Dynamic Non-Price Strategy and Competition: Models of R&D, Advertising and Location

Alison Hole

Submitted for the degree of PhD

London School of Economics

July 1997

Abstract

The dependence on past choices of present opportunities, costs, and benefits is pervasive in industrial markets. Each of the three chapters of this thesis considers a different example of such dependence affecting dynamic behaviour.

In the first chapter a single firm's present choices depend on what it has learnt from past experience. The firm is searching for the best outcome of many multi-stage projects and learns as stages are completed. The branching structure of the search environment is such that the payoffs to various actions are correlated; nevertheless, it is shown that the optimal strategy is given by a simple reservation price rule. The chapter provides a simple model of R&D as an example.

In the central model of the second chapter firms slowly build up stocks of goodwill through advertising. While many firms start to advertise in a new market, over time a successful set emerges and the others exit. The chapter explores the relative growth of firms and the determination of the number of successful ones. The chapter compares the results to those of a model in which a firm must complete all of a given number of R&D stages before being able to produce.

The final chapter considers one of the effects of urban bus deregulation in the UK: bus arrival times are changed very frequently. It is assumed that passengers do not know the timetable and once at a stop board the first bus to arrive. There can be no equilibrium in which an operator's bus arrival times are never revised: otherwise those of a rival would arrive just before and take all the waiting passengers. The chapter considers the pattern of revisions when they are costly. The chapter also shows that fares can be higher with two competing operators than with a single monopolist.

Contents

Acknowledgements

Surviving a PhD, the transition from student to independent researcher, means surmounting daunting intellectual, emotional and practical challenges. I have received incalculable support in respect of all of these: I could not have done it alone.

The first in any list of debts must be to my supervisor, Professor John Sutton, whose graduate lectures in Industrial Organisation meant that a whole generation of LSE trained economists will forever know the importance of "what is excluded by the model" and which instilled in me the firm belief that economics can and should contribute to our understanding of what is, and is not, the case. His intellectual integrity has constantly been a worthy standard to which to aspire. While this thesis does not live up to that standard, it is better for the attempt. He has provided sound practical advice at all levels: he has, moreover, done so with great kindness.

I have received invaluable encouragement and warm support during my long stay at LSE from many other members, amongst whom I must mention Max Steuer, Kevin Roberts, and Margaret Bray.

The whole experience has been made possible, fruitful and enjoyable by the interaction with colleagues: if I am ever again amongst so talented and kind a group of people I shall count myself extremely fortunate. I must especially thank Reinout Koopmans, with whom I spent many happy hours debating economics in general, and our own respective work in particular. As much as by anyone, my approach to economics and many of my ideas have been shaped by these debates. I owe a great debt to Godfrey Keller, my co-author for the first chapter. Not only would the chapter not have been written without him, this collaboration was the first occasion when research had been fun. I have also benefited enormously, both intellectually and just in making LSE a good time for me, from time spent with Mary Amiti, Mike Burkart, Abigail Fallot, Denis Gromb, Toni Haniotis, Donata Hoesch, Ana Lamo, Volker Nocke, Thomas Picketty, and Tomasso Valletti, among others.

Finally I must thank my friends and family who must often have despaired of my ever getting to the end, but who nonetheless gave freely of emotional and financial support; my husband, John, who has been through all the ups and downs with me, my parents and their partners, my grandparents, brother and sister in law.

Financial support from ESRC, FMG and STICERD is gratefully acknowledged, as is my time at Lexecon Ltd, where my belief that economics can be useful was renewed, and where I am now employed.

There are, in addition, particular acknowledgements for each chapter.

The first chapter reports joint work with Godfrey Keller. We received invaluable guidance and encouragement from our supervisors Patrick Bolton and John Sutton, and also from Margaret Bray, Denis Gromb, John Hardman Moore, Reinout Koopmans, Sven Rady and Kevin Roberts. Thanks are due also to participants at seminars at LSE, the University of Edinburgh, the University of East Anglia, the Economic Theory Workshop in Tel Aviv, the Royal Economic Society conference in Exeter, and the EARIE conference in Maastricht.

The second chapter contains some simulation results. I am indebted to

Ariel Pakes and his co-workers who developed the algorithms on which these are based, and made these algorithms freely available. I am also grateful to John Sutton for introducing me to Pakes' work, and for unstinting encouragement and advice throughout the project, to Volker Nocke and Tomasso Valletti for reading and commenting on numerous versions of the chapter, and to Reinout Koopmans and seminar participants at the EEA conference in Juan les Pins for valuable comments.

For the third chapter I owe thanks to John Sutton, who suggested I read the Select Committee Report into the effects of local bus deregulation, to Peter White, with whom I discussed the bus market and who steered me away from some blind alleys, to Pedro Marin, for earlier and relevant discussions on the airline industry, and to Volker Nocke and Tomasso Valletti for carefully reading the chapter and providing many useful comments.

Introduction

The three chapters of this thesis are concerned with the impact of the passage of time in models of Industrial Organisation.

The first considers a firm making sequential choices between a number of actions, each with a cost and an uncertain reward. The paradigm for studying such situations is Rothschild's paper on two-armed bandits [51]. Rothschild models the pricing decision of a monopolist facing an unknown stochastic demand. There are two possible prices (arms), each period the monopolist chooses which to charge (which arm to pull), and updates her beliefs about demand at the two prices in the light of the resulting sales. The central trade-off is between charging the price which maximises payoffs given current knowledge, and learning more about demand at the other price. Rothschild shows that, with positive probability, the firm will stop experimenting and charge the inferior price for ever more. A number of subsequent papers have also considered whether a firm, faced with a similar tradeoff, learns enough to act as it would have done had it known everything[1].

[1] As an example, see the paper by McLennan [42] on the persistance of price dispersion when agents must search for the best price. Aghion, Bolton, Harris and Jullien [1] have considered the conditions under which a decision maker, uncertain as to the shape of his payoff function, will obtain the true maximum payoff. This paper also contains detailed references to this branch of the literature.

A slightly different approach is taken by Weitzman [61], and by Roberts & Weitzman [50]. Although the central tradeoff considered by these authors is the same, they are concerned with the decision rule itself, with which action is optimal, rather than with whether optimal learning leads to a true maximum. Weitzman considers a problem where an agent must choose which project to implement. Before choosing, the projects can, at a cost per project, be sampled sequentially. Weitzman analyses the optimal order in which the projects should be sampled, and the circumstances in which it is optimal to stop sampling the projects, and implement one. Roberts & Weitzman look at an application to R&D in which there is a single multi-stage project. Benefits are received only at the end, and the choice facing the agent at each stage is whether to pay to resolve more of the uncertainty and bring the project closer to completion, or to abandon the project.

The first chapter of this thesis, on "exploring a branching structure", extends the work of Roberts and Weitzman. The innovation of the chapter lies in its description of the environment within which an agent makes her choices. This environment is intended as a model of the available choices during a phase of R&D. A firm which is engaged in R&D takes actions which both bring a particular avenue of research closer to completion, and provide information about the benefits of all possible products of that avenue. Moreover, an action may reveal a number of different directions in which an avenue can be pursued further: once a prototype has been completed, numerous possible improvements may become apparent. The search environment modelled in Chapter 1 captures aspects of the decision over which action to take through the way in which actions are related. Some actions can be taken only after another has been completed, and the result of one action will be informative about the benefits of completing any avenue of research proceeding from it.

The central result of Chapter 1 is that, despite the complexity of the decision environment, the optimal strategy is given by the same decision rule as is found in Robert's and Weitzmans' and Weitman's papers: namely to allocate an index[2] to each available action which depends only on what is known about avenues proceeding from that action, and to take the action with the highest index.

The second chapter turns to an issue closer to the heart of the literature on R&D: the amount invested, and its relationship to market structure. Dasgupta & Stiglitz [9], among others, have encapsulated the modern theoretical understanding of the relationship between the amount spent on R&D and market structure. Previously debate had centred on whether market structure caused R&D intensity, or vice versa. Dasgupta & Stiglitz suggested that the question was badly posed, and introduced a game theoretic framework in which both were simultaneously determined. In their model there are three stages[3]. In the first, firms decide whether or not to enter, on the basis of their (correct) expectations of the profits they will earn if they do. In the second stage those firms which entered in the first invest in R&D which reduces marginal costs, making the best decision given the number of firms which entered. In the final stage firms compete in the market place, making the best decision given the marginal costs which result from the previous R&D investment.

What matters here is that fixed investment in R&D leads to a higher

[2]the Gittins index. See the references to papers by Gittins, Glazebrook, Jones and Nash in Chapter 2, and in Whittle's paper [62]

[3]Many other authors have used 3-stage games to model market structure when firms can make some investment which affects short run competitiveness. See for example Shaked and Suttons' 1982 paper [53].

quality, or lower production cost, for each unit sold. Sutton, as part of his theoretical and empirical analysis of the relationship between advertising and market structure in "Sunk Costs and Market Structure" [54], has noted that advertising investments may behave in a similar way. Sutton's theoretical purpose, in this book, is to derive a constraint on the set of possible outcomes from consideration of a broad class of feasible models. In the case of advertising and market structure he shows that there is a lower bound to concentration which will not be violated no matter how large the market. He shows that there is strong empirical support for such a bound.

The second chapter is concerned with the dynamic behaviour of firms when investments are accumulated over time, rather than made all in one go. The *first* part of the second chapter explores a modification of one of Sutton's advertising models, in which firms accumulate quality according to a stochastic investment function. Each period firms revise their investment levels in the light of their own previous successes, and those of their rivals. There can be many firms, and the profits a firm earns depends on its own and its rivals' qualities.

Although there is no analytical solution available, simulated examples show that the most likely market structure is well defined, and depends on the model's parameters in a way which is consistent with the behaviour of Sutton's original model, so that Sutton's results are robust to at least this dynamic version. These examples also highlight some novel dynamic features. Early in a market's evolution many firms invest. Those which fall behind in the early stages stop investing and eventually drop out of the market. Later (should the number remaining in the market fall to the most likely number), firms converge: firms with lower qualities invest more. Finally, there are some states, in which many firms have high levels of accumulated quality, which are stable should they arise, but which are unlikely to do so.

This model is an example from a general class of dynamic oligopoly models described by Pakes & McGuire [47], as part of a program initiated by Ericson & Pakes [18]. This program was a response to the empirical findings of simultaneous entry and exit within markets, and of considerable flux in firms' relative positions, which sought to explain these phenomena in terms of stochastic research and exploration[4]. However, Chapter 2 also has strong links to a quite separate strand in the literature on R&D, one which has also found expression in dynamic oligopoly models: the literature on patent races.

In a patent race, more than one firm is trying to achieve the same patentable innovation. Market structure is not an issue: it is assumed that only one firm will win the patent. Patent races are interesting not so much because competition for a single patent is an important determinant of market structure in R&D intensive industries, but because races, or sequences of races, can be used as the basis for analysing the relationship between firm size and firm investment. In an early paper Harris & Vickers [33] assumed that in order to win a patent a firm must complete a number of stages[5]. The more is invested in R&D in a period, the more stages are completed that period. They considered two firms competing, and found that whichever firm had fewer stages left to complete would win the patent, while its rival would not invest at all. The winning firm invested at the rate it would have done had there been no rival.

Beath, Katsoulacos and Ulph [4] use the idea of a sequence of patent

[4]Other models addressing the same facts include those by Hopenhayn [35], Dixit [13], and Lambson [40].

[5]The stages are modelled as a continuum.

races to investigate industry evolution, specifically whether the leading or lagging firm wins the next patent[6]. As with the relationship between R&D and concentration, investment in R&D during a patent race can be used as a model of more general investment. As one example, Koopmans [39] considers a sequence of general investment opportunities and supposes that two firms compete by racing each other for each opportunity in turn. He derives results about whether the lagging firm catches up or falls increasingly behind in terms of the externality between firms from growth. Another generalisation addressing the relative growth rate of firms is that by Budd, Harris & Vickers [6]. They consider a somewhat abstract problem which is inspired by Harris & Vickers' original contribution but which differs from it in two important respects. First, a firm's state depends not on how many stages it has completed, but on how many more stages than its rival it has completed. This reduces the number of cases which need to be considered. Second, firms can earn profits whatever their current states, not just when one gets so far ahead of its rival that it can be declared the winner.

In the *second* part of the second chapter a generalised patent race is also used to explore market dynamics, and to consider the simulation results of the first part in a tractable setting. In the model each of a number of firms must complete a fixed number of staged before it is able to produce, but whereas in a patent race the first firm to finish would get the prize and all other firms would get nothing, in the third chapter any number of firms can finish and earn profits. The more firms that finish, however, the smaller their profits.

Equilibrium in this model has a number of features in common with the dynamic advertising model of the first part: there is a well defined expected

[6]See also Vickers' paper [60].

number of firms which complete all stages, some firms which complete the early stages stop investing and never finish, and for some realisations of the random variables in the model, more firms than expected finish. Model **B** is sufficiently simple for us to discern the intuition underlying these results. Firms must expect to recover costs sunk in establishing themselves in the market, and since profits fall as the number of rivals rises, while the amount spent on becoming established is fixed, the number of firms which can profitably establish themselves is restricted. Since costs are sunk incrementally and (in some equilibria) investment is stochastic, many firms may begin accumulating. If they are luckier than their rivals they will establish themselves in the market in the long run and make profits. If they fall behind early on they can stop investing without having lost too much.

The final, third, chapter considers a different sequential investment decision: choosing and changing location when location revision is costly. The empirical context is that of bus deregulation in the UK. It has long been accepted that where firms spend on R&D there will be change over time. The first theoretical task of the final chapter is to show that reversible choices of location can also give rise to change over time. Theoretical work by Foster & Golay [22] prior to the enactment of the 1980's legislation deregulating local bus services in the UK concluded that bus operators would probably adhere to stable timetables. This conclusion was in spite of evidence, for example that given by Chester writing 6 years after the 1930 Act which first introduced some control into London's bus routes [7], that before regulation had been introduced timetables were far from stable. Chester goes on to give a clear statement as to why instability would arise.

> If any operator fixed definite times, rival operators will seek to reach stopping places a few minutes earlier and take the traffic.

Subsequent experience of deregulation has confirmed that competition on urban local bus routes leads to timetable instability, and Chapter 3 of this thesis builds a formal model to show how it arises. The driving assumption is that passengers on urban bus routes tend to arrive at a bus stop at a convenient time and then board whichever bus arrives first. As noted by Chester, this means that if one company operated a fixed timetable a rival would choose to arrive just before, leaving no passengers waiting at the stop.

Once adjustment costs are introduced into the model a distinctive pattern emerges. The operator whose buses currently get fewer passengers is more likely to adjust its timetable than is its rival, and when it does so will tend to choose a new time which is just before its rival's current one, causing buses to arrive bunched together. Such bunching is another noted feature of deregulated bus routes.

The assumption on passenger behaviour also drives the second model of Chapter 3, which considers the impact on fares of competition on local bus routes. Contrary to expectations before deregulation, fares have not fallen much, and in some cases have risen [41]. In the fare model presented in Chapter 3, duopolists charge higher fares than would a monopolist. The intuition is straightforward. A duopolist which lowered its fare would increase the number of people choosing to travel by bus. Having chosen to go by bus, however, the additional passengers would board whichever bus arrived at the stop first, so that some would board the rivals bus, giving rise to a positive externality, and to underinvestment in fare reductions.

Chapter 1

Exploring a Branching Structure

1.1 Introduction

In many areas of human activity, an agent has to choose from a number of actions, each with a cost and an uncertain reward. Some of these actions are highly likely to produce a short-term gain, while others, such as gathering information to eliminate some of the uncertainty, may result in only a long-term benefit. The classic multi-armed bandit problem is a formalisation of such a situation: in each period the agent pays a unit cost to pull one of a fixed number of arms, different arms having different, unknown, and possibly interdependent pay-off probabilities; the agent's problem is to maximise the expected discounted sum of pay-offs.

In bandit problems currently in the economics literature, projects are equated with arms. There is no ambiguity about how to engage a project: with just one arm per project the only available action is to pull it. Further, taking an action leaves the number of possible actions unchanged: with still just one arm per project the only available action is to pull it again. However, many decision environments are more complex. Here we introduce a model of a more general sequential search process in which, when an action is taken in one period, several new actions become available in the next period. The set of projects and the actions available within them depend on the previous choices of the agent.

Even the classic multi-armed bandit problem resisted any general solution until Gittins and his co-workers showed, in a very general setting, that if the arms are independent (that is, pulling one arm is uninformative about other arms) then the optimal strategy is given by an index policy.[1] To each arm attach an index (known variously as a reservation price, dynamic allocation index or Gittins index) which depends on the current state of only that arm; the strategy is to pick the arm which currently has the highest index. Calculating the indices, however, can be a formidable task. In the economics literature, two notable applications of bandit problems with independent arms are by Weitzman [61][2] and Roberts & Weitzman [50],[3] in which the examples focus on cases where the reservation price is not so difficult to calculate.

Models in which the independence assumption is dropped have no simpli-

[1] See the references to papers by Gittins, Glazebrook, Jones and Nash here and in Whittle's papers [62].

[2] Weitzman considers a problem where there are several substitutable single-stage projects, which can be sampled sequentially. When the agent decides to stop searching, only one option is selected, namely the one with the maximum sampled reward.

[3] Roberts & Weitzman look at an application to R&D in which there is a single multi-stage project. Costs are additive (pay-as-you-go), benefits are received only at the end, and the choice facing the agent at each stage is whether to pay to resolve more of the uncertainty *and* bring the project closer to completion, or to abandon the project.

fying result comparable to that of Gittins to help in determining the optimal strategy. Nevertheless, the paper by Rothschild [51][4] which introduced bandit problems into the economics literature centres on an example of such a model, and he derives strong results on how much a monopolist learns about a stochastic demand function. Subsequent work on similar pricing problems[5] has abandoned the bandit terminology altogether, and indeed the usage of the term bandit *now* appears to be reserved for cases where the different arms are independent.

In this paper, we introduce a general sequential search process in which the possible actions belong to branching projects. This process generalises a standard multi-armed bandit in a number of significant ways: an action can reveal information about more than one reward; the pay-offs to various actions are correlated; and there is a natural way to talk about the diversity of rewards. We give a simple characterisation of when the independence assumption can be relaxed, but with the problem retaining the analytical convenience of the optimal strategy being determined by an index policy or reservation price rule.[6]

---

[4]In this well-known paper, Rothschild models the pricing decision of a monopolist facing an unknown stochastic demand as a two-armed bandit problem. No assumption is made that the parameters governing demand at the two prices are independently drawn and Rothschild does not derive the optimal strategies. The main result is that optimal experimentation may not result in adequate learning, that is, there is a positive probability that after some finite period the agent will settle for the inferior arm for ever more.

[5]See, for example, Aghion, Bolton *et al.* [1, section 6], and the references in their introduction.

[6]Gittins uses the example of job scheduling with precedence constraints to motivate an abstract model which is a finite horizon version of that which we present in this chapter, but without the information revelation aspects or the reward correlation which we have here [28]. Our proof of the optimality of the Gittins index policy in this set-up was arrived

A branching project is special case of a multi-action project, a project in which there may be several alternative actions which the agent can take at any one time, and where this set of available actions depends on the agent's previous choices. A Gittins index can be attached to a multi-action project in much the same way as to a single-action project. In an extension of his proof of the original result (see Gittins & Jones [29] and Gittins [27]), Whittle gives a condition under which the Gittins index policy is optimal for multi-action projects [63]; note that it does not specify the optimal action, only the project to which the action belongs. In the special case where the multi-action projects are branching projects we give a condition under which the Gittins index policy picks out not only the optimal project to engage but also the optimal action within that project. Essentially, this condition is that taking one action gives no information about actions which do not emanate from it.

The optimality of the Gittins index policy for a class of branching projects considerably reduces the problem of characterising the optimal search strategy. We use a simple model of R&D in order to demonstrate the usefulness of our result, deriving the optimal strategy in a generalised way and discussing some of its features.

In the next section, we present the example of R&D in order to illustrate some of the features which branching projects possess and introduce some notation. Then in Section 1.3 we give a formal description of the general model, and the central theoretical result as a corollary of Whittle's theorem. In Section 1.5, we apply it to the model of R&D and provide some results and examples. We conclude with a discussion and some remarks. Proofs of the main technical results are to be found in the appendices.

---

at independently and adopts what we believe is a self-contained and accessible approach.

1.2 A simple model of R&D

A simple branching project is represented in Figure 1.1 by a tree, with node 1 as its root and nodes 4 through 7 as its terminal nodes. When there is an arc from node p down to node q we say that node p is a parent of node q and that node q is a child of node p. The terms ancestor and descendant have the obvious meanings.

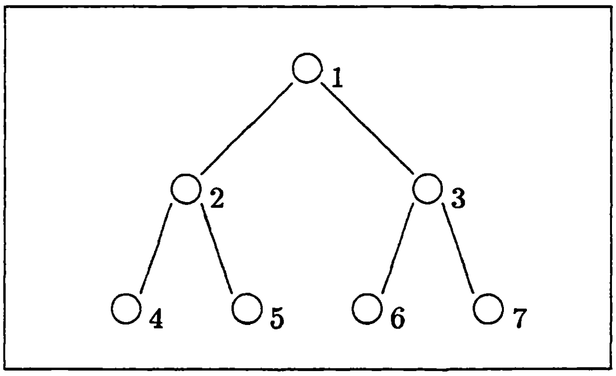

Figure 1.1: No nodes explored

The nodes correspond to possible actions, a subset of which are available in any given period. There are two sorts of possible action: one is to pay a cost c_n to explore node n and then continue; the other is to collect a prize whose value is y_n and which is located at an *explored* terminal node n, and stop. The actions which are available in any period depend on previous actions and can be summarised using the tree. We assume that initially no node has been explored, and now in any period the agent can (a) explore any node that has not yet been explored, provided that either it is the root or its parent has been explored, and then continue, or (b) collect the prize at a terminal node that has been explored and stop.

We shall often consider there being an additional fall-back option available in any period, and if it is chosen the agent collects a prize of value m and stops. For example, suppose that the situation is as illustrated in Figure 1.2, in which filled nodes have been explored and empty ones have not, and there is a fall-back. The available actions are: explore node 3, explore node 4, take y_5, or take the fall-back m.

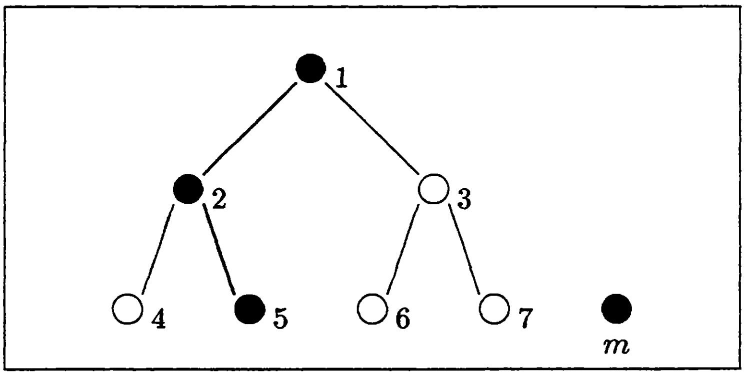

Figure 1.2: Some nodes explored, & a fall-back

In an R&D setting, node 1 might represent a feasibility study, and nodes 2 and 3 would represent two different avenues of basic research, each of which leads to two development opportunities. One would then think of nodes 4 through 7 as representing substitutable technologies to produce a product. To take the fall-back option is to use the existing technology, and abandon R&D. Note that 'production' is also a terminating action – it corresponds to stopping R&D and commercially exploiting the know-how that has been gained.

Exploring a node not only imposes costs on the agent and affects which actions are available in future periods, but also reveals information about the prizes at all its descendent terminal nodes: when the agent explores node n she receives a random signal z_n, which is independent for each node. The value of the prize at a terminal node is the sum of the signals at that node and its ancestors, so, for example, $y_5 = z_1 + z_2 + z_5$. (Because the signals contribute additively to the prize, we sometimes refer to them as *increments.*) The implication of this for the model of R&D is that each piece of basic research is informative only about products which embody that research, and that developing one product is uninformative about the value of other products. This means that, whenever the agent updates the expected value of any product, she uses only what has been learnt at its explored ancestors.

The agent's problem is to choose a strategy which maximises the expected value of the prize that she collects when she stops, net of the expected costs from exploring nodes before she collects the prize.

Note that the way in which actions become available leads to a natural measure of the diversity of prizes: those with a common parent are closer than those with only a common grandparent. Moreover, as a result of the specification of the prizes themselves, the values of closer prizes are more correlated.[7] Two features of this example worth stressing are that in any period each available action can be considered as the root of its own *separate* and *independent* sub-tree. Reconsider the situation illustrated in Figure 1.2.

We can in fact represent the agent's choice as between the projects shown in Figure 1.3 in which each project now contains only one available action: explore an unexplored root and continue, or collect a prize and stop. This representation is legitimate because all the ancestors of currently available actions have been explored, and we can use the state of each project to

[7] At the start, before the agent has received any signals, the values of all prizes are correlated random variables: they all depend on the realisation of z_1. The values y_4 and y_5 are closely correlated because $\mathrm{Cov}(y_4, y_5) = \mathrm{Var}(z_1) + \mathrm{Var}(z_2)$, and even when z_1 has become known they are still correlated. Contrast this with y_5 and y_6: $\mathrm{Cov}(y_5, y_6) = \mathrm{Var}(z_1)$, and once z_1 has become known they are uncorrelated.

effectively summarise the signals received at the ancestors of its root (and at the root itself, if it has in fact been explored). Further, these separate projects are independent: nothing that is subsequently learnt in one project reveals anything about the prizes available elsewhere, an inherited property that follows from the fact that the signals received at one node are informative about the prizes only at terminal nodes which descend from it.

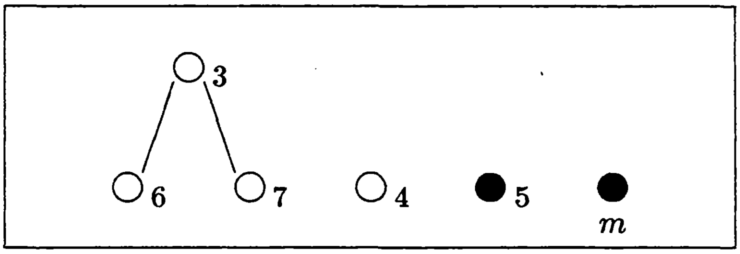

Figure 1.3: Separate and independent sub-trees

With regard to an index policy, were the agent to be in the above situation and treat the whole tree as a single project as in Figure 1.2, then a rule which selected the project with the highest index would simply tell the agent whether to proceed with the project or to take the fall-back. However, if she views the process with the perspective provided by Figure 1.3, and applies the rule to these separate projects, the strategy is completely characterised because just one action is picked out. Further, as we shall show, the fact that these separate projects are independent ensures that the Gittins index policy is optimal.

1.3 The General Model

In this section we develop more formally the central model of the paper: a sequential decision process in which the alternative projects are branching projects. We introduce our definition of a branching project and state our result (Claim 1) that if the agent is choosing an action from among a set of independent branching projects then the optimal *action* in each period is given by the Gittins index policy. This is shown to be a corollary of a more general result on stationary Markov decision processes (Theorem 1) which gives the conditions under which the Gittins index policy picks out the optimal multi-action *project* to engage in each period.

1.3.1 Branching Projects

Borrowing some notation from graph theory, we represent a branching project by an out-tree,[8] in which the number of nodes may be infinite, but such that the out-degree of any node is finite, i.e. the tree can have infinite depth but only finite branching. The nodes are the actions within the project, and the arcs represent precedence constraints: an action (other than the root) can be taken only if its parent action has previously been taken. An action is available if it has not previously been taken, and either it is the root or it is the child of an action which *has* previously been taken.

We shall consider a family of branching projects, and in each discrete period, a risk neutral agent chooses one project and an available action within it. We first note that the set of alternative projects need not be the same in

[8]Consider a directed graph, which is a set of nodes and a set of arcs, each arc being an ordered pair of nodes. An out-tree is a connected directed graph with a single root, no circuits and in which each node has no more than one parent.

every period.

Lemma 1 *Consider a family of N branching projects. In every period, there is a partition of the actions which have not yet been taken into a set of branching projects in which only the root action is available.*

PROOF: That such a partition exists initially is clearly the case, so assume that such a partition exists at time t. If the agent engages project k by taking its root action then each of the children of that root is an action available at time $t+1$ and is the root action of a distinct sub-tree, none of whose actions have been taken. Also, each of the projects which was not engaged at time t is still a branching project in which only the root action is available. Hence such a partition exists at time $t+1$, and the lemma is proved by induction. ■

When project k is engaged by taking action u, the agent receives a reward and observes a signal, the signal affecting what the agent knows about the rewards associated with actions that may be available in later periods. The state of the project, denoted by x_k, is a sufficient statistic for the observational history. It summarises what has been learnt from past signals about future rewards, availability of actions, etc. and both the reward, $R_k(x_k, u)$, and the signal, $S_k(x_k, u)$, depend on the current state and the action taken. The new state of a project depends only on the old state and the action taken, both directly and indirectly via the signal. If signals are informative only about the rewards at descendent actions,[9] then the branching projects are *independent*, i.e. the state of unengaged projects remains unchanged.

[9] Let u be the action taken, and u' be any action which is not a descendant of u. The agent's expectation of the reward to be obtained from taking action u' is unchanged by the signal received from taking action u.

Lemma 2 *Consider a family of N independent branching projects. If, after each period, the actions which have not yet been taken are repartitioned as in Lemma 1, then the branching (sub-)projects remain independent.*

PROOF: Consider the partition at time t, and observe that no action in one project is a descendant of the root action of another. So when taking an action is uninformative about actions which do not descend from it, engaging any project by taking its root action is uninformative about other projects, and the lemma follows. ■

The importance of the above two lemmas lies in the fact that when an action in a project is taken, the state of the project changes but thereafter the action does not affect the agent's choices or pay-offs, so that in each period we need consider only those actions which have not yet been taken. The lemmas then imply that, if we start with independent branching projects, in each period we can view the agent as choosing between actions in a family of branching (sub-)projects which are still independent and in each of which there is just one action available, namely the root action. This is at the heart of Claim 1 below.

1.3.1.1 The agent's problem

Rewards are additive and discounted in time by a factor β, so the agent's problem is to choose a strategy to maximise the expected discounted sum of rewards from this process, whose state at time t is written as $x(t) = \langle x_1(t), x_2(t), \ldots, x_N(t)\rangle$. The maximal expected reward over feasible policies π, denoted by the value function $F(x)$, is given by:

$$F(x(0)) = \sup_{\pi} \mathrm{E}_{\pi}\Big[\sum_{0}^{\infty} \beta^t R(x(t), u(t)) \ \mid \ x(0)\Big],$$

where $R(x, u)$ is the immediate reward realised when action u is taken in state x. When the rewards are uniformly bounded, standard assumptions from dynamic programming are sufficient to establish that the value function is the unique bounded solution to the associated dynamic programming equation and that an optimal policy exists.[10]

1.4 Optimality of the Gittins Index Policy

Following the approach of Gittins and his co-workers, it can be shown that, under certain conditions, all optimal policies for the general model are contained in a simple class of policies, and the optimal action is that recommended by the Gittins index.

1.4.1 Gittins Index Policy

Suppose that we can attach an index to any project k, that is a value $m_k(x_k)$ which is a function only of the project and its current state. When the agent selects the project with the currently highest index, she is said to be following an index policy. The specific index we shall look at is the Gittins index, whose definition makes use of a fall-back option. When there is a fall-back option m, then the agent has a stopping problem in which in each period (given that the fall-back option m has not yet been taken) the agent can either take the fall-back and stop, or continue the project for another period (the option of taking the fall-back remaining open in subsequent periods). The smallest value of m which makes the agent indifferent between stopping and

[10]Given that the rewards are additive, discounted in time by a factor β, and are uniformly bounded, the assumption that the agent is facing a stationary Markov process, for example, is sufficient.

continuation is the *Gittins index* of the project.

Denote the value function for the modified problem consisting of a fall-back M together with N projects by $\Phi(M, x)$. Since the rewards are bounded, we see that $\Phi(M, x) = M$ when M is large, and that $\Phi(M, x) = F(x)$ when $-M$ is large, and so the Gittins index is well-defined. The usefulness of this index is shown in the following result.

1.4.1.1 Optimality of the index policy for branching projects

Result 1 *Consider a family of N independent branching projects in which the rewards are uniformly bounded.*

Then the Gittins index policy selects not only the best project to engage but also the optimal action within that project.

PROOF: Using Lemmas 1 and 2, after each period we can repartition the actions which have not yet been taken into independent sub-projects in each of which just the root action is available. The claim then follows as a corollary of the more general result for super-processes which we present in the next sub-section, because the two sufficient conditions for the theorem hold. Essentially these are: (a) the state of unengaged projects remains unchanged (because signals are informative only about descendent actions); and (b) the optimal action within the engaged project is independent of the size of the fall-back (because repartitioning after each period ensures that there is only one action available in each sub-project). The theorem then tells us that the project to which the optimal action belongs is the one with the highest Gittins index, and so the optimal action is the root action of the sub-project picked out by the Gittins index policy. ∎

The above proof highlights the dual role of repartitioning actions into projects with only root actions available: it provides a key condition for the theorem, and it allows us to move immediately from 'best project' to 'optimal action'.

1.4.2 Bandit Super-processes

The proof of the above result relies on a theorem for super-processes which we present here.[11]

A super-process[12] is defined by the following collection:

(1) a set of projects, indexed by $k = 1, \ldots, N$;

(2.1) a state space, with generic element denoted by x;

(2.2) a set of available actions for each project when in state x, denoted by $U_k(x)$;

(2.3) a bounded real-valued reward function $R_k(x, u)$ which describes the instantaneous reward from taking action u in project k when in state x;

(2.4) a state transition rule giving the probability of next period's state, conditioned on this period's state, the action taken & the project it is in;

(3) a discount factor β.

The agent discounts the future by a factor β and aims to maximise the expected discounted sum of rewards from this process.

It is a *bandit* super-process when the state transition rule refers to each project rather than the process as a whole, and also when the action set and the reward are functions not of the *process* state but of the *project* state. (So, items (2.1) through (2.4) above would be for each project, and x should be

[11] For a fuller treatment, see the appendices and the references cited there.

[12] The terminology is due to Gittins [27], though the notion is due to Nash [43]. However, Glazebrook [30] uses 'super-process' to mean a multi-action project and so discusses a family of alternative super-processes.

replaced by x_k.)

Thus, given a bandit super-process, if project k is engaged in period t by choosing action $u \in U_k(x_k(t))$, the agent receives a reward of $R_k(x_k(t), u)$; states of unengaged projects do not change and the state of the engaged project changes by a Markov transition rule: if $j \neq k$ then $x_j(t+1) = x_j(t)$, and the value of $x_k(t+1)$ is conditioned only by $x_k(t)$, u & k.

We assume that the Markov process is *stationary* or time-homogeneous, i.e. the available action set, the reward, the state transition rule and the discount factor do not depend *explicitly* on time. (To give this some force, we do not allow time to be incorporated into the state.)

When the agent is maximising the expected reward from a super-process she must choose both which project to engage and which action to choose within that project. The theorem below shows that the Gittins index policy is optimal if two conditions are met: (a) projects are independent (i.e. it is a bandit super-process); (b) when there is a fall-back available, the optimal action within the engaged project is independent of the size of the fall-back.

Theorem 1 (Whittle) *Consider a super-process consisting of N alternative multi-action projects. Assume:*

(a) the projects are independent, i.e. the states of unengaged projects do not change;

(b) when there is a fall-back option available, the optimal action within the engaged project is independent of the size of the fall-back.

Then the Gittins index policy is optimal, in that it selects the best project to engage.

