated.

Machine I (or equivalently, the different values of p which they embody) are themselves drawn from some underlying distribution. If this is the case, then we are back in the simple setting of a single Machine I; for the agent can proceed by simply taking an expectation of p; and can proceed to make inferences from past values of x, in essentially the same way as before.

The second case is more interesting. Here, we posit that there is no underlying population of 'Machine Is' from which the machine used in any particular period is drawn. Rather, we just have a sequence of machines (i.e. values of p) which are 'different'. But this begs the question, for the issue now becomes: what do we mean by a 'sequence of values that are different'? Does this sequence have any ergodic property? Does it admit of a probabilistic description? The aim of the construction described above is to provide an explicit description of a set-up in which the answers to such questions is 'no'. Introducing the idea of 'different machines' in each period does not advance the discussion, and so we operate throughout with a set-up in which the same machine, of either type I or type II, is used in successive periods.[10, 11]

[10] That said, some readers may prefer to bypass the use of 'Machines I and II', and may be willing to appeal instead to the idea that outcomes in setting I are stochastic (probabilistic) but in setting II, agents face an 'arbitrary sequence of 1s and 0s'. Everything that follows can be built on this foundation, so long as it is accepted that this phrase implies that a rational agent has no basis for carrying forward any useful information from past periods, or of appealing to any probabilistic arguments in respect of future outcomes. The idea here is simply to retain the classical notion of the fully rational decision maker, while placing that decision maker in a fully specified setting in which recourse to any appeal to certain kinds of argument available in the classic setting of 'risk' or 'probability' is logically untenable.

[11] One further point is worth noting here, in relation to the 'Schumpeterian' question that motivated our introduction of the intrinsically non-probabilistic Machine II: why can we not replace Machine II with a one-shot machine that selects an infinite sequence of 1s and 0s *ab initio*? Since this is a one-shot machine, and since it attaches unknown probabilities to each possible sequence, the environment it generates is operationally equivalent from the point of view of the decision-maker from the repeated use of Machine I. In this sense, a 'hidden probability, one-shot device' can always provide an environment equivalent to that of Machine II. Unlike Machine I used repeatedly, however, successive draws have no ergodic property, and the Bayesian updating apparatus of the subjective probability approach is not applicable. (Gilboa and Schmeidler (1993), Dow and Werlang, (2006)).

Reasonable Actions

Recall that we may appeal to the inner and outer Lebesgue measures of S to place bounds on the fraction of occurrences of 1, and of 0. With this in mind, and writing p = μ_* (S) and q = 1- μ^*(S), we introduce the following definition:

Definition ('Reasonable Actions'): The action a′ is reasonable at period t, if there exists some ('justifying') sequence of 1s and 0s, denoted $(x_1, x_2, ...)$, such that for all actions $a \in A$, the agent's expected payoff, conditional on this given sequence, satisfies $E\pi\,(a' \mid p_t, p_{t+1}, \ldots) \geq E\pi\,(a \mid p_t, p_{t+1}, \ldots)$ where $p_t = p + (1 - p - q)\, x_t$.

Remark: In the special case p = q = 0 (pure uncertainty), this reduces to the requirement that there exists some ('justifying') sequence of outcomes $(x_t, x_{t+1}, ...)$ such that for all actions a ϵ A, $\pi(a' | x_t, x_{t+1}, ...) \geq \pi(a | x_t, x_{t+1}, ...)$. Here, we can avoid introducing any explicit reference to probabilities or beliefs, and work directly in terms of the availability of a justifying sequence of outcomes.

Two remarks are in order regarding this setup:

i. The key point of departure from Savage (1954) lies in dropping Savage's axiom 6 ('measurability'). This leaves intact Savage's notion of 'qualititative probability'. Indeed, we may construct an 'objective' counterpart of Savage's concept as follows, by taking advantage of the examples described above. For instance, we may choose two versions of Machine II, the first containing the set S defined in equation (1) and the second containing the set $S' = M \cap (E_1 \cup E_2)$, whence $S' \subset S$. An agent who initially faces the first machine, on being told it has been replaced by the second machine, will consider the outcome x=1 to be (weakly) 'less likely' than before, in the sense of Savage's 'qualitative probability'.[12]

[12] On later contributions to the literature on 'qualitative probability', see, for example, Fishburn (1969, 1996).

ii. In the conventional Schmeidler approach, and related approaches, the aim is to construct a 'theory of decision', and so the agent is equipped with some preference system such as 'uncertainty aversion' in Schmeidler et al. or 'maxmin behaviour' in Garlappi et al. (2004); the effect of introducing such a restriction in the model which follows is considered in Section 3. It is difficult, however, to find any compelling or a priori justification for imposing any single behavioural postulate on <u>all</u> agents. This problem becomes acute in contexts like the one explored in the present paper, in which the pattern of outcomes is not driven by some typical, average or predominant type of behaviour among agents, but is instead critically dependent on whether a certain action is chosen by any (possibly very small) fraction of agents.

   In contrast to these conventional approaches, the aim in what follows is not to pin down some unique 'best' action. Rather, all agents are equipped with the same preferences, but the setup is such that it will not in general be the case that there is a unique 'reasonable action'; in other words, the above criterion of 'reasonableness' will lead only to a partial ordering of actions. The agent may justify the action chosen by avowing a belief that the machine will generate some specific 'justifying sequence'. This interpretation of the above definition is intuitively attractive, and is convenient in thinking about the relationship between this approach and the conventional 'subjective probability' approach, but it is not essential to the present setup, which rests directly on the above concept of 'reasonable actions'. The key intuitive idea underlying the present approach is this: we aim to construct a fully specified ('closed' or 'objective') model of the environment faced by agents, within which all agents have equal and complete access to a full description of their environment, but in which an appeal to agents' rationality cannot be used to confine agents to a unique best action.

We now add the key behavioural assumption, which ensures that there will be a diversity of choices across agents:

Denote by $\hat{A}(\subseteq A)$ the (finite) set of 'reasonable actions'. We assume that there is some distribution of 'agent types' from which agents are drawn at random. Faced with a set $\hat{A}$ of reasonable actions, each type is mapped into one chosen action.

Assumption: ('Non-degeneracy'; 'Agent Types'): there is no proper subset of $\hat{A}$ such that a fraction 1 of all 'agent types' choose actions in this subset.

Remark 1: The above assumption is natural, given the definition of 'reasonable actions'; for if both a' and a'' lie in $\hat{A}$, then there is no rational argument for favouring a'', or vice versa; and so there is no rational argument that an outside observer might use to argue that all agents will (or should) avoid the use of a' in favour of a''.

