

Structural change in the presence of network externalities: a co-evolutionary model of technological successions^{*}

Paul Windrum^{1,2} and Chris Birchenhall³

¹ Manchester Metropolitan University Business School (MMUBS), Aytoun Street,
Manchester M1 3GH, UK

² MERIT, University of Maastricht, NL, Tongersestraat 49, PO Box 616,
6200 MD Maastricht, The Netherlands (e-mail: p.windrum@mmu.ac.uk)

³ School of Social Science, University of Manchester, Oxford Road, Manchester, M13 9PL, UK
(e-mail: Chris.Birchenhall@man.ac.uk)

Abstract. The paper uses a two-stage, multi-agent simulation model to examine the conditions under which technological successions can occur in the presence of network externalities. Data is used to identify a robust econometric model of the probability of succession. Four key factors are identified. First, the trade-off between higher direct utility from new technology goods and the network utility of old technology goods. Second, the relative innovative performance of new and old technology firms. Third, cost (price) differentials due to increasing returns in production. Fourth, the time old (new) firms have to develop their product designs prior to entry.

Keywords: Succession – Demand – Innovation – Externalities – Sail ship effect

JEL Classification: O30, C15

* The authors would like to thank Richard Nelson and Uwe Cantner for their helpful comments on an earlier draft of the paper, and for the comments provided by the anonymous referees. The usual disclaimers apply. The authors gratefully acknowledge supportive funding through SEIN-Project, European Commission's Framework 4 Programme (contract# SOEI-CT-98-1107) and through the Policy Regimes and Environmental Transitions (PRET) project funded by the Dutch Scientific Research Council (NWO).

Correspondence to: P. Windrum

1 Introduction: technological successions and network externalities

The paper investigates the conditions under which technological successions occur. It draws together two areas of research that have, by and large, previously treated as separate: competition between sequential technologies and network externalities. Interest in sequential technology competitions dates back to Schumpeter's proposition that new technologies are the fuel of long-run economic growth, and increasing welfare. Long-run economic development, he argued, occurs when an economy moves from one base technology to another. Empirical research by Grübler (1990, 1991) and Nakićenović (1986, 1991) on the changing relative market shares of sequential technologies in primary energy, shipping, and urban transport over long historical periods, clearly indicates that new technologies displace (however infrequently) established dominant technologies. In order for this to occur, a new technology must overcome the network externalities enjoyed by an established (old) technology.

Research on network externalities has tended to consider contemporaneous competitions between rival variants of the same technology (e.g. Arthur, 1989; Katz and Shapiro, 1986; Farrell and Saloner, 1985). A notable exception is David's empirical case study of the qwerty keyboard (David, 1985). The dvorak keyboard, a later and more efficient typing layout, was unable to displace the established qwerty technology due to the network externalities enjoyed by the latter. This has been challenged, theoretically and empirically by Liebowitz and Margolis (1990). They make two distinct arguments. First, they challenge the proposition that unsponsored competitions between rival variants of the same technology variants will result in the selection of an inferior variant. Second, they challenge the proposition that established technologies can then subsequently lock-out (in the manner described by David) later, superior technologies.

Arthur (1988), and David and Greenstein (1990) list a range of supply and demand side factors that may lock-out a new technology. In addition to pecuniary and non-pecuniary switching costs faced by users, scale economies in production, learning and competence creation costs of firms, and the marketing and advertising costs of establishing a market for a new technology may be significant. For each of these factors, increasing marginal returns accompanies growing market size and share. Taking these factors into account, the paper will consider the conditions under which supply and demand side externalities may be overcome by a later technology.

In order to proceed, we need to clarify the stage of technology development that we are considering here, and to distinguish between technological successions and technological substitutions. Here we are considering the factors relevant to the diffusion of a new technology. This is the third, and final, phase of technology development discussed by Schumpeter (1912, 1939). The first phase, invention, is usually related to some empirical or scientific discovery. In itself, an invention has no economic or social significance, and typically offers no hints about possible applications. For various reasons, many discoveries are never applied in the economic domain. This may be due to problems in attracting risk capital, or else a lack

of interest by entrepreneurs who may perceive the risks as too high or else see no clear market application. Alternatively, the discovery may conflict with social or religious norms (e.g. the current debate on genetic cloning). Hence, demand is low and the venture is risky.

The second phase, innovation, is the point at which the invention is actually applied for the first time, whether this be in the form of a product or a process (Mensch, 1979). Again, there are many reasons why a technology may not pass this particular stage, or else be retarded for a significant period of time. A new technology may require changes in other, upstream/downstream technologies. For example, the car required developments in materials (high quality steel sheets), in chemicals (oil refining), the development of a new way of organising production (Fordist mass manufacturing combined with Taylorist scientific management principles), and an extensive supporting infrastructure of roads, petrol stations, and repair facilities (Rosenberg, 1982; Grübler, 1998). If such changes had not occurred, then this technology would be retarded, or at best have had a localised and limited impact on mobility and productivity.

The first applications of a technology are invariably crude and inefficient. Not only is their performance usually poor compared to existing (alternative) technologies, but the (fixed) production costs are likely to be very high. Hence, innovations are not automatically capable of diffusing. They are rather like swan's eggs, requiring a (possibly long) gestation period in which further basic research and development is needed to develop them¹. If this is not possible, then many will perish. Rosenberg (1982) observes that survival of the new technology requires the establishment of a protected space in which further development can be achieved. This can take the form of distinct niche or sub-niche in the market, which may be complementary to the established technology, or else take the form of public sector support, where users are often also contributors to the R&D process. Examples of the former include railways, which did not initially compete with canals and water transport but were used as feeders to them, and steam ships, which (as tug boats) towed large sail ships into and out of port. Examples of the latter include the internet, which was invented and subsequently developed by the U.S. military, and the world-wide-web, which was invented and developed by a network of European and US universities.

In this paper we focus on the third phase of a new technology; its diffusion. Diffusion involves the widespread assimilation of a technology within a socio-economic setting (Grübler, 1998). Here we are specifically interested in the case where diffusion involves new technologies directly competing with, and successfully supplanting, existing practices and artefacts. If the development of a technology at the innovation stage can be likened to a swan's egg requiring further gestation, then it can be likened to an ugly duckling at the outset of this diffusion stage. The protection afforded by its niche has enabled the technology to be further developed and improved. However, if it is to continue developing into a swan, then it must directly compete with the dominant technology for resources.

¹ Mokyr (1990) has, rather less poetically, labelled technologies at this stage as 'hopeful monstrosities'. 'Hopeful' because they have particular features that are of interest, and 'monstrous' because of their initial crudeness and inefficiency.

Having clarified the phase of technological development under consideration, we next need to clarify the distinction between a technological substitution and a technological succession (the focus of our paper). In a technological substitution, a new technology is used in the same way as the old technology. It is adopted because it offers a superior quality/price performance in the same basic use. In terms of Lancaster's characteristics approach (Lancaster, 1971), the new technology has a superior performance in one or more service characteristics that are common to both it and the old technology. For example, the compact disk (CD) replaced the vinyl LP in domestic music systems in the 1980s. Consumers adopted the CD as a storage medium for music because it is more convenient (i.e. is smaller in size), requires less maintenance (no need to clean disks regularly), individual tracks can be accurately and easily selected (using a remote control unit), and is far less prone to degradation (i.e. scratches) than vinyl. Hardware manufacturers saw an opportunity to increase profits through sales of new CD players, while record companies realised that significant profits could be generated if, in addition to new material, consumers could be convinced to repurchase previously owned material, this time in the CD format.

In contrast to a technological substitution, a technological succession opens up new consumption possibilities. In terms of Lancaster's characteristics approach, the new technology offers one or more new service characteristics that were previously unavailable using the old technology. These may not necessarily have been evident in the invention or innovation stages, but become apparent in the diffusion stage as producers and consumers continue to experiment with the new technology, possibly combining it in novel ways with other technologies that appear in this phase. For example, the car initially competed with alternative forms of urban transport (predominantly horse-drawn vehicles such as trams)². By the 1930s it had become the dominant form of urban travel in the USA, while in Western Europe the transition occurred shortly after WWII. Subsequently, new markets emerged as a consequence of wider social, economic and political changes after the war. Notable amongst these were the rise of suburban living, and long-distance recreational and business travel. The public health movement, which had begun in the late 19th century as a reaction to pollution and disease in overcrowded cities, gave rise to the concept of suburban living as a morally and physically healthy alternative to the city (McShane, 1994). The development of suburban living after WWII saw the emergence of new type of commuter, the suburban car commuter.