Moreover, writing $\phi_k(m, x_k)$ as the analogue of $\Phi(M, x)$ when only project *k is available, the following identity holds:*

$$\Phi(M, x) = B - \int_M^B \prod_k \frac{\partial \phi_k(m, x_k)}{\partial m} \, dm$$

where B is the bound on the reward functions.

PROOF: The proof is outlined in the appendices. Appendix A.1.1 gives the proof for simple bandit processes (for which the second condition is vacuous), and Appendix A.1.2 generalises it to bandit super-processes for which the second condition is crucial. The approach is essentially due to Whittle [63] and the proof elaborates on that in Whittle [62]. ∎

It should now be clear from the definitions that a branching project is a super-process, and that a family of *independent* branching projects constitutes a *bandit* super-process, so the first condition for the theorem is met. Moreover, the lemmas show that it is legitimate to reorganise the available choices in a convenient way, so that not only is the second condition for the theorem met, but also the result is strengthened from the Gittins index selecting the best project in a general bandit super-process to it picking out the optimal action from a family of independent branching projects.

1.4.3 Discussion

The index result reduces the original problem significantly: the index is calculated without reference to any other outside option or project, and the optimal action emerges from a comparison of the indices $m_k(x_k)$ attached to the various projects; further, the index of any unengaged project does not change, and so need not be recalculated. We should stress that the index is used to determine which project to engage next when the other projects will still be available in the next period. *It is not the expected value of the project.* A brief example will illustrate this point.

Consider two projects A and B. You must decide which project to engage first, and then whether you want to stop, or to engage the other project and take the larger pay-off. The cost of project A is 20 and it results in a pay-off of either 200 or zero, each outcome being equally likely. The cost of project B is 10 and it results in a pay-off of 170 or 130, again with each outcome being equally likely. So, the net expected value of project A is 80, and that of project B is 140. However, the Gittins indices for the projects are 160 and 150 respectively, so it is optimal to engage project A first, and only then engage project B if the low outcome prevails.[13]

It is to the calculation of the indices, or reservation prices, that we turn in the next section, after a few remarks on processes consisting of projects with variable length project stages, and on finite versus infinite horizon problems with discounting.

1.4.3.1 Variable length project stages

If projects have stages whose length can vary, we assume that when the agent engages a project she is committed to it for a possibly random number of periods, that number being dependent on the current state of the project but not on the actual period in which the stage was begun. As is indicated in the appendices, the proof of the optimality of the Gittins index policy continues to hold.

[13]This also demonstrates the principle that you should engage the riskier project first – the down-side is unimportant because you will never end up taking the low outcome from project A. This is shown more formally in Result 2 of the next section.

1.4.3.2 Finite versus infinite horizon, and discounting

There are two ways of looking at the fall-back m. The first is: in any period, *either* select from the available projects, *or* settle for a once-and-for-all lump sum pay-off of m and abandon selection for ever. The second is: in any period, *either* select from the available projects, *or* take a fixed reward of $(1-\beta)m$ this period and continue selection next period. In the latter case, if it is optimal to take the fixed reward of $(1-\beta)m$ this period, the agent learns nothing about the other projects, and so it is optimal to take the fixed reward of $(1-\beta)m$ in all subsequent periods, and the total discounted reward from this period forward is just m. Thus, in the infinite horizon case with discounting, the two views are equivalent.

Similarly, in the case when some projects have a terminating action,[14] if the agent selects such a project which is in a terminal state, this can be viewed as either settling for the associated lump sum reward, say y, and abandoning selection for ever, or as taking a fixed reward of $(1-\beta)y$ now (with the state of *all* projects remaining unchanged) and continuing selection next period. If we take the former view, this may seem to imply that the selection of a project which is in a terminal state affects the state of other projects because they are no longer available. However, if we redefine the fall-back as the maximum of m and y whenever a project reaches a terminal state with an associated lump sum reward of y, then once more the choice is between selecting from the available projects which have not yet reached a terminal state and taking the fall-back.

In the finite horizon case when all projects have terminating actions and there is no discounting, we are forced to take the former view (i.e. to take

[14] This corresponds to the notion of *stoppable* super-processes in Glazebrook [30]. The simple model of R&D presented in Section 1.2 is an example of such a process.

the fall-back is to settle for a lump sum pay-off of m and abandon selection for ever) and the last remark (i.e. redefinition of the fall-back whenever a project reaches a terminal state) applies.

1.5 Reservation Prices – Results and Examples

This section returns to the example of the project that was introduced in Section 1.2 and employs the interpretation of it as a model of R&D. Using the results just derived, we characterise the optimal strategy, and then discuss some implications of this strategy. Figure 1.4 illustrates the project. It differs from Figure 1.1 in that, to be more consistent with the exposition of Section 1.3, the new figure also shows the actions of costless production (nodes $4'$ through $7'$). Also, although the figure only ever shows two branches, we may wish to assume that in the project itself there are more, and denote the number of branches by γ.

1.5.1 Characterising the Optimal Strategy – Gittins Indices

The project is clearly an independent branching project, in which the only action initially available is the root, and the out-tree which describes the structure is the set of arcs illustrated in Figure 1.4. As noted after Theorem 1 in the previous section, this means that a Gittins index policy selects the optimal action, and so to characterise the optimal strategy we need to determine the Gittins indices for the possible branching projects which may arise. Then, if the value of the best available product is greater than the

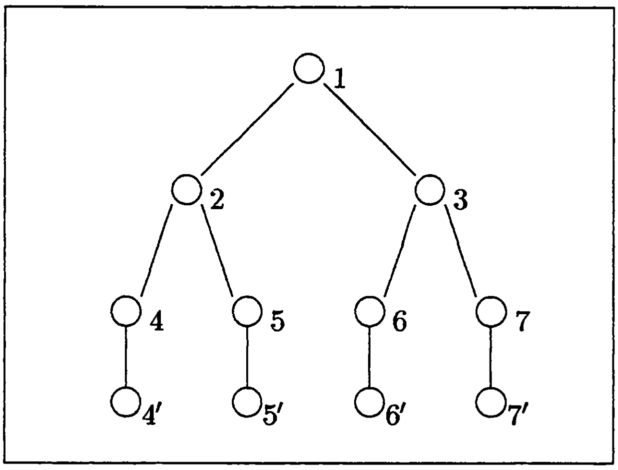

Figure 1.4: The project

highest Gittins index of the available (sub-)projects, the agent stops experimenting and makes that product; else she works on the (sub-)project with the currently highest index, and continues.

The possible projects can be classified into four types: either a project contains just a terminal action (making a product), or it is a branching project of depth 1, 2, or 3 (corresponding to a development project, a research project, and a feasibility study respectively). These are illustrated in Figure 1.5.

The rest of the analysis of this section concerns representative projects, and we adopt the convention that a representative project of type d corresponds to production if $d = 0$, and is a branching project of depth d if $d > 0$.[15] The *initial state* of a such project is the state when only the root action is available, and is a summary of everything known about the products which

[15] Subscripts on parameters, variables and functions, etc. will henceforth indicate the project depth and no longer the node, but when discussing generic properties we omit the subscript.

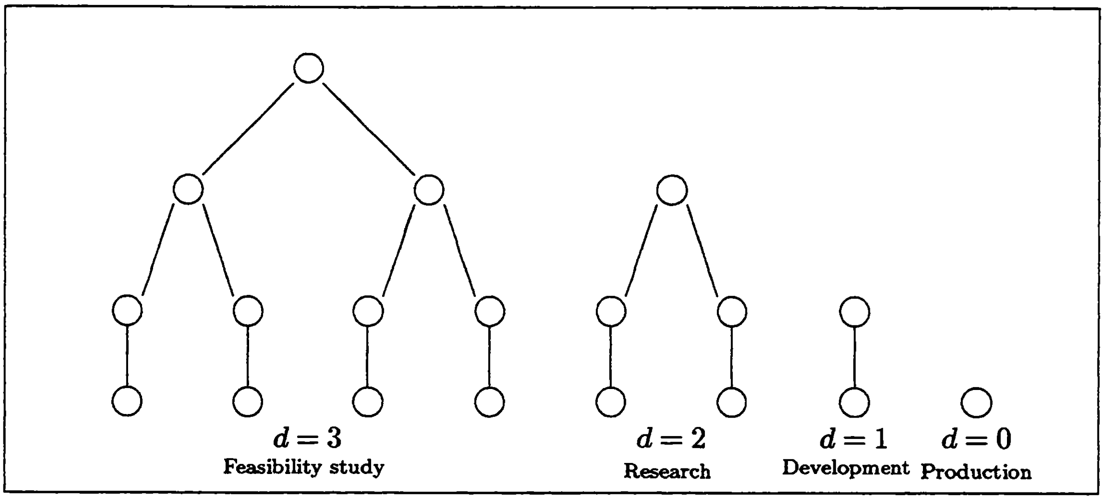

Figure 1.5: Four types of branching project

may emerge from that project. The sum of the signals received on taking actions which are ancestors of the root is such a summary, which we denote by y. Consider a project of type $d > 0$ and suppose that it is in its initial state y at time t. If the agent takes the root action then she learns z_d and updates the expected value of the products in the project accordingly. The root action can now be ignored, being no longer available, and the products can be considered as being in one of the γ (sub-)projects of type $d-1$, each of which is in its initial state $y + z_d$ at time $t+1$.

To find the Gittins index for a project, consider the process which consists of just that project and a fall-back m, and let $\phi(m, y)$ denote the value of this process when the initial state of the project is y. Denote the Gittins index, or reservation price, of the project by $r(y)$. By definition, if $m > r(y)$ the agent stops with the fall-back m, otherwise she pays c to learn the increment z and then continues. Denoting the continuation value by $\tilde{\phi}(m, z + y)$, we have the general formula for the continuation region:

$$\phi(m, y) = -c + \mathrm{E}\Big[\tilde{\phi}(m, z + y) \ \mid \ y\Big].$$

As the Gittins index is the minimal fall-back which makes the agent indifferent between stopping and continuation, we see that $r(y) = \phi(r(y), y)$, so $r(y)$ satisfies:

$$r(y) = -c + \mathrm{E}\Big[\tilde{\phi}(r(y), z + y) \ \mid \ y\Big].$$

For the rest of the section we will make the following simplifying assumption.

Assumption

(a) there is no discounting, i.e. $\beta = 1$;

(b) the number of branches emanating from the root of any project of type $d > 0$ is the same, namely γ_{d-1}, with $\gamma_0 = 1$;

(c) the cost of visiting the root of any project of type $d > 0$ is the same, namely c_d;

(d) the signal z_d received at the root of any project of type $d > 0$ is independently drawn from the same continuous distribution with support $[a_d, b_d]$, CDF $G_d(\cdot)$ and pdf $g_d(\cdot)$.

It will transpire that $r(y) = r(0) + y$, which is intuitively plausible: if the agent is indifferent between an project with initial value y and a fall-back of $r(y)$, she will also be indifferent between that project with initial value 0 and a fall-back of $r(y) - y$.

The implication of the above remark, together with the assumption, is that the optimal policy in our example will be fully characterised by just four quantities, namely r_0, r_1, r_2 and r_3, the index for each of the four types of project when the initial state is zero. We now derive expressions for these.

1.5.1.1 Production

As we have assumed that production is costless and its value is known, in this case c is zero, z is the degenerate random variable equal to zero, and so the continuation pay-off is simply the larger of m and y, i.e. $\tilde{\phi}(m, z+y) = m \vee y$. So, subscripting the variables and functions by 0:

$$r_0(y) = r_0(y) \ \vee \ y$$

and the minimal $r_0(y)$ which satisfies this is clearly given by $r_0(y) = y$. For consistency with what follows, we define r_0 as $r_0(0)$, and then we have:

$$\begin{aligned} r_0 &= 0 \\ r_0(y) &= r_0 + y. \end{aligned}$$

1.5.1.2 Development

In the continuation region for *production* ($m \leq r_0 + y$):

$$\phi_0(m, y) = m \ \vee \ y$$

and indeed in general:

$$\phi_0(m, y) = m \ \vee \ y.$$

For development, we subscript the variables and functions by 1. If the agent reveals z_1 she will be facing a single production project, the value of which will be $\phi_0(m, z_1 + y)$. So, in the continuation region for *development*:

$$\begin{aligned} \phi_1(m, y) &= -c_1 + \mathrm{E}\Big[m \ \vee \ (z_1 + y) \ \mid \ y\Big] \\ &= -c_1 + \int_{a_1}^{b_1} m \ \vee \ (z_1 + y)\, dG_1(z_1) \\ &= -c_1 + m + \int_{a_1}^{b_1} 0 \ \vee \ (z_1 + y - m)\, dG_1(z_1) \\ &= -c_1 + m + \int_{m-y}^{b_1} (z_1 + y - m)\, dG_1(z_1) \\ &= -c_1 + m + \int_{m-y}^{b_1} (1 - G_1(z_1))\, dz_1 \end{aligned}$$

the last line following from integrating by parts. So, from indifference:

$$\begin{aligned} r_1(y) &= -c_1 + r_1(y) + \int_{r_1(y)-y}^{b_1} (1 - G_1(z_1))\, dz_1 \\ c_1 &= \int_{r_1(y)-y}^{b_1} (1 - G_1(z_1))\, dz_1. \end{aligned}$$

This implicitly defines the value of $r_1(y) - y$ in terms of c_1 and the CDF $G_1(\cdot)$, and this value is therefore independent of y. As above, we define r_1 as $r_1(0)$, and then we have:

$$\begin{aligned} c_1 &= \int_{r_1}^{b_1} (1 - G_1(z_1))\, dz_1 \\ r_1(y) &= r_1 + y. \end{aligned}$$

1.5.1.3 Research

In the continuation region for *development* ($m \leq r_1 + y$):

$$\begin{aligned} \phi_1(m, y) &= -c_1 + m + \int_{m-y}^{b_1} (1 - G_1(z_1))\, dz_1 \\ &= m + \int_{m-y}^{b_1} (1 - G_1(z_1))\, dz_1 - \int_{r_1}^{b_1} (1 - G_1(z_1))\, dz_1 \\ &= m + \int_{m-y}^{r_1} (1 - G_1(z_1))\, dz_1 \end{aligned}$$

and in general:

$$\phi_1(m, y) = m \vee \Big(m + \int_{m-y}^{r_1} (1 - G_1(z_1))\, dz_1\Big).$$

In the case of a research project, if the agent reveals z_2 she will be facing several development projects, the value of each of which will be $\phi_1(m, z_2 + y)$. Let $\Phi_1(M, y)$ denote the value of these γ_1 projects when the fall-back is M and the state of each of them is summarised by y. Using the formula given in Theorem 1, we have:

$$\Phi_1(M, y) = B - \int_M^B \Big(\frac{\partial}{\partial m}\phi_1(m, y)\Big)^{\gamma_1} dm$$

where B is the bound on the reward functions. In the stopping region ($m > r_1 + y$), the partial derivative is 1, otherwise, in the continuation region, $\partial\phi_1(m,y)/\partial m = G_1(m-y)$. Thus:

$$\Phi_1(M,y) = M \vee \Big(M + \int_{M-y}^{r_1} (1 - G_1(z_1)^{\gamma_1})\, dz_1\Big).$$

So, in the continuation region for the research project:

$$\begin{aligned}
\phi_2(m,y) &= -c_2 + \mathrm{E}\big[\Phi_1(m, z_2 + y) \mid y\big] \\
&= -c_2 + \int_{a_2}^{b_2} m \vee \Big(m + \int_{m-y-z_2}^{r_1} (1 - G_1(z_1)^{\gamma_1})\, dz_1\Big)\, dG_2(z_2) \\
&= -c_2 + m + \int_{a_2}^{b_2} 0 \vee \Big(\int_{m-y-z_2}^{r_1} (1 - G_1(z_1)^{\gamma_1})\, dz_1\Big)\, dG_2(z_2) \\
&= -c_2 + m + \int_{m-y-r_1}^{b_2} \Big(\int_{m-y-z_2}^{r_1} (1 - G_1(z_1)^{\gamma_1})\, dz_1\Big)\, dG_2(z_2) \\
&= -c_2 + m + \int_{m-y-r_1}^{b_2} \Big[1 - G_1(m - y - z_2)^{\gamma_1}\Big](1 - G_2(z_2))\, dz_2
\end{aligned}$$

the last line again following from integrating by parts. Again using $r_2(y) = \phi_2(r_2(y), y)$, we obtain:

$$c_2 = \int_{r_2(y)-y-r_1}^{b_2} \Big[1 - G_1(r_2(y) - y - z_2)^{\gamma_1}\Big](1 - G_2(z_2))\, dz_2$$

This time, it is not as obvious that this equation uniquely determines the value of $r_2(y) - y$. However, having observed that, say, an increase in $r_2(y) - y$ would decrease both the integrand and the range of integration whilst leaving the LHS unchanged, we conclude as before that $r_2(y) - y$ is independent of y and so we define r_2 as $r_2(0)$ to give:

$$\begin{aligned}
c_2 &= \int_{r_2-r_1}^{b_2} \Big[1 - G_1(r_2 - z_2)^{\gamma_1}\Big](1 - G_2(z_2))\, dz_2 \\
r_2(y) &= r_2 + y.
\end{aligned}$$

1.5.1.4 Feasibility study

In the continuation region for *research* ($m \le r_2 + y$):

$$\phi_2(m,y) = -c_2 + m + \int_{m-y-r_1}^{b_2} \Big[1 - G_1(m - y - z_2)^{\gamma_1}\Big](1 - G_2(z_2))\, dz_2$$

$$\begin{aligned}
&= m + \int_{m-y-r_1}^{b_2} \left[1 - G_1(m - y - z_2)^{\gamma_1}\right](1 - G_2(z_2))\, dz_2 \\
&\quad - \int_{r_2-r_1}^{b_2} \left[1 - G_1(r_2 - z_2)^{\gamma_1}\right](1 - G_2(z_2))\, dz_2 \\
&= m + \int_{m-y-r_1}^{r_2-r_1} \left[1 - G_1(r_2 - z_2)^{\gamma_1}\right](1 - G_2(z_2))\, dz_2 \\
&\quad + \int_{m-y-r_1}^{b_2} \left[G_1(r_2 - z_2)^{\gamma_1} - G_1(m - y - z_2)^{\gamma_1}\right](1 - G_2(z_2))\, dz_2.
\end{aligned}$$

The derivation of the Gittins index for a feasibility study follows the same steps as above for a research project. As the calculations are somewhat laborious, we simply note that $r_3(y) = r_3 + y$, state the implicit formula for r_3, and collect the results together.

Reservation prices

0: $r_0 = 0$

1: $c_1 = \int_{r_1}^{b_1} (1 - G_1(z_1))\, dz_1$

2: $c_2 = \int_{r_2-r_1}^{b_2} \left[1 - G_1(r_2 - z_2)^{\gamma_1}\right](1 - G_2(z_2))\, dz_2$

3: $c_3 = \int_{r_3-r_2}^{b_3} \left(1 - \left[1 - \int_{r_3-z_3-r_1}^{b_2} (1 - G_1(r_3 - z_3 - z_2)^{\gamma_1}) g_2(z_2)\, dz_2\right]^{\gamma_2}\right)(1 - G_3(z_3))\, d$

1.5.2 Implications of the Optimal Strategy

Much of the intuition underlying the determinants of the index and so of the following result is illustrated by considering how r_1, the index for a development project, depends on the 'riskiness' of the pay-offs. In a development project (with an initial value of zero) there are two actions: the root action is to observe a signal z_1, and its child is to make the product whose value is z_1. The Gittins index is given by the formula $c_1 = \int_{r_1}^{b_1} (1 - G_1(z_1))\, dz_1$. Notice that the Gittins index does not depend on the distribution of low values of z_1, because when deciding how to proceed the agent always has the option, exercised if z_1 is low, of taking the fall-back rather than making the new product. The idea that the Gittins index depends just on the likelihood of high outcomes is captured by the result that r_1 increases if we consider a mean-preserving spread of the distribution of z_1.[16]

Result 2 *Let $H(\cdot)$ and $G(\cdot)$ be two CDFs such that $H(\cdot)$ is a mean-preserving spread of $G(\cdot)$ with the 'single-crossing property'. The Gittins index of the single stage project whose pay-off is distributed according to $H(\cdot)$ is greater than that of a similar project whose pay-off is distributed according to $G(\cdot)$.*[17]

PROOF: When H and G have the same mean:

$$\int_a^b (1 - H(z))\,dz = \int_a^b (1 - G(z))\,dz.$$

When $H(\cdot)$ is a spread of $G(\cdot)$ with the single-crossing property:

$$\int_a^x (H(z) - G(z))\,dz \geq 0$$

with equality at $x = a$ and $x = b$ and strict inequality for some $a < x < b$. Together,

$$\int_x^b (H(z) - G(z))\,dz \leq 0.$$

Denoting the two reservation prices by r_H and r_G, we have by definition:

$$c = \int_{r_H}^b (1 - H(z))\,dz = \int_{r_G}^b (1 - G(z))\,dz,$$

so

$$\begin{aligned} 0 &= \int_{r_H}^b (1 - H(z))\,dz - \int_{r_G}^b (1 - G(z))\,dz \\ &= \int_{r_H}^b (G(z) - H(z))\,dz - \int_{r_G}^{r_H} (1 - G(z))\,dz. \end{aligned}$$

[16]This is another illustration of the difference between the Gittins index of a project and its expected value.

[17]This point is explored by Weitzman [61].

The first integral is non-negative, and so $\int_{r_G}^{r_H}(1 - G(z))\,dz \geq 0$.

This implies that $r_H \geq r_G$, and if there is some difference in H and G towards the upper end of their support then the inequality is strict. ■

Thus if there is a choice between two development projects in which the expected value of the product from each project is the same, but with different variance, then it is optimal to do the more risky development first.

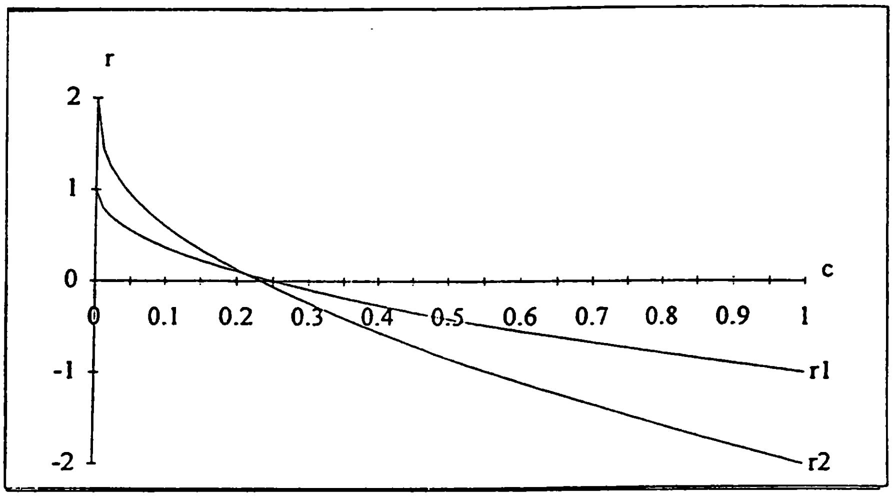

Figure 1.6: Reservation prices v. cost

Example 1 *Reservation prices as a function of cost*

The above result can be used to understand the relative behaviour of the Gittins indexes r_1 and r_2 as the cost of experimentation increases. For the case with two-way branching, equal costs of research & development, and where the distribution is uniform on $[-1, 1]$, the reservation prices vary with costs as shown in Figure 1.6.

The indexes r_1 and r_2 are calculated assuming that the initial states of the projects are zero.[18] Also note that the expected value of any signal is zero. Since the value of a product in a project is the sum of that project's initial state and the signals about the product that are subsequently observed, then initially the expected value of any product in both the research project and the pure development project is zero. However, the values of products in the research project have a higher variance. When the cost of search is low, this difference in the variance is the main consideration, and as we would expect from Result 2, the Gittins index for research is higher than that for development. As the search cost rises, however, a new consideration becomes increasingly important: the agent must spend more before production if the product is at the end of a research project than if it is in a development project. Thus as the cost rises, the Gittins index for development becomes higher than that for research. □

The main focus of this section is on how branching affects the way that agents pursue R&D. The example above shows that as costs rise, the balance tips in favour of pursuing development before engaging in more research, and this remains qualitatively the case if we allow the amount of branching to vary. In the example below, we shall see that as branching increases the agent tends to do more initial research before embarking on any development.

[18] It is easy to show that r_1 satisfies $c_1 = [(1-r_1)/2]^2$, giving $r_1 = 1-2\sqrt{c_1}$ for $0 \le c_1 \le 1$. Determining r_2 is a little more complicated. For $0 \le c_1, c_2 \le 1$, it is the positive root which is less than 2 of

$$m^4 - 24m^2 + 32(2 - 3c_1 + c_1\sqrt{c_1})m - 48(1 - 4c_1 + 4c_1\sqrt{c_1} - c_1^2) + 96c_2 = 0$$

and it is the negative root which is greater than -2 of

$$-24m^2 + 32(2 - 3c_1 + c_1\sqrt{c_1})m - 48(1 - 4c_1 + 4c_1\sqrt{c_1} - c_1^2) + 96c_2 = 0.$$

First, note that the expected value of a project consisting of γ_1 identical development opportunities is $r_1 - \int_{a_1}^{r_1} G_1(z_1)^{\gamma_1}\,dz_1$, which is increasing in γ_1, the number of branches from their common research parent.[19] Next, consider the effect of the amount of branching on the Gittins index for research (it has no effect on the Gittins index for development).

Result 3 *As the amount of branching increases, r_2 increases and r_1 is unchanged.*

PROOF: The expression giving r_2 implicitly is

$$c_2 = \int_{r_2-r_1}^{b_2} \left[1 - G_1(r_2 - z_2)^{\gamma_1}\right](1 - G_2(z_2))\,dz_2$$

If we hold r_2 fixed and increase γ_1, then the right-hand side increases. To restore the equality with c_2, we must increase r_2 thereby decreasing the range of integration and also the term $[1 - G_1(r_2 - z_2)^{\gamma_1}]$.

The expression for r_1 is independent of the amount of branching. ■

Now, what is the probability that, having explored one research avenue, the agent prefers to explore a second research avenue before pursuing any development of the first? Assume, without loss of generality, that the signal received from the feasibility study was zero. If the signal received from the first piece of research is z, then the Gittins index for developments of that research is $r_1 + z$. Thus the agent will undertake a second piece of research if $r_2 > z + r_1$ so that the probability of doing the second piece of research

[19]This leads to the final illustration of the difference between the reservation price for a project and its expected value. The reservation price for a project consisting of γ_1 identical development opportunities is simply the reservation price for just *one* development opportunity, namely r_1. This is strictly greater than the project's expected value noted above, which approaches the reservation price as the amount of branching tends to ∞.

first is given by $\Pr(r_2 > z + r_1)$, which is just $G(r_2 - r_1)$. As we would expect, this is increasing in the reservation price of research and decreasing in the reservation price of development. The final example presented here is a direct consequence of Result 3.

Example 2 *As the number of ways of developing a single piece of research increases, the agent is more likely to do a second piece of research before pursuing any development of the first.*

The intuition behind this is that the larger the number of development opportunities from a single research avenue, the higher are the expected rewards from after the development phase, and so it becomes more attractive to learn about these expected rewards before pursuing existing development opportunities. □

1.6 Conclusions

The central innovation of the chapter is the introduction of a sequential search process which can be represented as a family of trees, and the central theoretical result is that the optimal action to take in this process is given by a Gittins index policy. This result extends the existing work on multi-armed bandits in the economics literature in two important ways. In existing models, either projects are fully independent and the Gittins index policy is optimal, or they are not independent and the models have no such simplifying result. In our process the stochastic specification means that actions can have correlated rewards, so that independence is relaxed, yet the index policy remains optimal.

The second generalisation is that in existing multi-armed bandit models there is just one action available in each project in any one period, whereas in our process, the agent constantly faces choices about the direction in which to advance a project. The technical device which allows us to do this, while maintaining the result that the Gittins index policy identifies the optimal *action* and not just the optimal *project*, is to recognise that the way that actions are grouped into projects need not be the same in every period.

The final part of the chapter turns to economic applications. The representation of the process as a family of trees reflects the notion of precedence: some actions follow on from others; and it gives a measure of the diversity of rewards: close rewards have a nearer common ancestor than distant ones. The process also generates the feature that close rewards are more highly correlated than distant ones. This structure is clearly a natural one within which to study R&D and technological change[20] [59] and we investigate a very simple model of R&D in order to illustrate the main technical result. We find that as costs rise the agent expects to pursue development before engaging in more research, but that as the amount of branching increases, the agent expects to do more research before embarking on any development.

There are several ways in which this work could be extended. As mentioned in the introduction, modelling R&D as searching a branching structure provides a means of investigating the diversity of products that are developed and marketed, and how this depends on the nature of competition in R&D. Branching projects also provide a framework within which to examine

[20]Vega-Redondo [59] develops a similar model, though there the author's focus is on industry turnover rather than optimal search. Furusawa [25] also employs a branching structure to aid a game-theoretic analysis of the costs and benefits of Research Joint Ventures.

the dual role of patents as not simply conferring monopoly rights over some products, but simultaneously revealing information about related products not covered by the patent.

Chapter 2

Accumulation Games: Incremental Sunk Costs and Market Dynamics

2.1 Introduction

John Sutton in his book "Sunk Costs and Market Structure" [54] explores in some depth the idea that in many markets concentration can be understood by supposing that in a first stage firms choose to enter a market and, if technology and preferences allow, invest in a private stock of some non-price strategic variable which we call a state (Sutton focuses on advertising to increase customers' willingness to pay). In a final stage firms compete and receive profits which depend on the number of rivals, and on the state of each according to a profit function. Concentration is constrained by the need for final stage profits to cover costs incurred in the first stages. Since spending in the first stages is assumed to be unrecoverable in the last, this approach can be dubbed the sunk cost view of concentration. In his book Sutton [54] considers properties which hold over one of two very general classes of n-stage game and showed that the model predicted various features of concentration in 20 4-digit industries across 6 countries. Others have since found empirical support for his theoretical findings (see references in [55])

These sunk cost models of concentration form one starting point for the work presented in this chapter. The second starting point will be Pakes' response to the heterogeneity and variability that has been documented in a number of recent studies of firm level panels ([20], [21], [17], [45] and [12]). This response, developed in several papers with various co-authors, ([18], [45], [47]) has been to build an empirical program in which a model market with heterogeneous firm histories can be calibrated against a real market. This calibrated model might then be used to asses the expected impact of a policy intervention. The basic model on which the program is based is presented by Ericson and Pakes in "Markov-Perfect Industry Dynamics: A Framework for Empirical Work" [18]. They describe their work as considering

> the impact of uncertainty arising from investment in research and exploration-type processes. It analyses the behaviour of individual firms exploring profit opportunities in an evolving market place.

Their actual model takes a more restricted view of the actions available to firms, and the way the market can evolve. It does however allow a broad interpretation of what constitutes research and exploration. The present chapter observes that implicit in Ericson and Pakes' model is a particular hypothesis concerning the dynamic behaviour of the market, viz: each firm has a stock of some strategic variable, called its state, a firm's profit at any time depend on the state of each firm, changes to a firm's state in one period compared to the previous one depends on that firm's previous investment, and large expected changes in the strategic variable cost proportionately more than small changes.

According to this hypothesis the dynamic behaviour of a market and the level of concentration are related. Both derive from the dependence of one firm's profits on all firms' states, i.e. from the profit function. However, whereas concentration can be understood by thinking that the state is a stock accumulated in a single step, market dynamics derive from the accumulation of stocks over time. We dub this hypothesis the incremental sunk cost view of market dynamics, and begin an exploration of its implications using a number of particular examples.

The central sections of the chapter deal with two incremental sunk cost models, Model **A** and Model **B**, in which firms compete over an infinite number of periods. In both models a firm's state is measured by a stock which can be accumulated over time through investment. Firm's earn profits in each period, which are related to the state of firms in that period by the profit function. Each incremental sunk cost model has a corresponding sunk cost model of concentration, with which it shares the profit function. Both give similar results, summarised by the following claims.

Claim 1. There is a well defined number of firms in the incremental sunk cost model in the long run. This number of firms behaves in the same way as the number of firms in the corresponding sunk cost model of concentration.

Claim 2. Firms in the incremental sunk cost model converge in the long run: firms with lower stocks invest more.

Claim 3. There can be excess entry in the incremental sunk cost model, where more firms enter and accumulate early on, than persist in the long run. Firms which fall behind early on, invest more slowly or stop investing altogether.

Claim 4. There may be states in the incremental sunk cost model which are stable should they arise, but which are unlikely to do so. These can be related to market structures in the corresponding sunk cost model in which some equilibrium conditions are satisfied, but in which final stage profits do not cover investments made at earlier stages.

Model A describes a branded goods market where market shares depend on customers' perceptions of the quality of the various products on offer. The profit function is taken from the sunk cost model of concentration in advertising intensive industries given in Sutton's book [54, chapter 3]. The incremental sunk cost model embeds the profit function from the sunk cost model into a dynamic framework developed by Ericson, Pakes and McGuire ([18], [47]). The dynamic behaviour of firms in the model is very rich. However the mechanisms underlying this behaviour are somewhat obscure; no analytic solution is available and all results relate to numerically specified examples.

We propose that the dynamic behaviour is driven by two things. First the number of firms which can be supported in the long run is governed by the same structural features as equilibrium market structure in the corresponding sunk cost model. This corresponding market structure is symmetric, which has as its counterpart the convergence of firms in the incremental sunk cost model in the long run. Second, firms which establish themselves as long run participants in the market earn profits. Early in the market's history, when all firms have low qualities, many firms enter in the hope that they will be luckier than at least some of their rivals, will outlast them and so become one of the profitable firms supported in the long run. Model **B** is introduced in order to examine this proposal in a tractable setting.

Model **B** corresponds to the exogenous sunk cost model described by Sutton [54, chapter 2], in which firms pay an exogenously given cost in the first stage if they choose to enter. Model **B** displays all the behaviour discussed for Model **A**; convergence of firms in the long run, excess entry and divergence of firms early on and the possibility of stable fragmented states. Yet here the model is simple and the behaviour manifestly derives from the constraint placed on market structure by the corresponding sunk cost model.

The next two Sections deal with the two central models. Section 2.4 concludes.

2.2 Model A: Advertising

One approach to modelling the impact of advertising on market structure supposes that consumers perceive the quality of more heavily advertised goods as being higher. In these vertical differentiation models all consumers agree on the ranking of products by quality, and would buy the best one, all other things, including price, being equal. A firm's profits depend on the quality of every firm. This is the approach to advertising that Sutton adopts using sunk cost models, and tests with some success against data from the food and drink industry [54]. We consider one of Sutton's examples from this study and describe behaviour in an incremental sunk cost version where the profit function is the same, but where firms accumulate quality over time through investment.

It is the contention of this chapter that dynamic behaviour in an *incremental* sunk cost model is driven by the same factors constraining concentration in the sunk cost model with which it shares a profit function. We begin with an analysis of the sunk cost model. Sutton [54] has already given a treatment of the central result of this model , that no matter how large the market, the number of firms is bounded above. Our reason for analysing the model yet again is to consider four aspects of equilibrium that we will later compare to the behaviour of the incremental sunk cost model; equilibrium is symmetric, and when the market is large the number of firms does not depend on market size, the amount spent on advertising is proportional to market size, and the number of firms is lower at higher elasticities. We will show that all four are preserved as features of equilibrium in the incremental sunk cost model, and suggest that the same mechanisms are at work in both cases.

2.2.1 The sunk cost model of concentration

Consider a three stage sunk cost model between countably many firms, indexed by $i = 1, 2, \ldots$.