Remark 2: It is of interest to ask: what if we lower the qualitative probability of the outcome 1, either in some single period or in all periods, by replacing the set S by $S' \subset S$ in Machine II? Under the above definition of 'reasonable actions', the set of reasonable actions is thereby unchanged. This, however, leaves open the possibility that the fraction of agents choosing each action may now change. This provides an avenue for examining how changes in information impinge on outcomes in the present setup; but the development of this theme lies outside the scope of the present paper.[13]

3. Flexibility, Profitability and Survival

<u>The Model</u>

Consider a set of N identical isolated 'islands'. We will be concerned with the market, on each island, for a certain good, which can be supplied in many alternative 'varieties'

[13] In the model developed below, replacing S by S' will leave the final outcome unchanged, but may change the rate of convergence to that outcome.

(or ‘qualities’). The size of the market on each island is such that this good will be supplied, at equilibrium, by a single ‘active’ firm.[14]

Over successive periods, an exogenous shock may occur which makes feasible the production of a new variety which consumers will prefer to the variety currently offered.[15] When such a shock occurs, any firm offering the new variety completely displaces any firm offering the old variety. Within any given period, the (single) firm offering the currently favoured variety enjoys a profit of one unit.[16] A firm’s payoff is the net present value of its net revenue stream, discounted at a common discount factor δ.

The focus of interest lies in asking: will the new variety be offered by the incumbent (i.e. the single firm active in the preceding period), or by a new entrant? The answer to this question will turn on the form of investment in capability that the incumbent has made.

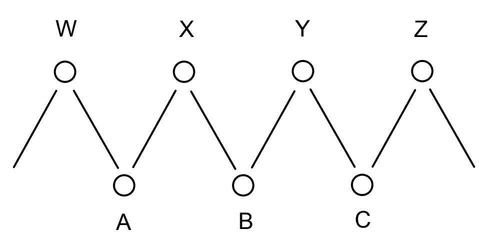

Figure 1: The relation between product varieties and underlying technologies. Learning to produce X (or Y, or Z) costs C. Mastering know-how element A (or B, or C) costs c, where $c<C<2c$.

In this example, we confine attention to a very simple mapping between the set of products (‘varieties’) and the set of underlying technologies used in their production.

[14] Alternatively, we may suppose that price competition is à la Bertrand, so that, given positive a sunk cost of entry, it is never profitable for a second supplier (or the same-‘current’ or ‘best’-variety) to enter.
[15] An alternative version of the model involves the introduction of new, higher quality, products that displace current offerings.
[16] i.e. There is a continuous flow of net revenue to the incumbent between t and t+1 whose net present value as of t equals unity.

Figure 1 illustrates the relationship. Product varieties are labeled X, Y, Z… and the underlying technologies are labeled A, B, C…

A firm aiming to produce any specific product, say X, can choose to master the know-how required directly in the production of X, at some fixed cost C, incurred over some 'learning period' τ. This is referred to in what follows as the FIX policy.

Alternatively, it can choose to master the underlying technologies ('elements of know-how') A and B. (This will be referred to as the FLEX policy). A firm choosing this latter route can master each of these elements of know-how (sequentially) in time $\tau/2$ each, and bring product X to market within the same time τ as required by a firm using the FIX policy , but the cost it incurs is c per element of know-how, or a total of 2c. We assume $c<C<2c$.[17] The return from this more costly investment accrues if and when demand switches to product Y (or W), since the incremental fixed outlay that must be incurred to produce Y is now the cost of mastering the know-how element C, i.e. a cost of c, and the time it takes to do this is $\tau/2 < \tau$. For a firm that followed the (low-cost) FIX policy in the production of X, the time taken to bring Y to market is τ and the fixed outlay it must incur to produce Y equals C (assuming it learns to produce Y directly; it may alternatively do so by mastering know-how elements B and C at cost $2c > C$ again in time τ, thereby moving to a FLEX policy).

Suppose that in some period $t > 0$ there is a switch to a new variety. Suppose an incumbent is in place who has followed the low investment FIX policy. Its incremental (minimum) fixed outlay in bringing the newly favoured variety to market (and the time required to bring it to market) is identical to that of a potential entrant. We assume that there is an infinite sea of potential entrants, and that each firm, whether incumbent or entrant, has an equal chance of being selected, so that the probability of the incumbent's remaining in place is zero. On the other hand, if the incumbent has followed the high investment FLEX policy, then its incremental cost of introducing the new product is c,

[17] If $C < c$, the FLEX strategy can never be worth using, while if $C > 2c$ the FIX strategy can never be worth using.

and it can introduce it in time $\tau/2 > \tau$. It will be feasible and profitable, under this scenario, for the incumbent to pre-empt entrants and offer the new variety.

It remains to describe the special case of period t = 0, when the market opens, i.e. there is no incumbent. Before the market opens, there is a demand announcement at time $-\tau$. The time taken to bring the product to the market is τ, irrespective of whether the firm follows a high or low investment policy. At time $-\tau$ a single firm is selected from among the potential entrants, and it may choose either a FIX or FLEX policy.

In order to focus on the main issues, we present the model in terms of this 'reduced form' description throughout, so that the only decision faced by a firm lies in choosing between the available actions FIX and FLEX, at any period when either (a) t = 0, and it is the firm selected or (b) a switch occurs, the incumbent is a FIX firm, and it is the firm selected.[18]

We now turn to the manner in which switches of demand may occur. The pattern is as follows: either demand switches from X rightwards to Y leftwards to W, or there is no switch (Figure 1). In the first case, demand may at any later date switch one step further to the right (i.e. to Z, etc.). In the second case it may in any future period switch one step further to the left. No 'reverse' switches occur in either case. We will without further loss of generality, confine the analysis to rightward switches (X → Y → Z…) in what follows.

<u>Machine I</u>

We begin by investigating the evolution of outcomes under Machine I as a function of the underlying parameters. Nothing of interest is lost by fixing the value of C to unity;

[18] A 'structural' version of the model can be formalised in a conventional way within the conventional (probabilistic) model below, by reference to pure strategy Nash equilibrium outcomes in the 'entry' sub-game beginning from any period t. (The entry sub-game is played between the incumbent and a (large) number of potential entrants. The actions open to potential entrants are 'FIX, FLEX, or Don't enter'. For the incumbent, the actions are 'insert C' or 'Don't invest'). Under the 'uncertainty' regime, 'expected profit' is not defined, however, and this extension is not straightforward.

the parameters of interest are then δ, c and p; and the range of interest of c is then $\frac{1}{2} \le c \le 1$.[19]

It is useful to think of the state of each island as the action taken by its current incumbent. As we have noted, once any island is occupied by a firm that has chosen FLEX, that firm remains the incumbent forever. We refer to such islands as being in the FLEX state; the other islands are in the FIX state. In this language, FLEX is an absorbing state.

We begin with the case where agents know the probability of a switch p, $0 \le p \le 1$. We define the expected payoff (N.P.V.) derived by entering at t=0 with a FIX, or with a FLEX, policy, as a function of c, δ and p.