Individual mobility was further heightened by the development of motorway systems after WWII. These brought the car into competition with railways for the first time. In addition to a rapid expansion in the number of road haulage companies operating nationally and internationally, another type of car-owning consumer class was born, the 'business rep'. This class has remained the staple purchaser of new car sales ever since. The post-war economic boom, which lasted from the 1950s to the 1970s, brought higher wages and leisure time to both the middle and working classes. The car became an important symbol of the emergent 'consumer culture',

² For detailed historical studies of the rise of the car see Flink (1990), Kirsch (2000), McShane (1994), and Mom and Kirsch (2001).

and motorways facilitated the boom in leisure travel. As the car example illustrates, a technological succession is associated with the emergence of new consumer classes with new preference sets. This contrasts sharply with technological substitutions, where users adopt a new technology because it better fulfils the same role, increasing consumer utility over an unchanging set of preferences.

Having clarified the focus of the current study, and having identified the supply and demand side factors that generate positive externalities for a dominant technology, we proceed, in Section 2, to consider the conditions under which a technological succession can occur. Of particular importance is the identification of factors, associated with a new technology, which either offset or outweigh the network externalities associated with the old technology. Here we take into account the insights of previous papers on the subject. Notable are Shy (1996), and Malerba et al. (1999). We also take into account the historical literature on technological displacements, particularly in land transport. This is for two reasons. First, sequential technological competitions in land transport in the US and Europe have been thoroughly researched by historians and economists. Second, it is an area of ongoing interest for economists, and understanding the dynamics of past successions may assist in debates regarding the development of more environmentally friendly and less congested urban transport.

Section 3 presents a model that incorporates the salient features of this discussion. Notably, it incorporates the increasing returns in production, and demand side externalities enjoyed by dominant technologies, and the factors that can offset these effects, thereby enabling a technological succession to occur. Important differences between this model and previous models lie in the treatment of product quality. Here quality is modelled as a multi-dimensional space of service characteristics. Another important difference is that technological diffusion is modelled as a co-evolutionary process in which both supply and demand conditions change over time as consumers and firms learn about the new possibilities associated with the production and consumption of a new technology. Given the large number of parameters involved, and the desire to examine these in a meaningful manner, Section 4 reports the findings of a large number of simulation runs conducted on the model. On the basis of these simulations we generate an econometric model for the prediction of succession. The results of this econometric modelling are detailed in Section 5, with a review and summary of the implications of the findings in Section 6.

2 Conditions for a technological succession

In this section we identify a set of potential conditions for a technological succession. An important starting point for the current discussion is Shy's 1996 paper on consumer substitution between quality and network externalities in sequential technology competitions. Shy develops a successive generations model in which there are *repeated* technology adoptions. Preferences between an 'old' consumer generation and a 'young' consumer generation can differ in the model (although preferences within each generation are assumed to be identical). It is assumed that

the new technology offers an additional characteristic not provided by the old technology, but that the two technologies are otherwise identical. In this way, the new technology is *de facto* assumed to be of higher quality. Further, it is assumed that the prices charged for the two technologies are the same. Shy finds that a technological succession can occur provided the new consumer generation treats quality and network externalities as substitutes. If, however, quality and network externalities are treated as complements, then the old technology will be selected. Shy identifies a number of necessary (but not sufficient) conditions for a technological succession. First, both a new consumer class (with a new preference set) and one or more firms producing the new technology must be present³. Logically, a succession cannot occur if only one or other is present. Second, the members of the new consumer class must positively value the additional characteristic(s) offered by the new technology. Third, the new technology must be competitive in those characteristics that are common to both technologies.

The Malerba et al. (1999) multi-agent model also emphasises the need for both a new consumer class (with a new preference set) and the presence of new technology producers. Their analysis is more sophisticated in its treatment of new and old consumer classes. Both old and new consumer classes are present, with the outcome of competition depending on how well these alternative classes are serviced by old and new technology firms. The analysis is conducted within a two-dimensional quality/price space. This can be divided into a number of compartments, each representing a different technology. The authors structure and order the quality/price space in a very specific manner. Namely, it is dividing it into two distinct compartments, one containing all the old technology designs, the other all the new technology designs. In this set-up, old technology designs are located in a compartment that is lower in quality and higher in price than new technology designs. Initially, only the old technology space is explored. Later, after a predetermined point in time, the new technology space is explored by firms and consumers. The set-up is such that, once firms and consumers search this new space a succession must occur *de facto*.

A technology space can be populated by a number of distinct niches, or 'sub-markets'. A niche is defined as containing one consumer group or 'class'. Since each class has a distinct preference set (here a particular point in quality/price space), the number of potential market niches is determined by the number of consumer classes that are initialised by the modeller. The utility functions of the classes are randomised within the overall quality/price parameters of each compartment (i.e. old and new technology), and remain fixed thereafter. Because user preferences are initialised as fixed points in quality/price space, a particular class will not become 'active' until a minimum level of quality/price performance has been reached. Once this threshold has been reached, the value that a consumer class places on a technology design becomes an increasing function of its performance and its cheapness. The utility of a design for a particular class is given by a Cobb-Douglas function in which the exponents are measures of the extent to the quality and price threshold requirements have been exceeded. Consumer utility also depends on the number

³ Also see Rosenberg (1982) and Rogers (2003) on this point.

of other users currently using the design (measured by current market share) and advertising. Together, these can generate strong lock-in effects.

Firms' profits are gross margins on production costs. Using these funds, firms can innovate and explore the quality/price space. Through innovation, the first set of firms search the pre-specified old technology quality/price space, and succeed or fail in developing designs that satisfy a consumer class. Later, at a given time in a simulation run, a new group of firms are created and the design space is opened up to include the new technology market. A new set of firms/designs is randomly distributed across a (pre-specified) new quality/price space. This new technology space differs to the quality/price space of the old technology by at least one characteristic, i.e. the characteristics are different in at least one dimension. The survival of the new technology firms also depends on the identification of a design that satisfies the price/quality preference of a consumer class. If the initial endowment of a firm is exhausted before its design meets a consumer class' minimum quality-price threshold, it will become bankrupt. Note that in this model old technology firms can survive by setting up a spin-off company to produce the new technology product.

As in the Shy model, a succession will only occur if welfare improvements due to the higher quality offers of the new technology outweigh the externality effects associated with the old technology. In addition, Malerba et al. handle user heterogeneity in a more sophisticated manner. A succession not only involves the emergence of new firms producing rival technological artefacts, but also the emergence of one or more consumer classes, with distinct preference sets, that rival the old consumer classes. Hence, a succession sees the displacement of an established production-consumption nexus by a new production-consumption nexus. Together with the discussion of niche and consumer class, this greatly clarifies our understanding of technological successions. The probability of a succession occurring is likely to depend on differences between the utility enjoyed by consumers of the new technology and those enjoyed by consumers of the old technology. This prompts the first of the research hypotheses to be tested in the model;

Hypothesis 1: the probability of a succession occurring depends on there being a positive differential between the direct utility of consuming new technology goods and the direct utility of consuming the old technology goods.

Unfortunately, a number of key issues are inadequately addressed in these models. First, there is the issue of quality differentiation between old and new technologies. Shy assumes that each technology is identical but for one service characteristic, which is only offered by the new technology. Malerba et al., meanwhile, treat quality as a simple integer value, and assume the potential quality/price performance of the new technology is always superior to the old technology. Yet, a universal measure of product quality does not exist. Perceived quality depends on the particular set of consumer preferences that exist in the market at any given moment in time, and so is subjective, temporal, subject to change (as preferences change over time), and is external to the firm. Second, innovation is treated in a highly problematic manner in the Malerba et al. model, and is completely omitted in the Shy model. Both aspects are likely to play a key role in determining whether or not a succession occurs. We

shall address these in detail in the remainder of this section. Further, they will form an important part of the formal model presented in Section 3.