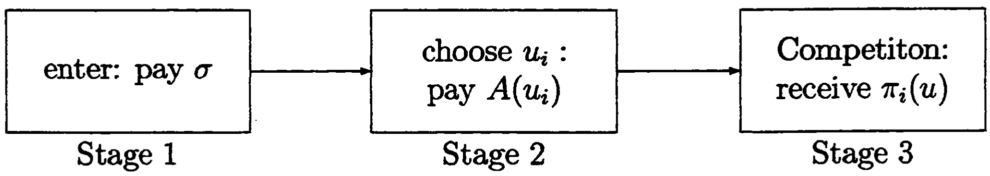

Figure 2.1: 3-stage model

In the first stage each firm chooses whether or not to enter. If firm i enters it pays a sunk entry cost σ, otherwise the firm plays no further part in the game and is ignored henceforth. In the second stage, firms observe the entry choices of the first stage. Let N be the number of entrants. Each firm

i chooses how much to spend on advertising, k_i. Advertising determines the perceived quality of the firm's product, u_i, according to:

$$k_i = A(u_i) = \frac{\alpha}{\gamma}(u_i^\gamma - 1), \quad \gamma > 1 \tag{2.1}$$

In the final stage each firm i supplies a chosen quantity q_i to the market, knowing the number of rivals and the quality of each firm's product. Each firm's marginal cost of production is c.

If a consumer buys an amount x of a product with quality u, and consumes z of an outside good, her utility is:

$$U = (ux)^\delta z^{1-\delta}$$

As a result each consumer spends δ of her income in the market. Denote the total value of sales in the market by S. If the market clearing price of product i is p_i, each consumer buys a product which maximises the quality price ratio u_i/p_i.

Subgame perfect Nash Equilibrium is found by solving Stage III first, taking the choices made in earlier stages as given. This leads to a characterisation of a firm's profits as a function of the vector of qualities of each firm in the market. Re-label firms so that they are in descending order of quality. The number of firms producing positive quantities is given by[1] the largest n for which:

$$\sum_{j=1}^{n} u_n/u_j > n - 1 \text{ and } u_n > 0$$

[1]Not all firms may find it profitable to produce. Firms which have a lower quality must charge a lower price if they are to make any sales. If a firm's quality is low enough the price it would need to charge may be lower than the marginal production cost, in which case it does better not to produce at all. In the equilibrium stated here a firm will not produce if any firm with a higher quality does not produce. There may be other equilibria corresponding to other orders.

And profits are given by[2]:

$$\pi_i(u) = \begin{cases} S\left(1 - \frac{n-1}{\sum_{j=1}^{n} u_i/u_j}\right)^2 & \text{if } i < n \\ 0 & \text{otherwise} \end{cases} \tag{2.2}$$

It is easy to confirm that if all firms had the same quality the final stage profits of each would simply be S/N^2.

Turning to the second stage, consider the choice of advertising level, taking the number of entrants N as given. In any equilibrium no firm must be able to improve its total payoff through marginal adjustments to its advertising level. In other words, the marginal increase in final stage profits a firm earns from a marginal change in perceived quality must equal the cost of that change. Call this the marginal profit condition. Restricting attention to symmetric equilibria, it defines a unique quality for each N. It is clearer to state the result in terms of the amount spent on advertising rather than the resulting level of perceived quality[3]:

$$\text{Marginal Profit Condition} \quad k/S = \frac{2}{\gamma}\frac{(N-1)^2}{N^3} - \frac{\alpha}{\gamma S}$$

The quality at which this marginal profit condition is satisfied is larger, and so firms must spend more on advertising, when the size of the market is larger. This is not surprising. Advertising a product makes all consumers more willing to pay for it. When there are more consumers, or each spends more, the same amount spent on advertising yields a greater final stage profit. Quality will now need to be higher before the net gains from marginal increases in advertising are exhausted. Note that when the market size is very

---

[2] A derivation of this equation is given in Sutton's book "Technology and Market Structure" [56].

[3] The result is found by setting the marginal profit, derived from 2.2 equal to the marginal cost of quality derived from the definition of $A(u_i)$, and then requiring that quality be the same for all firms.

large, the amount spent on advertising becomes proportional to S, so that k/S becomes independent of S.

The final step is to consider firms' entry decisions. In any equilibrium, each firm must get a non-negative overall payoff, otherwise it would have done better not to enter. Moreover, any additional firm must get a negative payoff if, instead of staying out of the market, it chose to enter. We saw that in a symmetric equilibrium, each firm gets final stage profits of S/N^2. The most that each of N entrants can therefore invest on advertising, and still get a positive payoff overall, is $S/N^2 - \sigma$. This defines a zero profit condition:

$$\text{Zero profit condition} \quad k/S = \frac{1}{N^2} - \frac{\sigma}{S}$$

The larger the market, the greater the final stage profits and so the more each firm can advertise, given N. The most that can be spent becomes proportional to S as the market gets large and k/S becomes independent of S.

The two limiting conditions when $\gamma = 2$ are shown in Figure 2.2.

The equilibrium number of firms is $N = 2$. For any integer $N > 2$, firms lie above the zero profit curve when the level of spending satisfies the marginal profit condition, and so their profits do not cover total spending on advertising. This equilibrium number is independent of the size of the market (for large enough market sizes). Although firms earn higher final stage profits when the market is larger, they also spend more on advertising, and the number of firms which find it profitable to enter remains the same.

Although the number of firms in the market does not depend on market size, it does depend on the parameter γ. This parameter is related to the elasticity of quality to spending on advertising:

$$\frac{d \ln u_i}{d \ln k_i} = \frac{1}{\gamma}\left(1 - u_i^{-\gamma}\right)$$

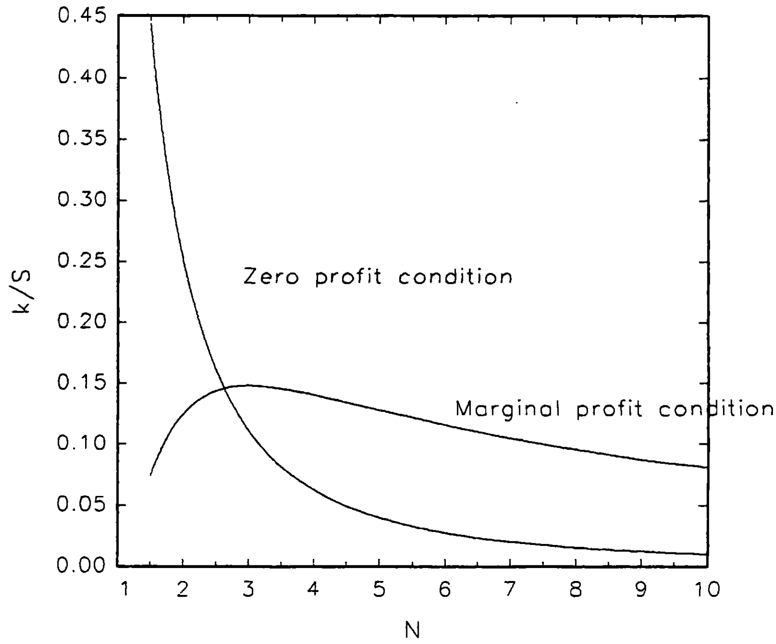

Figure 2.2: Marginal and zero profit conditions

We might expect that when $1/\gamma$ is high the number of firms in equilibrium is low. The reason for this is that if consumers' perceptions of the quality of a good are easy to change, advertising is an effective means of raising final stage profits. Firms' qualities will need to be correspondingly high before the net benefits of further marginal spending on advertising are exhausted. Certainly at large market sizes this intuition is correct. From the limiting marginal and zero profit conditions we can derive the limiting equilibrium number of firms as the largest integer N which satisfies:

$$N \leq 1 + \gamma/4 + \sqrt{(1 + \gamma/4)^2 - 1}$$

Here N falls as $1/\gamma$ rises.

2.2.2 The incremental sunk cost model

In the sunk cost model firms achieved their chosen level of perceived quality in a single step and in a final stage earned profits which depended on the quality of each firm. This was sufficient to analyse market concentration but in order to analyse market dynamics we need a model where firms can differ from period to period. Here we consider an incremental sunk cost model in which firms compete over an infinite number of periods, in each period they earn profits which depend on all firms' state according to the profit function (from the sunk cost model) in Equation 2.2, and can invest in accumulating additional quality. The rules governing the accumulation of quality are taken from a model which has been developed by Pakes and his co-authors [47], [46]. We describe the model in detail for clarity; it is only presented elsewhere as a special case of a more general model.

2.2.2.1 The Model

The model is an infinite period game in which a firm's current profits depend on the perceived quality of its own and its rivals goods. In each period, not only does a firm earn profits, but it has the opportunity to exit, sometimes to enter (if not already active), or invest, such investment giving a probability that its quality next period is one unit greater than the current level. Technically the model is a stochastic game and the solution concept we use is Markov Perfect Nash equilibrium[4].

The model is a stochastic game between I firms. Time is discrete. The state at time t, denoted k^t, is a vector listing the quality of each firm's product, $k^t = \{k_1^t, k_2^t, ...\}$. In period t each firm i takes an action a_i^t from a set which depends on the state of the industry. In each period a number of events and decision points take place. In the order in which they occur the are:

1 State realised
2 Exit
3 Competition
4 Investment
5 Entry
6 Quality depreciates

We take each event in turn and describe it in detail.

1. State Realised

   Firm i's quality in period t is a non-negative integer which is no more than some upper limit. $k_i^t \in \{0, 1, \ldots, \overline{k}\}$. If a firm is not currently

[4]Stochastic games are described by Fudenberg and Tirole in chapter 13 of "Game Theory" [24]. We adopt their notation.

active in the market its quality is zero and it is referred to as a potential entrant.

The industry's state, which we sometimes refer to as its structure, is observed by all firms at the beginning of the period. It is a random variable whose distribution depends on the state in the previous period, and on the actions taken by firms in the previous period. We will consider Markov strategies, so that firms' actions are time stationary functions of the state. As a result, state transition probabilities will also be time stationary, and the state will be a Markov Process [32]. We will return to describe these state transitions once we have looked at all of the events within a period.

2. Exit

After observing the state, each active (i.e. having positive quality) firm i chooses whether to exit. If so its action is denoted $a_i^t = X$. It sells its assets to get a current reward $g_i(k^t, X) = \phi$, and its state in the next period is $k_i^{t+1} = 0$.

3. Competition

Active firms which did not exit now compete in the marketplace, taking their own and their rivals' qualities as given. The rewards to firm i from this competition are given by a reduced form function $r_i(k)$, which depends only on the current state. Later we will enter the profit function given in Equation 2.2 as the reward function.

4. Investment

Any active firm which has not exited also invests in advertising. Investment levels are chosen simultaneously and affect the distribution of a firm's quality next period. If firm i chooses investment a_i^t, its total reward in the period is:

$$g_i(k^t, a_i^t) = r_i(k^t) - a_i^t, \quad a_i^t \in \mathcal{R}_+$$

Investment buys a firm specific random increment to quality, x_i, which has distribution (with parameter m):

$$x_i = \begin{cases} 1 & \text{with probability } p_x(a_i^t) = ma_i^t/(1+ma_i^t), \quad a_i^t \in \mathcal{R}_+ \\ 0 & \text{otherwise} \end{cases} \tag{2.3}$$

5. Entry

   At the end of the period a single potential entrant can choose whether or not to enter the market. Before deciding whether to enter, the firm observes the sunk entry cost σ^t, which is a random variable[5] and is uniformly distributed over the interval $[\sigma_l, \sigma_h]$. If the firm enters, its action is $a_i^t = E$, it pays the sunk cost immediately, though it will not be able to earn profits or invest until the next period, so its total reward in the current period is $g_i(k^t, E) = -\sigma^t$. The entering firm i's quality is set at $k_i^t = k_e$. Since the entrant has had no chance to invest, it cannot win a firm specific increment to quality and so $x_i = 0$. All potential entrants which did not enter have actions $a_i^t = \bar{E}$.

6. Market wide depreciation

   All firms which are active at the end of the period are subject to a

[5] The stochastic sunk cost is just a device to introduce noise and so help the numerical algorithms converge. See Pakes & McGuire [47] for details.

common negative stochastic increment to their qualities, y, whose distribution is given by:

$$y^t = \begin{cases} -1 & \text{with probability } p_y \\ 0 & \text{otherwise} \end{cases}$$

We interpret y as random depreciation. It is certainly reasonable to suppose that a branded good must be advertised continuously to maintain its high profile and perceived quality. It is not at all clear though that all brands within a product group should depreciate together. Pakes and McGuire interpret $y = -1$ rather as a random increase in the quality of some outside good which consumers use a benchmark [47]. Under this interpretation it is the difference between a good's quality and that of the benchmark which enters consumers' utility functions. Unfortunately this interpretation cannot be maintained[6]. We will refer to y as (brand) depreciation, since this seems the most natural reason why firms in branded goods markets must keep investing to maintain quality. However, nothing important is likely to turn on the interpretation of y: assuming that all firms share a common negative shock will affect the correlation of firm profitability, but should have little impact on the features of equilibrium we consider here.

---

[6]Index the outside good by 0. They claim that the real intrinsic quality of a good i is ν_i and what enters the demand functions is $u_i = \nu_i - \nu_0$, Furthermore when $y^t = -1$ then $\nu_0^{t+1} = \nu_0^t + 1$. However, having taken such care to derive the demand functions, it would be gratifying if this interpretation resulted from a well specified maximisation problem. It does not seem to in this case. Pakes and McGuire provide an example of a vertical differentiation model where demands *can* be written in these terms, but even here they cannot sustain their interpretation since the shocks are not to u_i, but to $k_i = c(u_i)$, c strictly convex. There is no ν_0^t, c strictly convex, such that $k_i^{t+1} = k_i^t - 1 \;\; \forall i, \; \forall k \in K$, when $y^t = -1$.

7. State transition

   Firm i's state in the next period will be $k_i^{t+1} = 0$ if i exited, or if i was a potential entrant which did not enter. Otherwise its state will be the sum of k_i^t and the specific and industry-wide increments x_i and y. The only exception is when this would specify a state higher than the upper bound $\overline{k}$. In this case the state is set to $\overline{k}$. This is summarised in a state transition function $q(k^{t+1} \mid k^t, a^t)$.

Firm i's payoff from period t onwards is the discounted sum of its rewards in each period:

$$U_i(k^t) = \sum_{\tau=t}^{\infty} \delta^\tau g_i(k^\tau, a_i^\tau)$$

Each firm i chooses a strategy s_i which we restrict to be functions only of the current state, i.e. to be Markov, and we look for Nash Equilibria, s^* of this game. Any Nash Equilibrium will necessarily be Markov. In addition we will restrict attention to those equilibria which are perfect, symmetric, and have the property that if, in a state, one firm finds it optimal to exit, then so do all other active firms with a lower quality.

At this point we can derive the incremental version of the sunk cost model analysed above simply by using the final stage profit functions $\pi(u)$ from Equation 2.2 for the reward function $r(k)$. To do so, however, it is important that the final stage profits are given as functions of advertising outlays, that is $r(k) = \pi\left(A^{-1}(k)\right)$. To see why, consider the cost of accumulating an additional unit of quality with a given probability in the incremental sunk cost model. This is derived from Equation 2.3 and is independent of the firm's current quality. In contrast, in the sunk cost model the marginal cost of quality, derived from Equation 2.1, is increasing. To improve the 'fit' to the incremental sunk cost model we will simply rename some quantities in the sunk cost models. Advertising outlays become quality, and the measure of quality which enters consumers' utility function becomes *effective* quality. This renaming has no impact at all on the sunk cost model. We have already anticipated this renaming by using the same symbol, k_i, for advertising outlays in the sunk cost model and quality in the incremental sunk cost model. Since in the latter game there is no deterministic relationship between the amount spend by i and the quality k_i, we are not free to rename quantities in the same way in the incremental sunk cost model.

2.2.2.2 The Example Markets

No analytic solution is available for the model. However Gowrisnakaran has encoded an algorithm which, if the user inputs the profit earned by firms in all possible states, and the values of the parameters of the accumulation process, will compute an equilibrium. The algorithm is described by Pakes, Gowrisankaran and McGuire [46] and the code is available (see Pakes and McGuire's paper for details [47]). The algorithm is computationally burdensome and convergence to equilibrium is slow, if it occurs at all[7]. This places practical restrictions on the examples we can consider. Ericson and Pakes showed that in this type of dynamic model, firms' equilibrium strategies may be independent of the maximum quality, $\bar{k}$, and the number of firms, I, so long as $\bar{k}$ and I lie above some bounds [18]. Ideally we would like to set $\bar{k}$ and I beyond these theoretical bounds in the examples; they are, after all, introduced to make numerical solution possible rather than as reasonable features of the model. In practice we need to restrict the state space to a manageable

[7]The algorithm essentially iterates the Bellman equations. In single agent problems this iteration is a contraction mapping and is garanteed to converge on the optimal solution. This is not the case in multi-agent problems.

size. However in all the examples reported firms only invested a little, if at all, in high states, which suggests, though it does not prove, that the results would be substantially the same if the maximum quality were raised. Also we have concentrated on cases where, loosely, the industry spent most of the time with at least one firm inactive. We would expect that in these cases behaviour would not be much different if the number of firms were increased.

We have not reported all, or even most, of the markets we began to examine. Those that looked as though practical values for $\bar{k}$ and I would constrain equilibrium too much were abandoned early on. Of the remaining ones we have included just those which relate to the comparisons on which the claims here are based. We cannot make direct comparisons between market structure in the sunk cost and incremental sunk cost models, but we can compare the comparative statics behaviour of market structure in both. The examples we have reported allow a comparison of equilibrium when market size changes, and when the effectiveness of advertising increases.

Clearly interpreting simulations is something of an art. However the claims support an intuitive story, and one that matches the results of the simple analytical models in the rest of the chapter. No example that we considered, reported or not, gave us reason to doubt that, were computer resources and time available, the claims we make here would prove substantially true over a wide class of possible markets. Table 2.1 lists the parameters for the numerical examples we discuss.

2.2.2.3 Results

Once we have found equilibrium strategies it is a simple matter to derive sample market histories[8] which chart the state of the market over time for

[8] The matter is simple not least because Gowrisankaran provides the basis of the routine.

Parameter	**Symbol**	**Common**	**A1**	**A2**	**A3**
No. of firms	I	3			
Highest quality	$\bar{k}$	20			
Quality of entrants	k_e	2			
Entry cost (low)	σ_l	0.95			
Entry cost (high)	σ_h	1.05			
Scrap Value	ϕ	0.1			
Parameter in $p_x(a_i)$	m	0.5			
Pr(depreciation)	p_y	0.7			
Effectiveness of advertising	γ		2.2	2.2	1
Parameter in $A(u_i)$	α	0.4			
Discount Rate	δ	0.925			
Market size	S		200	100	100

Table 2.1: Example Markets

some realisation of the random variables. Figure 2.3 gives one sample history of example **A1** in which all firms initially had zero quality. There are three time lines, each showing how the quality of one firm varies over the first 200 periods of the market's history.

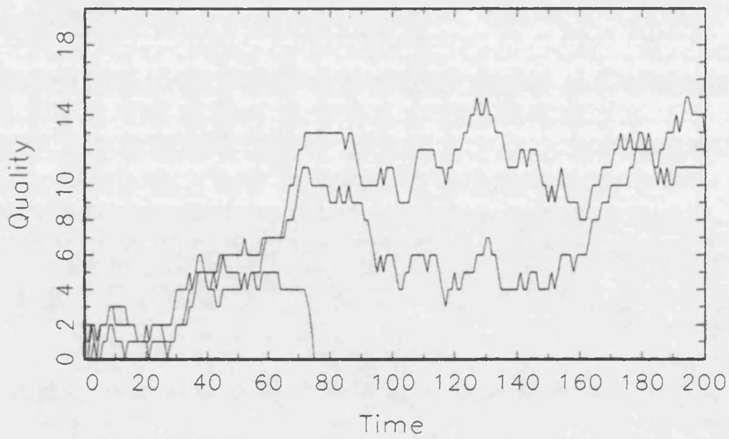

Figure 2.3: Sample history: example A1

There is variation in market shares, in the number of firms in the market, and their identities. Although Pakes developed the model specifically to exhibit this flux, we are more concerned with the stable features of the model, with what is likely to happen, and we develop some tools with which to examine various aspects of equilibrium, starting with market structure. We will organise the results under headings related to the four general claims set out in the introduction.

Claim 1(A). Market Structure

It is clear from Figure 2.3 that market structure changes over time. The first question is whether the industry spends more time in some states than others, and if so what can be said about the likely states. We take a long sample history (3 million periods) and find the proportion of all periods spent in each state during that history[9].

Result 4 *In each market there is a clear mode, and at the mode all active firms have the same quality.*

The modal state for example **A1** is shown in Figure 2.4. This shows the frequency distribution over states in which a nominated firm, the third firm, is not active.

This slice through the distribution is symmetric and single peaked. Moreover, 93% of the total distribution covers states in which at least one firm is not active. An indication of the extent to which this peak dominates the distribution is given by looking at a slice in which the third firm is active and has quality 16. This is shown in Figure 2.5. The distribution here is concentrated at the boundaries where one of the two remaining firms is inactive.

For most of the time this market is at or near the mode. This state has a strong claim to be a point description of market structure in the model and there is a strong sense in which the market supports 2 firms.

In example **A3** the frequency distribution is very different. Figure 2.6 shows the frequency distribution when one firm is inactive. Here the prob-

---

[9]This procedure approximates the time limiting distribution over states. Suppose that the Markov chain described by the state transition function has a unique irreducible subset. This will be finite and so have a unique stationary distribution $f(\cdot)$ which is the time limiting distribution. Let μ_k be the mean time between successive visits to state k. Then $f(k) = 1/\mu_k$. See for example Grimmett and Stirzaker's simple treatment of Markov chains [32, chapter 6].

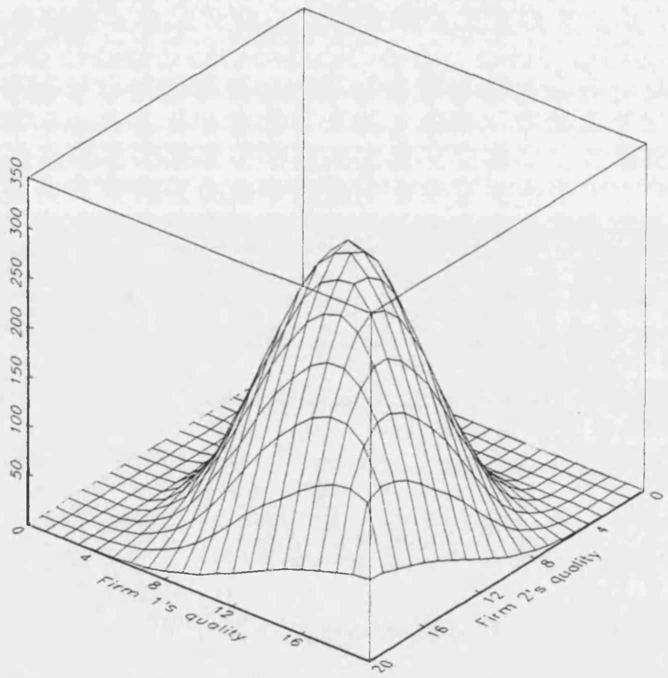

Figure 2.4: Frequency distribution when the third firm is inactive: example A1

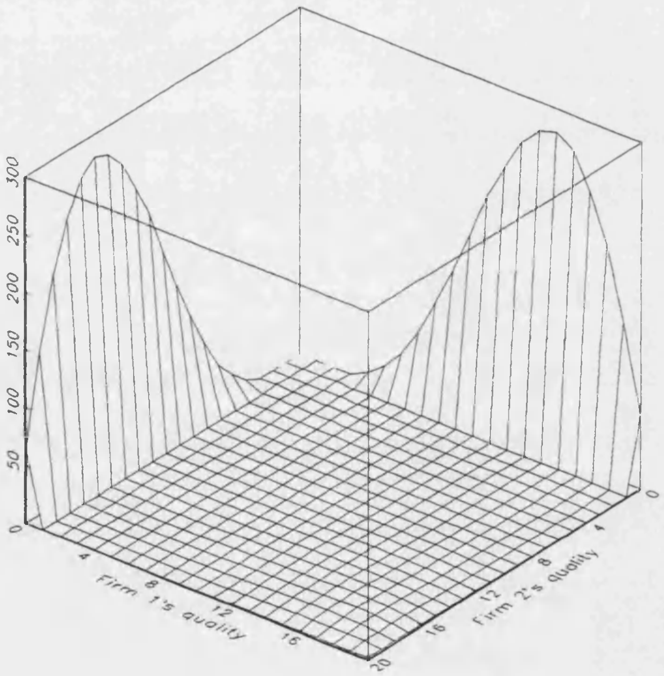

Figure 2.5: Frequency distribution when $k_3 = 16$: example A1

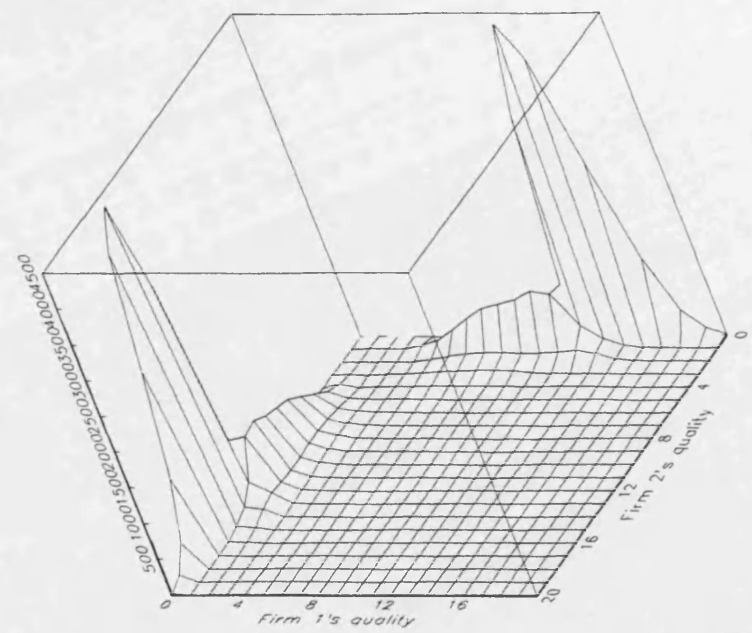

Figure 2.6: Frequency distribution when the third firm is inactive: example A3

ability is concentrated on those states in which just one firm is active: the market supports 1 firm.

There is no simple relationship between a modal state and the equilibrium of the sunk cost model for any particular example[10]. However, we have comparative statics results for the behaviour of concentration in response to changes in market size and the responsiveness to advertising. We can examine whether the modal state behaves in the same way. The result of this comparison is the important and reassuring one that the understanding of concentration developed using a 3-stage sunk cost model is preserved in the dynamic incremental sunk cost one.

Result 5 *(i) The number of firms in the modal state does not change with increases in market size, but quality in the modal state rises roughly in proportion to market size.*
(ii) If $1/\gamma$ rises, the number of firms in the modal state falls.
Table 2.2 gives the number of active firms and their quality in the modal state.

	A1	**A2**	**A3**
n	2	2	1
k	13	7	14

Table 2.2: The number and quality of firms in the modal state

[10]This is not at all surprising. For example Fudenberg and Tirole [24, chapter 13] describe work by Hanig who, in a differential capital accumulation game between two firms, found that steady state capital stocks reflected the replacement cost of depreciated capital, the discount rate, and the incentive of firms to increase current capital levels in order to reduce the rival's future levels. None of these factors play a part in a 2 or 3-stage sunk cost model.

From the above table we see that though the size of the market in **A1** is twice that in **A2**, the number of firms in the modal state is 2 in both cases. Moreover the quality of active firms in the modal state is (roughly) proportional to market size. In these two respects the modal state behaves in the same way as equilibrium in the sunk cost model. Markets **A2** and **A3** differ in the effectiveness of advertising, $1/\gamma$ and the number of firms in the modal state is greater when advertising is less effective (and so $1/\gamma$ is smaller). Again, this accords with the behaviour of equilibrium in the sunk cost model.

In the sunk cost model the number of firms is the largest number which can enter and still earn enough profits to cover their advertising costs. As market size is larger, final stage profits, all other things being equal, are larger. However, another firm would nevertheless not find it profitable to enter the larger market because the amount spent on advertising is also larger. The results here suggest that a basically similar mechanism dominates the determination of concentration in the dynamic model, despite the added complexity that costs are accumulated over time at a variable rate, and that firms earn profits in all states.

Claim 2(A). Convergence in the long run

Another way to look at equilibrium is in terms of the likely evolution of the state over time. We look first at states close to the modal one. Figure 2.7, illustrating example **A1**, shows all states in which the third firm is inactive.

The arrows show the expected change in the qualities of the remaining two firms. Specifically consider an arrow whose base is at (k_1, k_2). The horizontal (vertical) length of the arrow gives the difference between the probability that firm 1's (2's) investment is successful, minus the probability of depreciation. For clarity, the probabilities are shown scaled up by a factor

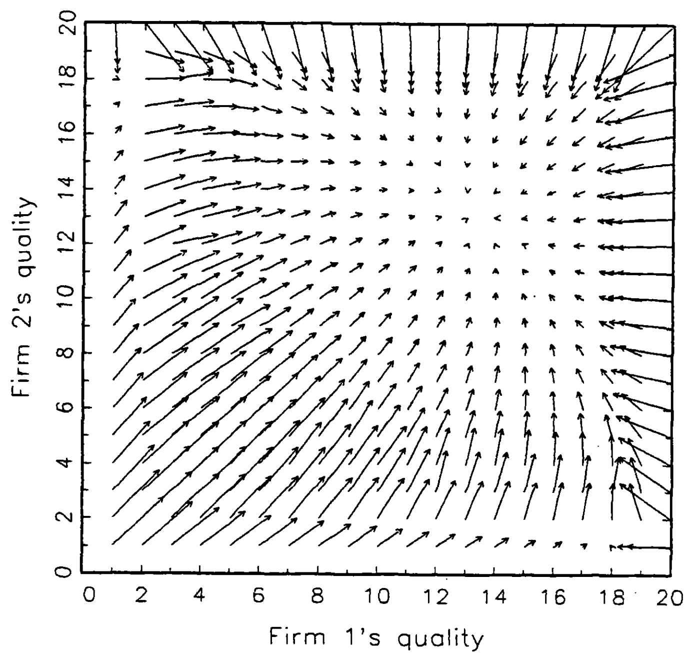

Figure 2.7: State transitions when $k_3 = 0$: example A1

10, if the distance between states is 1. The state's evolution does not depend just on investment and depreciation, but also on entry and exit decisions: we ignore these for now.

The arrows seem to be converging on a single point that lies somewhere between $(13, 13, 0)$ and $(14, 14, 0)$. What is more the arrows are shorter closer to this point. A state transition arrow with no length would indicate that the expected change in quality from an investment would exactly offset that from depreciation. The state would be stationary[11]. Though no actual state is stationary, it seems reasonable to say that the there is a stable stationary state between $(13, 13, 0)$ and $(14, 14, 0)$. The smaller of these two, $(13, 13, 0)$, is the modal state. Figure 2.8 shows the corresponding state transitions for example **A3**.

Result 6 *(i) The modal state is either stationary or one of the states closest to the stationary one.*
(ii) Active firms converge[12] *when the market is close to the modal one: the firm with the lowest quality invests more.*

Since the market is close to the modal state most of the time in examples **A1** and **A2**, most of the time firms are converging.

Finding that the modal state is stationary is no surprise. A state which persists is a good candidate as the most likely state. Neither is it a surprise

[11] There are states which are close to stationary at $(18, 1, 0)$ and $(1, 18, 0)$. However, they are not stable, and their appearance is probably due to such considerations as the positive scrap value for exiting firms and the fact that entrants start with quality 2.

[12] Convergence here is β-convergence: smaller firms grow faster. As pointed out by Quah, this does not necessarily imply that the difference between firm's states shrinks (it does not imply σ-convergence). If the two firms began in a symmetric state, we would be surprised, given the stochastic nature of investment, if they remained in a symmetric state forever. Quah has discussed these and other notions of convergence [48].

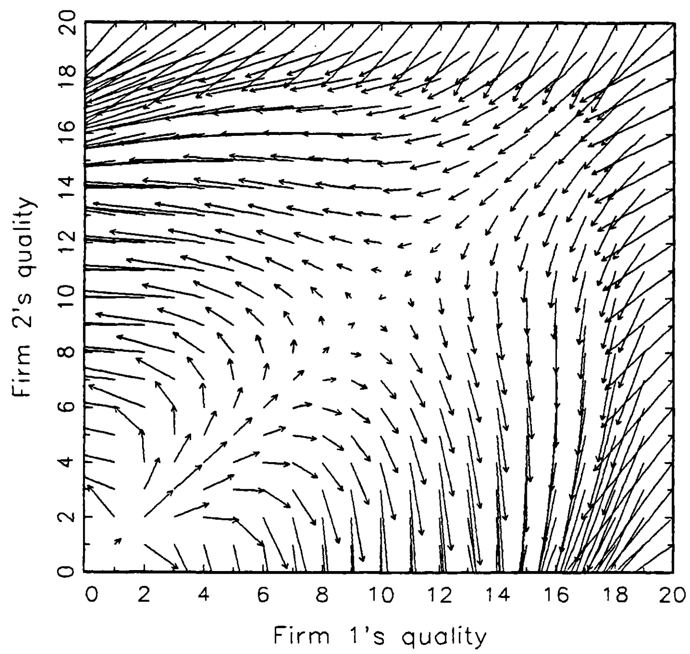

Figure 2.8: State transitions when $k_3 = 0$: example A3

to find that many trajectories in the state space end on the modal state, since this indicates that it is likely to arise. Stationary states in the dynamic game can be related to market structures satisfying the marginal condition in the sunk cost model: in both cases small changes to the state are unprofitable. States towards which the market moves, even when it starts far away, can be related to states in the sunk cost model which satisfy the zero profit condition: firms at the start of a trajectory must expect that their investment costs will be covered by their profits, otherwise they would exit. This reasoning, together with the finding that the comparative static behaviour of the modal state is the same as that of equilibrium market structure in the sunk cost model, suggests that in both cases the number of firms that the market will support is governed in the same way by structural features of the market: the ability of firms to increase their share of a market of fixed size through spending on advertising.

Claim 3(A). New markets: excess entry and divergence

Although firms converge most of the time in examples **A1** and **A2**, since in the long run these markets are mostly close to the modal state, there are situations where we have *a priori* reasons to think that the market starts out far from the modal state. One example is the new market. Suppose that a new market for branded goods opens. To examine the expected early history in such a case we took a large number (6000) of sample histories of example **A1**, each with an initial state $(0, 0, 0)$, and then found the average state of each of the highest, middle and lowest quality firm in each period. The result is shown in Figure 2.9.

What we see in the figure is excess entry. All firms initially enter and, initially, all begin to build up their quality. After a while, however, the lowest quality firm starts to fall increasingly behind the other two. It stops

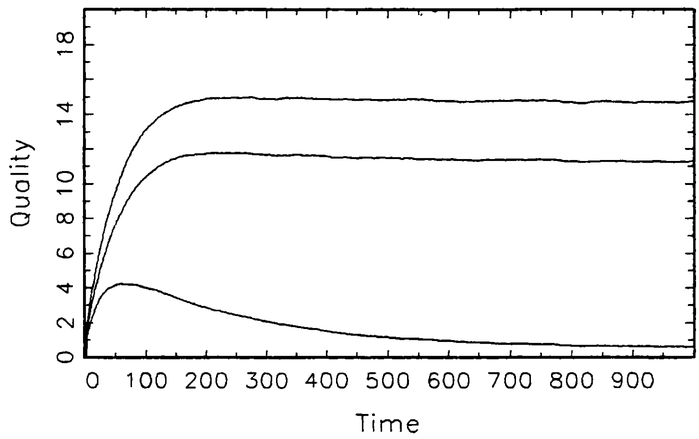

Figure 2.9: The expected early history of a new market: example A1

investing enough to counteract the depreciation of its brand and its quality falls. Eventually it exits: in the long run the average quality of the third firm is less than 1, which means it spends most of the time out of the market.