A FLEX firm invests 2c > 1 at t = 0. It receives a profit of 1 in each period from t = 0. In each period from t = 1 onwards it pays a switching cost of c with probability p. Hence its NPV as of t = 0 becomes

$$E\pi_{FLEX} = -2c + \frac{1}{1-\delta} - \frac{\delta}{1-\delta}\ cp \qquad (2)$$

A FIX firm invests C=1, and it receives a profit of 1, at t = 0. In each subsequent period it continues in operation with probability (1 – p), and conditional on doing so receives a profit of 1. Hence its NPV as of t = 0 becomes

$$E\pi_{FIX} = \frac{(1-p)\delta}{1-(1-p)\delta} \qquad (3)$$

We define a schedule in (c,p) space, along which $\pi_{FLEX} = \pi_{FIX}$, for each value of δ. In what follows, we show that this schedule, which we label c (p|δ), has the following properties (the proofs of these, and all subsequent results, are given in the Appendix):

[19] If c < ½ the FLEX strategy is superior on all trajectories, while if c > 1, the FLEX strategy no longer offers a cost saving relative to FIX when a switch occurs.

- At p = 0, c (p|δ) = ½, and is increasing with p
- There is a critical value of $\delta(=(3\sqrt{5})/2)$: for δ below this critical value, c(p|δ) is strictly increasing on p ∈ [0, 1]. For δ above this critical value, c (p|δ) is single peaked on p ∈ [0, 1].

Figure 2 (a) shows the curve c (p|δ) below which the N.P.V. of the FLEX policy as of t = 0 exceeds that of the FIX policy. It is for intermediate values of p that FLEX does relatively well.[20] The point to note from this figure is that there exists a maximal value of c, corresponding to the peak of the curve, above which FIX does better for <u>all</u> p. For small values of δ, this maximal value of c, which we labeled $\hat{c}$, above lies in the relevant region [1/2, 1]. Figure 2 (b) shows the curve in (c, δ) space which separates the region where FIX does better for <u>all</u> p, and the region where FLEX does better for some p. We label this curve $\hat{c}(\delta)$.

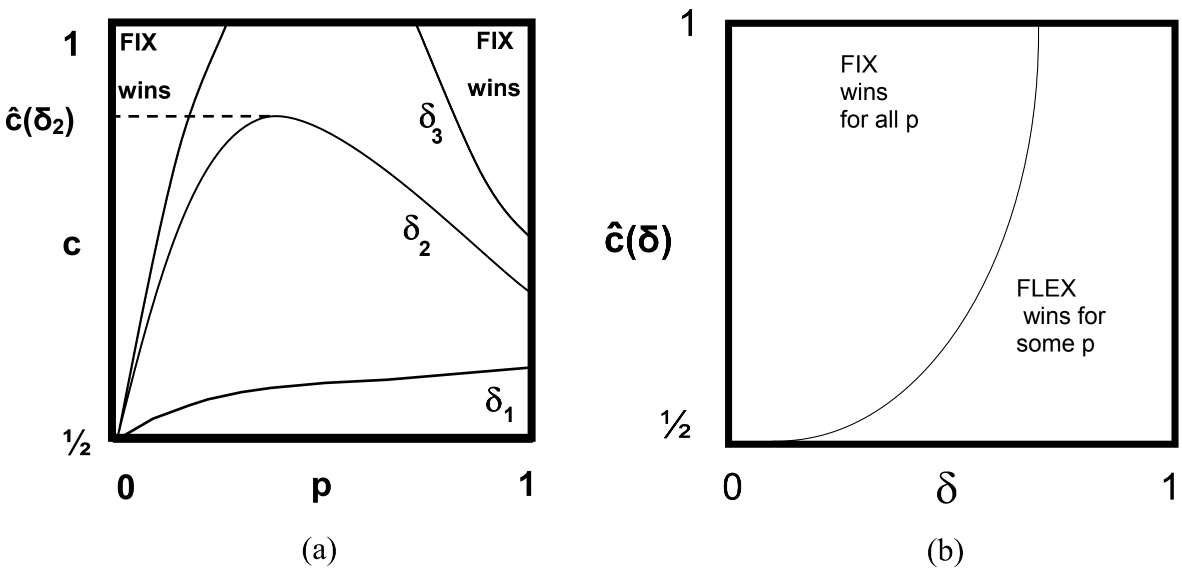

Figure 2: Optimal actions under the (probabilistic) Machine I: Panel (a) shows the curve in (c,p) space below which the N.P.V. of the FLEX strategy exceeds that of the FIX strategy, for various values of $\delta, \delta_1 < \delta_2 < \delta_3$. As δ increases, the area under this curve expands. Panel (b) shows the curve $\hat{c}$ (δ) above which FIX wins for all p, and below which FLEX wins for some p.

[20] When p is close to zero, FLEX does relatively badly as switches rarely occur, so the additional initial investment cost is not recovered. When p is high, FLEX does badly, since its continuing flow of profits is offset by continuing investment outlays.

The implications for the case in which p is known are as follows: depending on the probability of switches, and on the underlying parameters c, C and δ, either FLEX is optimal and every island is occupied in period 0 by a FLEX firm who remains the incumbent forever, or else FIX is optimal and every island is occupied in period 0 by a FIX firm, and is thereafter occupied by a continuing sequence of newly arriving FIX firms drawn from the pool of potential entrants.

So far, we have assumed that the agents know the value of p. The results illustrated in Figure 2, however, imply some immediate implications for the case where p is not (initially) known on the prior beliefs (subjective probabilities) of agents as to the value of p. If different firms have different priors, then within the right hand zone of Figure 2. (b), some firms may at least initially use the FLEX strategy, while other may use FIX. But in the left hand zone, <u>all</u> firms will use FIX, since FIX is superior to FLEX for <u>all</u> values of p. This is the only feature of the Machine I setting which is of interest in what follows. The key point is that this statement is true for the Machine I setup, independently of any assumptions that we might choose to make regarding priors and methods of updating (learning). In particular, equipping agents with non-common priors leaves this point unchanged. (For the non-common priors approach, see for example Kajii and Ui (2005), Harrison and Kreps (1979)).

<u>Machine II</u>

We now turn to the case in which Knightian uncertainty is present. Recall that we begin at t = 0 with all consumers demanding variety X. A <u>trajectory</u> (or sequence) of outcomes from t = 1 to t = T is a sequence of 0s and 1s, where 1 indicates that a switch occurred at period t. Thus (1, 0, 0…0) denotes a switch at t = 1 and no switch thereafter. Two trajectories are distinct if they differ in any element, and the number of distinct trajectories over a T-period horizon is 2^T.

An action is reasonable if there exists (at least) one 'justifying sequence' (p_t). It is convenient to write p_t in the form $p + (1-p-q)x_t$, $0 \leq x_t \leq 1$, and to refer to x_t as the 'justifying sequence'.

We begin with the characterization of a justifying sequence for the action FLEX. Consider the sequence (x_t^*) which maximizes the difference in expected profit $E\pi_{FLEX} - E\pi_{FIX}$. The action FLEX is reasonable if and only if $E\pi_{FLEX} \geq E\pi_{FIX}$ under this sequence:

Lemma 1: The sequence (x_t^*) that maximizes $E\pi_{FLEX} - E\pi_{FIX}$ takes the form (1,1,…1,0, 0, …), i.e. a finite sequence of 1s followed by an infinite sequence of 0s. For the 'pure uncertainty' case p = q = 0, it takes the form (1, 0, 0, …).