Quality is a complex phenomenon in its own right. In part, perceived quality is related to the technical performance that gives rise to a set of service characteristics that are consumed by the user. In part, perceived quality depends on the relative importance that users attach to the service characteristics that are offered (Lancaster, 1971). It is rare that a new technology is unambiguously superior to its rivals, even from a technical standpoint. Indeed, historical studies of technological change indicate that later technologies tend to be better in some aspects of performance but weaker in others. Further, while a new technology may offer one or more novel service characteristics, service characteristics provided by the old technology may be missing (Rosenberg, 1982; Grübler, 1990). Hence, even from a technical standpoint, one cannot state *ex ante* that a new technology is unambiguously superior to an older technology. Take, for example, the choice between the car and the bus. Both fulfil the same function, to transport people from one geographical point to another. The car has clear advantages over the bus in terms of the greater freedom it offers users in the timing and flexibility. Not only can a car user choose when to begin a journey (at any time, day or night), but the user can decide which route they prefer to travel. In addition, users derive certain benefits (e.g. status) from the fact that a car is a private consumption good. However, the car is a relatively inefficient and more expensive means of moving large numbers of commuters in urban areas, particularly in peak traffic times. Traffic congestion results in high levels of stress for the car driver. These are not endured if one uses the bus alternative.

It is important to explicitly recognise, in any discussion of technological successions, the extent to which perceived quality also depends on users' preferences. Different consumer classes attach different valuations to the range of service characteristics provided by alternative technologies, reflecting differences in their lifestyles, interests, and values⁴. For instance, a succession could occur if a new consumer group places a high value on a new service characteristic and is disinterested in the relatively inferior or missing service characteristics. The inherent difficulty in trying to assess whether a new technology will diffuse is further complicated by the fact that the population of consumer classes changes over time. Moreover, as the car example illustrates so well, this can happen within a short space of time as the technology is diffusing. In our model, the evolution of consumer groups is endogenised and made subject to evolutionary selection. This is achieved through a replicator algorithm, which allocates differential rewards to the competing consumer classes. Classes enjoying above-average levels of realised utility will grow in size over time. By contrast, classes with below-average levels of realised utility will decline in size.

Using the model, we analyse important aspects of price left unexplored in the Shy and Malerba et al. models. Higher levels of service characteristics tend to be more expensive to produce and will be reflected in the different prices of rival designs. Given that consumer welfare depends on the relative performance of each

⁴ See important contributions on the sociology of consumption by Simmel (1957), Bourdieu (1984), and Giddens (1991).

technology (direct utility) and the price paid for each technology (indirect loss of utility), a trade-off exists between quality and price, and this is likely to influence the probability of a succession occurring. Given this, an important potential barrier to succession is the economies of scale enjoyed by established firms prior to new firms' market entry. Over time, unit production costs tend to fall as a consequence of static and/or dynamic economies of scale. As noted in Section 1, Arthur (1988), and David and Greenstein (1990) have highlighted the importance of scale economies in determining the outcome of technological competitions. An established technology may enjoy static (level of production) scale economies and/or dynamic economies of scale (due to learning by doing) garnered over time. Both, of course, are initially unavailable to a new technology entrant. This prompts a second research hypothesis;

Hypothesis 2: economies of scale in production influences the probability of a succession occurring.

From the foregoing discussion of quality and price it is clear that, while new technology designs may not be fully-fledged swans at the outset of the diffusion phase, they must, at the very least, be attractive ducklings. In order to offset the demand side externalities and increasing returns in production enjoyed by the established technology, a new technology must offer a set of service characteristics that approximately meet the preferences of one or more new consumer classes. Otherwise, production levels will be insufficient for increasing returns in production, and revenues will be insufficient to fund further innovation. Of course, old as well as new technology firms can innovate in order to improve quality/price performance. Payson's (1994) study of a range of consumer products from Sears catalogues between 1928 and 1993 found significant improvements in product specifications, even in low technology goods such as sofas and shoes. Average improvement was 2% to 3% per year, with improvements in higher technology products ranging between 7% and 9% per year. In addition, average prices fell throughout the period as markets continued to increase in size and average unit costs fell.

Another important difference between our model and the Shy and Malerba et al. models is our consideration of the so-called 'sail ship effect'. Innovation by old technology firms may be stimulated by the arrival of new technology firms seeking to displace them⁵. Indeed, it is possible that the innovative activity of established firms may result in improvements in performance sufficient to see off the challenge posed by new technology entrants. This possibility is not considered in either the Shy or Malerba et al. model. As previously noted, there is no innovation in the Shy model, while the possibility of such an effect is explicitly omitted in the Malerba et al. model because it assumes the maximum quality/price performance of the old technology is at best equal to the minimum quality/price performance of the new technology. However it is clear that the relative innovation performance of old and new technology firms in sequential technological competitions is likely to be a key factor determining whether or not a succession occurs. This leads us to our third research hypothesis;

⁵ The term was first coined by Gilfillan (1935) when referring to the rapid spurt of technical improvement in sail ships that followed the introduction of steamships in the in the 1860s.

Hypothesis 3: the relative post-shock innovation performance of old and new technology firms influences the probability of a succession occurring.

A final factor we need to consider is time. The probability of a succession occurring is likely to be influenced by two time issues. First, the length of time old technology firms have to identify (through innovation) designs that meet the preferences of their target consumer classes, and to exploit increasing returns in production prior to the diffusion of the new technology. Second, the length of time new technology firms have to develop their designs and exploit increasing returns in production before yet another technological shock. Again, this issue is not addressed by Shy or Malerba et al. Yet, one would expect that the probability of a succession occurring will be lower the longer is the first time period, and the shorter is the second period. In this paper we test the following proposition;

Hypothesis 4: the probability of a succession occurring is negatively related to the length of time old technology firms have to innovate and identify successful designs, and positively related to the length of time new technology firms have to innovate and identify successful designs.

3 A model of technological successions

In this section we outline the simulation model that is used to investigate our four research questions. The model develops and extends the Windrum and Birchenhall (1998) model to consider competitions between sequential technologies. Most notable are the inclusion of user network externalities in consumer demand, the birth of new consumer classes and new technology firms, and the assignment of preferences to new consumers and of initial designs to new firms. Numerous, more subtle changes have been made which are elucidated below.

At the heart of the model is a market made up of a heterogeneous population of consumers that coevolve with a heterogeneous population of competing firms. There are a given number of individual consumers (specified at the outset by the modeller) in each period. These are distributed across a limited number of consumer groups or 'classes'. Associated with each class is a utility function. This embodies the particular criteria used by this class to select between the competing technology artefacts offered by firms. Total utility is the sum of three components: direct utility, indirect utility and the network utility. 'Network utility' is a function of the market share of the consumer's adopted technology design, in that the more units of the adopted good sold in the market, the greater the size of the network externality.

Following Lancaster (1971), 'direct utility' is the utility gained by consuming the set of *service characteristics* embodied in a good. Lancaster observes that a commodity is not desired for itself (i.e. as an engineering solution to a particular problem) but because it provides service characteristics that yield direct utility when consumed. For example, the basic function of a car is to transport people and goods from one place to another. Speed, fuel economy, comfort of ride, storage space, aesthetic design, and pride of ownership are the service characteristics. Alternative car designs will offer slightly different trade-offs amongst these common characteristics (e.g. family cars are less fast but provide better fuel economy and

storage space than sports cars). Firms offering different variants of the same generic technology are competing by offering consumers different levels of service within a common set of service characteristics. By contrast, alternative technologies differ with respect to the set of service characteristics offered. For example, users of a bus service derive no pride from ownership (it is owned by a third party). On the other hand, they are able to hand over driving stress to the bus driver. Hence, new and old technology firms offer consumers different sets of service characteristics. In our model, a key feature of the technological shock is the extended set of characteristics made available by the new technology firms.