To see the expected industry dynamics underlying this process we also looked at the average probability that investment is successful for each of the highest, middle and lowest quality firms in each period. These are shown in Figure 2.10. The time line which is the first to be positive is the high quality firm, the last is the low quality firm.

Certainly after 450 periods, when the expected quality of the third firm is less than 1, the remaining active firm with the highest quality invests least: this is the convergence in the long run discussed in the previous section. However, while three firms remain active, before the third has exited, the firm which invests least overall is the one with the lowest quality: early on there is divergence of low quality firms from those that will remain in the market in the long run. This early divergence can also be seen in the state

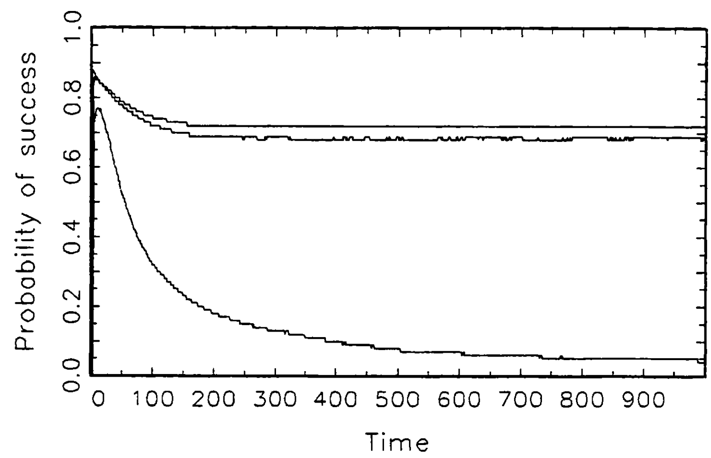

Figure 2.10: Expected probability that investment is successful over the early history of a new market: example A1

transitions. Figure 2.11 shows the state transitions of two of the firms in example **A1** when the third firm's quality is 8, which Figure 2.9 indicates is the expected quality of the leading firm when the trailing firm's expected quality reaches its maximum.

This clearly shows that when the two other firms both have low qualities they diverge. Similar analyses for the other markets yields:

Result 7 *Suppose that, initially, no firm has positive quality. In the expected industry history:*
(i) more firms accumulate early on than are in the modal state.
(ii) early on the lowest quality firm invests least and diverges from its rivals.

One explanation for this excess entry is that additional firms enter early on and exit later because profits are high when no firms have high qualities, and fall when qualities rise[13]. A second explanation, and the one we support

[13]This is the basic mechanism generating excess entry in the models developed by Klep-

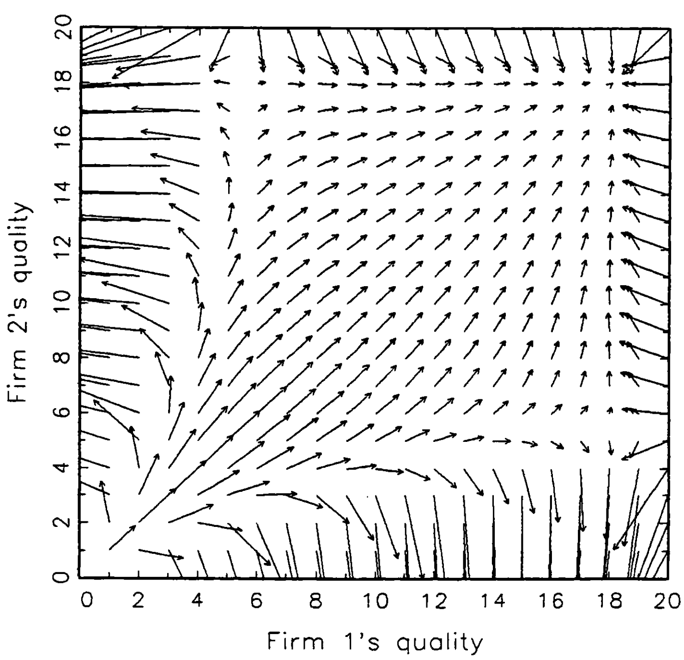

Figure 2.11: State transitions when $k_3 = 8$: example A1

here, is that an extra firm enters early on in the hope that it will be luckier than its rivals and will overtake and outlast at least one. In this view the number of firms which the market can support is restricted by structural features of the market, in much the same way as the number of firms in the concentration model. Excess entry represents rent dissipating competition to become one of the supported firms. Firms which fall behind are more likely to be among those that exit[14], and so have less time to recoup the benefits of their investments and invest less.

Claim 4(A). Fragmented stationary states

One possible consequence of there being more firms than the market will support is, as we have just seen, that the lagging firm diverges from the others and eventually exits. This result arose in the context of new markets where not only did too many firms enter, but they all began with low qualities. Another case where too many firms may be in the market is when an exogenous increase in the effectiveness of advertising makes advertising more attractive, increases a firm's willingness to invest, and so reduces the number of firms which can afford to be in the market. Again, "too many firms" is associated with "too low a quality", and we might expect much the same behaviour as firms compete to remain as one of the reduced number

---

per and his co-authors ([31], [38] and [37]) except that in his models it is capacity (and, in the paper with Gort, experience) which is accumulated over time rather than quality. Geroski has raised the related possibility that the reason many entrants do not establish themselves as long run players in the market may be that they are only exploiting finite lived profit opportunities, though he quickly dismisses this as unable to account for the fact that de novo entrants have higher failure rates than entrants who are diversifying into the market from elsewhere [26].

[14]This is practically guaranteed by the restriction of equilibria to symmetric ones in which, if a firm exits, so do those with lower qualities.

that the market will support, and invest less when it looks as though their efforts have been unsuccessful and that they will soon be forced to exit. Not all fragmented states, however, are unstable.

Figure 2.12 shows the transition probabilities of two firms when the 3rd has quality 16 in example **A1**. The two firms are stationary when each has

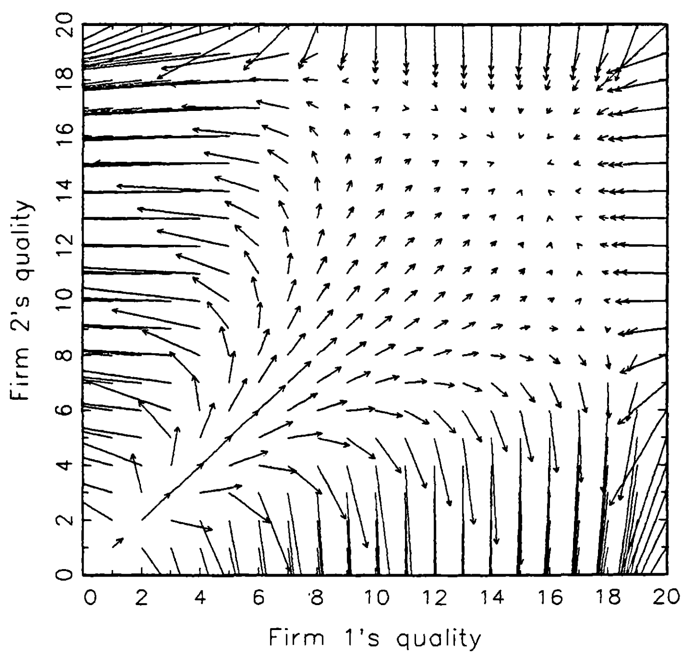

Figure 2.12: State transitions when $k_3 = 16$: example A1

quality 15. Since equilibrium is symmetric this suggests that all three are stationary at some notional state between (15, 15, 15) and (16, 16, 16). What is more, firms converge on this state from nearby. We might wonder whether the industry spends much time in this stable, locally convergent state. In fact there is a local mode close by as we see below. However, we have already seen that since all three firms are active for only 7% of the time, this local mode arises much less often than does the modal state discussed above. We have already seen the frequency distribution over states in which the third firm has quality 16, in Figure 2.5. The frequency is dominated by states where one remaining firm is inactive and to see the local 3-firm mode we must look at the Figure without these states. This is shown in Figure 2.13, and for clarity the same distribution is shown again in Figure 2.14, but using a contour map.

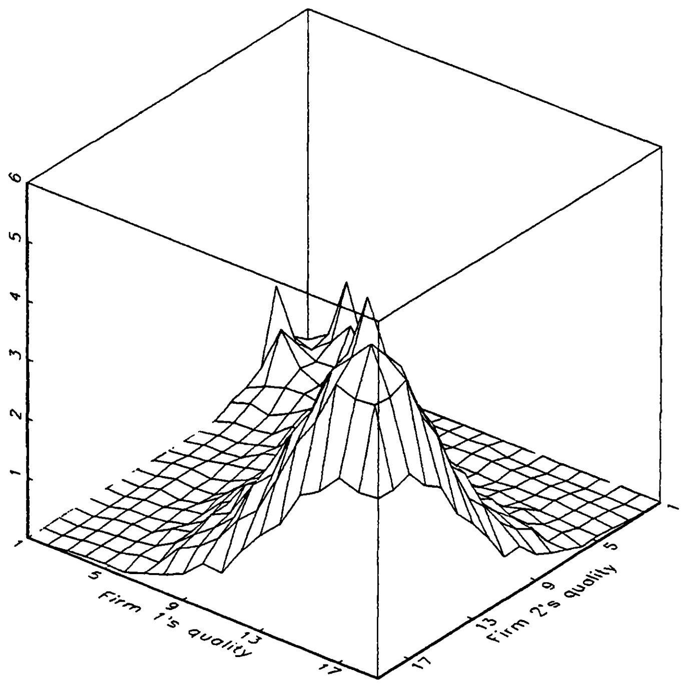

Figure 2.13: Frequency distribution when $k_3 = 16$, excluding k_1 or $k_2 = 0$: example A1

There is a stable fragmented state in each of the example markets, though it is not locally convergent in example **A2**.

Result 8 *There is a symmetric stationary state with more active firms than in the modal state. Table 2.3 shows the number of firms and their qualities*

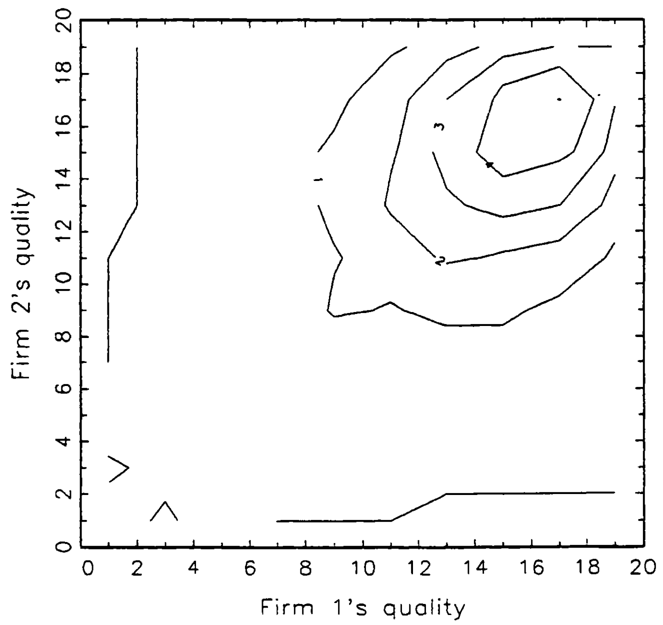

Figure 2.14: Frequency distribution when $k_3 = 16$, excluding k_1 or $k_2 = 0$: example A1

in these states.

	A1	**A2**	**A3**
n	3	3	2
k	$15-16$	$8-9$	$9-10$

Table 2.3: The number and quality of firms in the fragmented stationary state

We previously argued that the modal state can be related to equilibrium market structure in the sunk cost model. In both cases marginal changes to quality are unprofitable (the marginal profit condition of the sunk cost model) and in both cases firms expect to make enough profit to cover their investment costs (the zero profit condition of the sunk cost model). The fragmented stable states of Result 8 can be related to market structures in the sunk cost model in which the marginal condition, but not the zero profit condition, are satisfied. Once the state has arisen in the dynamic model no firm wants to make small changes to its state. We suggest, however, that firms never meant the state to arise in the first place; it just happened that, for example, while they were competing to see which of them would remain to be supported in the market no clear winner emerged until they had all accumulated high qualities. Since investment costs are sunk, once firms have high qualities, the requirement that firms earn enough profits to cover their total investment costs has no bite. Comparing examples **A1** and **A2** in Table 2.3 shows that when market size doubled, the quality of the three firms in the fragmented stable state (roughly) doubled, just as the quality doubles in the market structure which satisfies the marginal condition for n firms in the sunk cost model.

The market spends very little time in these fragmented states, and unless there are good *a priori* reasons to believe a market starts with too many firms, but with high enough qualities, they are not likely to be empirically relevant. The existence of these states, however, lends support to the view that market dynamics can be understood by referring to the sunk cost model.

2.2.3 Conclusions

There are severe constraints on the number of firms and set of possible qualities which it has been practical to consider. We can say, nevertheless, that the examples explored support the following claims.

Claim 1(A). There is a well defined modal state which is symmetric. The number of firms in the modal state behaves in the same way as the number of firms in the sunk cost model. The number of firms does not change as the market size is increased, but k_i rises proportionately, while the number of firms is lower as the elasticity is higher.

Claim 2(A). Most of the time the market is converging on the modal state, so that firms with higher qualities invest less.

Claim 3(A). There is a significant exception to convergence: the new market. Suppose that a new market for branded goods opens up. A related situation arises when there is a sudden increase in the effectiveness of advertising in an established market which means that existing levels of quality are too low. We find that there may be a competition to remain in the market during which more firms advertise than are in the market in the long run. Firms which fall behind diverge from their rivals: they invest less.

Claim 4(A). Finally there may be states which are stable, and so persist at least for a while should they arise, but which were unlikely to arise.

These correspond to market structures in the sunk cost model which satisfy the marginal profit conditions but in which firms earn negative profits overall.

We propose that these results can be understood as consequences of an accumulation race between firms, which compete by investing in quality, to remain as one of the firms supported in the market in the long run. Because the model is so intractable, the mechanisms are somewhat obscure, and we will examine our proposal in the context of a simple tractable model which shows the same features as this one.

2.3 Model B: Exogenous Sunk Costs

Equilibrium in the dynamic model in this section can have all the features we noted in model **A**: a well defined market structure, convergence of firms in the long run, excess entry, and unlikely but persistent, fragmented states. Here these features are related in a transparent way to the restriction that the underlying sunk cost model places on the number of firms which the market will support, as well as the need for firms to accumulate their strategic variable over time, rather than all at once. This transparency allows us to understand some of the forces at work in the advertising game in particular, and in incremental sunk cost models in general.

The profit function is taken from an exogenous sunk cost model in which firms pay to enter in a first stage. They compete in the product market in the second, earning profits which depend on the number of first stage entrants. The strategic variable here is just the ability to produce, and it is paid for all at once in the first stage, or not at all. We construct an incremental sunk cost model in which the ability to produce must be acquired in a number of stages, rather than just one. For concreteness consider the following two stage story.

> Market research has confirmed that there is a demand for a bicycle with automatic gears. Before being able to make such a bike, a firm would need to conduct R&D to discover an engineering principle which would allow gearing to alter as the torque applied to the pedals increased, and then to construct a prototype which would be cheap enough for consumers to buy. These tasks can only be done sequentially.

In this model there is only one reason to invest: firms want to reach the end of the project, and cannot earn any profits until they do. The number of stages they must complete is fixed and in particular does not depend on the number of rival firms which are also active in the market. In these two respects it is simpler than the advertising game where firms earnt profits in all states. Moreover in the advertising game the quality of firms in the stationary states, those states where firms stopped, was endogenous and depended on the number of active firms, as well many other factors. In the exogenous sunk cost model the restriction that profits cover sunk costs constrains the number of firms the market can support in an obvious way. The simplicity of this incremental exogenous sunk cost model allows us to consider the pure effects of this constraint on market dynamics.

We first confirm that there is always an equilibrium in which exactly the same number of firms begin and complete the project in the incremental sunk cost model as enter the sunk cost model, and typically these firms have a strictly positive payoff overall. However, in the case where the market supports 1 firm, we also analyse a second equilibrium in which there is excess entry. Here 2 firms compete to remain in the market in the long run. Both firms initially invest, but if early on one firm falls behind, it stops investing and never finishes. During the competition to remain in the market both firms invest and hope that they complete stages faster than does their rival. However, sometimes neither gets ahead of the other until both have completed so many stages that the cost of completing the remaining ones is bound to be less than the profits from producing, even if the rival produces as well. In this case both firms complete the project: deciding which firm will stay in the market through a competition to accumulate the strategic variable fastest leads in a natural way to the emergence of fragmented stable states if the competition does not produce a winner early on. We cannot say whether active firms converge or not if just one firm remains in the long run. For this reason we also analyse a particular 3-firm example in which at least 2 always finish. Even if the finishing firms have different numbers of stages to complete at some point, they are bound to end up with the same number, i.e. none.

The incremental sunk cost model does not directly model market dynamics in homogeneous goods industries, which are those industries where the exogenous sunk cost concentration game has been most successfully applied. The problem is that dynamics in the model relate to behaviour before firms are able to produce, whereas in the world firms will only enter our data sets once they are producing. In a later section we give an example which is more relevant for market dynamics in homogeneous goods markets.

2.3.1 The sunk cost model of concentration

The profit function is taken from a simplified version of what Sutton [54, chapter 2] characterises as the exogenous sunk cost model of concentration.

The market is one where firms offer an identical product and where, before being able to produce the good, firms must incur an unrecoverable cost σ. Once in the market firms engage in some form of competition which generate profits which are a function of the number of firms in the market, n, and are denoted $\Pi(n)$. Here we assume that the final stage competition yields symmetric profits. The situation is modelled as a 2-stage game. In the first stage a large number of firms each choose whether to pay σ and enter the market. If n choose to do so, then in the second stage each firm receives $\Pi(n)$. These profits are discounted in the first stage by a factor $\hat{\delta}$.

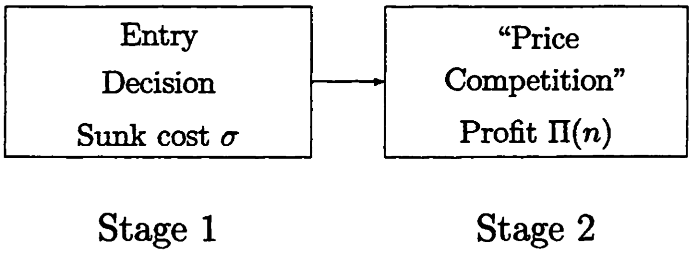

Figure 2.15: The Exogenous Sunk Cost Model of Concentration

In any pure strategy equilibrium both entrants and those that stayed out of the market must have chosen optimally. The final stage profits must cover the sunk entry cost for those that did enter, but would *not* have done so had an additional firm have entered. If $N(\sigma)$ is the number which enter in pure strategy equilibrium when the sunk cost is σ, $N(\sigma)$ is the largest n for which $\hat{\delta}\Pi(n) \geq \sigma$ and firms which enter earn profits $\hat{\delta}\Pi(n) - \sigma$.

There are also many equilibria in mixed strategies[15]. For example there

[15]Equivalently, suppose that firms can invest in a probability of successful entry, say by searching for suitable premises, and that when the probability of success is p, the cost is

is one where firms $i = 1, \ldots, N(\sigma)+1$ all choose a probability $q^* \in (0,1)$ and no other firms enter with positive probability. To see this, suppose that firms $i = 1, \ldots, N(\sigma)$ choose q and consider the optimal reply of firm $i = N(\sigma)+1$. Each firm's final stage payoff if it enters, net of entry cost, is (omitting the dependence of N on σ):

$$-\sigma + \sum_{n=0}^{N} \binom{N}{n} q^n (1-q)^{N-n} \hat{\delta}\Pi(n+1)$$

The definition of $N(\sigma)$ implies that this is negative when $q = 1$ and positive when $q = 0$. Since the payoff is continuous in q there will exist a q^* for which the payoff is zero and the final firm is indifferent over entering or not, and so entering with probability q^* is a best reply. Hence all $N(\sigma)+1$ firms entering with probability q^* is an equilibrium. The number of firms producing in the final stage is a binomial random variable which will sometimes be greater than $N(\sigma)$ and sometimes less. Of course we can easily imagine repeating the entry stage, so that if fewer than $N(\sigma)$ initially entered, others would enter later on, but even if we added an exit stage, once too many firms had entered they would not exit while their final stage profit was positive (gross of entry cost).

Even in this trivial example we see some of the features that arose in the advertising game: excess entry (more firms are initially active in the market than produce in the long) which dissipates the rents from the pure strategy equilibrium, and the possibility that fragmented stable configurations can arise. The basic mechanism in the dynamic model we present is essentially the same, but the dynamic features that it gives rise to are more strongly drawn.

σp. There are many equilibria where firms choose interior probabilities.

2.3.2 The incremental sunk cost model of market dynamics

In the story about developing a new bicycle randomness entered the model through the stochastic investment function. In the incremental sunk cost model we investigate here, we suppose that firms either invest or not, that investment yields a deterministic outcome, and randomness enters through the choice of mixed strategies. This supposition simplifies the exposition of the model and results, though much would remain unchanged were we to interpret the model in terms of a linear, stochastic investment function.

2.3.2.1 The Model

The model is an infinite period stochastic game [24] in discrete time between I firms. Firm i's state in period t, denoted k_i^t, is the number of stages i has yet to complete before it can produce, and lies in the set $k_i^t \in \{0, 1, \ldots, d\}$: firms produce in state 0. The industry's state at t lists the state of each firm, $k^t = (k_1^t, \ldots k_I^t)$.

In each period firm i's action a_i^t is either to invest, $a_i^t = 1$, or not, $a_i^t = 0$, and i receives a reward $g_i(k^t, a_i^t)$. If $k_i^t = 0$ firm i's only feasible action is to do nothing: $a_i^t = 0$. It earns profits which are non-increasing in the total number of producing firms. Let the number of producing firms be: $n(k^t) = \#\{k_i^t = 0\}_{i=1,\ldots,I}$, then each earns a profit $\pi(n(k^t))$. In other states i earns no profits. It can invest in completing a stage, however, in which case its action is $a_i^t = 1$ and it incurs an investment cost c. Otherwise $a_i^t = 0$. To summarise, the reward function is:

$$g_i(k^t, a_i^t) = \begin{cases} \pi(n(k^t)) & \text{if } k_i^t = 0 \\ -ca_i^t & \text{otherwise} \end{cases}$$

The evolution of each firm's state depends on its own investment. Specifically:

$$k_i^{t+1} = k_i^t - a_i^t$$

Firm i's payoff at t is the discounted sum of rewards in all subsequent periods:

$$u_i^t = \sum_{\tau=t}^{\infty} \delta^\tau g_i\left(k^\tau, a_i^\tau\right)$$

Each firm i chooses a strategy $s_i(k)$ which specifies the probability of investing as a time stationary function of the current state, i.e. strategies are Markov. The strategy is pure if $s_i(k) = 1$ or 0 for all k, and mixed otherwise[16]. We look for Markov strategies $s^* = (s_1^*(k), \ldots, s_I^*(k))$ which form a Perfect Nash Equilibrium[17].

We will consider a family of incremental sunk cost models parameterised by the number of stages d. Suppose a firm invests in every period until it completes the project, and thereafter shares the market with $n-1$ others. We require that the values of the profit and investment streams, evaluated at the start, be the same for all members of the family, and the same as in the sunk cost model. A family which satisfies this requirement, and the one we analyse, is:

$$\begin{aligned} \delta^d &= \hat{\delta} \\ \frac{\pi(n)}{1-\delta} &= \Pi(n) \\ \frac{c}{1-\delta} &= \frac{\sigma}{1-\hat{\delta}} \end{aligned} \tag{2.4}$$

[16]This is not standard notation, but there should be no ambiguity.

[17]Equilibrium in this game is identical to that in a related one where a firm's action specifies the probability that its investment is successful, the firm incurs cost ca_i^t, and the firm's state evolves according to: $k_i^{t+1} = k_i^t - 1$ with probability a_i^t, and is unchanged otherwise.

2.3.2.2 Results

A firm which had x stages yet to complete and which invested in every period until it completed the project, and thereafter shared the market with $n-1$ others would get a prize worth $\delta^x\pi(n)/(1-\delta)$ and would incur costs $c(1-\delta^x)/(1-\delta)$. Define the $u(x,n)$ as the net payoff in this case:

$$u(x,n) = \frac{\delta^x}{1-\delta}\pi(n) - \frac{1-\delta^x}{1-\delta}c$$

If the firm starts at d, this gives $u(d,n) = \hat{\delta}\Pi(n) - \sigma$, exactly what it would have received ~~if the concentration~~ had it entered when $n-1$ others did. This means in particular that $N(\sigma)$ is the largest n for which $u(d, N(\sigma)) \geq 0$. We will later use the properties that $u(x,n)$ is non-increasing in n (because $\pi(n)$ is non-increasing in n) and non-increasing in x. We first consider equilibrium in pure strategies. Here whether or not a firm finishes the project is deterministic, and the number which do so is the number which the market can support.

Claim 1(B). Market Structure

Suppose that initially no firms had completed any stages, and so $k_i^0 = d, \forall i$. Since the outcome is deterministic no firm will invest unless it finishes, otherwise it is bound to make a loss. Intuition suggests that along the equilibrium path those firms that finish at all will do so by investing in each of the first d periods, since waiting one period means that the payoff will be discounted. If this is so, and n firms finish, each gets a payoff $u(d,n)$, and since $N(\sigma)$ is the largest number of firms for which this is positive, exactly $N(\sigma)$ will finish, the same number as entered in the sunk cost model. This is indeed what happens in any equilibrium. What is ruled out in equilibrium is that, in the continuation from any state k, one firm will eventually complete the project while another, which at k was closer to completion than the first one, will not. Essentially, if only one of two firms can profitably finish, the firm which is ahead can guarantee that it is the one by investing until a state is reached in which it would continue to invest even if the other did[18].

For each state k we define a boundary distance $x^b(k)$ which is the furthest a firm can be from the end and still make a profit if it and all firms closer to the end invest fully in every period until they have completed the project. More formally, the largest number of firms which would be able to profitably complete the project, if all firms were at a distance x, is $\max\{n \mid u(x,n) \geq 0\}$. The boundary $x^b(k)$ is the furthest distance x from the end for which this number is greater than the number of firms closer than x. The only exception to this is when all firms would be able to profitably complete the project when none have yet completed any stages, i.e. when $N(\sigma)$ is greater than the number of firms in the market, in which case $x^b(k)$ is set to $d+1$, $\forall k$. Formally:

$$x^b(k) = \begin{cases} d+1 & \text{if } N(\sigma) > I \\ \max\{x \mid \max\{n \mid u(x,n) \geq 0\} > \#\{i \mid k_i < x\}\} & \text{otherwise} \end{cases} \tag{2.5}$$

Result 9 *If $N(\sigma) > I$, all firms invest in every period unless they have completed the project. Otherwise, in any equilibrium, if the state at time t is k^t, the continuation satisfies the following characterisation, for all possible k^t:*
(i) all firms which are closer to the end than $x^b(k^t)$ invest in every subsequent period until they have completed the project
(ii) if there are any firms at $x^b(k^t)$, some or all will invest in every subsequent period until they have completed the project. The number which invests is the

[18]This sort of reasoning is familiar in the literature on patent races. In particular Fudenberg et al. discuss the conditions under which a firm with a small advantage can be sure of winning the race [23].

smaller of all firms at the boundary, and sufficient firms to bring the total number of firms which complete the project, and which are at $x^b(k^t)$ or closer, up to the largest number of firms which could profitably complete the project were all firms at $x^b(k^t)$. In other words the number of firms at the boundary which invests is:

$$\min\left\{\#\left\{i \mid k_i^t \leq x^b(k^t)\right\}, \max\left\{n \mid u(x^b(k^t), n) \geq 0\right\}\right\} - \#\{i \mid k_i^t < x^b(k^t)\}$$

(iii) no other firm invests in any subsequent period.

PROOF: See Appendix A.2.1 ■

We illustrate equilibrium for an example.

Example 3 *$d = 3$, $I = 2$ and the signs on $u(x, n)$ are as in Table 2.4*

		x			
		3	2	1	0
n	1	+	+	+	+
	2	−	−	+	+

Table 2.4: The sign of $u(x, n)$ in Example 1

Equilibrium for Example 3 is given in Figure 2.16. The Figure shows the state space, and arrows between states show the transitions implied by firms' investment behaviour. Consider, for example, the state $k = (3, 2)$. Recall that the boundary distance $x^b(k)$ ~~which~~ is the furthest a firm can be from the end and still make a profit if it and all firms closer to the end invest fully in every period until they have completed the project. The boundary when $k = (3, 2)$ must be $x^b(k) = 2$. The boundary cannot be at 3: if the firm at 3 invested in every period until it completed the project, and so did the firm at 2, the firm at 3 would get a payoff of $u(3, 2)$, which from

Table 2.4, is negative. However, if just the firm at 2 were to invest until it had completed the project, it would get $u(2,1)$, which is positive. As a result, in the equilibrium path subsequent to $k = (3,2)$, firm 2 invests for 2 periods, while firm 1 never invests, as illustrated in Figure 2.16.

Not all strategies which satisfy Result 9 form an equilibrium, however. As an example, consider the strategies illustrated in Figure 2.17. The continuation from any state satisfies Result 9. Nonetheless, in state $(3,3)$ firm 2 can profitably deviate by investing for one period, causing the market to enter a state where it, rather than 1, finishes the project.

Our original concern was with the number of firms that would finish the project and we now turn to the equilibrium path when all firms begin with d stages to complete. A simple implication of Result 9 is that exactly $N(\sigma)$ firms finish, the same number that entered the sunk cost model, and these earn the same profits as did entrants in the sunk cost model.

Result 10 *Suppose initially $k_i^0 = d \, \forall \, i$ and that $I \geq N(\sigma) > 0$. In any pure strategy equilibrium exactly $N(\sigma)$ firms invest in each of the first d periods. Each of these firms has an initial value of $\hat{\delta}\Pi(N(\sigma)) - \sigma \geq 0$. No other firms invest at all.*

PROOF: This follows straightforwardly from Result 9. The boundary $x^b(k^0)$ is the furthest from completion a firm can be and still be sure of a profit if just it and all firms closer to completion invest in every period until they have reached the end. Since no firms are closer than d, and since $N(\sigma) > 0$, we must have $x^b(k^0) = d$. Result 9 then gives that exactly $\max\{n \,|\, u(d,n) \geq 0\} = N(\sigma)$ firms invest in each of the first d periods, and no other firms invest at all. The initial value of each of the investing firms is $u(d, N(\sigma)) = \hat{\delta}\Pi(N(\sigma)) - \sigma \geq 0$ ■

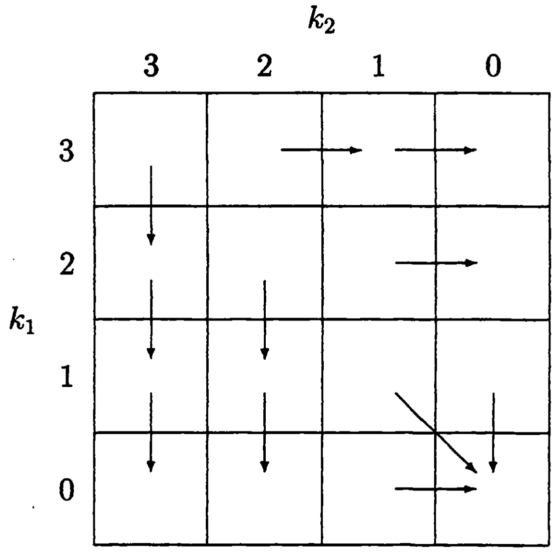

Figure 2.16: An example of pure strategy equilibrium

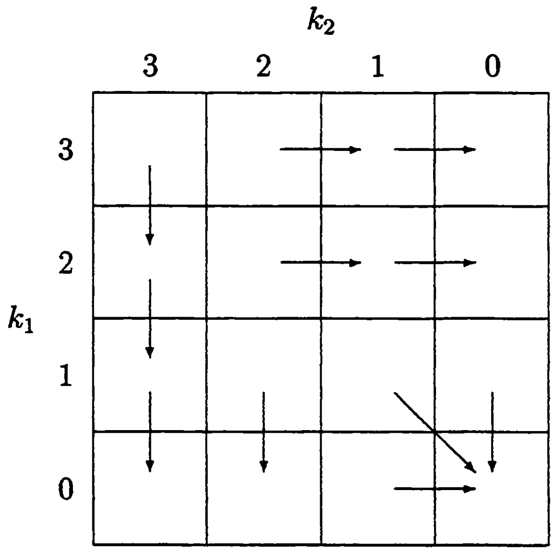

Figure 2.17: An example of strategies which satisfy Result 9 but which do not form an equilibrium.

This result confirms that concentration in the dynamic model is governed by the same need to cover entry costs as in the sunk cost model and, as in the sunk cost model, firms typically have positive values in pure strategy equilibrium.

Claims 2(B), 3(B), and 4(B). Convergence in the long run, excess entry and stationary fragmented states.

Also as in the sunk cost model there are other equilibria where firms sometimes choose mixed strategies. The general case is hard to analyse: later we consider a particular 3 firm example which, amongst other things, shows the kinds of problem we can expect to arise in the I-firm case. Most of the results here will relate just to the 2-firm case where the market supports only 1 firm. Here there is a unique symmetric equilibrium which is the same as the pure strategy one except in those symmetric states where, in pure strategy equilibrium, just one firm invests. In the symmetric equilibrium both firms choose mixed strategies in these states. Thus if one firm is strictly closer to completing the project it invests. If both firms are so close to completing the project that investing to the end is profitable, even if the market must subsequently be shared with the rival, then both invest. Finally if investing to the end is profitable only if the rival does not finish, firms choose mixed strategies.

To make the following expressions clearer we define $u^m(x) := u(x, 1)$ and $u^d(x) := u(x, 2)$. Also denote i's rival by j. We are able to give the value function for this case, which greatly simplifies the proof that what is proposed is indeed an equilibrium, though the proof that it is the only symmetric one is just as laborious as in the pure strategy case and is relegated to an appendix.

Result 11 *Suppose that $N = 1$ and $I = 2$ and consider the symmetric strategies $s^*(k)$ and value function $V(k)$ given by:*

$$s_i^*(k) = 0 \text{ if } k_i = 0, \text{ otherwise}$$

$$s_i^*(k) = \begin{cases} 1 & \text{if } u^d(k_i) \geq 0 \text{ or } k_i < k_j \\ \frac{1}{1-u^d(k_i)/u^m(k_i)} & \text{if } u^d(k_i) < 0,\ u^d(k_i-1) \geq 0 \text{ and } k_i = k_j \\ \frac{1}{1+c/u^m(k_i)} & \text{if } u^d(k_i-1) < 0 \text{ and } k_i = k_j \\ 0 & \text{otherwise} \end{cases}$$

$$V_i(k) = \begin{cases} u^m(k_i) - \delta^{k_j}\left(u^m(0) - u^d(0)\right) & \text{if } u^d(k_j) \geq 0 \text{ and } k_i < k_j \\ u^d(k_i) & \text{if } u^d(k_i) \geq 0 \text{ and } k_i \geq k_j \\ u^m(k_i) & \text{if } u^d(k_j) < 0 \text{ and } k_i < k_j \\ 0 & \text{otherwise} \end{cases}$$

(i) $s^*(k)$ is a perfect Nash equilibrium, and $V(k)$ is the corresponding value function.

(ii) $s^*(k)$ is the unique symmetric equilibrium.