Proof: Appendix

In what follows we confine attention to the case of pure uncertainty. We begin by constructing a curve in (c, δ) space, labelled $\tilde{c}(\delta)$, along which the FLEX strategy is equi-profitable with the FIX strategy along the specific trajectory (1, 0, 0…) on which a single switch occurs in period 1 (Figure 3). It is shown in the Appendix that, to the left of this schedule, FIX wins on all trajectories, while to the right, FLEX wins on some trajectories. (Note that FIX always wins on at least one trajectory, viz. (0,0,0,…)).

The schedule $\tilde{c}(\delta)$ is constructed by reference to the profits earned along the specific trajectory (1, 0, 0…). On this trajectory,

$$\pi_{FIX} = 0 \text{ and } \pi_{FLEX} = -2c + \frac{1}{1-\delta} - \delta c$$

whence, equating $\pi_{FIX} = \pi_{FLEX}$ to define $\tilde{c}(\delta)$ we have

$$\tilde{c}(\delta) = \frac{1}{(1-\delta)(2+\delta)} \qquad (4)$$

We now state:

Lemma 2: The curve $\tilde{c}(\delta)$ lies strictly to the left of the curve $\hat{c}(\delta)$ over the range $½ < c \le 1$.

Proof: Appendix.

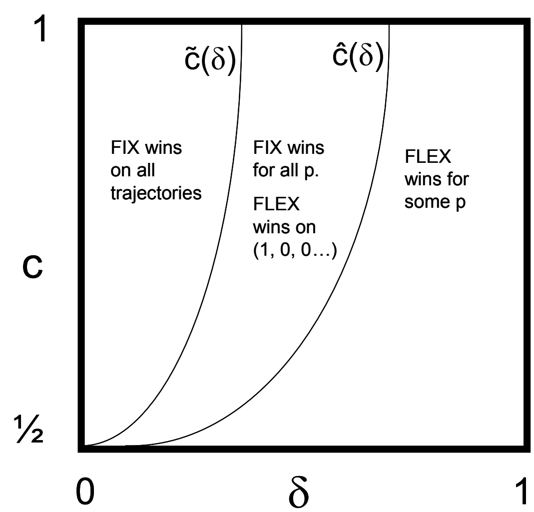

Figure 3: Comparing reasonable actions under Machine II with optimal actions under Machine I. The left hand curve $\tilde{c}(\delta)$ shows the values of (c, δ) along which FIX and FLEX yield equal N.P.V. on the trajectory (1, 0, 0…). To the right of this curve, FLEX wins on <u>some</u> trajectories. The right hand curve $\hat{c}(\delta)$ is taken from Figure 2 (b), and shows the boundary between the zone where FIX wins for all p, and that in which FLEX wins for <u>some</u> p.

The zone of interest in what follows is the middle zone, lying between the two curves. It follows that both the FLEX and FIX policies constitute reasonable actions inside this middle zone. (To see this, note that FIX always wins on some trajectories, since it necessarily wins on (0, 0, 0…), while FLEX wins on (1, 0, 0…) by construction).

The following result now follows immediately from the above 'Non-Degeneracy' assumption. Note that the 'Non-Degeneracy' assumption implies that there will be some probability $p > 0$ with which the (randomly chosen) entrant chooses FLEX at each period when a switch has occurred; and this probability is independent of the (uninformative) history of past outcomes.

Proposition 1: Under Machine II, let (c, δ) lie to the right of $\tilde{c}(\delta)$. Then, for any $\eta, \varepsilon > 0$, there exists T, such that for a fraction exceeding $1 - \eta$ of trajectories, the probability that any island is occupied by FLEX firms at any $t > T$ exceeds $1 - \varepsilon$.

Proof: Appendix

Proposition 1 implies that FLEX firms must, for almost all trajectories, eventually dominate the set of island markets. Recall that, under Machine I, in the middle zone of parameter space, only FIX firms enter the market, and Proposition 1 does not hold.

We now turn to the relative profitability of FIX and FLEX strategies. As noted already, this can only be judged on an ex-post basis.

The relative profitability of the two strategies varies as we move across the zone, and across trajectories. We first note that, in a neighbourhood of the horizontal axis (i.e. where the (c, δ) pair lies close to the line $c = ½$), then FLEX is more profitable on almost all trajectories:

Proposition 2: For any $\eta > 0$, there exists a zone in (c, δ) space such that $\pi_{FLEX} > \pi_{FIX}$ on a fraction $(1 - \eta)$ of all trajectories.

Proof: Appendix

On the other hand, as we move close to the curve $\tilde{c}(\delta)$, the FLEX strategy becomes <u>less</u> profitable on all trajectories:

Proposition 3: For any $\eta > 0$, there exists $\varepsilon > 0$ such that for all (c, δ) satisfying

$$c > \tilde{c}(\delta) - \varepsilon, \; \pi_{FLEX} < \pi_{FIX} \text{ on a fraction } (1 - \eta) \text{ of all trajectories.}$$

Proof: Appendix

The idea that emerges from Propositions 1, 2 and 3 is as follows:

i. The FLEX strategy tends to dominate FIX over time, so long as we are in the middle or right hand zones of parameter space;
ii. The ex-post profitability of the two strategies varies across these zones. Profitability does not go 'in step' with survival, but is quite a separate issue. Depending on where we are in (c, δ) space we may find that dominant FLEX firms score either higher or lower than FIX firms on overall lifetime profitability.
iii. In particular, a sharp divergence occurs in a neighbourhood of $\tilde{c}(\delta)$. Here FLEX dominates all islands for almost all trajectories; but FLEX is less profitable ex post on almost all trajectories. This extreme case does not arise in the probabilistic model. Irrespective of the value of p, the information possessed by agents, the priors they hold on p, or the way in which they update their priors, only FIX is ever played in this zone of parameter space under Machine I.

<u>Discussion</u>

In the model examined here, there is a sharp divide between profitability and survival. The FLEX firm, by investing in capability, enhances its chances of survival; to do this may or may not be profitable ex post. This point is not dependent on the distinction between machines I and II; it could be developed by reference to either of these setups so long as the parameters c and δ lie in the right hand zone of Figure 2 (b), by introducing a wide range of suitably chosen priors. Moreover, this idea extends beyond the present example ('investing in capability'; see below).

So how do machines I and II differ? The difference is that the divide between profitability and survival emerges in a much stronger form under machine II. Proposition 2 does not hold in any model of the hidden probability kind, even if we introduce such features as bounded rationality, exogenous differences in firms' attitudes to risk or incomplete information. The key result special to machine II is that FLEX firms enter and dominate even when it is <u>almost</u> <u>everywhere</u> less profitable to do so.