'Indirect utility' is the utility the consumer obtains from spending residual income – income minus price of good bought in the modelled market – on other goods. The higher the price of a good, the less money the consumer can spend in other markets and the lower is indirect utility. Note that a consumer may decide not to purchase any of the technology goods currently offered but instead use the money to purchase goods in other markets. In other words, expenditure in the modelled market has an opportunity cost. The utility obtained from making zero purchases in the modelled market (i.e. spending all disposable income in other markets) is the *null utility*. Together, direct utility and indirect utility make up the total 'intrinsic utility' gained through consumption. Intrinsic utility (relative to null utility) measures the net gain from purchasing the good. If the consumer does not purchase any of the competing goods in this market, then realised intrinsic utility is equal to the null utility.

Formally the utility function of each consumer class is modelled as

$$U(x, p, s) = d(x) + V(m - p) + e(s(x)) \quad (1)$$

where $d(x)$ is the direct utility provided by the design vector x , $V(m - p)$ is the indirect utility having purchased the good (m is initial income of the consumer and p is the price of the good), and $e(s(x))$ is the network utility. Here $s(x)$ is the current market share of design x . We assume linearly increasing network externalities with $e(s) = \gamma s$, γ being a positive constant.

In Shy's model, succession depends on how consumers treat network utility and intrinsic utility. This is also the case in our model. Whether or not a succession occurs depends on a number of other factors. First, it depends on how consumers treat the two components that make up intrinsic utility: direct utility $d(x)$ and indirect utility $V(m - p)$. Within direct utility there may be a trade-off between the novel characteristics that are only offered by the new technology, and superior service characteristics offered by the old technology (for characteristics that are common to both the old and new technologies). Second, succession is likely to depend on the rates at which old and new technology firms improve, through innovation, the quality of the service characteristics of their goods. Note that if old technology firms are more successful innovators than new technology firms, then the rate of improvement in direct utility for old technology products will be greater than for new technology products (the sail ship effect), increasing the probability that a succession will not occur. As discussed in Section 2, this possibility is excluded in the Shy and the Malerba et al. models. Third, within indirect utility, differential costs (e.g. due to scale economies) will lead to differences in the relative prices of

the alternative technology goods. Finally, a trade off exists between the quality of service characteristics (direct utility) and price (indirect utility).

In our model, consumer classes appear in the market in a random order, and each class uses its utility function to select the best combination of currently available offers⁶. The average realised utility of a consumer class depends on the distribution of offers made by competing firms at the time of purchase. The average realised utility of each class can be assessed and compared. This is the basis on which the distribution of the consumer population evolves. Using a replicator dynamic algorithm, classes enjoying above average-levels of realised utility grow in size over time, while classes with below-average levels of realised utility decline in size. Thus, utility plays a dual role in the model. It not only embodies preferences over firms' alternative technology offers, it is also the basis on which the vitality of alternative consumer classes is assessed.

Initially, all firms in the model are endowed with identical levels of capacity and wealth. Firms are heterogeneous with respect to the set of service characteristics that make up their product designs, and the consumer class that they target. These are randomly generated at the outset. In each period, every firm has a current design, a productive capacity (setting an upper limit on output), and a non-negative inventory of stock carried over from the previous period. The price of its design is determined by a fixed mark-up on the unit cost of production (i.e. prices do not adjust to clear the market). Fixed mark-up pricing is a common feature in a number of evolutionary models. Probably the best-known research in this area is Hall and Hitch (1939). Their study of 38 businesses found that the most common pricing procedure was average cost with a 'normal' mark-up. The same finding has appeared in more recent studies in the US and UK. More than half the 72 US firms (with annual revenues of \$10 million and over p.a.) interviewed by Blinder (1991) reported that cost-based pricing was a moderate or very important factor in explaining price adjustment, while 37% of respondents in the Hall et al. (1997) study of 654 UK companies use a cost-based pricing rule.

Coordination of market supply and demand occurs through quantity adjustments in this model. Firms adjust output and capacity in light of past demand. Firms compete by offering a combination of service characteristics, with a consequent price, that they believe will be more attractive than those offered by their rivals. In this way, a firm effectively offers consumers a distinct point within a multi-dimensional service characteristic/price space. Product innovation is the means by which firms search this space. Unit cost is the sum of an average fixed cost (a common fixed cost Φ which includes a fixed cost for innovation, divided by the firm's level of production y) and an average variable cost that is a function of the good's design (the vector of service characteristics offered by the design)⁷. The average variable cost of a design is thus independent of the level of production. The

⁶ Available quantities are net of previous sales to other consumer classes.

⁷ In order to simply, this average variable cost function (mapping designs on to unit variable cost) is assumed to be a fixed convex separable function that is common to all firms. The marginal cost of each service characteristic is positive and increasing. The partials of the average cost function are positive, and the diagonals of the Hessian are positive.

average total cost C is given by

$$C = (\Phi/y) + (\sum_k \gamma_k c_k(x_k)) \quad (2)$$

where γ_k are constants and the c_k are monotonically increasing, convex functions of the k th service characteristic. Firms set prices according to a simple mark up rule,

$$p_{jt} = (1 + \eta_{jt}) \times C_{jt} \quad (3)$$

where C_{jt} is the j th firm's average total cost in period t and η_{jt} is the j th firm's mark up in period t . In the current version of the model we assume there is a common and constant mark up so that $\eta_{jt} = \eta$.

At the beginning of each period, each firm offers a quantity (stock q_{jt} plus production y_{jt}) of goods with design x_{jt} at a price p_{jt} that reflects both the variable cost of producing the current design and an average fixed cost. Given sales s_{jt} and the level of production y_{jt} a firm's net revenue is

$$\Pi_{jt} = p_{jt}s_{jt} - C_{jt}y_{jt} \quad (4)$$

This profit is added to its monetary wealth M_{jt} , which follows the dynamic $M_{jt+1} = M_{jt} + \Pi_{jt}$. Successful firms, with high levels of sales and production, gain a direct advantage from their lower average fixed costs and (in turn) lower prices, making their goods more attractive to consumers. In this model, the growth of productive capacity is financed from initial wealth or profits, so a firm with relatively high levels of sales, and relatively high profits, will be able to finance a higher growth of capacity. Loss making firms, by contrast, will initially use up their monetary wealth and, once this is exhausted, thereafter finance themselves by reducing, that is selling, capacity. Once capacity becomes exhausted, the firm is bankrupt and exits the market.

As noted, each firm is randomly assigned a target consumer class. Its design strategy is to maximise the utility function of this target class. To simplify, we assume that each firm knows what this function is. We also assume that firms cannot switch to other consumer classes, but multiple firms may target the same class. Firms targeting the same consumer class clearly compete with each other. In addition, competition between firms occurs across classes because consumers are *not* restricted to buying from firms that target that particular class, but are free to purchase the goods of any firm. Having said this, the evolution of consumer classes (with relatively well serviced classes growing as a proportion of the total population over time while poorly serviced classes decline) is such that firms whose target consumer class are declining will lose market share over time.

Given that all firms face the same underlying technology and cost function, and that firms have fixed mark-ups, the only means by which firms can improve their competitive position is through innovation. Product innovation involves the creation and evaluation of new designs in each period. New designs are created through a combination of imitation (of the service characteristics of successful rivals) and through the firm's own R&D activities. These are modelled using a

modified genetic algorithm⁸ (see Windrum and Birchenhall, 1998; Windrum, 2004). As a consequence of performing R&D, there is a random mutation in one or more service characteristics. In the evaluation process, the firm uses its knowledge of the utility function of its target class to determine whether the proposed design should be put into production or whether, alternatively, the existing design should be retained. In other words, a firm will only implement the proposed design if it yields a higher utility to the target class than the current design.

It is worth emphasising, once again, that a ‘design’ is a particular point in the service characteristic space and *not* a point in an engineer’s technical space. Thus, even though we have simplified by assuming each firm knows the utility function of its target consumer class and that it can, in principle, identify an optimal design for that class⁹, firms do not necessarily know how to implement an optimal design. To be more precise, the filtered process of imitation and mutation means firms are climbing hills in design space, but there is no guarantee that designs will converge on to optima. Furthermore, there is no guarantee that the socially optimal combination of consumer class and design will emerge from the process. The socially optimal consumer class can be defined to be that class with the highest utility. If all firms produced optimal designs for their respective target class at the outset, and each class has at least one firm targeting it, then selection (here modelled using a replicator algorithm) would work to select the socially optimal combination of class and firm design over time. However, if there are no firms that target this class, or else the initial set of designs are very poor (such that the class is poorly serviced and starts to decline), then the system can converge to a sub-optimal firm/consumer class combination.