PROOF: For part (i) we only need to confirm that the expressions form a fixed point of a set of Bellman equations[19] as follows. Let:

$$U_i^*(k, a_i) = -a_i c \qquad (2.6)$$
$$+s_j^*(k)\,\delta\, V_i\,(k_i - a_i, k_j - 1)$$
$$+\left(1 - s_j^*(k)\right)\,\delta\, V_i\,(k_i - a_i, k_j)$$

then the Bellman equations are, for each state and each firm:

$$V_i(k) = \max_{a_i}\ U_i^*(k, a_i)$$
$$s_i^*(k) \in \left\{\arg\max_{a_i}\ U_i^*(k, a_i)\right\}$$

We just give a couple of examples. First suppose $k_i = k_j = k'$ and $u^d(k' - 1) < 0$. The two firms are in the same state, but if both

[19] Tirole gives a simple justification for the one period deviation criterion underlying this approach [58, Page 265].

invested in each period until the project was complete, each would make a loss. According to Result 11, the only possible successor state in which firm i's value is not zero is $k = (k' - 1, k')$, where i has completed another stage but j has not. The value of i in this state is given as:

$$V_i(k'-1, k') = u^m(k'-1)$$

The Result also gives j's strategy as to invest with probability:

$$\frac{1}{1 + c/u^m(k')}$$

Substituting into Equation 2.6:

$$\begin{aligned} U_i^*\left((k', k'), 1\right) &= -c + \left(1 - \frac{1}{1 + c/u^m(k')}\right) \delta\, u^m(k'-1) \\ &= -c + \left(\frac{1}{1 + u^m(k')/c}\right) \delta\, u^m(k'-1) \end{aligned}$$

and

$$U_i^*\left((k', k'), 0\right) = 0$$

Since $u^m(k') = -c + \delta\, u^m(k'-1)$ this implies that $U_i^*(k, a_i) = 0$ whatever action i takes. In particular i can choose any mixed strategy, including the equilibrium one, and its value will be 0 as stated.

As a second example consider a state $k = (k_i, k_j)$ and suppose $u^d(k_j) \geq 0$ and $k_i < k_j$, so i has fewer stages left to complete, but both would have a positive payoff if both completed all stages. According to Result 11, j's equilibrium strategy is to invest. Substituting j's strategy and i's value in possible future states, as given in Result 11, into Equation 2.6, gives:

$$\begin{aligned} U_i^*(k, 1) &= -c + \delta u^m\left(k_i - 1\right) - \delta^{k_j}\left(u^m(0) - u^d(0)\right) \\ U_i^*(k, 0) &= \delta u^m\left(k_i\right) - \delta^{k_j}\left(u^m(0) - u^d(0)\right) \end{aligned}$$

Since $u^m(k') = -c + \delta\, u^m(k'-1)$ and $\delta < 1$, $U_i^*(k, a_i)$ is maximised when $a_i = 1$. In other words, i's optimal strategy is to invest, as stated in Result 11

The proof of part (ii) is the same as the proof of Result 9, except in those states where firms have the same state but only one invests in the pure strategy equilibrium. In these states firms choose mixed strategies. The full proof is in Appendix A.2.2. ∎

Figure 2.18 illustrates the state transitions implied by the symmetric equilibrium strategies for Example 3 on Page 83. A full arrow indicates that the transition is certain, a dotted one that the transition occurs with some probability. Firms choose mixed strategies in two states, when both have 3

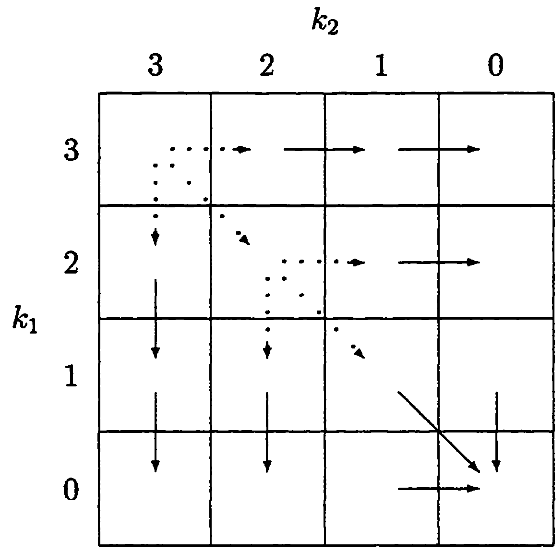

Figure 2.18: Mixed strategy equilibrium

and 2 stages yet to complete. The probabilities chosen in these states are:

$$s_i^*(3,3) = \frac{1}{1 + c/u^m(3)}, \quad s_i^*(2,2) = \frac{1}{1 - u^d(2)/u^m(2)}, \quad i = 1, 2$$

k_1 \ k_2	3	2	1	0
3	0	0	0	0
2	$u(2,1)$	0	0	0
1	$u(1,1)$	$u(1,1)$	$u(1,2)$	$u(1,2)$
0	$u(0,1)$	$u(0,1)$	$u(0,2)$ $+\pi(1)$ $-\pi(2)$	$u(0,2)$

Figure 2.19: Firm 1's value function in mixed strategy equilibrium

Figure 2.19 shows the value function for firm 1. This equilibrium has the excess entry and the possibility of stationary fragmented states that we noted were features of equilibrium in the dynamic advertising game in the previous section. It also has convergence of firms in the long run, though in a very trivial sense.

Claim 3(B). Excess Entry

In equilibrium we can see excess entry, in which both firms initially invest, though only one completes the project. Consider the symmetric equilibrium of Example 3, discussed above. With probability $(s_i^*(3,3))^2 > 0$ both invest in the first period, though the probability that both finish is less than unity. Denote the probability that both complete the first stage by $p(3)$. If both complete this stage one of two things must happen. Either both complete the stage in the first period, which happens with probability $(s_i^*(3,3))^2$, or neither complete the stage in the first period, and they subsequently both complete the stage, which happens with probability $(1 - s_i^*(3,3))^2\ p(3)$. On rearranging this gives the probability that both firms complete the first stage as:

$$\begin{aligned} p(3) &= (s_i^*(3,3))^2 + p(3)\ (1 - s_i^*(3,3))^2 \\ &= \frac{1}{1 + \frac{2c}{u^m(3)}} \in (0,1) \end{aligned}$$

Likewise the probability that both firms complete the second stage together, given that they complete the first together, is:

$$p(2) = \frac{1}{1 - \frac{2u^d(2)}{u^m(2)}} \in (0,1)$$

If both firms complete the second stage together, the state will be $k = (1,1)$, and subsequently both will certainly invest and complete the project. The probability that both finish is therefore $p(3) \cdot p(2) < 1$.

The excess entry in this example is associated with rent dissipation: whereas in pure strategy equilibrium the firm which entered had an initial value of $\hat{\delta}\Pi(1) - \sigma > 0$, here both firms have an initial value of 0. Excess entry arises as the two firms compete to be the one firm that the market will support. Early on a firm invests in the hope that it will become the only firm which completes all stages and produces: as we saw in the pure strategy equilibrium. Were a firm certain that its rival would finish, it would not invest. In the symmetric equilibrium a firm which falls behind early on cannot catch up to its rival, and so cannot be the sole producing firm. A lagging firm will therefore stop investing and fall increasingly behind, as in asymmetric states when at least one firm has two or more stages to complete in Figure 2.18.

Claim 4(B). Stationary fragmented states

Although both firms enter in the hope that their rival exits, there is some probability that their competition never produces a winner and both finish the project. This probability was calculated above as $p(3) \cdot p(2)$. Of course once both firms are producing they stay producing forever. We see here that it is possible for stationary fragmented states to arise: they are a natural outcome when firms compete to remain in the market by hoping to build up stocks of a strategic variable faster than rivals. The possibility that both firms finish does not just result from the fact that firms complete the project in a finite number of discrete stages. In the limit when the number of stages d is countably infinite, the probability that both finish is bounded away from zero. In order that both finish, both must complete each of the stages for which $u^d(x)$ is negative at the same time as its rival (since otherwise the lagging firm would cease investing), and so the number of events which must all be true gets large. However, the cost of completing each stage becomes very small as each stage covers less ground, and the probability that both complete a stage together tends to 1 at a rate which exactly offsets the increase in the number of stages.

Result 12 *Let ϕ be the probability that both firms finish the project.*

$$0 < \lim_{d \to \infty} \phi$$

PROOF: See Appendix A.2.3 ■

Claim 2(B). Convergence in the long run

The symmetric equilibrium has convergence of active firms in a somewhat forced sense. If one firm stops investing leaving a single firm active we cannot say whether there is convergence or otherwise of active firms. If both complete enough stages then each completes all remaining stages no matter what its rival does and here we can say firms are in a convergent region. In Figure 2.18 showing equilibrium strategies in Example 1 this convergent region is the one where both firms have completed at least 2 stages. If $k = (1, 0)$ the lagging firm will invest and catch up. However, although in the long run either just one firm is active or firms are in a convergent region, firms never actually converge: if the initial state is $k = (3, 3)$, $k = (1, 0)$ can never arise. Actual convergence here could arise only in response to unforeseen shocks or mistakes.

A more meaningful case of convergence in the long run is seen in an example where the market supports 2 firms and 3 compete to be among them, as in Example 2 below.

Example 4 *$d = 2$, $I = 3$ and the signs on $u(x, n)$ are as given in Table 2.5*

		x = 2	1	0
	1	+	+	+
n	2	+	+	+
	3	−	+	+

Table 2.5: The sign of $u(x, n)$ in Example 2

Result 13 *The unique symmetric equilibrium in Example 2 is given in Table 2.6*

PROOF: See Appendix A.2.4. ■

In the first period all firms invest with probability q_2. If, when actions are realised, just one firm invests, the state next period will be $k(1, 2, 2)$. The leading firm is now sure that it will complete the project, and invests and completes the project with certainty, while the lagging firms spend time resolving which of them will also complete the project, and each invests with

k	$s^*(k)$	
(0 0 0)	(0 0 0)	
(0 1 1)	(0 1 1)	
(0 1 2)	(0 1 0)	
(0 2 2)	(0 q_1 q_1)	$q_1 = \frac{1}{1-u(2,3)/u(2,2)}$
(1 1 1)	(1 1 1)	
(1 1 2)	(1 1 0)	
(1 2 2)	(1 q_1 q_1)	$q_1 = \frac{1}{1-u(2,3)/u(2,2)}$
(2 2 2)	(q_2 q_2 q_2)	$0 < q_2 < 1$

Table 2.6: Equilibrium strategies in Example 2

probability $q_1 < 1$. At least one of them is bound to finish, and so to converge on the state of the firm which has already finished. Convergence follows straightforwardly from the fact that once they are sure they will remain in the market all firms invest until they reach the exogenously determined end at which they can earn profits, and so, even if at some point finishing firms have different numbers of stages to complete, they are bound to end up with the same number of stages left, viz none.

Result 13 not only confirms that we can see firms catching up in equilibrium. The proof also brings to light a curious consequence of the possibility that stationary fragmented states can arise: a firm's value may increase if a rival has fewer stages to complete. The reason is that there is some probability that two distant firms with the same number of stages to complete both end up finishing the project, even if a closer rival is certain to finish and there is only room for two firms. If one of the distant firms had fewer stages to complete, just it and the close firm would finish. Consider the value of the close firm. If the profit when two share the market is not much less than when one firm is producing alone, the close firm will be unconcerned that a rival finishes sooner if it has fewer stages to complete. However if the profits when three firms share the market are low, the possibility that both rivals may finish can have a large negative impact on the close firm's profits. The proof of Result 13 confirms that the value of the close firm in state $(1,2,2)$ can be lower than that in state $(1,1,2)$. This possibility means that we need to confirm explicitly that at the start a firm will find it optimal to invest if neither rival invests, which in this case is sufficient to confirm that firms do initially choose non-zero probabilities. It is this complication which makes it difficult to generalise the symmetric equilibrium to more than one firm.

2.4 Conclusions

All of the features of equilibrium in the incremental sunk cost model of advertising discussed in Section 2 are also features of the incremental exogenous sunk cost model here: a well defined market structure, convergence of firms in the long run, excess entry and divergence of lagging firms early on, and the possibility that stable fragmented states arise. In the exogenous sunk cost case these features are clearly consequences of two factors. First firms must expect to recover costs sunk in establishing themselves in the market. Their profits fall as the number of rivals rises, yet the amount spent on becoming established is fixed. The result of this is that the number of firms which can profitably establish themselves is restricted. Second, since costs are sunk incrementally and, in some equilibria, investment is stochastic, many firms may begin accumulating. If they are luckier than their rivals they will establish themselves in the market in the long run and make profits. If they fall behind early on they can stop investing without having lost too much. The advertising model is more complicated than this exogenous sunk cost model. In the advertising case there is no exogenously given end. Firms earn profits in all states, and the states in which firms stop net accumulation of quality, which we can take as states in which firms are established, are endogenous. The similarity of the results in the two cases however, suggests that the same factors dominate dynamic market behaviour in both.

Chapter 3

Reversible Location Choice: Bus Deregulation in the UK

3.1 Introduction

The 1930 Road Traffic Act created a bus and coach market in which all aspects of service were tightly regulated. In order to run a service, an operator had to meet prescribed standards of vehicle safety and driver competence and, more restrictively, acquire a Road Service Licence from the Traffic Commissioners. A licence would only be issued if the applicant could show that its service was in the public interest. In practice permission would often not be granted if existing licence holders, or British Rail[1], objected, creating a barrier to the entry of independent operators. Permission was also required for changes to fares or timetables for existing services, and again the onus was on the applicant to prove that such changes were in the public interest. Since British Rail could raise objections, this restricted the ability of the

[1]Or, prior to 1948, the various railway companies.

incumbents themselves to compete effectively against rail.

The 1980 Transport Act abolished the need for operators to obtain a Road Service Licence for express services (defined as those carrying passengers a minimum distance of 30 miles in a straight line), thereby allowing entrants to compete against what had effectively been the protected monopoly incumbent, National Express, a marketing arm of the publicly owned National Bus Company. Since the Act, operators of express services have simply needed to notify the Traffic Commissioners 28 days before starting a new service. The 1985 Transport Act which followed privatised the incumbent operators of express coaches, and deregulated the local bus markets. Subsequent to the Act an operator needed only to register its timetable and satisfy basic safety requirements in order to run a local bus service. The main remaining restriction was that the Traffic Commissioners had to be notified of all new services, and changes to existing services, 42 days in advance.

A survey of the secondary literature discussing the effects of these Acts suggests that while some effects are common to both local bus and express coach markets, for example both have emerged with a surprisingly concentrated market structure, there are also some striking differences. When comparing conditions in those urban local bus markets in which there was competition with the situation prior to deregulation, or to those areas where one operator was dominant, observers have found that there was little, if any, fall in price, and that the frequency with which timetables were revised was very high. With express coaches on the other hand, when there was competition on the road then observers found a marked reduction in prices, and they made no mention of timetable instability.

This paper puts forward the hypothesis that these two differences in the effects of competition stem from a basic difference in passenger behaviour.

Passengers boarding express coaches know the timetable, indeed they often book their trip in advance, and they arrive at the terminal in time to board their most preferred coach. In the local bus markets where the instability and high fares are most marked, the metropolitan areas where bus frequencies are particularly high, passengers intending to travel by bus arrive at the stop independently of bus arrival times and wait for a bus to arrive. They may have some idea about what fare they expect to pay, and how long they expect to wait, but they do not know the exact arrival times.

The paper constructs simple models of a local urban bus market and an express coach market which embody these contrasting assumptions on passenger behaviour. There are two main results. First, arrival times change from day to day when there are two competing operators in a local bus market. If it is costly to revise the timetable, the later bus is the more likely to change arrival time, and so choose a time just before the early bus. One consequence of this is that bus arrival times tend to be bunched together. Competition is vital to this result: if all buses are run by a single firm, arrival times are never changed between periods. The reason for the instability is straightforward. Neither operator will stick to a fixed timetable. If it did its rival would start to arrive just before it and would take all the passengers.

Turning to the fare level set by buses, we show that in the express coach market fares are lower when there are two firms rather than one. In stark contrast the paper presents an example model of a local bus market in which competing firms charge a higher price than would a monopolist. If one firm lowers its fare more passengers will decide to travel by bus. However, since they arrive at the stop independently of the arrival times of the buses, and then board the next bus, only some will board the bus which lowered its fare; the rest will board its rival. If all buses were operated by a single firm, that firm would gain all the additional passengers that made journeys as a result of a fall in fares. The monopolist therefore has a greater incentive to cut fares than a duopolist does.

3.2 Bus Deregulation: the Stylised Facts

3.2.1 Local Buses

In 1995 the select committee on transport produced a report on the effects of bus deregulation [8]. Among the many issues raised by the committee are four relating to the impact of competition between buses, operating the same route, on the performance of the market; on the stability of the timetable, on fare levels, on the frequency of service, and on the quality of buses used.[2]

It is clear from the report that the frequency with which bus timetables are updated is a considerable source of irritation to bus users, and that this instability is a feature of on-the-road competition. The rise in timetable instability was the first point raised by the National Federation of Bus Users (NFBU) in its evidence to the committee and the Road Traffic Commissioners agreed that this was the problem which was of most concern to passengers.

[2] There are many other issues raised in the report, which we do not discuss here. For example, rival bus operators on the same network do not always allow through ticketing. The committee spent a lot of time considering whether or not bus markets are contestable, the issue which dominated the theoretical debate while the 1985 legislation was being drawn up, and which is relevant to the policy question of whether the government should view consolidation in the industry with equanimity. Another feature of the deregulated market which had surprised observers is the relative failure of high quality minibus services to develop alongside full size bus services. Such differentiated markets are common in South East Asia. The committee also considered the vexed question of whether or not there has been predation in local buses.

Indeed those Commissioners questioned by the committee would have liked powers to restrict the frequency with which changes to the timetable could be made. In many places the cost of keeping passengers informed about timetables is borne by the local authority and those councillors from the metropolitan areas, where on-the-road competition is most common, testified to the committee that the frequent timetable changes were very expensive to them. Councillor McLellan from Strathclyde, where four large operators and numerous small ones competed, testified that whereas before deregulation timetables were produced at 3 or 6 month intervals, since deregulation 5 timetable changes had been notified to the Strathclyde Traffic Commissioners every working day. The NFBU observed that frequent changes to the timetable are associated with another feature: bus departure times tend to be bunched together, and they added that two operators running practically identical timetables hardly increased customer choice. In its final report the committee made a similar point, stating that entrants into bus markets typically registered times just before those of the incumbents.

There is anecdotal evidence from previous periods of unregulated bus services of similar behaviour. Glaister, in his evidence to the committee, cites the example of the horse bus Associations in 19th century London, which had to make great efforts to enforce service regularity on their members in order not to alienate their passengers. Another example is given by Chester [7][3] in a book written 6 years after the 1930 Act which first introduced some control into London's bus markets. There he described the ills of "unfettered competition", and argued that such competition means

> the running of vehicles to a regular timetable will become impossible.

[3] Cited by Mackie & Preston [41]

A number of witnesses to the select committee commented that fares in the deregulated bus industry have not fallen, as might have been expected, but have risen(for example see Glaister's evidence). In a recent book on the effect of deregulation on the local bus market Mackie & Preston [41], using data supplied by the Department of Transport, confirm the rises in fares and give the following table comparing the rise between metropolitan and shire counties. On-the-road competition is more common in metropolitan

	Fares
English metropolitan counties	+49
English shires	+8.8

Table 3.1: %change 1985/86 - 1993/94

areas, so that the fare increases are greatest where competition is greatest. Both Glaister and the NFBU, in their evience to the select committee, noted that fares have been affected not just by the changes in entry conditions, but also by other changes in market conditions; the removal of subsidies at the time of deregulation and subsequent increases in the duty on diesel. However, the evidence from previous periods of unregulated competition in local buses supports the view that competition does not result in low fares. Barker & Robbins, in their "History of London Transport" [2] discuss the fierce competition on London's roads in the 1920's and note that:

> Perhaps the most interesting feature of the bus competition... is the fact that there was very little attempt at competition in fares..

Instead of cutting fares, various commentators have noted that competing local bus operators put on more buses, leading to very low bus loads and a great deal of city congestion.

A central concern of the Select Committee was safety, and in particular whether there was adequate investment in new buses. This concern can be seen as part of the wider question of the impact of deregulation on the quality of service. In much of the existing literature on bus deregulation, quality was identified with service frequency [14]. However, the actual behaviour in the deregulated markets suggests that service frequency and bus quality behave differently. Frequency has been mentioned earlier; here we just consider bus and service quality. There seems to be a basic disagreement about the effects of competition on the quality of service. On the one hand in MMC report Cm2423 we find

> competition and potential competition, in our view, are the main safeguards against ... lower services.

And on the other hand the view of the NFBU, given in its evidence to the select committee [8], is that

> on the whole, competition is wasteful. Bristol, with a dominant operator, probably has the best service...

a view which White endorses in his evidence.

Looking specifically at the issue of investment in new buses, much of the evidence suggested that investment was highest where on-the-road competition was absent. The NFBU noted that the lack of modern buses is most marked in Sheffield and Manchester, where competition is very fierce, and Glaister pointed out that the original 1930's legislation introducing regulation was brought in specifically to create barriers to entry in order to improve the then poor safety standards[4]. A dissenting voice is that of Norris, then a

[4]It should be noted that this poor record was due in part to bad behaviour by drivers, and not just to inadequate maintenance.

junior Minister for Transport. He claims that there is no evidence that new buses are only being bought by de facto monopolies. Finally, White considers a wider notion of quality and notes that service innovation has occurred in areas of relatively little inter-operator competition, and has probably arisen as bus companies compete against other methods of transport, such as walking or the car, rather than competing with other operators. In summary, the majority view seems to be that quality is higher when there is no competition.

3.2.2 Express Coaches

One of the purposes behind the 1980 legislation was to introduce institutional changes that would lead to improvements in transport services [11] and it does seem that since 1980 there have been significant changes to prices, quality, frequency and the coverage of the coach network, and these changes have on the whole benefited passengers. Deregulation has led to changes in both the level and structure of prices. The structures of concessionary fares and of different types of ticket available from National Express have seen changes that Thompson & Whitfield [57] characterise as making prices simpler and more customer friendly. Price levels have also changed: they have fallen. The initial fall was dramatic, and even though prices began to rise again from 1982 onwards, in 1990 they remained below the 1980 levels in real terms [57]. Prices could have fallen for a number of reasons, among them that the pre-deregulated prices were higher than even a protected monopolist would want, that coaches were trying to take market share from rail, that operators kept prices low so as not to attract entry, or that low prices resulted from competition between operators. Of course all of these, and other, mechanisms could have contributed something to the fall in prices, and various writers, particularly Thompson and his co-authors [57] [36], have attempted to disentangle their effects. The main conclusion to emerge from this work is that the fall in prices is greater when there is actual competition.

As well as its effects on prices, deregulation led to a dramatic increase in frequencies on routes, and an extension of the coach network. Before the 1980 Act the coach network had been declining [3], yet 4 years after the Act Barton & Everest [3] found that an extra 48 million vehicle kms were being operated by new express services. The frequency of services has also seen important changes. Although Robbins & White [49] found, in 1986, that frequencies had fallen on some minor provincial routes, frequencies have on the whole seen dramatic increases. Thompson & Whitfield [57] examined how the frequency rise depended on the amount and success of entry. They found that frequencies rose by over 700% on routes within Scotland, where entry was most successful, whereas on UK trunk routes, where National Express has maintained dominance, frequency rose by just 179%. The picture is not entirely clear, but it does seem that competition pushes frequencies up.

There have been a number of innovations in the type of service offered, including a large expansion of airport linked services, and the introduction of luxury services offering such facilities as videos, refreshments, hostesses and the like [11]. Many of these innovations were introduced first by entrants, and indeed those entrants which did survive often did so by offering upmarket services [57]. Journey times, another indicator of passenger convenience, have also fallen [57].

Thus far the behaviour of deregulated express coach markets is roughly in line with what was expected: removing barriers to entry and allowing operators to compete against British Rail led to lower prices, higher quality (and vertically differentiated products) and more travel. However, the market has provided one major surprise. It was widely expected that the dominance of National Express would dissolve on deregulation [11], yet despite an initially high rate of entry it has remained dominant, and has thrived.

3.2.3 Comparing Local Buses and Express Coaches

Some effects are common to both markets, most notably that both have emerged with a surprisingly concentrated market structure, and that in both cases, when operators competed on the road, then the number of buses/coaches rose dramatically. There are also some striking differences. Briefly, when comparing conditions in those urban local bus markets in which there was competition with the situation prior to deregulation, or to those areas where one operator was dominant, observers have found that there was little, if any, fall in price, that buses were old and badly maintained, that, despite the prediction of many writers that deregulation would lead to the appearance of services of differing qualities, no vertical differentiation emerged, and that there was a very considerable amount of timetable instability. With express coaches on the other hand, when there was competition on the road then observers found a marked reduction in prices, they made no mention of timetable instability, or of a fall in the cleanliness and newness of coaches, and they commented that entrants often introduced higher quality services than those of the incumbent, although in some cases the incumbent itself offered a differentiated service, even in the absence of competitors.

Two particularly striking and clear cut differences in the effects of on the road competition between urban local buses and express coaches stand out from this summary: first, competition leads to substantial instability in the arrival times of local buses, but not of express coaches, and second that prices remain surprisingly high under competition in local bus markets, but fall markedly when express coaches compete. The rest of this paper seeks to support the hypothesis that both differences arise from the same difference in passenger behaviour between the two markets. Before turning to this hypothesis a few words are in order to clarify the phenomena the paper seeks to explain.

?

It is clear what is meant by timetable instability: it arises when operators frequently change the arrival times of their buses. However, there are different senses in which buses could be bunched. It seems clear that congestion and random variations in journey times can cause buses to bunch together, even if their published arrival times at a stop are well spaced. Also, bus arrivals may be concentrated at particular times of the day simply because there are peaks in demand. This chapter will be concerned with neither of these types of bunching. Rather it will seek to explain the kind of bunching which the NFBU complained of when it pointed out that two operators running timetables which were almost identical, for example where both ran an hourly service, with one arriving on the hour and one arriving 5 minutes past, hardly increased passenger choice.

3.3 Explaining the Differences

There are some similarities in the impact of competition on market performance between express coaches and local buses. In both cases competition is associated with higher frequency, though this is perhaps more marked in local buses, and in both the level of concentration is surprisingly high. There are also some differences between the two markets. Competition is associated with frequent changes to the timetable on local bus, but not express coach, routes, and competition leads to reduced fares on express coach, but not local bus, routes. There is also a difference in the impact of competition on the quality of bus services. Competition in local bus markets seems to reduce the quality of buses used, while competition on express coach routes leads to innovation and improved service. Timetable instability, high fares and low quality buses are particularly marked features of local bus routes in metropolitan areas and for the rest of the chapter we focus on high density urban local bus routes.

This paper puts forward the hypothesis that these differences arise from a basic difference in passenger behaviour. On the one hand the time at which passengers on urban local bus routes arrive at the stop is taken to be independent of the arrival times of buses, while on the other hand passengers on express coach routes are assumed to arrive just in time to board their most preferred coach.

The assumption that on urban local bus routes passengers arrive at a bus stop independently of the arrival time of buses has independent support. Savage [52] cites work on passenger waiting times in Greater Manchester. It was found that when the intervals between buses are comparatively short there is a random element in the arrival patterns of potential passengers at stops, so that as frequency increases the average waiting time falls. In fact if the intervals are less than around 12 minutes, then arrivals become totally random. Savage in his own empirical work on competition on selected bus routes assumes that a bus arriving just before its rival will get all the market. This behaviour on the part of passengers has also been put forward by others as important to understanding the effect of competition on timetable stability and low reductions in fares on local bus routes. Chester, who, as we saw above, argued that 'unfettered competition' on local bus routes would undermine regular timetables [7] went on to give the reason for this as that if any operator fixed definite times, rival operators will seek to reach stopping places a few minutes earlier and take the traffic.

Despite the obvious appeal of this mechanism, early attempts to formally model timetable instability were structured so as to exclude the mechanism's operation. Foster & Golay [22] used a Hotelling framework in which passengers have an ideal departure time. One component of the cost of making a journey is an item which increases as the difference between the actual departure time and the ideal gets larger. To the passenger it is unimportant whether a bus arrives before or after its ideal time. It follows that the benefit to be gained by pre-empting the rival is offset by the loss incurred as a result of the increased separation from the preceding service. They identify instability with lack of equilibrium, and since an equilibrium *does* exist in this location model[5], they conclude that there will not be instability. Subsequent work on the choice of arrival time has used a similar framework[6]. A slightly different perspective on timetable choice is provided by Glaister in his evidence to the Select Committee. He suggests that irregularity arises because there are revenue benefits from service regularity which are external to the individual operator but internal to the market as a whole. Presumably there is a market benefit because demand is higher for a regular service, and this

[5]The authors assume sequential entry. Moreover in proving the existence of equilibrium the authors rely on the modified zero conjectural variation introduced in Novshek [44], so that they do not show that a pure strategy Nash Equilibrium necessarily exists. Their result can be seen as part of the debate about the conditions under which a pure strategy Nash Equilibrium exists in Hotelling location games, when firms choose both price and location. d'Aspremont et al [10] pointed out that when transport costs are linear there is not neccessarily a pure strategy price equilibium when locations are too close together, but there is an equilibrium when costs are quadratic.

[6]See the papers by Foster & Golay [22], Evans [19], Dodgson et al [15], [16].

demand benefits all operators, not just the one whose choice of timetable led to a more stable and regularly spaced service. This observation will only imply an underprovision of regularity, however, if there is a private gain to creating irregularity. Glaister leaves the source of this gain unexplained.

The hypothesis examined in this chapter is that timetables are unstable when there is competition on local bus routes for exactly the reason set out by Chester. The formal model we consider also predicts that bus arrivals will be bunched together. The select committee concluded that this bunching did occur on local bus routes. Finally, the model predicts that the operator most likely to update its timetable is the one whose buses currently collect least passengers. Many witnesses[7] testified to the select committee that passengers on high density bus routes take the first bus to arrive, regardless of price differentials, and that this undermines attempts by bus operators to win market share through cuts in fares. This explanation was also put forward by Mackie & Preston [41] as a reason why fares remained so high. Many go on to note that competition on local bus routes focuses on being first rather than cheapest, and as a consequence operators put many buses on a route.

This chapter explores the fares set on local bus routes. Passenger behaviour in the model we develop has two important features: passengers arrive at the stop independently of bus arrival times, and passengers will board the first bus to arrive even if its fare is a little higher than that charged on the next bus. We show that when passenger behaviour has these features, not only is the incentive to cut fares reduced, competition between two operators can lead to higher fares than a monopolist would set. The reason for this is two-fold. First, as noted by previous commentators, if a bus company

[7] See the evidence of White and the TGWU

cuts its fares it does not increase its share of passengers. However, the the number of passengers which make bus journeys will increase. Some of these additional passengers will board the rival bus, giving a positive externality which leads to under-investment in fare reductions.

The reasons put forward in the Select Committee report to explain why investment in new buses is low when there is competition rely on financial constraints: essentially the profits which firms earn when there is competition are too low to cover the investment costs. However, in the concluding section of the chapter we suggest that when passengers board the first bus to arrive competition will lead to under-provion of investment in much the same way as it leads to under-provision of fare reductions.

3.4 Timetable Instability and Bunching

3.4.1 Local Buses

In the model in this section frequent changes to the timetable are a natural result of on-the-road competition when passengers arrive evenly throughout the day and board the first bus to arrive. This behaviour means that two competing operators running alternate buses will try to schedule their buses to arrive as late as possible after their rival's, as the later after one bus the next one arrives, the more passengers will be waiting. Both operators cannot simultaneously choose arrival times just before those of its rival. Each operator will keep its rival guessing as to its arrival time since, if it chose any time with certainty, its rival would arrive just before it and leave it with no passengers.

That each bus operator will keep its rival guessing would, on its own, induce bus operators to choose all possible arrival times with equal probability. We also assume, however, that there is a cost to revising the timetable. This is in fact likely to be the case. Timetable changes must be registered, giving rise to at least some administrative costs. Other costs arise from the managerial time needed to decide on a change, and on the form of the new timetable. Moreover, since timetables do not come into effect until 42 days after initial registration, deciding on a change will involve planning and research into the rival's planned actions. The cost has a striking effect on the pattern of timetable changes: the bus operator whose bus, yesterday, arrived just before its rival's, and so had most passengers, is more likely not to revise its timetable today at all, while the other is more likely to change so that its buses arrive just before the time its rival's arrived yesterday. The tendency is for buses to leapfrog each other in order to arrive earlier and earlier. One result of this behaviour is that bus arrivals tend to be bunched together as each bus operator, if it revises its timetable at all, will choose a new time just before its rival's old one.

The model is highly stylised in order to draw out the effects on timetable stability of the twin assumptions that passengers board the first bus to arrive, and timetable revision is costly. In particular we do not endogenise passenger boarding behaviour, it is just taken as a primitive of the model. Also we will treat a day as circular in order to focus attention purely on the question of whether firms want their buses to arrive before or after those of their rivals, without the complications caused by end effects.

We first give a simple discrete example which shows instability, bunching and leapfrogging. The full continuous model draws out the underlying mechanisms more clearly.

As in the earlier work on timetable choice, instability in the model here will arise when a pure strategy equilibrium does not exist. However, the lack of such an equilibrium here is more fundamental than that discussed in early formulations of the Hotelling location game. In those cases the lack of equilibrium arose because of problems in optimal pricing when firms were located too close to each other, and the problem could be resolved through a suitable choice of cost function[8]. Here there is no pricing problem. Instability results directly from a lack of pure strategy equilibrium in the choice of location.

3.4.1.1 An Example

Two buses, A and B, compete to pick up passengers during each of an infinite number of days. Each day has 4 minutes arranged around a circle, so that minute 3 is just before minute 0. Figure 3.1 illustrates a day.

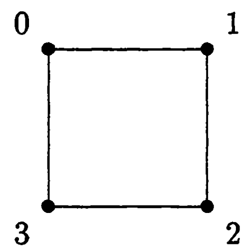

Figure 3.1: A day

In period t bus i, $i = A, B$, picks an arrival time a_i^t, $a_i^t \in \{0, \ldots, 3\}$. The state in period t, denoted k^t, is the arrival times of buses in the previous period, $k^t = a^{t-1} = (a_A^{t-1}, a_B^{t-1})$.

One unit of passengers arrive every minute and if a bus arrives in the same minute they board it, otherwise they board the first bus to arrive. Two buses arriving at the same time share the waiting passengers equally.

[8]See Footnote 5.

A bus gets a gross profit of 1 for every unit of passengers which board. Let $\pi_i(a_A, a_B)$ be the gross payoff of bus i when arrival times are a_A, a_B. The complete gross profit matrix is given in Table 3.2 which shows the pair π_A, π_B for each possible combination of arrival times.

		a_B							
		0		1		2		3	
	0	2	2	3	1	2	2	1	3
a_A	1	1	3	2	2	3	1	2	2
	2	2	2	1	3	2	2	3	1
	3	3	1	2	2	1	3	2	2

Table 3.2: Table of values of π_A, π_B

Buses also incur a cost of c if they revise their timetable. Let $C_i(a_i^t)$ be the cost bus i pays if it arrives at a_i^t in period t. Then:

$$C_i(a_i^t) = \begin{cases} 0 & \text{if } a_i^t = k_i^t \\ c & \text{otherwise} \end{cases}$$

For simplicity we assume that each bus i is myopic: it seeks only to maximise the expected current profit net of any revision cost. However, we will see that this assumption is not as restrictive as it appears: even if buses maximised the sum of discounted future net profits the equilibrium strategies would be the same as the ones we find here. The problem is essentially a one period one (though the past influences the present through the state) and we now drop the time superscript. A strategy for bus i, denoted s_i, specifies the probability that i chooses each arrival time, so that $s_i(m)$ is the probability that $a_i = m$. Strategies can be conditioned on the state, but on nothing else. In any equilibrium $s^* = (s_A^*, s_B^*)$ any arrival time chosen with positive probability must maximise expected profit net of cost, given the rival's strategy, i.e. if $s_i^*(m) > 0$ then m maximises $E\left[\pi(m, a_j) - C(m) \mid s_j^*\right]$.