The intuition underlying the difference between the two models is as follows: if our underlying device is probabilistic (i.e. if we are using machine I), then there is an implied link in the likelihood of different trajectories, and as the key trajectory (1, 0, 0 …) becomes more (or less) likely to occur, so too do other trajectories less favourable to the FLEX strategy. Under machine II, we can have an 'optimistic' FLEX firm who believes that (1, 0, 0…) will be the outcome. Under machine I, the most optimistic view favouring FLEX is that of an agent who believes some 'intermediate' value of p will drive outcomes; but the attractiveness of FLEX is lower in this setting. Hence there will be a ('middle') zone in parameter space in which some firms play FLEX under machine II, but not under machine I; and once FLEX is played by some fraction of firms, it will necessarily come to dominate on almost all islands. Moreover, this remains the case even in the neighbourhood where FIX is almost everywhere more profitable ex-post.

<u>Welfare</u>

It is natural to ask which strategy is superior in terms of welfare? Given the setup chosen here, a social planner facing the (probabilistic) Machine I, and having access to the true value of p, would choose the strategy that yielded the highest N.P.V. of net revenue accruing to the full population of firms. (Consumer surplus is unaffected by the choice of FIX versus FLEX). Now the net present value of a FLEX strategy is the same for the planner as for the firm, but this is not true of a FIX strategy. For the planner, an exit event requires the incurring of a cost C = 1 by the 'replacing' firm, but the profit flow of one unit per period is maintained. It follows that for any parameter configuration (c, δ, p), if the FIX strategy is optimal for the firm, then it is a fortiori optimal for the planner. In other words, the zone of dominance of the FIX strategy for the planner is a strict subset of the zone of dominance of the FIX strategy for the firm, shown in Figure 2 (a).

The payoff from FLEX under the (probabilistic) Machine I is the same for the firm, and for the planner. The firm's payoff from FIX is given by expression (2) above (where c = 1). The planner's payoff from FIX is

$$W_{FIX} = -1 + \frac{1}{1-\delta} - \frac{\delta}{1-\delta} p = \frac{(1-p)\delta}{1-\delta} \tag{4}$$

where the last term in the middle expression represents the planner's need to pay C=1 when a switch occurs, in order to maintain the unit flow of profit per period. A comparison of (3) and (4) for any (δ, p), shows that FIX is superior in welfare terms ex ante if it is more profitable ex ante, but not vice versa.

Under machine II, the only form of comparison available is an ex-post comparison along a specific trajectory. Since the welfare score equals the profit score plus the profit (net of entry costs) earned by subsequent entrants, it follows that on any trajectory except (0, 0, 0 …), the welfare score from FIX exceeds the profit score; while the welfare score from FLEX coincides with the profit score on all trajectories. It follows that, on all trajectories except (0, 0, 0 …), FIX is strictly preferred to FLEX on welfare grounds if it is more profitable, but not vice versa. In particular, in a right-neighbourhood of $\tilde{c}(\delta)$, FLEX is (slightly) more profitable than FIX, on, and only on trajectories of the form (1, 0, 0, … , 0, ˙ , ˙, …). Here FLEX is strictly less preferred to FIX on welfare grounds.[21, 22]

---

[21] To see this, note that on the trajectory (1, 0, 0, …), $\pi_{FIX} = 0$ but the welfare score of FIX is $W_{FIX} = -\delta + \frac{\delta}{1-\delta} > 0$ where the first term in the middle expression denotes the discounted unit entry cost of the new entrant and the second term denotes its profit stream.

[22] The bias identified here disappears in the special setting of 'zero profit' environments, such as 'perfect competition with constant returns, identical firms and free entry', or 'monopolistic

What these results illustrate is that there is a bias in favour of over-investment in capability, that derives from the wedge between the firm's private return, as against the combined return accruing to the population of firms as a whole. The exit of any one firm creates a positive externality for (potential) entrants, and the planner's calculation internalizes this effect, thus leading to an increase in the planner's return from the FIX policy. The market, in this sense, generates too much 'investment in flexibility'.

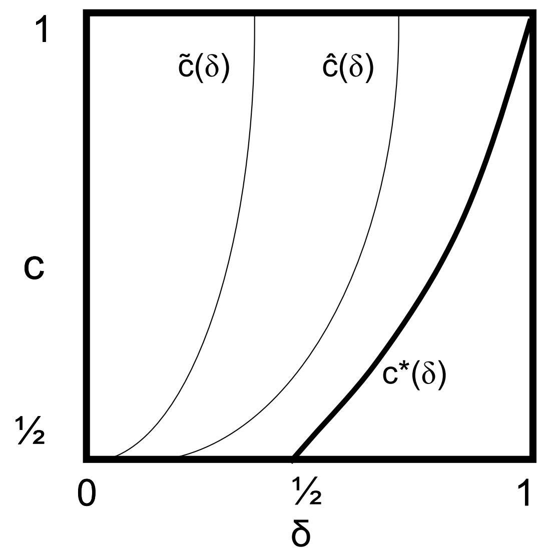

Figure 4 The minimax agent plays FLEX (resp. FIX) to the right of the schedule $c^*(\delta)$, which in turn lies to the right of $\hat{c}(\delta)$. The zone of interest is where $0 < \delta < \frac{1}{2}$, and c is close to $\frac{1}{2}$. Here FLEX is welfare optimal on almost all trajectories, but is not played.

A key feature of the present approach is that no uniform restriction of the conventional kind is placed on agents' 'attitude to uncertainty' (see remark (ii), page 13). It is of some interest, therefore, to ask: how would these results be altered if all agents followed the conventional maximin rule? Here, the agent assigns a 'worst case' payoff to each action. For FLEX, the lowest expected profit occurs on the trajectory (1, 1, 1, …) and

competition with identical firms in large markets'. It is present in such standard settings as Hotelling-type models, or markets with vertical product differentiation.

takes the value $\pi = W = -2c + (1-c)/(1-\delta)$. For FIX, the worst case is reached on any trajectory of the form (1, ·, · , …) and takes the value zero. It follows that a maximin agent plays FLEX only to the right of the schedule $c^*(\delta) = 1/(3-2\delta)$, which lies to the right of $\tilde{c}(\delta)$, (Figure 4). The zone of interest is that for which $0 < \delta < 1/2$ and c lies close to 1/2. Here FLEX is not played, yet a simple calculation shows that FLEX is superior both in terms of profitability and in terms of welfare on almost all trajectories (see Appendix). In a 'maximin' setting, the welfare bias towards excessive FLEX play observed in the present model is replaced by a bias towards excessive FIX play.

4. Summary and Conclusions

The aim of this exercise has been to isolate the relationship between flexibility, profitability and survival in a setting of Knightian uncertainty. The central conclusion is that profitability and survival do not in general go together: investing in flexibility always favours survival, but it may or may not be the case that it will turn out to be more profitable in an ex-post sense. The second conclusion is normative: firms in this model place no weight on survival per se; they are straightforward profit maximizers. Nonetheless, the market outcome is biased in favour of an over-investment in flexibility; and so in favour of a greater degree of longevity among firms than would be optimal on welfare grounds.