As implemented in this model, the replicator governs the distribution of consumers across classes. Let $\rho_{it} = G_{it} \div G$, where G is the total consumer population and G_{it} is the number in class i at time t , so that ρ_{it} is the proportion of the consumer population belonging to the i th consumer class in period t . Given the average utilities w_{it} for the i th class in period t , the new distribution ρ_{it+1} is calculated as

$$\rho_{it+1} = \rho_{it} \times (w_{it} \div W_t) \quad (5)$$

where $W_t = \sum_k (\rho_{kt} \times w_{kt})$ is the average level of utility across all current classes. Classes with above-average utility grow as a proportion of the total distribution, while groups with below-average utility decline (i.e. they have a negative growth rate). A fundamental feature of human development is the adoption of criteria used to shape decisions. Preferences and values are not genetically determined but are the product of social development. Use of this replicator dynamic on consumer

⁸ Some may suggest our algorithm is too different from a GA to be called a modified GA. We use this description to clearly signal the origins of the algorithm. The key feature of the algorithm is that cross-over and selection are replaced by a selective transfer; one firm adopts part or all of a second firm’s design, but the second firm does not have to adopt the matching design, or part design, of the first.

⁹ The model assumes that the whole of the service characteristic space is technically feasible, and that this optimal design is the one that maximises the intrinsic utility of the target consumer class. Recall that improving a service characteristic increases direct utility but also increases cost and price, thereby reducing indirect utility.

classes is a relatively simple means of building this observation into our model. Part of human adaptation is adapting preferences and values in light of individual experience and the experiences of our social peers.

In the absence of shocks, this replication process almost always reduces the number of classes to one in the limit. In turn, this dominant consumer group will consistently select out the design that yields the greatest total utility. Further, since firms do not compete through variable mark-ups in this model, the process almost always leads to a single dominant design¹⁰ and, typically, to a monopolistic firm¹¹.

To investigate the possibility of succession – the breaking away from an established dominant firm/consumer class combination – we introduce a technological shock. This introduces new consumer classes, firms and designs into the model. Picking up on the issues raised in Section 2, a key feature of new technology designs is their ability to offer new characteristics that are not available with the old technology designs. This is implemented in the model as an extension of the service characteristic space. The other defining features of the shock are the arrival of new consumer classes that put a positive weight on these new characteristics, and the simultaneous arrival of new firms that develop the new technology designs and target the new consumer classes.

Following the creation of the new classes, firms and designs, the dynamics of the model runs as before, i.e. the distribution of consumers across classes is governed by the replicator and firms grow and decline as before. New firms and new consumer classes are disadvantaged by the fact that old technology firms have low average fixed costs and old consumer classes enjoy high network utilities. However, a new firm may start with a design that offers sufficiently high utility to a new consumer class. Further, the new firms can innovate to improve their designs. Success depends upon making the total utility of their target classes greater than average population utility, inducing a virtuous circle of growth and eventual succession. Still, it should be remembered that old technology firms are also able to innovate, and thereby able to see off the challenge of the new entrants. One possibility is that new firms start well but eventually wither away as old firms re-establish dominance through ongoing innovation. Consequently, the probability of a succession occurring not only depends on the initial set of new firm designs and new consumer class preferences, but also the relative innovative performance of new and old firms. To put it another way, the outcome is not determined by the initial ‘shock’ itself. A succession is the outcome of a stochastic process of innovation, conditioned by the initial shock and by the state of the old firms and consumer classes at the time of the shock.

This raises a problem. While it is possible to identify *sufficient* conditions for succession¹², a full analytical specification of the functional relationship between the immediate post-shock state and the probability of succession is non-trivial. The best one can do is to investigate the relationship between the *probability* of

¹⁰ Competition through design and price allows two design-price combinations to coexist.

¹¹ The probability of one firm perfectly replicating a second is zero and (in the long run) the replicator dynamic selects on any non-zero difference.

¹² For example, if a new firm-consumer class pairing yields an initial total utility that exceeds the optimal total utility over all the old firm-consumer class pairings, then succession will occur.

succession and the conditions found at the time of a technological shock. In the next section of the paper we formulate a *predictive* econometric model of a succession that maps the variables observed at the time of a shock on to the probability of a succession occurring. We then present the predictive model derived from a large number of runs of our simulation model. In this way we can generate robust models that meet the requirements of sensitivity analysis (see Windrum, 2004).

4 Simulations and an econometric model of succession

Our simulation model is dependent on a set of parameters that are assigned values at the beginning of each simulation run. For example, all consumers in the model have the same income at the beginning of each period. This common level of income is set before each simulation run. Let θ represent the vector of these parameters; note this includes a specification of the shock. Given θ , the probability of succession is, in principle, determined; let $p(\theta)$ be this probability. Furthermore, we can in principle calculate the probability of a succession occurring, conditional on these parameter values and the state ω of the model at the time of the shock; let $p(\theta, \omega)$ be this conditional probability. If we wish to predict succession immediately after the shock, the best possible predictor is $p(\theta, \omega)$. Ideally we would find the analytical structure of $p(\theta, \omega)$, but given the difficulty in identifying $p(\theta, \omega)$ we adopt a different approach. First, we perform simulation runs at a number N of points in the parameter space. For each simulation run we observe the values of the parameters, a set of “summary statistics” z at the time of the shock as well as observing whether or not succession occurs at the end of the run.

On the basis of these observations, we next estimate a logistic model that maps a *selected subset* of the observed values (θ, z) on to a probability of succession. The econometric procedures used here are those developed for forecasting business cycle regimes in Birchenhall et al. (1999)¹³. There are two key ideas driving this methodology. When constructing a predictive model it is important to identify a robust relationship. In choosing the econometric model, we do not aim to maximize statistical fit with a given sample of observations, as this can lead to distortion of the model by fitting to noise or sample specific variations. Researchers using artificial neural networks will be familiar with this issue of over-fitting. We need to proceed with some caution in the search for robust relationships. Unless we are confident that we have a model that encompasses the ‘true’ model, we must treat standard errors and standard t -tests with caution. The procedure adopted here, as in Birchenhall et al. (1999), is based on a useful result by Sin and White (1996). If we choose our model so as to maximize an appropriately penalized likelihood function then, asymptotically, we will select the ‘best’ model *even if all models are misspecified*¹⁴. The specific penalized likelihood function we use is the negative of the Schwarz Information Criterion (SIC), i.e. we choose that subset of parameters and summary statistics (θ, z) that gives the smallest SIC value. Although this procedure has

¹³ Also see Birchenhall et al. (2001) and Sensier et al. (2004)

¹⁴ The selected model is best in the sense that the selected model will be the ‘nearest’ to the true model, i.e. it will have the smallest Kullback-Liebler distance from the true model.

desirable asymptotic properties, we further test the robustness of the ensuing models by inspecting out-of-sample predictions.

The reader should be aware of a nuance of the procedure that makes it less than mechanical. With a penalized likelihood function, every extra free parameter (i.e. free parameter on an included variable) incurs a penalty, and thus free variables are only retained if the increase in the likelihood exceeds the penalty. These penalties tend to favour parsimonious models and prior restrictions on parameters can improve the performance of selected models. To date, we have not attempted to specify the space of possible restrictions, and the choice of restrictions is based on prior judgement as much as statistical evidence. For example, two summary statistics that we use in our prediction model are the utilities of old and new consumers. Rather than leave them as independent variables, we have chosen to use their difference on the understanding that this is important in driving the replicator dynamic (see Sections 2 and 3 above).