It is easy to see that if the revision cost is less than 1 firms will never choose a pure strategy in equilibrium, no matter what the state. Suppose A arrived in minute 0 with certainty. If B arrived at 3 it would collect 3 units of passengers, and pay a maximum cost of c, giving a net profit of more than 2. If B arrived at any other time it could collect at most 2 units, and so it will certainly arrive at 3. But if B is arriving at 3, A does best to arrive just before at 2, and so on. When $c < 1$ the only equilibrium is in mixed strategies. Each bus randomises to keep its rival guessing as to exactly when it will arrive. Equilibrium strategies are shown in Tables 3.3 and 3.4 below.

k_B	$s_A^*(0)$	$s_A^*(1)$	$s_A^*(2)$	$s_A^*(3)$
0	$\frac{1}{2}$	–	$\frac{1}{2}-\frac{c}{2}$	$\frac{c}{2}$
1	$\frac{1}{2}+\frac{c}{4}$	–	$\frac{1}{2}-\frac{c}{4}$	–
2	$\frac{1}{2}$	$\frac{c}{2}$	$\frac{1}{2}-\frac{c}{2}$	–
3	$\frac{1}{2}-\frac{c}{4}$	–	$\frac{1}{2}+\frac{c}{4}$	–

Table 3.3: A's equilibrium strategy

k_B	$s_B^*(0)$	$s_B^*(1)$	$s_B^*(2)$	$s_B^*(3)$
0	$\frac{1}{2}$	–	$\frac{1}{2}-\frac{c}{2}$	$\frac{c}{2}$
1	–	$\frac{1}{2}-\frac{c}{4}$	–	$\frac{1}{2}+\frac{c}{4}$
2	$\frac{1}{2}-\frac{c}{2}$	–	$\frac{1}{2}$	$\frac{c}{2}$
3	–	$\frac{1}{2}-\frac{c}{4}$	–	$\frac{1}{2}+\frac{c}{4}$

Table 3.4: B's equilibrium strategy

In both tables we assume that A arrived in minute 0 yesterday, i.e. $k_A = 0$. The state is therefore summarised just by k_B. We can always ensure that $k_A = 0$ simply by relabelling the minutes at the start of the current period.

Before examining what these strategies imply for the pattern of timetable revisions, we first confirm that they do form an equilibrium. To do this we need just show that each firm only chooses an arrival time with positive probability if arriving at that time maximises its expected net profit given its rival's strategy. Table 3.5 sets out the expected net profit of A when B plays the strategy given in Table 3.4. Denote A's expected profit when it arrives at m and B plays s_B^* by $\Pi_A(m)$. It is clear by inspection that, given

k_B	$\Pi_A(0)$	$\Pi_A(1)$	$\Pi_A(2)$	$\Pi_A(3)$
0	$2-\frac{c}{2}$	$2-\frac{3c}{2}$	$2-\frac{c}{2}$	$2-\frac{c}{2}$
1	$2-\frac{c}{2}$	$2-c$	$2-\frac{c}{2}$	$2-c$
2	$2-\frac{c}{2}$	$2-\frac{c}{2}$	$2-\frac{c}{2}$	$2-\frac{3c}{2}$
3	$2-\frac{c}{2}$	$2-c$	$2-\frac{c}{2}$	$2-c$

Table 3.5: A's expected profit, given s_B^*

the state k_B, the strategy given in Table 3.3 only assigns positive probability to those arrival times which maximise A's expected profit. A similar table to Table 3.5 could readily be found for B and this would show that the strategy given in Table 3.4 likewise only assigns positive probability to those arrival times which maximise B's expected profit. This confirms that Tables 3.3 and 3.4 do specify an equilibrium in mixed strategies. Moreover this is the only equilibrium when $0 < c < 1$, though to check this requires an exhaustive search of other possibilities and the results of this search are not repeated here. One final general feature of the equilibrium is that a bus operator's expected net profits do not depend on the state: they are always $2 - c/2$. Since the current period can only affect the future through the state, this means that the current period has no effect on future net profits and firms would not change their behaviour if they were not myopic[9].

Turning to the implications of these equilibrium strategies for the pattern of timetable revision, the case that is of particular interest is when the buses arrived in two successive minutes yesterday, so either $k_B = 1$ if B arrived just

[9]Suppose buses maximise the discounted sum of net profits, the discount rate is δ and the revision cost $c\delta$, then there is a perfect eqilibrium in which strategies are identical to those found here.

after A, or $k_B = 3$ if it arrived just before. The strategies of the two buses are shown diagrammatically in Figure 3.2 for the case where $k_B = 1$. According

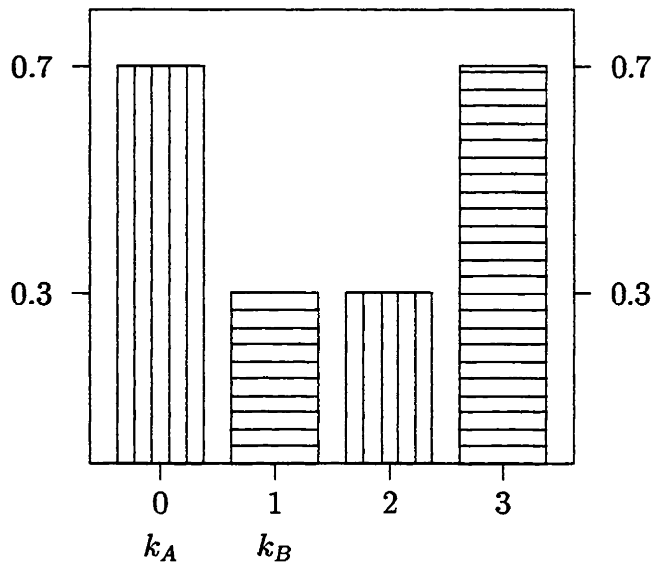

Figure 3.2: Strategies when $k_B = 1$ and $c = 4/5$

to equilibrium strategies either A arrives at a particular time with positive probability, or B does, but not both. In the Figure there is a bar at each minute whose height is proportional to the probability that a bus arrives at that minute: if it is bus A the bar has vertical stripes, if B horizontal. We assume that $c = 4/5$.

The reason why this case is the most important is that whatever the actual realisations of firms' random strategies, they will never arrive at the same time, and neither will they arrive evenly spaced: if buses arrived one after the other in the previous day, they are bound to arrive one after the other in the current day, and so in all future days. When $k_B = 0$ or 2, the only other possible cases, firms randomise over three possible arrival times, and so with positive probability arrive one after the other in the current day, and if not in the current day, then with positive probability in the next day, and so on. In the long run, buses will always arrive one after the other every day. This phenomenon resembles the bunching of bus arrivals that many commentators have noted is a feature of deregulated local bus markets.

In the long run not only do buses always arrive one after the other, but we see a tendency for buses to leapfrog each other backwards round the day. In Figure 3.2 bus A, which arrived just before B and collected most passengers yesterday, was most likely to arrive at the same time today, whereas B was most likely to revise its timetable in order to arrive just before A's arrival time yesterday. The continuous time model in the next section explores the mechanisms underlying this leapfrogging and bunching more fully.

3.4.1.2 The Model

Two buses, A and B, compete to pick up passengers during each of an infinite number of days. Each day has length 1, and is circular, with later times being further clockwise round the circle. In each period t each bus i, $i = A, B$, picks an arrival time a_i^t, $a_i^t \in (0, 1]$. At the start of period t all times are relabelled so that A's arrival time in period $t-1$ is at 0, which just has the effect that all times in t are measured in terms of minutes later than A's arrival time in the previous period. The state in period t, denoted k^t, is the (relabelled) arrival time of B in the previous period. We will assume henceforth that A arrived 'before' B in the sense that $k \leq 1/2$. By symmetry, this is without loss of generality.

Passengers arrive at a uniform rate throughout the day, with a total mass of 1 per day, and board the first bus to arrive after they do, unless both arrive at the same moment, in which case half board each bus.

A bus gets a gross profit in the day equal to the mass of passengers which boards and the mass boarding a bus is just the minutes after the previous bus that this bus arrives. Let $\pi_i(a_A^t, a_B^t)$ be the gross profit of bus i when arrival times are a_A^t, a_B^t. Then:

$$\pi_A(a_A^t, a_B^t) = \begin{cases} a_A^t - a_B^t & \text{if } a_A^t > a_B^t \\ 1/2 & \text{if } a_A^t = a_B^t \\ 1 + (a_A^t - a_B^t) & \text{if } a_A^t < a_B^t \end{cases}$$

and similarly for π_B.

If buses must pay when they update their timetable different arrival times will entail different costs. Denote the updating cost incurred by i should it arrive at x when the state is k by $C_i(x, k)$.

We assume that buses are myopic and seek only to maximise the expected current profit net of any revision cost. The problem is essentially a one period one (the past is summarised by the state), and we now drop the time superscripts. We restrict attention to Markov strategies which depend only on the state. A pure strategy for bus i is a function $s_i(k)$ which gives the arrival time chosen when the state is k. A mixed strategy is a distribution function $F_i(x, k)$ which gives the probability of arriving in the interval $[0, x]$. We consider Nash Equlibria where each bus chooses a strategy which maximises its expected net profit given the strategy chosen by its rival.

The first point is that for updating costs sufficiently low there is no equilibrium in pure strategies. Consider the extreme case where the updating cost is everywhere zero. In this case the state does not affect current payoffs and Markov strategies will not depend on it. The best reply function is not even defined here. If B chooses $s_B = a_B$, A will maximise the profit from boarding passengers by arriving as late as possible while still arriving before a_B, i.e. by setting its arrival time as the largest a_A such that $a_A < a_B$. When time is continuous there is no a_A which satisfies this. However, even without this technical problem there would be no equilibrium in pure strategies. Consider whether an ε-equilibrium (s_A^*, s_B^*) exists, where if j arrives at s_j^*,no arrival time gives i a payoff of ε more than $\pi_i(s^*)$, for an arbitrarily small ε. No such ε-equilibrium exists. To see why, simply note that in any ε-equilibrium A will arrive no more than ε minutes before B, and B will arrive no more than ε minutes before A. When ε is small, these conditions cannot both be met.

From now on we will consider only mixed strategies. Denote the average arrival time of bus i by $\bar{a}_i$. Let $\Pi_i(x, F_j)$ be the expected gross profit of bus i when it arrives at x and its rival's strategy is F_j. This will be given by:

$$\begin{aligned}\Pi_i(x, F_j) &= \lim_{\epsilon \to 0} [x - (\bar{a}_j | a_j < x)] F_j(x - \epsilon) \\ &\quad + [1 + x - (\bar{a}_j | a_j > x)] (1 - F_j(x)) \\ &\quad + \frac{1}{2} (F_j(x) - F_j(x - \epsilon))\end{aligned}$$

Let $\lim_{\epsilon \to 0} (F_j(x) - F_j(x - \epsilon)) = \Pr_j(x)$ (this will be zero when there is no atom in the distribution at x). Then:

$$\begin{aligned}\Pi_i(x, F_j) &= [x - (\bar{a}_j | a_j < x)] (F_j(x) - \Pr{}_j(x)) \\ &\quad + [1 + x - (\bar{a}_j | a_j > x)] (1 - F_j(x)) \\ &\quad + [x - (\bar{a}_j | a_j = x)] \Pr{}_j(x) \\ &\quad + \frac{1}{2} \Pr{}_j(x)\end{aligned}$$

which rearranges to:

$$\Pi_i(x, F_j) = 1 - \bar{a}_j + x - F_j(x) + \frac{1}{2} \Pr{}_j(x) \tag{3.1}$$

Using these expressions we first find equilibrium when timetable revision is costless.

When there is no cost to revising the timetable, so that in the absence of other considerations all arrival times are equally attractive, there is an equilibrium in which both operators choose arrival times according to a uniform probability distribution:

$$F_i^*(x) = x, \; i = A, B$$

To confirm that these strategies do form an equilibrium, substitute into the expression for expected profit above to give:

$$\Pi_i^*(x, F_j^*) = 1/2 \quad \forall x, \quad i = A, B$$

Since, when its rival chooses arrival times according to a uniform distribution over all times, an operator earns the same expected profit no matter what time it chooses, it is indifferent over all possible strategies, including arriving according to a uniform distribution. Here we see a radical instability in the timetable. Buses choose any arrival time with equal probability, independently of their rival's or their own previous arrival time. This instability arises from the desire on the part of both buses to arrive just before their rival, when there will be many passengers waiting at the stop.

Once we assume that it is costly to adjust the timetable more structure on the probability distribution chosen by firms emerges.

For technical reasons we assume that the cost of choosing different arrival times changes continuously. In particular we assume that if the arrival time is the same as in the last period there is no updating cost, and that the cost rises linearly at a rate m with the absolute change in the arrival time until a maximum updating cost of c is reached, at which point the updating cost remains constant. When m is large this function will approximate the situation where a firm pays c for every arrival time except that at which it arrived in the previous period, for which it pays nothing, and henceforth we assume $m > 1$. The updating cost functions are:

$$C_A(x,k) = \begin{cases} mx & \text{if } x \in [0, c/m] \\ c & \text{if } x \in [c/m, 1-c/m] \\ m(1-x) & \text{if } x \in [1-c/m, 1) \end{cases} \tag{3.2}$$

$$C_B(x,k) = \begin{cases} c & \text{if } x \in [0, k-c/m] \\ m(k-x) & \text{if } x \in [k-c/m, k] \\ m(x-k) & \text{if } x \in [k, k+c/m] \\ c & \text{if } x \in [k+c/m, 1) \end{cases} \tag{3.3}$$

An example of these costs are illustrated in Figure 3.3 below. In writing and illustrating $C_B(x,k)$ we have assumed that the interval between bus arrivals in the previous period, k, was not too small, specifically that $k \geq c/m$. In the limit where $m \to \infty$ this will almost always be true, but in any case what is at issue is notation rather than results. Looking at the illustration of C_B in the second panel of Figure 3.3, note that the circular day has been mapped to a line in the diagrams by cutting it at the point where 1 and 0 meet up and placing 0 at one end and 1 at the other. If $k < c/m$ the diagram is essentially the same, but the two ends of the line will lie in a region where the updating cost is less than c. The exact expression for C_B would differ from the one given in Equation 3.3, though the function is, in essence, the same. We will ignore this notational complication in what follows. The reader should be able to construct the exact expressions relevant to the case $k < c/m$ from the results that follow.

Now that the cost of arrival times varies, there can no longer be a completely mixed strategy equilibrium in which A chooses to arrive according to a uniform probability distribution. If A did so, B would choose to arrive at the same time in one period as it did in the previous one, i.e. at k, since all times give the same expected profit from boarding passengers, and by

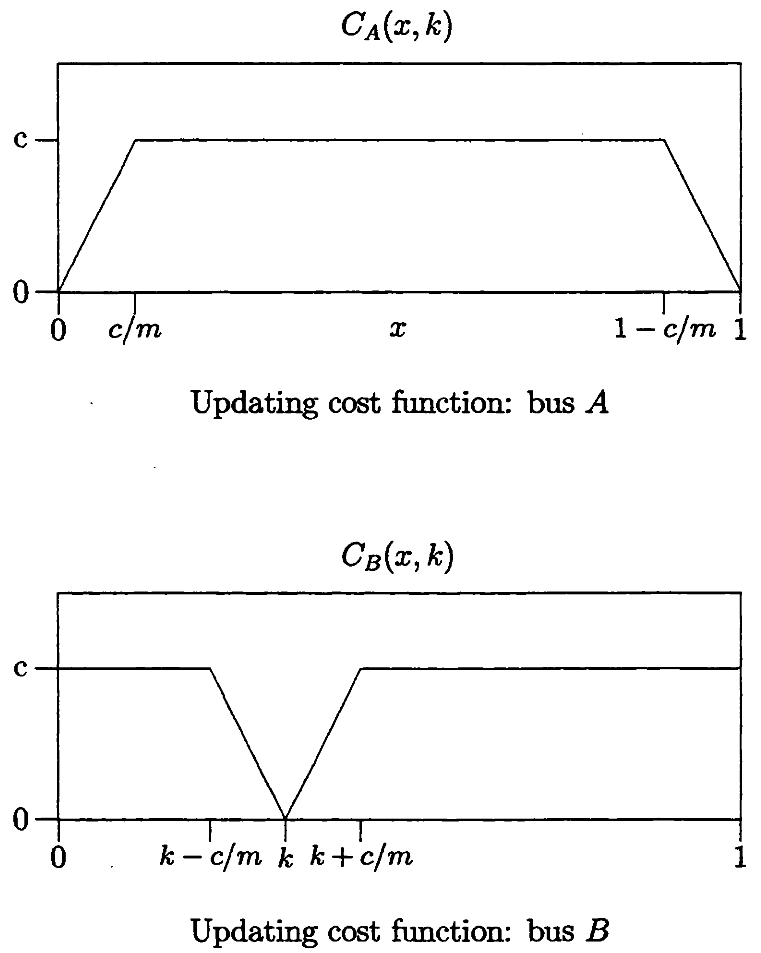

Figure 3.3: An example of the updating cost as a function of arrival time x for a given k

arriving at k bus B avoids all updating costs.

There are two cases, depending on whether the buses were bunched together in the previous period or were evenly spaced, specifically on whether $k > c$ or not. We consider the simpler case first where buses were fairly evenly spaced and $k > c$ (this case is only possible when $c \geq 1/2$). We will first formally state and prove the result before describing its implications and providing some intuition as to why it is true. We have:

Result 14 *There is an equilibrium (F_A^*, F_B^*) in which for $k \in (c, 1/2]$:*

$$F_A^*(x,k) = \begin{cases} c & x \in [0,c] \\ x & x \in [c, k - c/m] \\ (1+m)x - mk + c & x \in [k - c/m, k] \\ k + c & x \in [k, k+c] \\ x & x \in [k+c, 1) \end{cases}$$

$$F_B^*(x,k) = \begin{cases} 0 & x \in [0,c] \\ x - c & x \in [c, k) \\ k & x \in [k, k+c] \\ x - c & x \in [k+c, 1-c/m] \\ (1+m)x - m & x \in [1-c/m, 1) \end{cases}$$

Proof. To show that these strategies form an equilibrium, we need to show that the net profit a bus operator expects to earn is the same, no matter what time in the support its bus arrives, and that this net profit is no less than that from arriving at any time not in the support. Substituting B's strategy into the expression for A's profit from boarding passengers gives:

$$\Pi_A(x, F_B^*) = 1 - \bar{a}_B + \begin{cases} x & x \in [0, c] \\ c & x \in [c, k) \\ c/2 & x = k \\ x - k & x \in (k, k + c] \\ c & x \in [k + c, 1 - c/m] \\ m(1 - x) & x \in [1 - c/m, 1) \end{cases}$$

which gives an expected net profit for A as a function of its arrival time of:

$$\Pi_A(x, F_B^*) - C_A(x) = 1 - \bar{a}_B + \begin{cases} x(1 - m) & x \in [0, c/m] \\ x - c & x \in [c/m, c] \\ 0 & x \in [c, k) \\ -c/2 & x = k \\ x - k - c & x \in (k, k + c] \\ 0 & x \in [k + c, 1) \end{cases}$$

Similarly the expected net profit of B as a function of its arrival time is:

$$\Pi_B(x, F_A^*) - C_B(x) = 1 - \bar{a}_A + \begin{cases} x - 3c/2 & x = 0 \\ x - 2c & x \in (0, c] \\ -c & x \in [c, k] \\ (x - k)(1 - m) - c & x \in [k, k + c/m] \\ x - 2c & x \in [k + c/m, k + c] \\ -c & x \in [k + c, 1) \end{cases}$$

This net profit is shown in Figure 3.4 below. Inspection of the expressions and the Figures reveals that the expected net profit is $1 - \bar{a}_B$ if A arrives at any time in the support of $F_A^*(x)$, and is less than this should A arrive at any other time. This confirms that F_A^* is a best response to F_B^*. Similar reasoning confirms that F_B^* is a best response to F_A^* and so that these strategies are an equilibrium. ∎

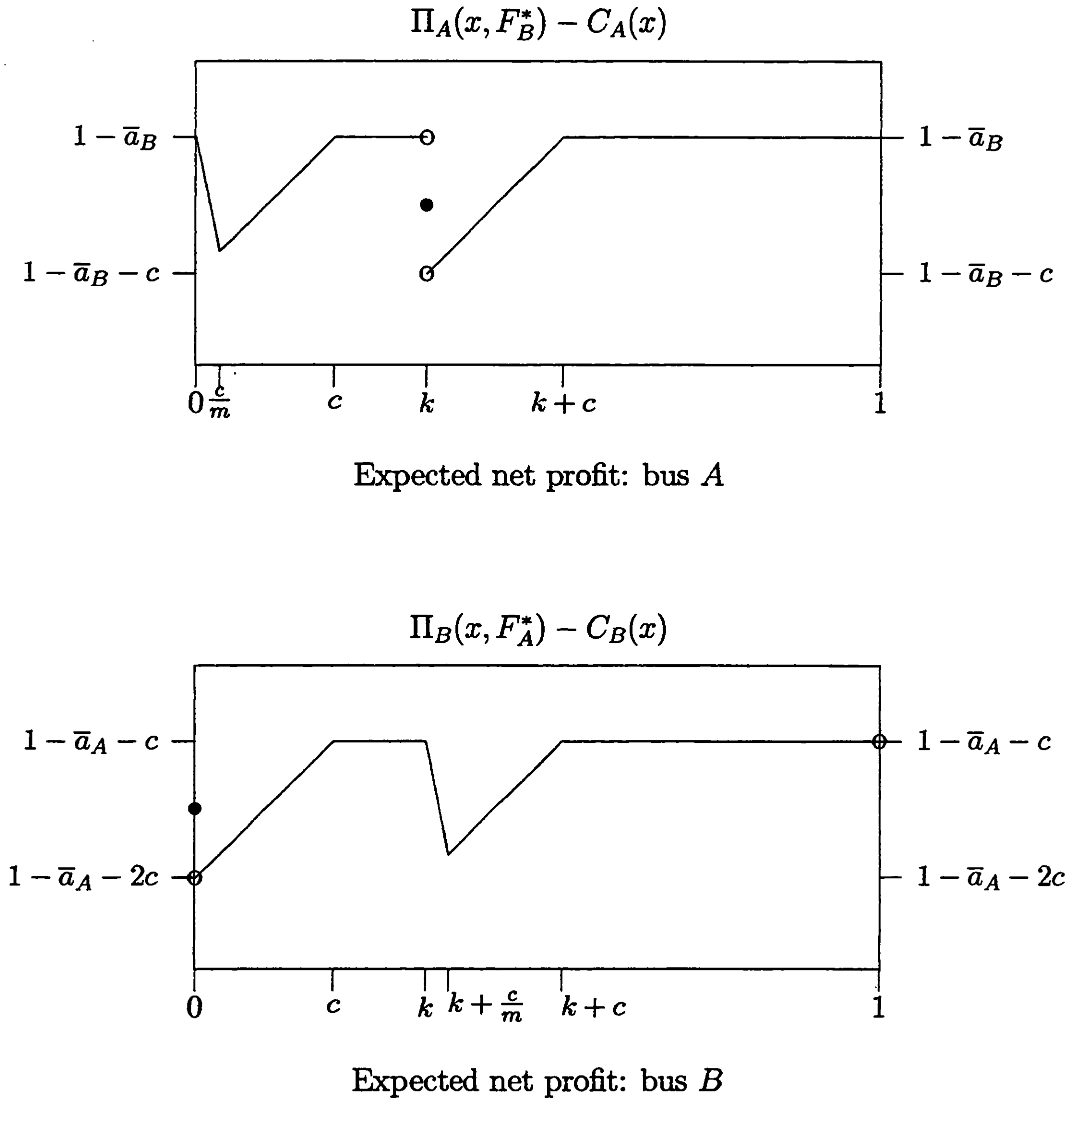

Figure 3.4: Expected net profit as a function of arrival time: $k < c$

Figure 3.5 below shows the equilibrium strategies. Since marginal prob-

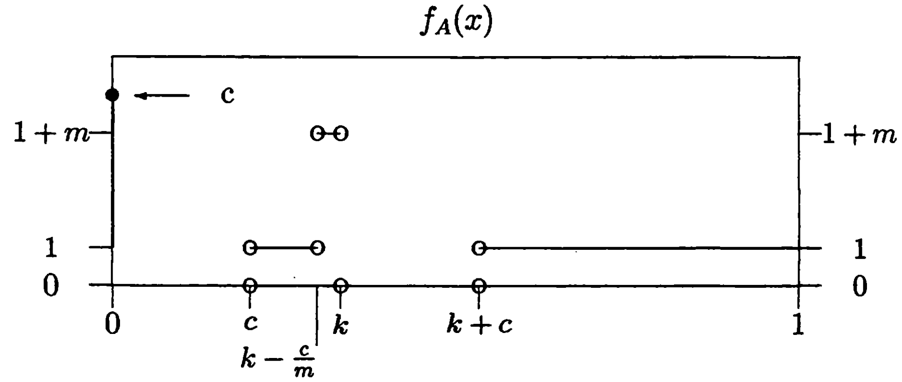

Marginal probability of arriving: bus A

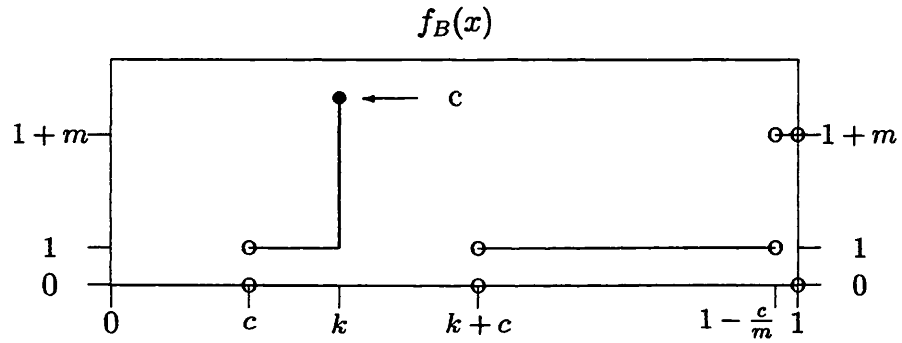

Marginal probability of arriving: bus B

Figure 3.5: Equilibrium marginal probabilities of arrival time for given k

abilities are simpler to interpret than the related distribution function, the figures give the marginal probability chosen by each firm in equilibrium, where this is defined. A filled square at the top of a line means that there is a probability mass at that point, and the probability in that mass is marked. The probability mass of c in the distribution means that each bus arrives at the same time as it did in the previous period with probability c. Thus, not surprisingly, the higher the updating cost, the higher the probability that a bus chooses not to incur it. If a bus does update its arrival time, it never chooses to arrive a little later than previously, but may arrive little earlier. Also we see that a bus never arrives a little later than its rival's previous arrival time, but may arrive a little earlier. In particular with probability $c+c/m$ it will arrive in an interval of width c/m immediately before its rival's arrival time in the previous period[10]. The implication of these strategies is that buses tend either not to update their arrival times, or if they do, to arrive just before their rival's previous arrival time. They never arrive later than either their own or their rival's previous arrival time. Since a lot of probability is concentrated close to the same two arrival times for each bus, this behaviour will cause a tendency for buses to choose close arrival times this period and so to bunched arrival times. Moreover we see some leapfrogging to earlier and earlier times as bus operators avoid arrival times later than their rival's previous time, but sometimes choose an arrival time before.

As stated in the proof, to show that these strategies do form an equilibrium we need to show that the expected net profit of each bus is maximised by its arriving at any time in the support of its equilibrium distribution function. Suppose A arrives according to F_A^* and consider B's expected payoff. All other things equal B would arrive at k and avoid all updating costs. However, with a relatively high probability A arrives just before k which increases B's expected payoff if it arrives a little earlier still. A's distribution function is such that this inventive exactly offsets the disincentive from having to pay an updating cost. Also the atom in A's distribution at 0 makes the expected profit from arriving just before this higher so that these times also lie in B's

[10] $c + c/m = (k - (k - c/m))\,(1 + m)$

support.

So far we have just considered the case where the buses were not too bunched together in the previous period. Now we turn to the case where their arrival times were separated by less than c in the previous period and so $k < c$. In this case:

Result 15 *There is an equilibrium* (F_A^*, F_B^*) *which, when* $k < c$*, has the form:*

$$F_A^*(x,k) = \begin{cases} k+c & x \in [0, k+c] \\ x & x \in [k+c, 1) \end{cases}$$

$$F_B^*(x,k) = \begin{cases} 0 & x \in [0,k) \\ k & x \in [k, k+c) \\ x-c & x \in [k+c, 1-c/m] \\ (1+m)x - m & x \in [1-c/m, 1) \end{cases}$$

Proof. To confirm that these form an equilibrium we can calculate the expected net profit from arriving at different times, assuming the rival's times are given by these distributions, the same way as above. This gives:

$$\Pi_A(x, F_B^*) - C_A(x) = 1 - \bar{a}_B + \begin{cases} x(1-m) & x \in [0, c/m] \\ x-c & x \in [c/m, k] \\ k/2 - c & x = k \\ x-k-c & x \in (k, k+c] \\ 0 & x \in [k+c, 1) \end{cases}$$

$$\Pi_B(x, F_A^*) - C_B(x) = 1 - \bar{a}_A + \begin{cases} -k/2 - 3c/2 & x = 0 \\ x - k - 2c & x \in (0, k - c/m] \\ (x-k)(1+m) - c & x \in [k - c/m, k] \\ (x-k)(1-m) - c & x \in [k, k + c/m] \\ x - k - 2c & x \in [k + c/m, k + c] \\ -c & x \in [k + c, 1) \end{cases}$$

These expected net profits are illustrated in Figure 3.6 below. Inspection of the Figure and the expressions confirms that each bus' expected net profit is maximised at any point on the support of its equilibrium distribution function. ■

These equilibrium strategies are illustrated in Figure 3.7: When bus arrivals were close together in the previous period, the bus which arrived just before its rival and so had more passengers, i.e. bus A, is less likely to have its timetable updated this period than is bus B which had fewer passengers. This is shown by the fact that the atom at 0 in A's equilibrium strategy has mass $k + c$, whereas the atom in B's equilibrium strategy at k only has mass k. Moreover if bus A does have its timetable updated, it will arrive earlier than its own previous arrival time, but will avoid time either a little earlier or later than B's previous arrival time. Bus B on the other hand will, with relatively high probability, arrive in the interval c/m just before A's previous arrival time. The leap-frogging to earlier and earlier times first seen for the case when buses were fairly evenly spaced previously, $k > c$, is a much stronger feature of the equilibrium when buses were bunched together previously. The later bus is both more likely to have its timetable revised than its rival, and if it is revised at all, is relatively likely to arrive just before the previous arrival time of the early bus.

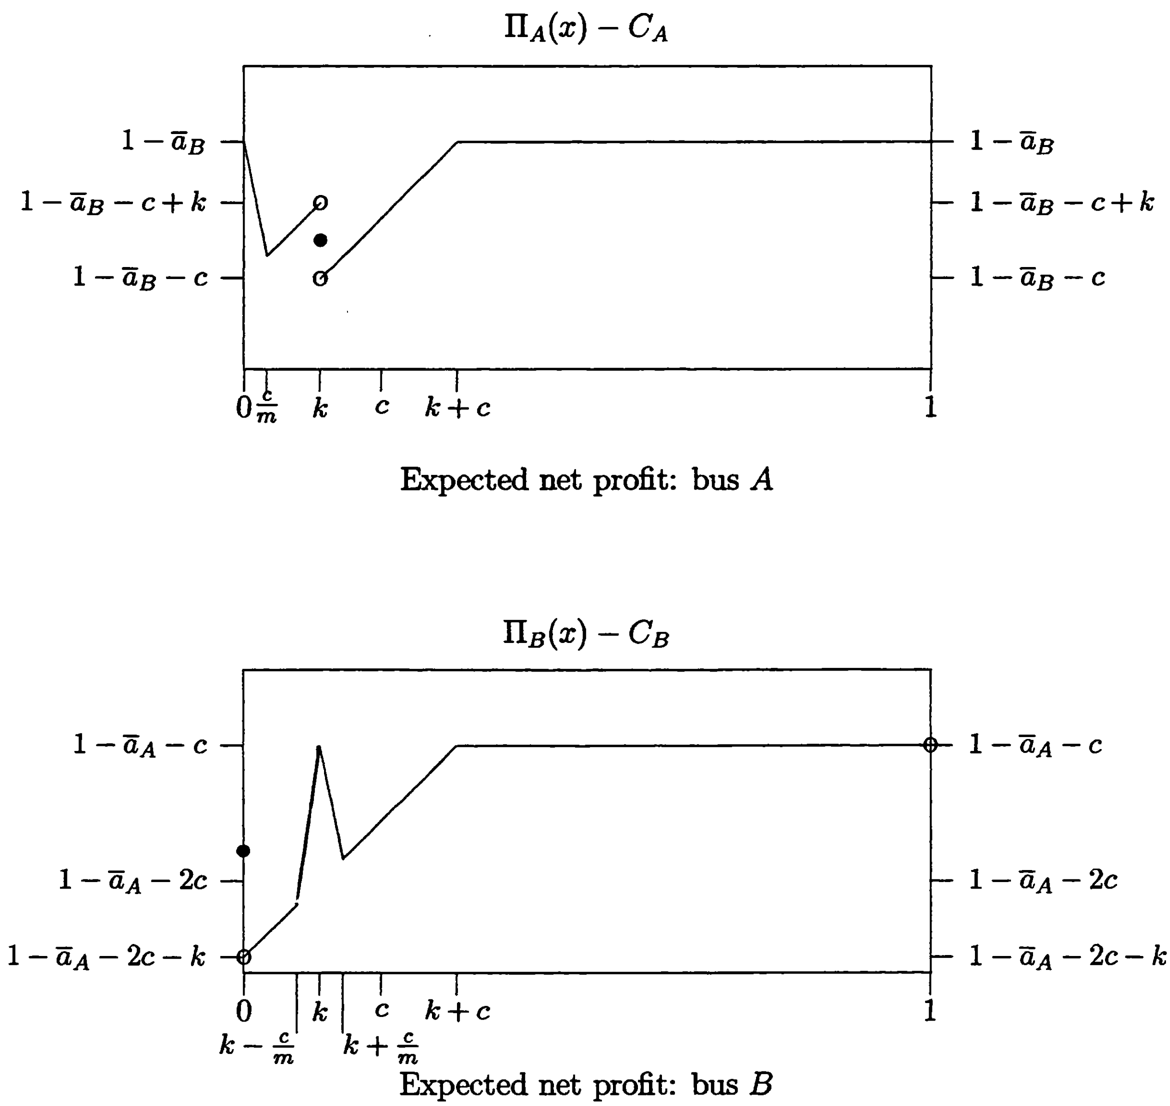

Figure 3.6: Equilibrium marginal probabilities of arriving: $k < c$

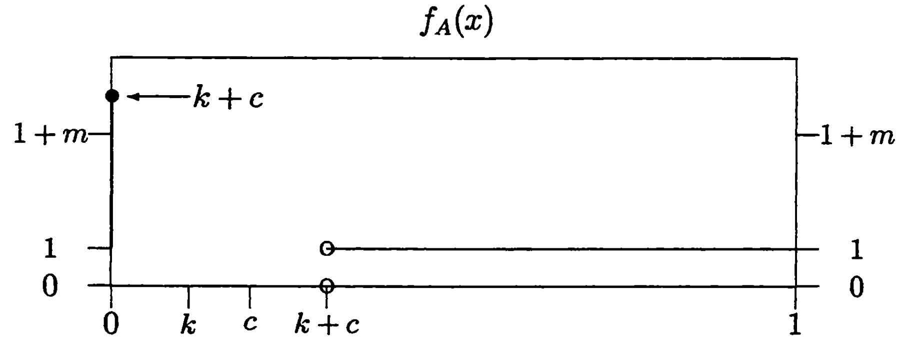

Marginal probability of arriving: bus A

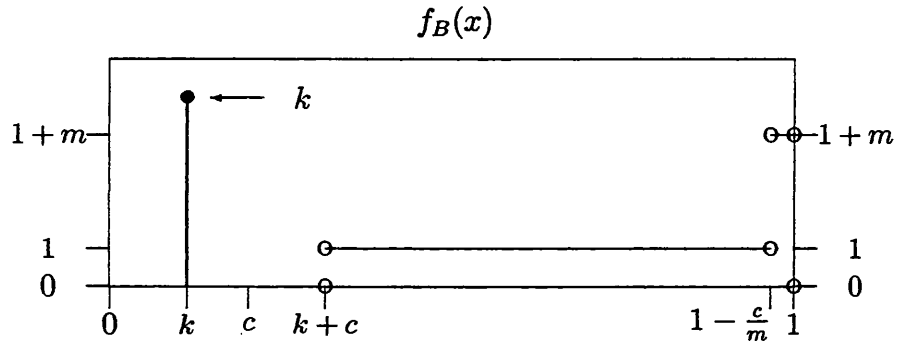

Marginal probability of arriving: bus B

Figure 3.7: Equilibrium marginal probabilities of arriving: $k < c$

3.4.2 Express Coaches

Passengers travelling by express coach know the timetable and travel on that coach which they most prefer. Although preferences are not modelled explicitly in this section, the assumption underlying the demand function we will use is that what matters to passengers is the absolute difference between the arrival time of the bus they board and the passenger's most preferred time.