A caveat is in order regarding these implications. The present model has not been designed with empirical applications in mind.[23] Rather, it has been designed to isolate, in the simplest possible setting, two aspects of the profitability-survival nexus that run counter to the mechanisms and effects that have been emphasized in the literature: the idea that profitability and survival do not as a general rule go in parallel, and the idea that a bias towards over investment in capability is necessarily present whenever survival is positively affected by such investment. Empirical investigations in this area pose serious challenges in terms of research design, but it would appear appropriate to

[23] One particular empirical feature which is abstracted from in the present model relates to 'radical shocks', i.e. the introduction of new products, the know-how for whose production is non-overlapping with that of existing products. 'Rare' shocks of this type have played a key role in certain industries (Sutton, 1998).

attack such issues within a framework that is general enough to incorporate the effects noted in the present model.

A second theme of this paper has been to address the 'Schumpeterian' claim, that there is a fundamental difference between a probabilistic setting, as against one in which the firm faces a set of outcomes driven by 'unique, unrepeatable circumstances'. To address this claim, it is necessary to work in a setting in which the firm is equipped with a fully specified model of its environment, and in which we can relax the 'measurability' assumption, and so the probabilistic character of the setup.

This leads to an approach to the modeling of Knightian uncertainty that differs from the now-standard approach of Schmeidler et al. in a number of respects, and it may be worth recapitulating the points of difference:

i. The present model is closed ('objective') in the sense that it rests on the notion of an agent facing a fully specified environment, relative to which actions are chosen. The environment is such that it cannot be described probabilistically; in particular, no complete set of <u>subjective</u> probabilities can be attached to the set of draws w that lead to outcome 1 (resp. 0) without running into internal contradictions. An (arbitrary) subjective probability can, of course, be attached directly to the outcome 1 (resp. 0), but this cannot be updated (via Bayesian or other means) in a way that leads to some asymptotic limiting value).

ii. A second distinction can be made between the two approaches by noting that there are two routes we might in principle take in departing from a probabilistic setup. Any probabilistic setup requires (i) some listing or classification, and some ordering or grouping of states of the world, and (ii) the attachment of probabilities to (groups of) states. In Schmeidler et al., either certain states may be omitted at step (i) or else certain sets of states leading to different outcomes may not have probabilities attached to them at step (ii). In the present setup, the departure lies in introducing a description of the list of scenarios in step (i) that precludes the possibility of attaching probabilities to the sets of states corresponding to each outcome. Thus the intuition behind the Schmeidler et al. approach is that the agent knows that there are some possible states that he or she cannot envisage; while the intuition behind the present approach is that the set of scenarios is large, and admits of no tidying up which allows a probabilistic description. An exploration of a countable infinity of possible scenarios by an agent can lead to no settled view of the likelihood of different outcomes (see the discussion in Section 2).

iii. The final difference is one of standpoint, as noted in section 2: while the standard subjective approaches focus on decision-theoretic issues, with a view to pinning down a specific action for a particular agent by reference to the agent's objective function and 'attitude to uncertainty,' the standpoint of the present approach lies in using appeals to rationality by (expected) profit-maximizing agents to exclude candidate actions whenever possible, while noting that there then remains a set of actions that can be chosen. The focus lies, not in showing how a particular agent selects an action, but rather in exploring the way in which the operation of the market impinges on a population of agents whose actions span the set of 'reasonable actions'.

This last point is related to the point made in Section 1 above regarding the division in game-theory between the standard 'Nash' approach, and the 'evolutionary game theory' approach. The present approach is in the spirit of the standard 'Nash' approach in using profit (payoffs) as a criterion of success; yet the focus of the model investigated here is concerned with questions of survival. In evolutionary game theory, the payoff (profit) in period t is the driver of survival in period (t + 1) by construction. It is difficult, in this setting, to distinguish between the issues of profitability and survival. Tackling 'profitability and survival' within the present setup (or some alternative 'Nash' type setup) offers the advantage of separating 'survival' from 'success'; and this seems appropriate, given the dichotomy between the two ideas which the present model has been designed to highlight.

Appendix: Proofs

1. Characterization of $\hat{c}(\delta)$

To establish the properties of $c(p\,|\,\delta)$, equate expressions (2) and (3) of the text and solve for c to obtain

$$c(p\,|\,\delta) = \frac{1-\delta(2-\delta)+(2-\delta)\delta p}{(2-2\delta+\delta p)(1-\delta+\delta p)} \quad = \quad \frac{(1-\delta)^2+\delta(2-\delta)p}{[2-(2-\delta)p][(1-\delta)+\delta p]}$$

Differentiating with respect to p, we find that $c(p|\delta)$ is increasing at p = 0 for all δ and is strictly increasing on $p \in [0, 1]$, if and only if $\delta \leq (3-\sqrt{5})/2$,. If this condition holds, $c(p|\delta)$ takes its maximum over $p \in [0, 1]$ at p = 1, whence its maximum value $\hat{c}(\delta) = c(1|\delta) = 1/(2-\delta)$. For $\delta > (3-\sqrt{5})/2$, it is easily verified that $c(p|\delta)$ takes a unique maximum at some $p \in [0, 1)$. Denoting this maximum as $\hat{c}(\delta)$, this defines the schedule shown in Figure 3b.

2. Proof of Lemma 1:

Without loss of generality we label the time at which the choice is made as t = 0. Let p_t denote the probability of a switch at time t, where p_t takes the value $p_t = p + (1-p-q)\,x_t$. We write 1 – p – q as θ in what follows.

Since

$$E\pi_{FLEX} = -2c + \frac{1}{1-\delta} - c\left\{ \sum_{t=1}^{\infty} \delta^t p_t \right\} \qquad (*)$$

it follows that

$$-\frac{1}{\delta^t}\frac{d}{dp_t}E\pi_{FLEX} = c$$

Since

$$E\pi_{FIX} = \delta(1-p_1) + \delta^2(1-p_1)(1-p_2) + \quad ... \quad + \delta^t(1-p_1)...(1-p_t) + \quad ... \qquad (**)$$

it follows that

$$-\frac{1}{\delta^t}\frac{d}{dp_t}E\pi_{FIX} = (1-p_1)\,...\,(1-p_{t-1})\left\{1+\delta\,(1-p_{t+1})+\delta^2\,(1-p_{t+1})\,(1-p_{t+2})\right\} \qquad (***)$$

and the maximizing sequence (x_t) satisfies x_t = 1 or 0 according as the r.h.s. expression is greater or less than c respectively.