While our aim is to approximate the theoretically best predictor $p(\theta, \omega)$, we recognise that this probability will be very small in a good part of the (θ, ω) -space. That is to say, for many – if not most – parameter values θ and immediate post-shock states ω , the probability of succession will be small. Approximating $p(\theta, \omega)$ in these regions is highly problematic since successions will constitute a small fraction of any sample, i.e. information on successions in these regions is likely to be low. To largely circumvent these issues, we have restricted the analysis to a region in the parameter space where successions are not rare. Hence, the estimated logit model is an approximation of $p(\theta, \omega)$ within a restricted domain. Some 1000 simulation runs on the model were conducted. Of these, succession occurred in 456 cases. We suggest that this approach is the same as that found in empirical studies. These concentrate on cases where successions occur, and attempt to understand the factors that have influenced this outcome. Clearly, one should be explicit about this methodology and temper one's conclusions in the light of these restrictions, i.e. that the results are valid within a restricted domain.

One detail is usefully noted before turning to the results. Prior to estimation of the parameters in the logit model, the variables are standardised by subtracting the sample mean and dividing by the sample standard deviation. This makes parameters comparable, in the sense that their size reflects the impact of a standard deviation change in the associated variable.

The logit models take the form

$$P = lf(\alpha + \beta'v)$$

where P is the probability of succession, $lf(r) = \exp(r)/[1 + \exp(r)]$ is the logit function, α is a constant and $\beta'v$ is a linear combination of the selected variables v . Alternatively, we can express this relationship using the following log-odds ratio

$$\ln\left(\frac{P}{1-P}\right) = \alpha + \beta'v.$$

Recalling that the selected variables v are standardised forms of the simulation parameters and the post-shock state variables, we can see that the β_i associated

with the i th variable v_i measures the impact of a standard deviation increase in v_i on the log-odds ratio.

5 Results

The simulation model was run 1000 times. In the logit model reported here the initial, pre-selection, list of explanatory variables were 22 in number, and were a subset of the 36 variable parameters and post-shock state variables recorded at the end of each simulation. The remaining 14 simulation variables were redundant or had consistently been selected out in previous logit analysis. At the beginning of each run, the random number generators used in the model were initiated with new, random seeds and all variable parameters in the model were assigned random values. For example, in each run both the number of periods T_1 before and the number of periods T_2 after the technological shock are randomly set. The common fixed cost is set randomly at the start of each run. A full list of included variable parameters and state variables is provided in the Appendix.

To clarify the structure of the selected model, we recall that a consumer's utility function has three components,

$$U(x, p, s) = d(x) + V(m - p) + e(s)$$

where $d(x)$ is the direct utility of design x , $V(m - p)$ is the indirect utility dependent on income m and price p and $e(s)$ is the utility arising from network externality and s is the market share of the good. The principal summary statistics we use to assess the technological shock is the difference between the utility of old and new consumer classes immediately after the shock¹⁵. This difference can be decomposed into three parts; namely, the difference in direct utilities, the difference in indirect utilities, and the difference in network utilities between old and new classes. These differences measure the extent to which the consumer population as a whole is enjoying higher levels of utility immediately after the technological shock. New consumer classes will do relatively well if the service characteristics (direct utility) they enjoy are relatively higher than old consumer classes and/or the cost and price (indirect utility) of the new goods are relatively lower than for old consumer classes. Given that the new technology has been just introduced, their network utility will, by definition, be lower.

Table 1 provides information on the variables that were selected out of the original 22 variables, together with the estimated coefficients (α, β) of the selected logit model. Note that, rather than using the independent time variables T_1 and T_2 , we imposed a restriction by using the time ratio T_2/T_1 ¹⁶. Of the 22 original variables, only 4 variables were selected by the SIC criterion. These are: the time

¹⁵ To be more precise we find the old consumer class with the highest within-class-average utility and the new consumer class with the highest within-class-average utility and subtract the old utility maximum from the new utility maximum.

¹⁶ Prior experiments using independent time variables tended to generate coefficients that were approximately equal in size and of opposite sign, suggesting the use of the ratio. By contrast, the difference in network utility and the difference in indirect utility were not found to be significant.

Table 1. Logit results

Coefficient	Value	Statistic	Value
Constant	-0.3071	SIC	848.5
Δ DirectU	1.7610	Error1S	26%
NetExt (γ)	-0.6534	Error2S	17%
Time ratio	0.3981	Error1O	16%
Fixed cost	-0.3309	Error2O	17%

ratio, the difference in direct utility between new and old consumer classes, the network utility associated with the old technology, and fixed costs. Examining Table 1, we see that the constant (α) is -0.3071 . Hence, if all other variables in the model are zero (i.e. pre-standardised variables are equal to their sample mean) then the predicted log-odds ratio would be -0.3071 . This corresponds to a probability of succession equal to 0.4238¹⁷. The coefficient on the difference in direct utility (Δ DirectU) is 1.7610 and larger (in absolute value) than the others. If all variables are initially at their (raw) sample mean, so that the predicted probability of succession is 0.4238, then an increase in Δ DirectU by one sample standard deviation will raise the log-odds ratio to $-0.3071+1.7610=1.4539$, or a predicted probability of succession of 0.8105. The reader is invited to do similar calculations for the other three variables in Table 1.

Let us consider the signs of the selected variables. A positive Δ DirectU implies new consumers give a higher assessment of the quality of the available offers than the assessment old consumers give to the available offers. Thus, the new consumer groups will tend grow faster (or decline less quickly) as a proportion of the population. Since new consumers tend to prefer the offers from new firms, new firms have a higher probability of gaining dominance. The positive sign on Δ DirectU is consistent with this expectation, and with hypothesis 1 in Section 2.

If the network externality effect were stronger (i.e. γ were higher) than Δ DirectU, then we would expect the probability of succession to be low. Further, the fact that the estimated coefficient on γ is negative (-0.6534) is consistent with our prior expectation that the network externality enjoyed by the old technology has a negative impact on the probability of a succession occurring. Consequently, not only do consumers in the model treat intrinsic utility and network utility as substitutes (in accordance with Shy's theorem), but successions are likely to occur when the gain in intrinsic utility Δ DirectU outweighs the network utility γ .

If new firms have a relatively longer time to achieve succession – that is to say, the time ratio (T_2/T_1) is higher – then we would expect the probability of succession to be higher as new firms have greater relative time to exploit the potential of the new technology. The estimated coefficient (0.3981) is consistent with this expectation, and hypothesis 4 in Section 2. Finally, if fixed costs are higher, we would expect new technology firms to be relatively disadvantaged since they begin with a relatively

¹⁷ To calculate the predicted probability given the log odds ratio, evaluate the logit function at the log-odds ratio.

Table 2. Successions and utility differences

Succession	Positive utility difference	Negative utility difference	Sum
Yes	306	150	456
No	103	441	544
Sum	409	591	1000

low level of output. The estimated coefficient of this variable (-0.3309) is consistent with this expectation and with hypothesis 2 in Section 2.

All in all, the signs in Table 1 are consistent with our prior expectations. What we learn from these estimates is that, within the context of this particular sequence of simulations, changes in ‘quality’ are by far the most important predictor – the coefficient on $\Delta\text{DirectU}$ is almost three times larger than the next highest coefficient. Put simply, the best way for new technology firms to succeed is to hit the road running with a high quality good. ‘Technological potential’ is, by itself, insufficient. A technology that prematurely enters the market may not be given time to realise its potential. If firms do not immediately satisfy their target buyers, by producing and selling a competitive quality product, the chance of succession is greatly reduced. This is further confirmed by Table 2. From Table 2 we see that, of the 456 occasions in which a succession occurred, there was a positive difference between the utility offered by the new designs over the old designs in 306 cases.

Examination of Table 2 highlights another interesting finding. An initial positive difference between the direct utility of new over old technology designs does not guarantee a succession will occur. This is because, as noted previously, old technology firms can respond to new entrants by improving the quality of their own designs (the sail ship effect). If old technology firms innovate more successfully, i.e. increase the utility of their target consumer classes faster than new technology firms, then they will see off the challenge posed by the new entrants. Equally, it is to be noted new technology designs may succeed even if they start with a lower initial utility. The reason for this is the same: subsequent innovation by the new technology firms, if sufficiently effective and timely, can offset an initial disadvantage. Nevertheless, it should be stressed that Table 2 reinforces the proposition that the probability of succession is significantly enhanced when new technology firms enter the market with *relatively* good quality/price designs. This supports hypothesis 3 in Section 2.