The results of the model express coach market are much more straightforward than those for urban local buses, and the whole analysis can be dealt with informally. Consider a model identical to the one for urban local buses above except that the specification of demand differs. Assume that passengers are located evenly over the day and that the mass of passengers boarding a bus is equal to the mass which is closer to that bus than to its rival. For both buses this will be $1/2$ no matter what the arrival times chosen. So long as the updating cost is at a minimum when buses do not update their arrival times, there will be an equilibrium in which both buses arrive at the same time from day to day.

3.4.3 Conclusions

This section has considered the pattern of timetable revisions that result when passengers just turn up at a bus stop and board the first bus to arrive, and when there is a cost to updating the timetable, and compared this pattern to that when passengers choose their most preferred bus. The results provide an explanation for the difference in timetable stability, and in the bunching of bus arrival times, which has been noted as between urban local bus markets and express coach markets. The next section considers the impact of the same passenger behaviour on fare levels.

3.5 Fares

The difference in the impact of market structure on fare levels between local urban bus and express coach markets can be investigated by comparing fares under monopoly and duopoly in each of the two market settings. In the basic model buses of a fixed quality arrive at fixed, regular times. Operators have no choice over quality and frequency. At the end of the analysis we discuss relaxing this assumption and argue informally that the results may well be robust. Neither do operators have a choice over exact arrival times: work in previous sections suggests that in local bus markets competition can lead to a fundamental instability in choosing arrival times which would quickly make the model intractable. We sidestep the issue here by imposing regular arrival times. Varying the assumption on market structure in the model is effected by imposing different ownership patterns on the buses.

The results of this section rely on passenger responses to different fare levels and, unlike in the previous section, here we explicitly model the preferences underlying their behaviour.

Passengers differ in the cost of walking to the stop/station, which is exactly equivalent in the model express coach market to supposing that they differ in the value they put on a coach trip. Waiting at the stop/station is also costly and, *all other things equal*, passengers have a preferred time at which they would like to travel. In standard location games the cost of travelling at a time different from the most preferred one plays a central role, rather than the marginal one ascribed to it here. After the analysis we discuss the implications of a significant 'inconvenience' cost for the results, and suggest informally that the results would largely carry through.

Finally, passengers are assumed always to know the frequency of buses and the fares charged by different operators. However, in the model urban local bus market we assume that they do not know the exact bus arrival time, whereas in the model express coach market they *do*[11]. This difference generates very different effects of competition on price. In the express coach market the duopoly fare is *lower* than the monopoly one, but in the urban local bus the duopoly fare is *higher*. The basic mechanisms underlying this difference are straightforward. Starting with the more familiar case, with express coaches, where passengers know in advance when buses arrive and what fares they charge, those passengers choosing to travel at all choose that coach from the set charging the lowest fare which leaves at the most convenient time. The number of passengers travelling by coach is higher the lower is the lowest fare. A monopolist, running all coaches, is sure to carry all passengers deciding to travel and is able to set fares to trade off the numbers travelling against the revenue earned from each. Competing duopolists on the other hand get an additional benefit from cutting fares: they get half the passengers if they charge the same fare as their rival, but will capture all of them if they charge marginally less, giving rise to an over-incentive to cut prices compared to the monopolist. In equilibrium both operators set fares to zero, which is surely lower than the monopolist's fare.

In urban local bus markets where passengers do not know the exact arrival times of buses, passengers will walk to the stop at the most convenient time. Once at the stop, however, they will board the first bus to arrive, even if

[11] In the context of express coaches, it makes more sense to talk about "departure times" than "arrival times". However, we will continue to use "arrival times " to highlight the comparison with local buses.

it charges (not too much) more than the next one, since waiting at a bus stop is costly. Consider then a duopolist lowering its fare. More passengers will walk to the stop on the expectation of lower fares on average. However half of the extra passengers will then board the rival bus, so the duopolist gets only half the benefits of the lower fare. A monopolist on the other hand always picks up all passengers and so has a greater incentive to cut fares, leading to lower fares under monopoly.

3.5.1 The Model: Common Assumptions

Buses and Operators

The first part of the description sets down the assumptions on the arrival pattern of buses. This is taken to be fixed; what the model investigates is the effect of varying the ownership pattern of the buses.

During an infinitely long day, one bus arrives at time t_1 and thereafter one bus arrives every Δ minutes. The arrival time of the first bus, t_1, is drawn from a distribution F_t, which represents a uniform distribution on $[0, \Delta)$, before the day begins. If an agent knows t_1 and Δ then it knows the arrival time of all buses and, conversely, if the agent knows Δ but not t_1 then it does not know the exact arrival times of any buses. The identity of the operator of the s'th bus is given by m_s. Different market structures can be investigated through varying the assumptions made on m_s. The two with which the analysis of this paper will be concerned are first, a duopoly in which each of two operators runs alternating buses, and second a monopoly in which all buses are run by the same operator.

In addition we need an assumption on the action set of operators. The fare charged on the s'th bus is given by p_s. For most of the section we will assume that operators set their fares before the day begins (though after

t_1 is drawn). This is, in fact, what happens: bus drivers do not have the authority to start bargaining with passengers waiting at the stop over the level of the fare. However, there are no obvious formal constraints which dictate why this is so, and at the end of the section we will instead assume that prices are set only when a bus arrives at a stop. We find that under this alternative assumption prices are so high under all market structures that no one ever travels by bus. The main body of the section assumes that before the day begins, each operator i chooses the fares on its buses, that is operator i chooses $p_s \; \forall \, s \in \{s \,|\, m_s = i\}$. Operators choose simultaneously. This action set is further restricted by requiring that a duopolist charge the same fare on all its buses, and a monopolist charge the same fare on every other of its buses. The first restriction is standard and it will become clear that the fares chosen under the restriction would remain optimal, given passenger behaviour, were the restriction removed. The reason for requiring that a monopolist charge the same fare just on every other of its buses is that this makes comparisons of monopoly and duopoly behaviour more transparent. It means that the monopolist's strategy set contains all the options that would be available to a colluding duopoly. In fact, and not at all surprisingly, we find that it is optimal for the monopolist to charge the same fare on *all* its buses.

Assumptions 1 and 2 below set out the two alternative assumptions for the ownership of buses and the additional restrictions on the action sets of operators.

Assumption 1 *Monopoly:*

(i) $m_s = 1, \; \forall s$

(ii) $p_s = p_1, \; s$ *odd, and* $p_s = p_2, \; s$ *even*

Assumption 2 *Duopoly:*
(i) $m_s = 1$, *s odd, and* $m_s = 2$, *s even*
(ii) $p_s = p_1$, *s odd, and* $p_s = p_2$, *s even*

Under both assumptions, operator behaviour is fully specified by the fares on the first two buses, p_1 and p_2, and we can refer to two types of buses, odd and even, indexed by 1 and 2 respectively.

There are no costs to accepting a passenger on board and, if necessary, a bus could accommodate any number of passengers. Bus operators aim to maximise their average revenues per bus. If a bus charges p and n passengers board, the revenues earned by the bus is just pn.

Passengers

The second part of the description concerns passenger behaviour.

Each passenger is identified by a pair (τ, κ), $\tau \in \mathcal{R}$, $\kappa \in [\underline{\kappa}, \overline{\kappa}]$. The first variable is the time at which the passenger would most like to leave the house. The second is the cost a passenger pays should he walk to the bus stop/station. The mass of passengers with preferred leaving times in any interval δt is constant and is just δt. Sampling passengers with preferred leaving times in any time interval, the proportion with walking cost κ or less is given by the distribution function $F(\kappa)$ which represents a uniform distribution over $[\underline{\kappa}, \overline{\kappa}]$, and so:

$$F(\kappa) = \begin{cases} 1 & \text{if} \quad \kappa > \overline{\kappa} \\ \frac{1}{\overline{\kappa}-\underline{\kappa}}(\kappa - \underline{\kappa}) & \text{if} \quad \underline{\kappa} \leq \kappa \leq \overline{\kappa} \\ 0 & \text{if} \quad \kappa < \underline{\kappa} \end{cases} \tag{3.4}$$

A passenger (τ, κ) chooses whether to walk to the bus stop, or to stay at home and exit the game on the basis of its information set Φ_a, which is the same for all passengers. Assume that passengers know the prices set by firms, and the interval between firms, and so $\Phi_a \supset \{p_s, \Delta\}$, but they do not necessarily know the exact arrival times of buses. If he chooses to walk to the stop, he must also choose a time t. More precisely, each passenger chooses an action a from the set $\mathcal{A} = \{X\} \cup \mathcal{R}$, where passengers may choose different actions, and:

$$a = \begin{cases} t & \text{if the passenger walks to the stop at } t \\ X & \text{if the passenger stays at home and so exits the game} \end{cases}$$

Here we make what is largely a technical assumption. Passengers suffer a small inconvenience cost if they leave the house at a time different from τ. Specifically, the cost if a passenger leaves the house at t is:

$$I(t, \tau) = \varepsilon|t - \tau|$$

It is reasonable to suppose that passengers incur a cost through either rushing to be ready before they would want, or from waiting around with nothing to do because they are ready to leave before they need to be. However, such a cost is unlikely to be of first order importance and it would be a problem for the model if the results rested on this assumption. In fact the assumption is made for technical reasons: it means that when a passenger $(\tau, \cdot)$ is otherwise indifferent over leaving the house at any time t in some set T, he will prefer the time in T closest to τ. We will be looking at limit economies as this inconvenience cost ε goes to zero, its effect in picking between otherwise indifferent leaving times will be its only impact.

If passengers stay at home then they get zero payoff. Passengers choosing to walk to the bus stop arrive there instantaneously. Assume that they learn nothing new at the stop and that once there they wait for the next bus. A more realistic assumption might be that passengers learn the arrival times t_1 once they arrive at the stop, since bus stops often carry timetables, and arriving passengers may infer when the last bus left from the number of passengers waiting at the stop. The assumption that, on the contrary, passengers learn nothing at the stop is made to simplify the analysis: even if some passengers arrive at the stop so long before a bus is due to arrive that the passenger would do better to go straight home again, the passenger is no better informed about this than when he walked to the stop in the first place, so that he has no basis on which to update his decision to be at the stop, and we can legitimately ignore such considerations. Waiting at the stop is costly, with instantaneous cost c, and costs are paid as they arise. Therefore a passenger waiting at the stop for t minutes will incur a waiting cost $C(t)$ given by:

$$C(t) = ct$$

When a bus arrives, passengers at the stop learn the arrival times of all buses, and so the information set becomes $\Phi_b = \{p_s, \Delta, t_1\}$. Waiting passengers then decide whether to board, go home, or continue waiting. More specifically each passenger at the stop chooses an action b_s when the s'th bus arrives, from the set $\mathcal{B} = \{B, S, X\}$ where:

$$b_s = \begin{cases} B & \text{if the passenger boards} \\ S & \text{if the passenger stays at the stop} \\ X & \text{if the passenger goes home} \end{cases}$$

We make the additional restriction that passengers always behave in the same way when the same type of bus arrives. That is we impose that:

$$\begin{aligned} b_s &= b_1 \text{ if } s \text{ odd} \\ b_s &= b_2 \text{ if } s \text{ even} \end{aligned}$$

A passenger boarding bus s payoff associated with the bus trip of:

$$g_s = u - p_s,\ s = 1, 2, \ldots$$

Note that under either of Assumptions 1 or 2 trip payoffs on all buses are either g_1 or g_2 depending on whether the bus in question was an odd or an even one. The variable u can be thought of as the value to a passenger of a trip on a bus.

In summary, a passenger can receive four payoffs; the trip payoff, and the costs (which are negative payoffs) associated with leaving at an inconvenient time, walking to the stop and waiting for a bus. These payoffs are assumed to be undiscounted and additive.

Equilibrium

If operators choose their fares before the day begins, then there are no decisions for them to make during the day. The only decisions are those made by passengers over whether to walk to the stop, and whether to board a bus. We look for strategies;

$$\begin{aligned} g_i &\in \mathcal{R}, \quad i = 1, 2, \\ a(\tau, \kappa \,|\, \Phi_a) &: \mathcal{R} \times [\underline{\kappa}, \overline{\kappa}] \rightarrow \mathcal{A}, \\ b_i(\Phi_b) &\in \mathcal{B}, \quad i = 1, 2 \end{aligned}$$

which form a PNE. One such equilibrium is derived below. Equilibrium is derived by first finding how passenger behaviour depends on the fares set by operators, and then finding optimal operator choices of these fares.

3.5.2 Local Buses

Passenger behaviour will be dictated just by the fares offered by buses, not on bus ownership, and so the questions of whether passengers board arriving buses, and how many of them walk to the stop in the first place, can be dealt with independently of market structure. Passenger behaviour is accordingly dealt with first, and the findings are summarised in Result 16 on Page 149.

The analysis then turns to compare fare setting under monopoly and duopoly, and the findings are given in Result 18 on Page 157.

Passenger Behaviour

Here we examine the implications for passenger behaviour of the assumption that they do not know the exact arrival time of buses. The section shows first that lower fares on alternate buses will increase the total number of passengers walking to the stop. Second it shows that so long as the difference in fares between alternate buses is not too large then passengers board the first bus to arrive, and the share of passengers boarding each bus is just 1/2. Cutting fares on alternate buses, then, will mean that more passengers board *all* buses, and it is this fact that will lead to competing duopolists charging higher fares than a monopolist.

In markets for urban local bus services we take Assumption 3 below to hold. This asserts that passengers know the interval between buses, Δ, the fares charged on all buses, p_s, $\forall s$, but not t_1 and so not any actual arrival times. Passengers believe that t_1 could be any time in the interval $[0, \Delta]$ with equal probability

Assumption 3 *Urban Local Buses:*
$\Phi_a = \{p_s, \Delta, F_t\}$

We also make Assumption 4, which is a technical assumption on parameter values that ensures both that even if all buses were free, not every passenger would find it worthwhile walking to the stop, and that if every other bus were free, at least some passengers *would* find it worthwhile walking to the stop. It means that there are no reasonable circumstances under which all passengers walk to the stop, and that there is always a fare which is sufficiently low that, if charged on alternate buses, some passengers walk to the stop. This reduces the number of different cases that need to be considered during the following analysis.

Assumption 4 $\underline{\kappa} - c\Delta < u < \overline{\kappa} - c\Delta/2$

Passenger boarding behaviour b_i

Boarding behaviour is restricted to depend on whether the bus number is odd or even (though we confirm below that passengers would behave no differently were they able to condition on additional variables). This means that a passenger at a stop will either wait at the stop forever, if $b_1 = b_2 = S$, or will wait at the stop for no more than two buses to arrive, one from each operator. It is easy to see that waiting at the stop forever cannot be optimal since a passenger which did so would never get a trip payoff, but would incur an infinitely large waiting cost. Moreover, if a passenger stays at the stop when a bus of type i arrives, i.e. if $b_i = S$, then he will surely board a bus of type j, i.e. $b_j = B$. This follows because, as we have already seen, it cannot be optimal to also stay at the stop when both i and j arrive, and if the passenger goes home, i.e. $b_j = X$, then he has incurred an unnecessary wait, and would have done better to leave when i arrived. The only remaining choice is to board j.

Passengers' boarding behaviour depends, then, on a simple comparison of the payoffs accruing from; exiting, boarding the bus, and staying at the stop and boarding the next bus. These depend on the trip payoffs, the waiting costs, and the fall-back payoff of zero that the passenger gets from exiting. As an example, it will be optimal to board a bus i when it arrives if this gives a higher payoff over the rest of the game than exiting or waiting and boarding the next bus. This will be so when the trip payoff on i is positive, and greater than the trip payoff on j net of the cost of waiting for the next bus to arrive, which will be in Δ minutes The full description of passengers' optimal behaviour is given in Equation 3.5. Note that optimal behaviour is uniquely defined over almost all combinations of trip payoffs, the exceptions being just those boundary cases where more than one action gives passengers the highest payoff[12].

$$\begin{aligned} b_i(\Phi_b) &= \begin{cases} B & \text{if} \quad g_i \geq 0 \quad & \text{and} \quad g_i \geq g_j - c\Delta \\ S & \text{if} \quad g_j - c\Delta \geq 0 \quad & \text{and} \quad g_i < g_j - c\Delta \\ X & \text{if} \quad g_i < 0 \quad & \text{and} \quad g_j - c\Delta < 0 \end{cases}, \\ i &= 1,2, \qquad j \in \{1,2\}, \quad j \neq i \end{aligned} \tag{3.5}$$

A number of cases arise, depending on the fares of the two types of buses. Passengers at a stop may board the first bus which arrives, they may board those of just one type and either stay at the stop or go home should one of the other arrive first, or they may go home whichever bus arrives first. These cases are shown in Figure 3.8 which splits up the space of trip payoffs (g_1 and g_2) according to passenger boarding behaviours. The boundaries of the Figure are given by the equations in Equation 3.5.

Behaviour of arising passengers, $a(\tau, \kappa)$

Next consider the choice arising passengers make over whether or not to walk to the bus stop. Whether this is in a passenger's interest depends on

---

[12] Although we have restricted possible strategies to depend just on the identity of the operator of an arriving bus, it is worth pointing out that even if passengers were free to condition behaviour on other observable factors, such as time or the number of passengers at the stop, it is not optimal for them to do so. To see this, note that if a passenger eventually boards a bus i then it is always better to board the first bus i that arrives, so if passengers board a bus it will always be one of the first two to arrive, and optimal behaviour will be given by just the comparisons set out in Equation 3.5.

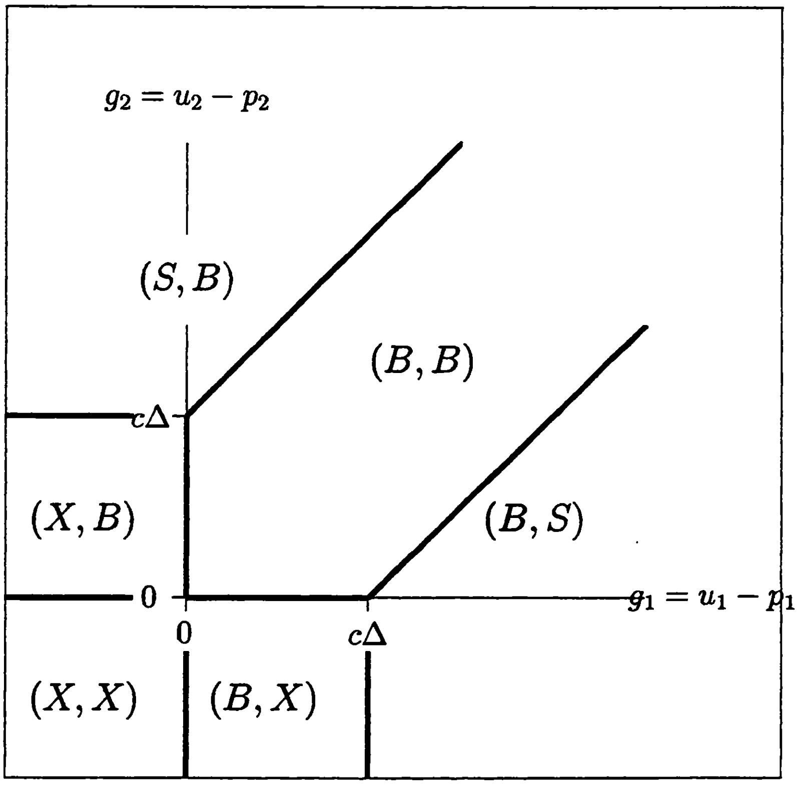

Figure 3.8: Regions of passenger boarding behaviour, (b_1, b_2), depending on trip payoffs. Regions are delineated by the heavy lines.

the expected trip payoff, net of the cost of waiting at the stop. Denote the expected trip payoff net of waiting cost of a passenger who walked to the stop by κ^*. This κ^* will depend on the prices, qualities and arrival intervals of the buses, but whatever its value the optimal behaviour of arising passengers is simply to walk to the stop at their preferred time if κ^* is greater than their particular walking cost, and to stay at home otherwise. When $\varepsilon > 0$ this optimal behaviour is uniquely defined for all passengers except those for which $\kappa^* = \kappa$. However, when there is no cost to leaving at an inconvenient time, although leaving at the most preferred time is still optimal, so is leaving at any other time, given that the net trip payoff exceeds the walking cost.

In the rest of the analysis we take optimal walking behaviour as given in Equation 3.6 below, which assumes that passengers leave at their preferred time, whatever ε.

$$a(\tau, \kappa \mid \Phi_a) = \begin{cases} \tau & \text{if } \kappa \leq \kappa^* \\ X & \text{if } \kappa < \kappa^* \end{cases} \tag{3.6}$$

The number of passengers walking to the stop, $F(\kappa^*)$

Bus operators choosing prices and qualities before the day begins are more interested in the total number of passengers walking to the stop than in the behaviour of individual passengers, that is they are concerned with $F(\kappa^*)$ and how their choices will affect this. The way that κ^* is affected by price and quality will depend on which buses passengers may board once they get to the stop: for example if passengers do not board odd buses then small changes in p_1 will have no impact on the number of passengers who find it worthwhile walking to the stop. In other words, the expression for $\kappa^*(g_1, g_2)$ will depend on the region of boarding behaviour in Figure 3.8 above that operators' choices place passengers in. As an example, consider region (B, B), where passengers always board the first bus to arrive. A passenger walking to the stop will arrive before bus of type i with probability $1/2$. Given that he arrives before i he expects to wait $\Delta/2$ minutes, and to get a trip payoff of g_i. Writing out the expected trip payoff net of waiting cost gives $\kappa^* = 1/2(g_1 - c\Delta/2) + 1/2(g_2 - c\Delta/2)$, which simplifies to $\kappa^* = 1/2(g_1 + g_2) - c\Delta/2$. Consider next region $(b_1, b_2) = (B, S)$, where passengers will always wait for bus 1 to arrive before boarding. The trip payoff is always

g_1 and the expected wait is Δ minutes, giving the expression above.

$$\kappa^*(g_1, g_2) = \begin{cases} 1/2(g_1 + g_2) - c\Delta/2 & \text{if} \quad (b_1, b_2) = (B, B) \\ g_i - c\Delta & \text{if} \quad (b_i, b_j) = (B, S) \\ 1/2(g_i) - c\Delta/2 & \text{if} \quad (b_i, b_j) = (B, X) \\ -c\Delta/2 & \text{if} \quad (b_i, b_j) = (X, X) \end{cases} \tag{3.7}$$

Note that if any passengers who were at the stop found it optimal to go home, i.e. if $(b_i, b_j) = (B, X)$ or (X, X), then no passengers would find it worthwhile walking to the stop in the first place. It is easy to see why. Suppose first that passengers at the stop go home whatever bus arrives, so $(b_i, b_j) = (X, X)$. In this case passengers never get a trip payoff, and cannot recover the cost of walking to the stop. Suppose now that $(b_i, b_j) = (B, X)$ so that passengers at the stop go home when bus j arrives. The trip payoff from i must therefore be less than $c\Delta$, or it would be better to wait for i. However, a passenger walking to the stop expects to wait at the stop for $c\Delta/2$ minutes, and yet only gets a trip payoff of g_i with probability$1/2$ (otherwise he gets zero). Consequently the expected trip payoff, net of waiting cost, is negative, and it is never worthwhile walking to the stop. This result is shown formally in Lemma 3.

Lemma 3 *If* $(b_i, b_j) = (B, X)$ *or* (X, X) *then* $\kappa^* < \underline{\kappa}$.

Proof. From Equation 3.7, $\kappa^* < 0$ if $(b_i, b_j) = (X, X)$ and since $\underline{\kappa} > 0$ by assumption, $\kappa^* < \underline{\kappa}$. Also, from Equation 3.7, $\kappa^* = 1/2g_i - c\Delta/2$ if $(b_i, b_j) = (B, X)$. But from Equation 3.5, boarding behaviour will only be (B, X) when $g_i \leq c\Delta$ and so $\kappa^* \leq 0$, and hence $\kappa^* < \underline{\kappa}$. ∎

The boundaries where $\kappa^* = \underline{\kappa}$ and $\overline{\kappa}$ will depend on the boarding behaviour of passengers. As an example, suppose that $(b_i, b_j) = (B, B)$, so that trip payoffs are in the region where passengers take the first bus to arrive. In this case the expected trip payoff, net of waiting at the stop, is $\kappa^* = 1/2(g_1+g_2)-c\Delta/2$. Then $\kappa^* = \overline{\kappa}$, when $1/2(g_1+g_2) = \overline{\kappa}+c\Delta/2$. Similarly, $\kappa^* = \underline{\kappa}$ when $1/2(g_1+g_2) = \underline{\kappa}+c\Delta/2$. Figure 3.9 shows the boundaries between the different regions of κ^*, superimposed on the boundaries of different boarding behaviour. The Figure is for parameters where the support of the passenger distribution is large relative to the waiting cost, specifically $c\Delta < \overline{\kappa} - \underline{\kappa}$, although it looks much the same for other parameter values: the only difference being in whether or not point X is at a higher value of g_2 than point Y.

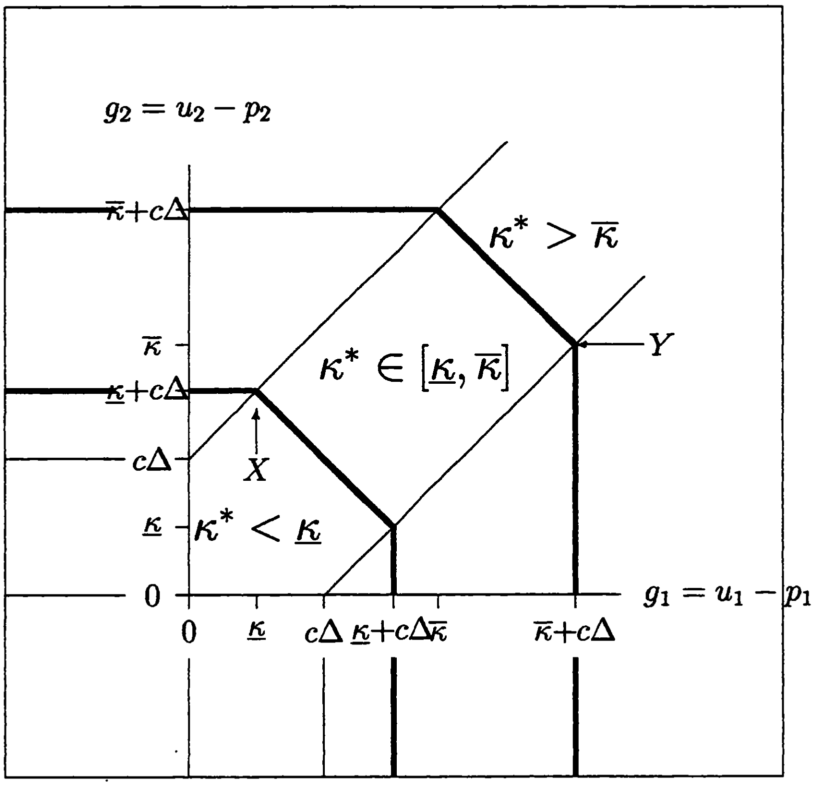

Figure 3.9: Regions of $\kappa^*(g_1, g_2)$. Region boundaries are given by the heavy lines.

The proportion of passengers walking to the stop is given by $F(\kappa^*)$ which is zero whenever $\kappa^* < \underline{\kappa}$ and equal to 1 whenever $\kappa^* > \overline{\kappa}$, and given by $(\kappa^* - \underline{\kappa})/(\overline{\kappa} - \underline{\kappa})$ otherwise. The full dependence on trip payoffs is given below in Equation 3.8. As before, in reading the equation, expressions such as "if (B, B)" should be read as "if $(b_i, b_j) = (B, B)$". The more compact notation has been used to make the various cases clearer.

$$F(\kappa^*) = \begin{cases} \frac{1}{\overline{\kappa}-\underline{\kappa}}\left(\frac{1}{2}(g_1+g_2) - \frac{c\Delta}{2} - \underline{\kappa}\right) & \text{if } \kappa^* \in [\underline{\kappa}, \overline{\kappa}] \text{ and } (b_i, b_j) = (B, B) \\ \frac{1}{\overline{\kappa}-\underline{\kappa}}\left(g_i - c\Delta - \underline{\kappa}\right) & \text{if } \kappa^* \in [\underline{\kappa}, \overline{\kappa}] \text{ and } (b_i, b_j) = (B, S) \\ 1 & \text{if } \kappa^* > \overline{\kappa} \\ 0 & \text{if } \underline{\kappa} > \kappa^* \end{cases} \tag{3.8}$$

Summary of passenger behaviour in local bus markets

The central findings of this section on passenger behaviour, which will be used below, are gathered into Result 16. The result states first, that if the trip payoffs do not differ by too much then passengers board the first bus to arrive. Second, when passengers are boarding the first bus to arrive, then increasing the trip payoff on alternate buses increases the number of passengers walking to the stop.

Result 16 *Passenger behaviour in local bus markets.*
(i) If $|g_i - g_j| \leq c\Delta$ *then* $(b_i, b_j) = (B, B)$.
(ii) If $(b_i, b_j) = (B, B)$ *and* $\kappa^* \in [\underline{\kappa}, \overline{\kappa}]$ *then* $F(\kappa^*) = \frac{1}{\overline{\kappa}-\underline{\kappa}}(\frac{1}{2}(g_1 + g_2) - \frac{c\Delta}{2} - \underline{\kappa})$

The next section turns to the choice of trip payoffs by operators.

3.5.2.1 Monopoly

The analysis begins by considering the fares offered by a monopolist and so Assumption 1 will hold throughout this section.

Characterising Optimal Fares

When choosing the optimal fares to set in order to maximise average revenues per bus, the monopolist must trade off the increase in the number of passengers boarding when the fare is reduced against the cut in the amount some passengers pay.

Denote by n_i the number of passengers which board buses of type i, given the trip payoffs offered on the two types of bus. Then, except for the first bus, which will have no impact on average revenues:

$$n_i = \begin{cases} \frac{1}{\Delta}F(\kappa^*) & \text{if } (b_i, b_j) = (B, S) \\ \frac{1}{2\Delta}F(\kappa^*) & \text{if } (b_i, b_j) = (B, B) \\ 0 & \text{otherwise} \end{cases} \tag{3.9}$$

The monopolist will then choose a pair (g_1^*, g_2^*) to solve:

$$(g_1^*, g_2^*) = \arg\max_{(g_1,g_2)} [u - g_1]n_1 + [u - g_2]n_2$$

The task of characterising optimal fares is greatly simplified by an initial observation, set out in Lemma 4, which states that the monopolist's optimal fares will always put boarding behaviour in (B, B). If this were not so, passengers would board some of its buses, but not others. The operator would do better to cut fares on the expensive buses to the level of those on the cheap ones so that passengers boarded all buses. This would not affect the revenue per passenger, and would encourage more passengers since the expected wait at the stop would be cut.

Lemma 4 *There exists a solution to the monopolist's problem and it satisfies* $(b_1, b_2) = (B, B)$

Proof. Suppose first that a solution exists. By Assumption 4 there is a positive fare at which the monopolist can attract some passengers to the stop. Clearly it cannot be optimal to set fares so that no one walks to the stop and so by Lemma 3, boarding behaviour will be either (B,B) or (B,S). Suppose it is $(b_i,b_j) = (B,S)$ and so all boarding passengers pay p_i. if the operator lowered p_j so that $p'_j = p_i$ then boarding behaviour would be (B,B), all boarding passengers would still pay p_i, and more would board since the expected waiting cost would be lower. Thus (B,S) cannot result from optimal fare setting. A solution, if it exists, will thus lie in (B,B). But the set of trip payoffs for which boarding behaviour is (B,B) is compact, from Equation 3.5, and $[u-g_1]n_1 + [u-g_2]n_2$ is continuous over this region, so that a solution exists. ∎

The importance of Lemma 4 is that optimal fares can be characterised using a modified problem in which the monopolist maximises per bus revenues as though passengers always boarded the first bus to arrive. We will then confirm that at this modified optimum the monopolist chooses fares which do indeed cause passengers to board the first bus, which is enough to confirm that if the optimum exists, it will be the modified one. In the modified monopolist's problem revenues are calculated as though passengers always board the first bus and make their walking decisions accordingly. Specifically we can define a modified function n'_i which will be identical to n_i when boarding behaviour in the original problem is (B,B) and is given by:

$$n'_i = \frac{1}{2\Delta}\frac{1}{\overline{\kappa}-\underline{\kappa}}\left(\frac{1}{2}(g_i+g_j) - \frac{c\Delta}{2} - \underline{\kappa}\right), \qquad g_i, g_j \in [0,u] \tag{3.10}$$

Then let $R'(g_j)$ be i's best response in the modified problem, so it is that trip payoff of i's which would maximise its per bus revenues were the numbers boarding given by n' rather than n.

$$(g'_1, g'_2) = \arg\max_{(g_1,g_2)} [u-g_1]n'_1 + [u-g_2]n'_2, \qquad g_1, g_2 \in [0,u] \tag{3.11}$$

It is easy to confirm that (g'_1, g'_2) will satisfy the first order conditions for

Equation 3.11 which are found to be, by standard calculations:

$$\begin{aligned} g_1' &= \frac{1}{2}\left(2(u - g_2') + c\Delta + 2\underline{\kappa}\right) \\ g_2' &= \frac{1}{2}\left(2(u - g_1') + c\Delta + 2\underline{\kappa}\right) \end{aligned}$$

These will be solved simultaneously at a unique symmetric pair $g_1' = g_2' = g^*$, which implies an equilibrium fare $p^* = u - g^*$ where:

$$\begin{aligned} g^* &= \frac{1}{2}\left(u + \frac{c\Delta}{2} + \underline{\kappa}\right) \\ p^* &= \frac{1}{4}\left(2u - c\Delta - 2\underline{\kappa}\right) \end{aligned}$$

Since, in the modified problem, it is optimal for the monopolist to set fares equal on all buses, passengers will indeed board the first bus to arrive, so that optimal fares are in the region in which the modified and original problems are identical. Therefore, if optimal fares exist, they will be the modified ones.