We first note that the maximizing sequence (x_t) must satisfy $x_1 = 1$.[24] To see this, note that the r.h.s. expression evaluated at t = 1 becomes $1+\delta\,(1-p_2)+... > 1 > c$. We next note that the sequence (1, 1, 1, …) can not be maximizing, for if $x_1 = x_2 = ... = x_T = 1$, then the r.h.s. must be less than $(1 - p - \theta)^{T-1}/(1 - \delta)$, which for T sufficiently large is less than c. Finally, we establish that if for any t, $x_t = 0$ and $x_{t+1} = 1$, then (x_t) is not an maximizing sequence. To show this, we consider the modified sequence in which $x_t = 1$ and $x_{t+1} = 0$, and we show that <u>either</u> this leads to an increase in $E\pi_{FLEX} - E\pi_{FIX}$, <u>or</u> that reducing x_{t+1} to 0 increases $E\pi_{FLEX} - E\pi_{FIX}$. To see this, we proceed as follows: from (*) and (**) it follows that the changes associated with this switch to the modified sequence $x_t = 0$ and $x_{t+1} = 1$ are:

$$\Delta E\pi_{FLEX} = -\theta c\delta^t\,(1-\delta)$$
$$\Delta E\pi_{FIX} = -\theta\delta^t\,(1-p_1)\,...\,(1-p_{t-1})$$

whence this switch is strictly improving unless

$$(1-p_1)\,(1-p_2)\,...(1-p_{t-1}) \le c\,(1-\delta) \qquad (****)$$

But in this case, reducing x_{t+1} to 0 raises $E\pi_{FLEX} - E\pi_{FIX}$; to see this, note that the associated changes are:

[24] For the special case of pure uncertainty, this implies p_1 = 1, and so it follows immediately from inspection of (***), that the optimal sequence is (1, 0, 0, …).

$$\Delta E\pi_{FLEX} = \theta\,\delta^{t+1}\ c$$

$$\begin{aligned}\Delta E\pi_{FIX} &= \theta\delta^{t+1}(1-p_1)\ldots(1-p_{t-1})(1-p_t)\ \{1+\delta(1-p_{t+2})+\delta^2(1-p_{t+3})+\ldots\}\\ &\leq \theta\delta^{t+1}\,c(1-\delta)(1-p_t)\ \{1+\delta(1-p_{t+2})+\ldots\} \qquad \text{using } (****)\\ &\leq \theta\delta^{t+1}\,c \quad \text{since } (1-p_t)<1 \text{ and } \{\bullet\}\leq 1/(1-\delta).\end{aligned}$$

2. Proof of Lemma 2:

To show that the schedule $\tilde{c}(\delta)$ lies to the left of $\hat{c}(\delta)$ it suffices to show that for any pair $(\delta, \tilde{c}(\delta))$,

$$\pi_{FLEX}(p|\tilde{c}(\delta),\delta) < \pi_{FIX}(p|\hat{c}(\delta),\delta)$$

for all p. To see this, note that $\tilde{c}(\delta)$ is defined by

$$\frac{1}{1-\delta} - 2\tilde{c}(\delta) = 0$$

whence for c = $\tilde{c}(\delta)$,

$$\pi_{FLEX}(p\ |\ \tilde{c}(\delta),\delta) = -2c + \frac{1}{1-\delta} - \frac{\delta}{1-\delta}cp < -2c + \frac{1}{1-\delta} = 0 < \frac{(1-p)\delta}{1-(1-p)\delta} = \pi_{FIX}(p\,|\,\tilde{c}(\delta),\delta)$$

4. Proof of Proposition 1:

Each ‘island’ can be regarded as a 2-state Markov process, with one absorbing state (FLEX). The ‘Non-Degeneracy’ assumption implies that FLEX will be played with probability $p > 0$ whenever an island is in the FIX state, and a switch occurs. The probability that the island remains in the FIX state following η switches is not greater than $(1-p)^n$, which is less than ε if $\eta > |\ln\varepsilon| / |\ln(1-p)|$. By choosing T sufficiently large, we ensure that the number of switches on a fraction $1-\eta$ of trajectories exceeds $|\ln\varepsilon| / |\ln(1-p)|$ (see, for example, Billingsley (1995) Chapter 1).

5. Proof of Proposition 2:

We first define a critical period T, such that at least a fraction $(1-\eta)$ of trajectories have at least one switch at or before period T; to do this, choose the smallest T satisfying $\left(\frac{1}{2}\right)^T < \eta$. We note that FLEX incurs an initial cost, relative to FIX, of 2c – 1; and, letting T denote the time of the first switch, then it earns a profit premium, relative to FIX , of $\delta^T(1\text{-}c)$. Thus a sufficient condition for FLEX to be more profitable, conditional on a switch occurring by time T, is $2c\text{-}1<\delta^T(1-c)$ or $\delta^T > (2c\text{-}1)/(1\text{-}c)$. This last inequality defines the required zone (which is a neighbourhood of the horizontal axis in Figure 2(b)).

6. Proof of Proposition 3:

We identify a trajectory with a vector Δ_t of 1s and 0s where $\Delta_t = 1$ corresponds to a switch at period t. (The initial period is labeled t = 0 and Δ_t is defined on t = 1, 2, 3…, T). Denoting the period of the first switch by T′, we note that

$$\pi_{FLEX} = \frac{1}{1-\delta} - 2c - \sum_{t=1}^{\infty} \delta^t c\Delta_t \qquad \text{(A1)}$$

$$\pi_{FIX} = \delta \frac{1\text{-}\delta^{T'\text{-}1}}{1\text{-}\delta} \text{ where } T' = \text{argmin}(\Delta_t = 1) \qquad \text{(A2)}$$

We first consider trajectories of the form (0, ·, ·), i.e. those with no switch in period 1. The fraction of trajectories of this type is one-half. Here, we define an ε such that π_{FLEX} < π_{FIX} in an ε-neighbourhood of $\tilde{c}(\delta)$. We have from (A1), (A2), that

$$\pi_{FIX} \geq \delta \text{ and } \pi_{FLEX} \leq \frac{1}{1-\delta} - 2c$$

whence

$$\pi_{FIX} - \pi_{FLEX} \;\geq\; 2c + \delta \;-\; \frac{1}{1-\delta} \tag{A3}$$

We note that the curve $\tilde{c}\,(\delta)$ is defined by the equation

$$\tilde{c}(\delta) = \frac{1}{(1-\delta)(2+\delta)}$$

whence

$$2\tilde{c}(\delta) \;+\; \delta\tilde{c}(\delta) \;=\; \frac{1}{1-\delta} \tag{A4}$$

Combining (A3) and (A4) we have

$$\pi_{FIX} - \pi_{FLEX} \geq \; 2c + \delta - [2\tilde{c}(\delta) + \delta\tilde{c}(\delta)]$$

$$\geq \; 0 \;\text{ if } \tilde{c}(\delta) - c < \frac{1}{2}\delta\,(1 - \tilde{c}(\delta))$$

We set $\varepsilon_1 = \frac{1}{2}\delta\,(1 - \tilde{c}(\delta))$, which on substituting for $\tilde{c}\,(\delta)$ and rearranging becomes $\frac{\delta}{2} \cdot \frac{1-\delta(1+\delta)}{(1-\delta)(2+\delta)}$, whence for all c,δ satisfying $c > \tilde{c}(\delta) - \varepsilon_1$ it follows that $\pi_{FLEX} < \pi_{FIX}$ on all trajectories of this form.