Let us next consider the right-hand side of Table 1. While the Schwarz Information Criterion (SIC) value provides little information in itself, the error counts are of interest. Given the estimated logit model, one can calculate the value of the logit function for any observation, and interpret this value as the predicted probability of succession. If this probability is greater than 0.5 we can predict succession will occur. If it is less than 0.5, we predict there will be no succession. If we observe a succession but the model predicts no succession, then we have a type 1 error. Alternatively, if we observe no succession but the model predicts succession, we have a type 2 error. To test the robustness of our selected model we divided the 1000 simulations into two parts. The first 900 simulations were used to select variables and estimate the logit model. The remaining 100 observations were used to test for

out-of-sample predictive errors. Error1S is the percentage (26%) of in-sample successions the model failed to predict (% of in-sample Type 1 errors), while Error2S is the percentage (17%) of in-sample non-successions the model failed to predict (% of in-sample Type 2 errors). Error1O (16%) and Error2O (17%) are the corresponding %-errors for out-of-sample observations. We suggest the in-sample error count is respectable in that no naïve model would out perform this model. Furthermore, the out-of-sample error count is consistent with a claim that the model is robust. Indeed, the out-of-sample performance is better than the in-sample performance.

6 Conclusions

The paper contributes to our understanding of key issues in technological successions. In the previous section we focused on the identification of factors that assist in *predicting* a succession. Specifically, we modelled the probability of a succession given the state of the market immediately following a technological shock. Incremental innovation – exploiting a set of old and new technologies – follows a technological shock and thus the best predictor of succession just after a shock is a probability. Our focus was thus directed to identifying those factors which influence the probability of succession. Drawing on existing theoretical analysis and historical studies of technological succession, four hypotheses on the process of succession were identified and a rich computational simulation model of succession was used as the basis of an econometric model to investigate those hypotheses.

The simulation model presented in the paper is able to address the complex phenomenon of quality and to break down utility into two components: direct utility gained through consuming service characteristics of a particular quality, and indirect utility associated with the price paid for a particular product. This enables us to investigate questions that could not be addressed in earlier models by Shy and Malerba et al. With regards to the first of our hypotheses, the findings support the proposition that succession depends on there being a positive differential between the direct utility of consuming new technology goods and the direct utility of consuming old technology goods. Interestingly, this finding supports Liebowitz and Margolis' (1990) argument that, over the longer term, it is possible for goods offering higher direct utility to displace established technology goods¹⁸. In our

¹⁸ The findings do not, however, support the Liebowitz and Margolis suggestion that an optimal technology will be selected in each technology regime. In simulations generated by our model, it is perfectly possible for sub-optimal consumer classes and technology design variants to be selected. This is true for both the old technology regime and the new technology regime, i.e. in line with the position of Arthur (1989), Katz and Shapiro (1986), and Farrell and Saloner (1985). Hence, while a shift from an old technology to a new technology (offering higher consumer utility) leads to an increase in overall welfare in our model, there is no guarantee that the optimal product variant (amongst either the old or the new product variants) will be selected. A different path may have led to higher utilities. This is not to say that state intervention could do better. Indeed, no agent or group of agents may be in a position to identify which path is best. Recall that underlying the service characteristic space is an implicit engineer's space which is not fully understood by agents. It is not possible to map all service characteristic vectors on to feasible engineering processes that will produce those service vectors. This mapping is *partially* explored through experience and experimentation. It is possible that the issue here reflects different definitions of 'optimal'. It could be argued that the outcome in our model is optimal

model, *ceteris paribus*, successions are more likely to occur when the gain in direct utility from the new technology is high. For this to happen, consumers must treat intrinsic utility and network utility as substitutes (in accordance with Shy's theorem). Under these conditions a high intrinsic utility can substitute for low network utility and a technology succession can occur.

Exploring this finding in greater detail, the quality differential between the initial set of new technology designs and the existing old technology designs is far the most important predictor for a technological succession, with the coefficient on $\Delta\text{DirectU}$ being almost three times larger than the next highest coefficient. This indicates that the most important factor in determining the probability of succession is the relative quality of new designs *at the point of entry*: the greater the relative direct utility of initial new designs, the higher the probability of succession. In the presence of network externalities and returns to scale in production, the initial set of new technology products must offer service characteristics that closely meet the preferences of (one or more) new consumer classes. If this is not the case, then production levels will be insufficient for increasing returns in production to arise, and revenues will be insufficient to fund innovation and, hence, build market share.

This highlights an important difference between the 'innovation phase' and 'diffusion phase' of the technology life cycle. A technology may be able to survive in the innovation phase – even if it is initially inferior in many respects to the old technology – provided it shows sufficient 'promise' or 'potential' to a key group of supporters. But in the diffusion phase (modelled here) new technology goods must be able to directly compete on quality and price with old technology goods. Consumers, as in our model, are not interested in a technology's potential in this phase of the life cycle. Hence, while the initial set of product offerings may not be fully fledged swans, it is important they are, at the very least, attractive ducklings. In order to offset the network utility enjoyed by the established technology (the second most important predictor for a technological succession),

A good initial product design may be insufficient to ensure a succession occurs, however. The findings indicate 3 factors which may prevent a succession, even though the direct utility effect of the initial new designs outweighs the network utility effect of the old technology. The first of these factors is the relative post-shock innovation performance of the old and new technology firms. A strong 'sail ship effect' was identified in our findings. The entrance of new, competing firms stimulates old technology firms to innovate in order to improve the quality of their goods and, thereby, to see of the challenge posed by the new entrants. Consequently, the probability of a succession depends on the relative rates at which new and old technology firms successfully innovate. Importantly, this possibility was precluded in earlier models of Shy and Malerba et al.

given the state of knowledge of agents in the market and the financial framework we have imposed on firms. But such a response would obscure the key argument that the market may have taken a different historical path with a different end state of knowledge and a higher well being for consumers. Surely, if there is any incentive for firms to seek new knowledge and generate new technological regimes, then it must be the case that things can be better than they are, in which case we must concede that a different historical path may have led to a better outcome than the current state of affairs.

Turning to our remaining hypotheses, we have identified two further variables which affect the probability of a succession: fixed costs and time. As with the relative innovation performance of new and old firms, these help explain why a succession can fail to occur, even when the direct utility effect of the new technology outweighs the network utility effect of the old technology. Firstly, with regards to fixed costs, we note that the complex interplay between quality, cost and price is another issue that is not open to investigation in either the Shy or the Malerba et al. model. In our model, a trade-off exists between quality (direct utility of characteristics) and price (indirect utility) in consumers' utility functions. We find that new (higher quality) technology products will not displace an old (lower quality) technology if there is a significant price differential. The assumption of mark-up pricing in our model means prices reflect average costs. These are partly determined by quality (higher quality goods cost more to produce) and partly by fixed costs of production, and the costs of conducting incremental innovation.

Increasing returns to scale, as measured by fixed costs, are found to play an important role in predicting a succession. In the simulations, a fixed cost component is randomly initialised and is the same for all firms (new and old). Given that unit cost is the sum of an average fixed cost and an average variable cost, which is a function of the good's design (the vector of service characteristics offered by the design), the probability of a succession occurring is, *ceteris paribus*, negatively related to fixed cost. This is because old technology firms are able to build their market share, increase output and, hence, reduce average fixed costs prior to the new firms entering the market. This is an important finding in the paper, and one that is supported by a significant body empirical research which indicates that high start-up costs are an important factor deterring new market entrants (cf. Klepper, 1996). Consequently, there are both demand-side network externalities and supply-side production externalities that inhibit successions occurring in our model.

Time is the remaining variable identified as affecting the probability of a succession in this set of simulations. Specifically, it is found that a succession is more likely to occur (i) the shorter the time for the old technology firms to innovate and develop designs that closely match the preferences of their target consumers, and (ii) the longer that new firms are given to innovate and turn their attractive ducklings into fully fledged swans. As a final observation, drawing our conclusions to a close, we observe the extent to which our co-evolutionary model highlights the fundamental structural impact of a technological succession, when it occurs, on an economic system. A succession is more than the displacement of old technology products with new technology products. It also involves the displacement of existing consumer classes and preferences with new consumer classes with different preference sets, the displacement of established market firms by new market firms, and established structures of production by new production structures. Schumpeterian 'gales of creative destruction' are widespread in reach and deep in their impact.

Appendix

Appendix 1: Parameters of the simulation model and post-shock state variables

In Logit	Symbol	Description	Value range
	T_1	Number of periods before technological shock	RandIndex[100,200]
Yes	T_2	Number of periods after shock	RandIndex [100,600]
	$T\text{-ratio}$	T_1 / T_2	
	N	Number of characteristics in full (new) design space	8
Yes	N_1	Number of characteristics shared by old and new	RandIndex [4,7]
Yes	r	Average proportional reduction in costs for new firms	Uniform[0,0.5]
Yes	x_{\max}	Maximum value of all x_i before T_1	Uniform[0.3,1.0]
Yes	ΔDU	New/Old difference in direct utility	Post Shock State Var
Yes	ΔIU	New/Old difference in indirect utility	Post Shock State Var
Yes	ΔNU	New/Old difference in network utility	Post Shock State Var
Yes	γ	Network externality parameter: as in $e(s) = \gamma s$	Uniform[0,1.5]
Yes	M	Monetary budget of consumers	Uniform[10,50]
	G	Total population of consumers	60
		Number of consumer groups	8
Yes	S_c	Old consumer groups with a population share less than this value are replaced by new consumer groups after shock.	Uniform[0.1,0.2]
Yes	$Bias_c$	Probability that utility coefficient on x_i is non-zero	Uniform[0.3,0.8]
Yes	$Bias_o$	Probability that x_i in initial design is non-zero	Uniform[0.3,0.8]
Yes	K_0	Initial capital = capacity for all firms	Uniform[5,10]
Yes	W_0	Initial monetary wealth of all firms	Uniform[10,20]
Yes	F_c	Partial adjustment factor to temper adjustment of capacity to target	Uniform[0.25,0.5]
Yes	F_p	Partial adjustment factor to temper adjustment of production target to recorded excess demand	Uniform[0.1,0.2]
Yes	F_s	Partial adjustment factor to temper adjustment of target market share to actual share	Uniform[0.1,0.2]
Yes	S_f	Old firms with a capacity share less than this value are replaced by new firms after shock.	Uniform[0.1,0.2]
Yes	μ	Probability of design mutation	Uniform[0.2,0.4]
Yes	κ	Mutation size factor	Uniform[0.05,0.1]
Yes	μ_X	Probability of selective transfer	Uniform[0.5,1]
Yes	Φ	Fixed cost of production in each time period	Uniform[1.0,3.0]

Uniform [a,b] indicates the parameter is uniformly distributed across the interval [a,b], while RandIndex [M,N] indicates the parameter is a random integer uniformly distributed across the range from M to N.

References

- Arthur B (1988) Competing technologies: an overview. In: Dosi G, Freeman C, Nelson R, Silverberg G, Soete L (eds) *Technical change and economic theory*. Pinter, London
- Arthur B (1989) Competing technologies, increasing returns and lock-in by historical events. *Economic Journal* 99: 116–131
- Birchhall CR, Jessen H, Osborn DR, Simpson P (1999) Predicting US business cycle regimes. *Journal of Business and Economic Statistics* 17(3): 313–323
- Birchhall CR, Osborn DR, Sensier M (2001) Predicting UK business cycle regimes. *Scottish Journal of Political Economy* 48: 179–198
- Blinder AS (1991) Why are prices sticky? Preliminary results from an interview study. *American Economic Review* 81(2): 89–96
- Bourdieu P (1984) *Distinction: a social critique of the judgment of taste*. Routledge & Keegan Paul, London
- David P (1985) Clio and the economics of QWERTY. *American Economic Review Papers and Proceedings* 75: 332–336
- David P, Greenstein S (1990) The economics of compatibility standards: an introduction of recent research. *Economics of Innovation and New Technology* 1: 3–41
- Farrell J, Saloner G (1985) Standardisation, compatibility and innovation. *Rand Journal of Economics* 16: 70–83
- Flink JJ (1990) *The automobile age*. MIT Press, Cambridge, MA
- Giddens A (1991) *Modernity and self-identity*. Polity, Cambridge, UK
- Gilfillan SC (1935) *Inventing the ship: a study of the inventions made in her history between floating log and rotorship*. Folliet, Chicago
- Grübler A (1990) The rise and decline of infrastructures. *Dynamics of evolution and technological change in transport*. Physica-Verlag, Heidelberg
- Grübler A (1991) Diffusion: long-term patterns and discontinuities. *Technological Forecasting and Social Change* 39(1): 159–180
- Grübler A (1999) *Technology and global change*. Cambridge University Press, Cambridge UK
- Hall RL, Hitch CJ (1939) Price theory and business behaviour. *Oxford Economic Papers*, 2: 12–45. Reprinted in: Wilson T, Andrews PWS (eds) (1951) *Oxford studies in the price mechanism*. Clarendon, Oxford
- Hall S, Walsh M, Yates A (1997) How do UK companies set prices? Bank of England Working Paper No. 67, available at < http://papers.ssrn.com/sol3/papers.cfm?abstract_id=114948>
- Katz M, Shapiro C (1986) Technology adoption in the presence of network externalities. *Journal of Political Economy* 94: 822–841
- Kirsch DA (2000) *The electric car and the burden of history*. Rutgers University Press, New Brunswick, NJ
- Klepper S (1996) Entry, exit, growth and innovation over the product life cycle. *American Economic Review* 86(3): 562–583
- Lancaster K (1971) *Consumer demand: a new approach*. Columbia University Press, New York
- Liebowitz SJ, Margolis SE (1990) Fable of the keys. *Journal of Law and Economics* 33: 1–25
- Malerba F, Nelson RR, Orsenigo L, Winter SG (1999) History friendly models of industry evolution: the computer industry. *Industrial and Corporate Change* 8(1): 3–41
- McShane C (1994) *Down the asphalt road: the automobile and the american city*. Columbia University Press, New York
- Mensch G (1979) *Stalemate in technology: innovations overcome depressions*. Ballinger, Cambridge
- Mokyr J (1990) *The lever of riches*. Oxford University Press, Oxford
- Mom GPA, Kirsch DA (2001) Technologies in tension: horses, electric trucks, and the motorization of American cities, 1900–1925. *Technology and Culture* 42: 489–518
- Nakićenović N (1986) The automobile road to technological change: diffusion of the automobile as a process of technological substitution. *Technological Forecasting and Social Change* 29: 309–340
- Nakićenović N (1991) Diffusion of pervasive systems: a case of transport technology. *Technological Forecasting and Social Change* 39: 181–200
- Payson S (1994) *Quality measurement in economics: new perspectives on the evolution of goods and services*. Edward Elgar, Aldershot

- Rogers EM (2003) *Diffusion of innovations*. 5th edn. Simon & Schuster, London
- Rosenberg N (1982) *Inside the black box: technology and economics*. Cambridge University Press, Cambridge UK
- Schumpeter JA (1912) *The theory of economic development: an inquiry into profits, capital, credit, interest, and the business cycle*. Oxford University Press, Oxford
- Schumpeter JA (1939) *Business cycles: a theoretical, historical and statistical analysis of the capitalist process*. McGraw-Hill, New York
- Sensier M, Artis M, Osborn DR, Birchenhall CR (2004) Domestic and international influences on business cycle regimes in Europe. *International Journal of Forecasting* 20: 343–357
- Shy O (1996) Technology revolutions in the presence of network externalities. *International Journal of Industrial Organisation* 14: 785–800
- Simmel G (1957) Fashion. *American Journal of Sociology* 62: 541–558
- Sin CY, White H (1996) Information criteria for selecting possibly misspecified parametric models. *Journal of Econometrics* 71: 207–225
- Windrum P, Birchenhall CR (1998) Is life cycle theory a special case? Dominant designs and the emergence of market niches through co-evolutionary learning. *Structural Change and Economic Dynamics* 9: 109–134
- Windrum P (2004) Neo-Schumpeterian simulation models. Merit Research Memoranda; available at <<http://www.merit.unimaas.nl>>. Forthcoming in: Hanusch H, Pyka A (eds) *The Elgar companion to neo-Schumpeterian economics*. Edward Elgar, Cheltenham