We now consider trajectories of the form (1, · , ·), i.e. those with a switch in period 1. Here the FIX firm exits after the initial period and earns net profit zero; whence we have from (A1) that

$$\pi_{FLEX} - \pi_{FIX} = \frac{1}{1-\delta} \;-\; c(2 + \delta + \sum_{t=2}^{\infty} \delta^t \Delta_t)$$

Denote by T" the period in which the second switch (if any) occurs, i.e. the next switch after that occurring in period 1. Note that

$$\pi_{FLEX} - \pi_{FIX} \le \frac{1}{1-\delta} - c(2 + \delta + \delta^{T''})$$

Note from the defining equation of $\tilde{c}(\delta)$ that

$$\frac{1}{1-\delta} = \tilde{c}(\delta)(2+\delta)$$

and so

$$\pi_{FLEX} - \pi_{FIX} \le (\tilde{c}(\delta) - c)(2+\delta) - c\delta^{T'''} \qquad \text{(A5)}$$

We note that the fraction of trajectories which have a switch at t = 1 but no switch during $t = 2, \ldots, T^0$ constitute a fraction $1/2^{T_0}$ of all trajectories. Choose T^0 as the smallest integer satisfying $1/2^{T_0} < \eta$, and set $\varepsilon_2 = c\delta^{T_0}/(2+\delta)$, whence it follows from (A5) that $\pi_{FLEX} - \pi_{FIX} < 0$ for all c, δ s.t. $c > \tilde{c}(\delta) - \varepsilon_2$. We now choose $\varepsilon = \min(\varepsilon_1, \varepsilon_2)$. This completes the proof.

7. The Welfare Properties of the Maximin Model

It suffices to show that FLEX is superior in welfare terms on almost all trajectories. To see this, note that a necessary condition for FIX to be superior is that $2c - 1 > \sum \delta^t (1-c)$ where the sum is taken over all t where $x_t = 1$. Define an ϵ-neighbourhood of $c = \frac{1}{2}$, viz. set $2c-1 = \epsilon$, and note that if T denotes the earliest period at which $x_t = 1$, then FIX is inferior on all trajectories except those in which no switch occurs in the first T periods, where T is implicitly defined by $\frac{1}{2}\delta^{T} = \epsilon$. (Since $\sum \delta^{t}(1-c) > \delta^{T}(1-c)\frac{1}{2}\delta^{T}$. It follows that FIX is superior on a fraction of trajectories not exceeding $\frac{1}{2^{T}}$.

References

Al-Nagar, Nabil I., Luca Anderlini and Leonardo Felli, (2006), 'Undescribable Events', Review of Economic Studies, forthcoming.

Anscombe, F. J. and Aumann R. J. (1963), 'A Definition of Subjective Probability', Annals of Mathematical Statistics, vol. 34, pp. 199-205.

Aumann, Robert (1987), 'Correlated Equilibrium as an Expression of Bayesian Rationality', Econometrica, vol. 55, pp. 1-18.

Aumann, Robert (1962), 'Utility Theory without the Completeness Axiom', Econometrica, vol. 30, pp. 445-462.

Bewley, Truman F. (1986), 'Knightian Decision Theory: Part I', Cowles Foundation Discussion Paper, no. 807, Yale University.

Bewley, Truman F. (1987), 'Knightian Decision Theory: Part II', Cowles Foundation Discussion Paper, no. 868, Yale University.

Billingsley, P. (1995), Probability and Measure, (3rd edition), New York: Wiley.

Blume, Lawrence, Adam Brandenburger and Eddie Dekel (1991), 'Lexicographic Probabilities and Choice under Uncertainty', Econometrica, vol. 59, pp. 61-79.

Burrill, C. (1972), Measure, Integration and Probability, New York: McGraw-Hill.

Dow, James and Sergio Rebeiro da Costa Werlang, (1994), 'Nash Equilibrium under Knightian Uncertainty: Breaking Down Backward Induction', Journal of Economic Theory, vol. 64, pp. 302-324.

Dubra, Juan and Federico Echenique, (2001) 'Measurability is Not About Information', Cowles Foundation Discussion Paper no. 1296, Yale University.

Fishburn, Peter C. (1996), 'Finite Linear Qualitative Probability,' Journal of Mathematical Psychology, vol. 40, pp. 64-77.

Fishburn, Peter C. (1969) 'Weak Qualitative Probability in Finite Sets,' Annals of Mathematical Statistics, vol. 40, pp. 2118-2126.

Garlappi, L., R. Uppal and T. Wang (2004), 'Portfolio Selection and Model Uncertainty: A Multi-Prior Approach,' unpublished.

Gelbaum, Bernard R. and John M.H. Olmsted (1964), Counterexamples in Analysis, San Francisco : Holden-Day.

Gilboa, Itzhak and David Schmeidler (1995) 'Case-based Decision Theory,' Quarterly Journal of Economics, vol. , pp. 605-639.

Gilboa, Itzhak and David Schmeidler (1993), 'Updating Ambiguous Beliefs,' Journal of Economic Theory, vol. 59, pp. 33-49.

Halmos, Paul R. (1950), Measure Theory, New York: Van Nostrand.

Harrison, J. Michael and David M. Kreps, (1979), 'Martingales and Arbitrage in Multiperiod Securities Markets', Journal of Economic Theory, vol. 20, pp. 381-408.

Kajii, Atsushi and Takashi Ui, (2005), 'Incomplete Information Games with Multiple Priors,' Japanese Economic Review, vol. 56, pp. 332-350.

Machina, Mark and David Schmeidler (1992), 'A More Robust Definition of Subjective Probability,' Econometrica, vol. 60, pp. 745-780.

Machina, Mark and David Schmeidler (1995), 'Bayes without Bernoulli: Simple Conditions for Probabilistically Sophisticated Choice,' Journal of Economic Theory, vol. 67, pp. 106-128.

Nelson, R. and S. Winter (1982), An Evolutionary Theory of Economic Change, Cambridge: Mass.: Belknap Press.

Savage, Leonard J. (1954), The Foundations of Statistics, New York: John Wiley and Sons Inc.

Schmeidler, David (1989), 'Subjective Probability and Expected Utility without Additivity,' Econometrica, vol. 57, pp. 571-587.

Schott, Peter and Andrew Bernard (2006), 'Survival of the Best Fit: Exposure to Low Wage Countries and The (Uneven) Growth of US Manufacturing Plants,' Journal of International Economics, forthcoming.

Sutton, John (2001a) 'Rich Trades, Scarce Capabilities: Industrial Development Revisited,' Keynes Lecture in Economics 2000, Proceedings of the British Academy, vol. 111, pp. 245-273.

Sutton, John (2001b), Marshall's Tendencies: What Can Economists Know,? Cambridge, Mass.: MIT Press.

Sutton, J. (2006), 'Competing in Capabilities', in preparation.