The reader can easily confirm that these are exactly the results that would have been found from imposing that the monopolist charge the same fare on all buses.

3.5.2.2 Duopoly

Now consider the trip payoffs offered on buses when alternate buses are run by competing operators, i.e. when Assumption2 holds. The analysis will closely follow that for monopoly.

Characterising Equilibrium

As under monopoly, when choosing the optimal fare to set in order to maximise average revenues per bus, operators must trade off the increase in the number of passengers boarding when the fare is reduced against the cut in the amount each passenger pays.

The number of passengers which board i's bus, given the trip payoffs offered by the two operators, will be the same as the number of passengers boarding buses of type i under monopoly, viz. the function n_i given in Equation 3.9.

Operator i's problem is to choose its best response $R(g_j)$ to the trip payoff offered by its rival, that is to solve the problem:

$$R(g_j) = \arg\max_{g_i}[u - g_i]n_i$$

As with monopoly, the task of characterising equilibrium is made significantly easier by the observation that if there is an equilibrium then it must be the case that passengers board the first bus to arrive. The reason is simple: any operator which finds itself with no passengers is earning no revenues and could have done better by cutting price until passengers wished to board. Lemma 5 states this more precisely.

Lemma 5 *If a price equilibrium in pure strategies exists, in which $F(\cdot) > 0$, then $(b_i, b_j) = (B, B)$.*

Proof. Note first that there can be no price equilibrium in which one operator earns positive revenues while the other does not. Suppose the contrary, and suppose, w.l.g., that $p_1 > 0$, $\kappa^* > \underline{\kappa}$ and $(b_1, b_2) = (B, S)$ which implies $p_1 \leq p_2 - c\Delta$. Operator 2 is getting no revenue. If it charged $p_2' = p_1$ then the number of passengers would not fall (κ^* is monotonically non-decreasing in g_i), and half of them would board bus 2, giving operator 2 positive revenues. Thus it cannot have been an equilibrium to have $(b_1, b_2) = (B, S)$. By Lemma 3 the only other region giving positive flows of passengers is $(b_i, b_j) = (B, B)$. This proves the first part. Now note that if $F(\cdot) > 0$ then prices must be strictly positive. If not then neither operator is earning positive returns, while at least the one charging the lower price is making positive sales. This one could raise price a little, still get positive sales ($F(\cdot)$ continuous, and the firm gets at least half the market so long as it charges a fare no greater than its rival's plus $c\Delta$) and so get strictly positive returns.■

Again, as with monopoly, Lemma 5 allows us to characterise equilibrium using a modified problem in which operators maximise fares as though passengers always boarded the first bus. We then confirm two things. First we show that in this modified equilibrium operators choose fares which do indeed cause passengers to board the first bus. This is enough to confirm that if equilibrium exists, it will be the modified one. Second, however, we show that there are parameters for which, even if operators assumed the correct functions for passenger boarding behaviour, their choice of fare would be the same as in the modified equilibrium, which means that there are parameters for which equilibrium exists.

In the modified problem operator i aims to choose its best response $R'(g_j)$ to the trip payoff offered on j's buses as though passengers always boarded the first bus, and make their walking decisions accordingly. If n' be as given in Equation 3.10, the modified problem is to solve:

$$R'(g_j) = \arg\max_{g_i}[u - g_i]n'_i \tag{3.12}$$

It is easy to confirm that $R'(g_j)$ will satisfy the first order condition $\partial/\partial g([u-g]n') = 0$ and is found to be:

$$R'(g_j) = \frac{1}{2}(u - g_j + c\Delta + 2\underline{\kappa}) \tag{3.13}$$

If p' is the fare corresponding to a trip payoff of g' then

$$p' = \frac{1}{2}(2u - p_j - c\Delta - 2\underline{\kappa})$$

Since the modified best response functions are the same for both operators, are linear and have slope different from 1, the modified equilibrium

(g_1^*, g_2^*) at which $R'(g_i^*) = g_j^*$, $i = 1, 2$, $j \neq i$, is unique and symmetric, and so let $g_1^* = g_2^* = g^*$, and let $p^* = u - g^*$. Substituting $R'(g^*) = g^*$ into Equation 3.13 and rearranging gives the modified equilibrium trip payoffs and fares:

$$\begin{aligned} g^* &= \frac{1}{3}\left(u + c\Delta + 2\underline{\kappa}\right) \\ p^* &= \frac{1}{3}\left(2u - c\Delta - 2\underline{\kappa}\right) \end{aligned}$$

Since operator fares in the modified equilibrium are equal, passengers will board the first bus to arrive, i.e. the equilibrium lies in (B, B), the region in which the functions n' and n are identical. Moreover, if equilibrium exists in the original problem it will lie within this region, and so if equilibrium exists it will be the modified one.

Having characterised equilibrium, the analysis turns to consider when an equilibrium exists.

Existence

This section examines the conditions under which equilibrium exists. The reason why equilibrium sometimes fails to exist is that at the modified equilibrium outlined above, each operator may have an incentive to cut fares to such a low level that all passengers waiting at the stop board its buses. As an example Figure 3.10 illustrates the number of passengers boarding 1's bus as a function of g_1 for the case $g_2 > c\Delta$ and $\overline{\kappa} - \underline{\kappa} > c\Delta$.

Three distinct regions are clear from Figure 3.10, corresponding to different regions of boarding behaviour. At low trip payoffs, specifically $g_1 < g_2 - c\Delta$, no passengers board bus 1. The trip payoff is so low that either none walks to the stop, if g_2 is also low, or if some do walk to the stop then all wait for a bus 2 to arrive, rather than board bus 1. At intermediate

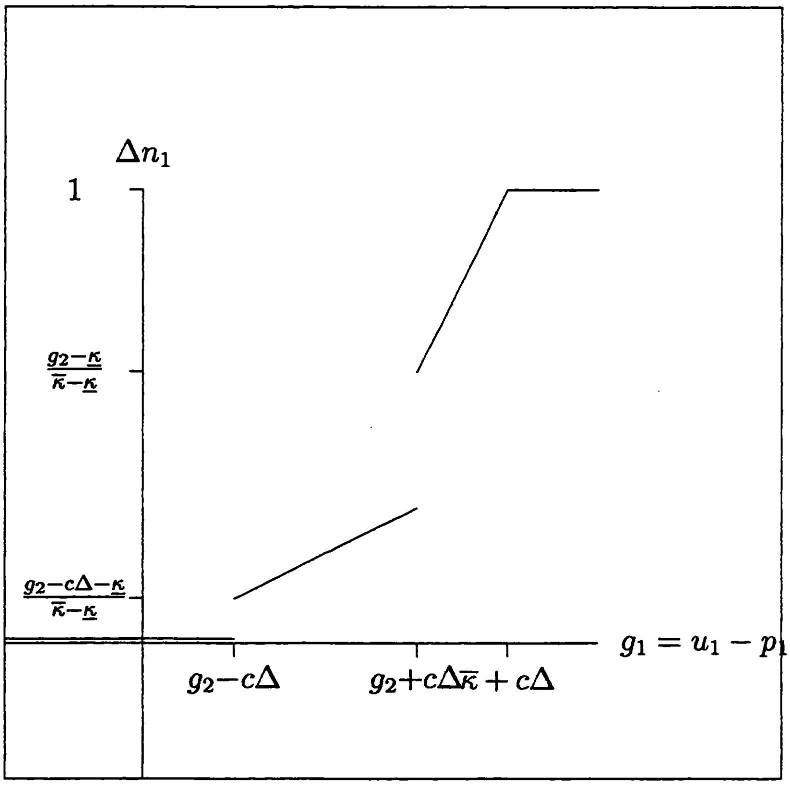

Figure 3.10: n_1 when $g_2 < c\Delta$ and $\overline{\kappa} - \underline{\kappa} > c\Delta$

prices, when $g_2 - c\Delta < g_1 < g_2 + c\Delta$, the trip payoff on 1's buses is sufficiently high that passengers at the stop when a 1 arrives will now prefer to board than wait for the next bus, and so boarding behaviour is (B, B). At $g_1 = g_2 + c\Delta$ there is a sudden doubling of the number of passengers boarding bus 1 since once the trip payoff exceeds $g_2 + c\Delta$ then all passengers at the stop prefer to wait for a bus 1 to arrive than board a bus 2.

There are two candidate best responses. There is a local maximum which maximises profit given that boarding behaviour is (B, B), and so the operator gets just half of the passengers walking to the stop. This is the modified best response analysed above. However, there is also a local maximum at a discretely higher trip payoff, and so lower fare. This maximises profits given that boarding behaviour is (B, S), so that the operator has undercut its rival and all passengers now board its buses.

Equilibrium may fail to exist, therefore, because of the possible incentive of operators at the modified equilibrium to undercut their rival by a large margin, thus winning all waiting passengers. However, it is only possible to undercut if fares in the modified equilibrium are sufficiently high, since in order to win all passengers one operator must charge at least $c\Delta$ less than its rival. Clearly then, undercutting is impossible if $p^* < c\Delta$. This will happen when:

$$\begin{aligned} \frac{1}{3}(2u - c\Delta - 2\underline{\kappa}) &< c\Delta \\ \text{or } u &< \underline{\kappa} + 2c\Delta \end{aligned}$$

Result 17 *If $u < \underline{\kappa} + 2c\Delta$ then equilibrium exists.*

Note that this is a strong condition guaranteeing existence: it guarantees that it is not possible to undercut and steal all passengers at the modified equilibrium. A weaker condition would rule out only those cases where undercutting was profitable.

3.5.2.3 Fares in local bus markets.

For convenience, we summarise the results on fares in local bus markets in Result 18 below. which shows that fares are lower under monopoly than duopoly.

Result 18 *Fares in local bus markets.*

(i) Optimal fares under monopoly exist, are the same on all buses, and are:

$$p^* = \frac{1}{4}(2u - c\Delta - 2\underline{\kappa})$$

(ii) Equilibrium fares under duopoly exist if $u < \underline{\kappa} + 2c\Delta$, are the same on all buses, and are:

$$p^* = \frac{1}{3}(2u - c\Delta - 2\underline{\kappa})$$

(iii) Fares are lower under monopoly than duopoly

3.5.2.4 Comparing Monopoly and Duopoly when fares are set at the bus stop.

There is an obvious institutional reason why we can take arrival times as fixed before the day begins: arrival times are fixed by the legal requirement that operators register their timetables. There are no corresponding formal constraints which dictate that fares on buses are set before the day begins, although in practice this is what happens. Here we drop this assumption and no longer require fares to be equal on alternate buses, so that we no longer restrict $p_s = p_1$, s odd, and $p_s = p_2$, s even. Instead we examine fare setting when fares are set when a bus is at the stop and look for strategies for each type of bus, $p_i(\cdot)$, $i = 1, 2$. In analysing behaviour in the model we will look for an equilibrium described by strategies for all players which satisfy a number of conditions. First, a player's strategy will specify an action, at each time at which the player can make a decision, which maximises the players expected payoff during the rest of the game, given the strategies of its rivals. In other words equilibrium is Perfect and Nash. Second, strategies are conditioned only on the current environment, and not on the past history of the game, or on time, i.e. strategies are Markov. Let x be the number of passengers at the stop when a bus arrives. We seek a set of Markov strategies;

$$\begin{aligned} b(p) &: \mathcal{R} \to \mathcal{B}, \\ a(\tau, \kappa) &: \mathcal{R} \times [\underline{\kappa}, \overline{\kappa}] \to \mathcal{A}, \end{aligned}$$

$$p_i(x) \; : \; \mathcal{R} \rightarrow \mathcal{R}, \quad i = 1, 2,$$

which form a Perfect Nash Equilibrium (PNE). The following Result characterises such an equilibrium.

Result 19 *There exists a MPNE in which no one travels by bus. Specifically, the strategies;*

$$\begin{aligned} b(p) &= \begin{cases} B & \text{if } p \leq u \\ X & \text{if } p > u \end{cases} \\ a(\tau, \kappa) &= X, \quad \forall \tau, \kappa \\ p_i(x) &= u, \quad i = 1, 2 \end{aligned}$$

form a MPNE.

Proof. Consider first $b(p)$, the choice made by waiting passengers when a bus arrives and offers fare p. If all buses in the future charge a price equal to their quality, the trip payoff in the future is zero. Therefore, passengers always do better to exit now rather than stay at the stop and incur additional waiting costs before exiting later or boarding a later bus. Passengers prefer to board now than exit if boarding gives a positive trip payoff, confirming $b(p) = B$ when $p < u$. If the trip payoff is negative, so $p > u$, then it is better to exit, so $b(p) = X$ in this case, and, finally, if the trip payoff is zero then passengers are indifferent between boarding and exiting, so that $b(p) = B$ when $p = u$ is one best response. Now consider the price offered by a bus of type i. All passengers board so long as $p \leq u$ and since, if there are any passengers at the stop, revenues rise with price, the bus charges $p_i(x) = u$, $\forall x > 0$. If there are no passengers then revenues will always be zero and $p_i(0) = u$ is a best response. Next consider $a(\tau, \kappa)$. Since all passengers earn zero trip payoff, the cost of walking to the stop is never recovered, and so it is always better for passengers to stay at home. ■

In equilibrium no passengers ever walk to the bus stop. Bus fares are always set at u, which is the full value of a bus trip to passengers. Passengers, knowing this, will board any bus with a fare of u or less, rather than wait at the stop and incur additional waiting costs. Given this behaviour on the part of passengers buses maximise revenues by charging the highest fare consistent with all passengers boarding. However, if passengers pay u to board a bus, then it is never worthwhile incurring the cost of walking to the stop.

The result rests on the type of hold-up problem that is explored in the literature on property rights[13]. By the time passengers and bus drivers come to trade over the level of bus fares, passengers have already sunk the costs associated with walking to the stop and waiting for the bus. Bus drivers can make a take it or leave it offer over the fare level and extract the full value of the ride from passengers, leaving them with no surplus to cover their sunk costs. Of course operators would be better off if they could assure passengers that they would set a lower fare, thus enticing them to walk to the stop and providing positive sales, albeit at the lower fare. However, in a Perfect equilibrium they cannot do so.

In fact, of course, people *do* travel by bus. It might be argued that reality differs from the model's result because opportunistic behaviour on the part of bus operators drives only some, and not all, potential passengers to stay at home: it might be thought for example that the starkness of the result arises from the simplifying assumption that all passengers make the same choice so that an operator which raises its fare will either keep all passengers or lose them all, or that the bargaining assumption, that buses make a take it or leave it offer to passengers, might also contribute to the starkness of the

[13]See, for example, the paper by Hart and Holmstrom on "The theory of contracts" [34]and the references therein.

result. More importantly, it might be observed that operators and passengers in fact interact over many days, and that current prices somehow affect the number of future passengers: if passengers regretted walking to the stop on one day, they may stay at home the next. In this case, operators would be concerned with the effect of their fares on the number of *future* passengers. What is certain is that in the real world the operational structure of bus companies rules out any opportunistic behaviour in the first place. Fares are in fact set centrally, and individual bus drivers do not have the authority to start renegotiating the fare with waiting passengers[14]. This operational structure may be chosen as a device to allow operators to credibly set their fares before the day begins, though it more likely results from factors outside those considered here. All that is important here is that the more appropriate assumption for modelling bus behaviour would seem to be that prices are set before the day begins, so that operators consider the effect on passenger numbers of their pricing decisions.

3.5.3 Express Coaches

The simpler market to analyse, and the more familiar, is that for express coaches in which passengers know the arrival time of buses. As before, passenger behaviour can be analysed before moving on to consider the fares set by operators under different market structures.

Passenger Behaviour

Passengers now know the timetable and so there is no reason for them to spend costly time waiting at the bus stop: they can time their departure to

[14]This is in contrast to the case of minicabs, where individual drivers *do* have some discretion over the fare. Minicabs, note, drive up to the door so passengers have not sunk a walking cost by the time they negotiate over fare.

arrive at the same time as a bus. Moreover, they know bus fares, and so if one bus is cheaper than the other they can make sure that they always catch the cheap one.

The distinguishing assumption for express coach markets is Assumption 5 below that passengers' information sets contain the exact arrival times of coaches.

Assumption 5 *Express Coaches*

$$\Phi_a = \{p_s, \Delta, t_1\}$$

We also make a technical assumption on parameters: passengers would always prefer to leave the house at an inconvenient time and arrive just in time to catch a coach than wait at the bus station for a coach to arrive.

Assumption 6 $\varepsilon < c$

In fact for much of the analysis, we take $\varepsilon = 0$, but we decide between multiple optimal strategies by taking the limit $\varepsilon \to 0$.

Passenger boarding behaviour, b_i

Once at the stop, boarding behaviour is as with local buses, and so is given by Equation 3.6. The difference in behaviour between express coaches and urban local buses lies in the decision to walk to the stop, and the time at which to leave the house.

Passenger travelling behaviour, $a(\tau, \kappa)$

In the case of express coaches, where passengers know bus times in advance, and where it is more costly to spend time at the coach station than to leave home at an inconvenient time, they will, in equilibrium, never wait at the stop, but will arrive at the same time as the coach they intend to board.

Passengers choosing to travel will therefore always choose a leaving time to coincide with one of the coaches.

Note that, for any positive inconvenience cost, passengers will only ever board either the closest odd coach or the closest even one (if any), since any other choice would involve unnecessary inconvenience. This means that a travelling passenger (τ, κ), where τ lies between the arrival of the s'th and $s+1$'th coach, will travel at either t_s or t_{s+1}. When the inconvenience cost is zero, travelling on one of theses two coaches is still optimal, but so is travelling on many others.

When contemplating which, if any, of the two closest coaches to board, passengers are aware of the fares. If the two sorts of coach charge the same fare, passengers do best to board the closest one, and so minimise the cost of travelling at an inconvenient time. This behaviour is uniquely optimal for any positive inconvenience cost (though when $\varepsilon = 0$, catching any other coach is also optimal). If the two sorts of coaches charge different fares, however, passengers will board that coach charging the lowest fare, if the fare difference is sufficiently high to compensate for the inconvenience. If the inconvenience cost is zero, *any* fare difference is large enough (though now it is also optimal to board any coach charging the lower fare). For example, a passenger (τ, κ) where τ lies between the arrival of the s'th and $s+1$'th coaches, who intends to travel, will use the s'th coach if:

$$\begin{aligned} u - p_s - \varepsilon(\tau - t_s) &> u - p_{s+1} - \varepsilon(t_{s+1} - \tau) \\ \text{or} \qquad p_{s+1} - p_s &> \varepsilon(t_{s+1} + t_s - 2\tau) \end{aligned}$$

Suppose that the above inequality is true, so that it is better for the passenger to travel at t_s than t_{s+1}. It will be worth making the journey at all if trip payoff on s, net of inconvenience cost, is greater than the agent's walking cost, i.e. if:

$$u - p_s - \varepsilon(\tau - t_s) \geq \kappa$$

Writing the optimal travelling behaviour of passengers when the inconvenience cost is ε as $a^\varepsilon(\tau, \kappa)$ gives:

$$\begin{aligned} \text{if } \tau &\in (t_s, t_{s+1}] \\ a^\varepsilon(\tau,\kappa) &= \begin{cases} t_s & \text{if } p_{s+1} - p_s > \varepsilon(t_s + t_{s+1} - 2\tau) \\ & \text{and } u - p_s - \varepsilon(\tau - t_s) \geq \kappa \\ t_{s+1} & \text{if } p_{s+1} - p_s \leq \varepsilon(t_s + t_{s+1} - 2\tau) \\ & \text{and } u - p_{s+1} - \varepsilon(t_{s+1} - \tau) \geq \kappa \\ X & \text{otherwise} \end{cases} \end{aligned}$$

There are a few points to note here. First, when fares are different and when the inconvenience cost becomes small, passengers travel on the cheaper coach, regardless of which is closer. The second point is that as the fare differential gets very small behaviour depends only on whether $(t_s + t_{s+1} - 2\tau)$ is positive or negative, i.e. on whether τ is closer to t_s or t_{s+1}. In other words, as fares on alternate buses become equal, passengers need take into account only the cost of travelling at an inconvenient time. Here we take the view that the cost of travelling at an inconvenient time is not a driving force underlying whether or not fares are higher under monopoly than duopoly, and accordingly we henceforth assume that $\varepsilon = 0$ and take passenger travelling behaviour $a(\tau, \kappa)$ to be the limit of $a^\varepsilon(\tau, \kappa)$ as $\varepsilon \to 0$.

Assumption 7 $\varepsilon = 0$

Let s satisfy $\tau \in (t_s, t_{s+1}]$ for passenger (τ, κ). Then if $\max\{g_1, g_2\} < \kappa$:

$$a(\tau, \kappa) = X$$

otherwise:

$$a(\tau,\kappa) = \begin{cases} t_s & \text{if } p_s > p_{s+1} \\ & \text{or } p_s = p_{s+1} \quad \text{and } |t_s - \tau| < |t_{s+1} - \tau| \\ t_{s+1} & \text{otherwise} \end{cases} \tag{3.14}$$

The number of passengers walking to the stop, $F(\kappa^*)$

As described above, passengers always board a coach offering the highest trip payoff and so the payoff of a passenger travelling by coach is $\kappa^* = \max\{g_1, g_2\}$ and the number of passengers walking to the stop is then:

$$F(\kappa^*) = \frac{1}{\overline{\kappa} - \underline{\kappa}}(\max\{g_1, g_2\} - \underline{\kappa}) \tag{3.15}$$

3.5.3.1 Monopoly

Suppose now that Assumption 1 holds and so all coaches are run by a single monopolist. The number of passengers boarding its coaches is $F(\kappa^*)$ and depends only on the largest trip payoff offered. Moreover by Equation 3.14 passengers will board only those buses with the highest trip payoff and so revenues per passenger also depend just on this. Therefore the monopolist's problem is to choose a pair of trip payoffs (g^*, g') to solve:

$$\begin{aligned} g^* &= \arg\max_g (u - g)\frac{1}{\Delta(\overline{\kappa} - \underline{\kappa})}(g - \underline{\kappa}) \\ g' &\leq g^* \end{aligned}$$

Standard calculations show that any solution satisfies:

$$\begin{aligned} g^* &= \frac{1}{2}(u - \underline{\kappa}) \\ g' &\leq \frac{1}{2}(u - \underline{\kappa}) \end{aligned}$$

and so the lowest fare offered, $p^* = u - g^*$, will be:

$$p^* = \frac{1}{2}(u + \underline{\kappa}) > 0$$

3.5.3.2 Duopoly

Now suppose that Assumption 2 holds so that alternate coaches are run by two different operators. From Equations 3.14 and 3.15, the number of passengers boarding operator i's coaches given the trip payoffs offered by both operators, n_i, will be:

$$n_i = \begin{cases} \frac{1}{\Delta}\frac{g_i - \underline{\kappa}}{\overline{\kappa} - \underline{\kappa}} & \text{if } g_i > g_j \\ \frac{1}{2\Delta}\frac{g_i - \underline{\kappa}}{\overline{\kappa} - \underline{\kappa}} & \text{if } g_i = g_j \\ 0 & \text{if } g_i < g_j \end{cases}$$

If its rival charges a positive fare, it can never be optimal for an operator to charge a higher fare: no one would board the operator's coaches, whereas charging the same or a lower fare as its rival gives positive numbers boarding at a positive revenue per passenger. Moreover, charging the same fare cannot be optimal since charging infinitesimally less leads to a negligible fall in per passenger revenues, but at least doubles the number of passengers boarding. The best response, then, must be to charge less than the rival, if the rival charges a positive fare. Clearly there is no set of fares (p_1^*, p_2^*) at which either is positive and both are charging less than their rival. If the rival charges a zero fare, then charging anything gives zero revenue and is a best response and so $p_1^* = p_2^* = 0$ is the unique equilibrium fare, and $g_1^* = g_2^* = u$ the unique equilibrium trip payoff.

Result 20 *Fares in express coach markets.*

(i) Optimal fares under monopoly exist, are unique, are the same on all coaches, and are:

$$p^* = \frac{1}{2}(u + \underline{\kappa})$$

(ii) Equilibrium fares under duopoly exist, are unique, are the same on all coaches, and are:

$$p^* = 0$$

3.5.4 Robustness

The model analysed here has been designed to make transparent the primary impact of assuming that potential bus passengers do not know the exact arrival times of buses. This transparency leads to stark results. When passengers do not know the timetable, fares are higher under duopoly than under monopoly. When, on the other hand, they do know when buses arrive, fares are higher under monopoly. Here we consider, informally, the likely effect on the results of relaxing some of the assumptions.

First, suppose that frequency, instead of being given as is assumed so far, were endogenous. In fact, when there is competition both in real express coach markets and local bus markets, frequency appears to be higher, and so the interval between buses, Δ, smaller. In the results for the model local bus market, the duopoly and monopoly fares are higher as Δ is smaller. To see why note that if Δ is high, not many passengers walk to the stop because the waiting cost is so high. But then, cutting fares is not very costly in terms of loss of revenue from existing passengers, while the gain in numbers walking to the stop is the same. If endogenising frequency in the model did result in higher frequency under duopoly, this would only serve to reinforce the finding that duopoly fares are higher than monopoly ones in local bus markets. In the model express coach market the interval between buses has no effect on prices in either market structure, so that endogenising market structure will not affect the finding that monopoly fares are higher than duopoly ones. It is worth noting, incidentally, that a monopolist in the express coach market would only choose to run one bus over the whole day, so frequency could only rise under duopoly (so long as there was some positive cost to running a bus).

Another stark assumption that has been maintained in the express coach market is that the inconvenience cost ε is zero. It is certainly plausible that passengers suffer a cost if they travel at a time which differs from their most preferred one, and introducing this would give duopolists some market power over close passengers, so that they would likely be able to charge a positive fare. It would remain the case, however, that if an operator cut fares on one type of coach, those coaches would gain market share from the other type, giving each duopolist excessive incentives to charge low fares compared to a monopolist.

It was noted earlier in the model local bus market that passengers arriving at the stop learnt nothing about the timetable. Suppose, on the contrary, that they did. There may then be some passengers who walked to the stop and, on learning how long was the wait until the next bus, and how much the fare would then be, decided to go straight home again. Cutting fares on one type of bus would then affect that type's market share, as well as the number of passengers walking to the stop. However, the additional market share would be gained not from the other type of bus, but just from affecting the number of passengers going straight home from the stop, and so it would still be the case that cutting fares on one type of bus increased the number of passengers boarding the other type, creating an under-incentive for duopolists to cut fares compared with a monopolist.

The final question concerns the robustness of the results in the situation where quality is a choice variable. This must remain an open issue: there seems to be no compelling way to model quality choice. Different ways could be appropriate in different markets. In local bus markets it seems that the main determinant of quality is the newness and maintenance levels of buses, and there is little evidence of differentiation in quality. In express coach markets operators on both duopoly and monopoly routes have introduced vertically differentiated services, with quality being such factors as journey time and whether or not drinks are available. Also marketing has been cited as playing a large role in determining market shares in express coach markets, which can be taken as a measure of spending on perceived quality and is a fixed cost with respect to the number of coaches.

3.6 Conclusions

In this chapter we have made a simple, well motivated assumption on the way that passenger behaviour differs between local bus routes on which the frequency of buses is high, and express coach routes where there are fewer departures. This assumption is that passengers travelling on high density routes tend to arrive at the stop independently of bus arrival times, and to board the first bus to arrive, whereas passengers using express coaches will arrive in time to board their preferred coach. We have shown that this assumption can provide an explanation for two different aspects of the experience of competition in the deregulated bus and coach markets of the 1980's and 90's. In the first model of a local bus market, there are frequent changes to bus timetables, bus arrival times bunched together, and the bus with least passengers being the most likely to make a revision. The first two of these results have been observed in local bus markets in which there is competition. In the second model local bus market we found that two operators competing set a higher fare than did a monopolist. Again this accords with what has been observed in local bus markets, where fares have risen since deregulation, and risen by most in those metropolitan areas where competition is most common.

Finally we will make a few remarks about quality. We have already noted in the discussion of the robustness of the results on fare levels that quality should properly be modelled in a different way in the two types of market. Here we will mention local bus markets where the main determinant of quality is the newness and maintenance levels of buses. The dominant view of those giving evidence to the Select Committee was that quality was lowest where there was most competition. This is not particularly surprising: competition reduces a firm's market share, and so the number of customers from which it can recoup any investments in advertising. The contrary would not have been surprising either: if advertising affects only market share, and not the size of the market, a monopolist will not advertise at all, though competing firms may well do so. However, there is additional evidence in the Select Committee Report that part of the explanation for the low quality is that passengers arrive independently of bus arrival times, and board the next bus to arrive. This additional evidence is the experience of Stagecoach in Manchester, related by the Traffic Commissioners in their testimony to the Committee. In Manchester, Stagecoach entered a market with new, high quality buses, while the incumbents used old poorly maintained ones. The Commissioners were surprised that the quality of Stagecoach's buses did not enable them to win greater market share (Stagecoach eventually withdrew from the market), and they attributed this failure to customer loyalty. While some customers may have been locked into using the incumbent's services through discount cards, if passengers turn up at the stop and board the first bus we would not expect higher quality buses to get higher markets.

This in turn suggests that one reason for the lower quality of buses on those local bus routes with competition is essentially the same as that for high fares. The mechanism would operate somewhat as follows: as the average quality of buses rises, consumers make more journeys by bus. However, once a consumer has made the decision to take a bus, he then arrives at the stop and takes the first bus to arrive, regardless of its quality. Firms investing in newer or better maintained buses do not, therefore, gain all the extra customers that result from this investment, leading to under-investment in quality.

Appendix A

Proofs

A.1 Exploring a branching structure

A.1.1 Optimality of the Gittins index policy for simple bandit processes

We give an outline of the proof of the optimality of the Gittins index policy for multi-armed bandits; it is essentially from Whittle's [63] and also used by Berry and Fristedt [5]. It is included here for accessibility, and contains some notational changes and expository material due to the present authors.

There are N projects[1] and in each discrete period you can work on only one project. The state of project k at time t is denoted by $x_k(t)$, and the project engaged at time t is denoted by $k(t)$. The state variable at time t is written as $x(t) = \langle x_1(t), x_2(t), \ldots, x_N(t) \rangle$, and the information at time t, namely past and current states and past actions, is written as $I(t)$. If

[1]Here, we are dealing with single-action projects. At any stage, each of the fixed number of projects has a single action (i.e. there is no branching) so that the notions of engaging a project and selecting an action are interchangeable.

project k is engaged at time t then you get an immediate expected reward of $R_k(x_k(t))$. Rewards are additive and discounted in time by a factor β. States of unengaged projects do not change and the state of the engaged project changes by a Markov transition rule: if $k(t) \neq k$ then $x_k(t+1) = x_k(t)$, and if $k(t) = k$ then the value of $x_k(t+1)$ is conditioned only by k & $x_k(t)$.

Assume that rewards are uniformly bounded:

$$-\infty < -B(1-\beta) \leq R_k(x) \leq B(1-\beta) < \infty.$$

Writing $R(t)$ for the reward $R_{k(t)}(x_{k(t)}(t))$ realised at time t, the total discounted reward is then $\sum_0^\infty \beta^t R(t)$ with a maximal expected reward $F(x)$ over feasible policies π given by:

$$F(x(0)) = \sup_\pi \mathrm{E}_\pi \Big[\sum_0^\infty \beta^t R(t) \ \mid \ I(0) \Big].$$

F will be the unique bounded solution to the dynamic programming equation:

$$F = \max_k L_k F$$

where L_k is the one-step operator if k is the project engaged:

$$L_k F(x) = R_k(x_k) + \beta \,\mathrm{E}\Big[F(x(t+1)) \ \mid \ x(t) = x, k(t) = k \Big].$$

Introduce a fall-back M, where the option of taking the fall-back remains open at all times. The maximal expected reward of the modified process, conditional on $x(0) = x$, is $\Phi(M, x)$ and solves

$$\Phi = M \ \vee \ \max_k L_k \Phi. \tag{A.1}$$

Let $\phi_k(m, x_k)$ be the analogue of $\Phi(M, x)$ when only project k is available; ϕ_k solves

$$\phi_k = m \ \vee \ L_k \phi_k. \tag{A.2}$$

(L_k changes only x_k, so $L_k\phi_k$ is well-defined.)

The Gittins index, denoted by $m_k(x_k)$, is the infimal value of m such that $m = \phi_k(m, x_k)$, namely the alternatives of stopping with m_k and of continuing project k (with the option of taking the fall-back staying open) are equitable, and so $m_k = L_k\phi_k$.

It is fairly easy to show that $\Phi(M, x)$, as a function of M, is non-decreasing & convex (convexity following from the fact that we are dealing with the supremum of expressions which are linear in M), and that $\Phi(M, x) = M$ when $M \geq B$. Also $\Phi(M, x) = F(x)$ when $M \leq -B$.

Similarly, $\phi_k(m, x_k)$, as a function of m, is non-decreasing & convex, and $\phi_k(m, x_k) = m$ when m is large, certainly if $m \geq B$, and more precisely for $m \geq m_k$, so $m_k \leq B$. Note that, since $\phi_k(m, x_k)$, as a function of m, is convex, the derivative $\partial\phi_j(m, x_j)/\partial m$ exists almost everywhere.

We "guess" the form of the value function:

$$\Theta(M, x) = B - \int_M^B \prod_k \frac{\partial\phi_k(m, x_k)}{\partial m}\, dm,$$

and proceed to verify it by showing two things:

- Θ satisfies (A.1), that is $\Theta = M \vee \max_k L_k\Theta$;
- the action recommended by the Gittins index maximises the RHS of the above equation, i.e. when $M > \max_k L_k\Theta$ it selects the fall-back, and when $M < \max_k L_k\Theta$ it selects the project which maximises $L_k\Theta$.

$$\begin{aligned} \text{So, define } P_k(m, x) &= \prod_{j \neq k} \frac{\partial\phi_j(m, x_j)}{\partial m} \\ \text{and } m_{\neg k} &= \max_{j \neq k} m_j. \end{aligned}$$

$P_k(m,x)$, as a function of m, is non-negative and non-decreasing, and $P_k(m,x) = 1$ for $m \geq m_{\neg k}$. (These follow directly from the properties of ϕ_j.) Note that

$$d_m P_k(m,x) \geq 0 \; \forall \; m, \text{ and } d_m P_k(m,x) = 0 \text{ for } m \geq m_{\neg k}.$$

Rewrite $\Theta(M,x)$ as

$$\Theta(M,x) = B - \int_M^B \frac{\partial \phi_k(m,x_k)}{\partial m} P_k(m,x)\, dm$$

and use integration by parts to obtain

$$\begin{aligned} \Theta(M,x) &= B - \Big[\phi_k(m,x_k)P_k(m,x)\Big]_M^B + \int_M^B \phi_k(m,x_k)\, d_m P_k(m,x) \\ &= \phi_k(M,x_k)P_k(M,x) + \int_M^B \phi_k(m,x_k)\, d_m P_k(m,x) \end{aligned}$$

noting that $\phi_k(B,x_k) = B$ because $m_k \leq B$, and $P_k(B,x) = 1$ because $m_{\neg k} \leq B$. Also, $d_m P_k(m,x) = 0$ when $m \geq B$, so we can amend the range of integration:

$$\Theta(M,x) = \phi_k(M,x_k)P_k(M,x) + \int_M^\infty \phi_k(m,x_k)\, d_m P_k(m,x). \tag{A.3}$$

Now fix x, so we can focus on the dependence of various function on m or M. We want to show that:

$$\begin{aligned} \Theta(M) &\geq M \text{ for any } M, && \text{(A.4)} \\ \text{and } \Theta(M) &= M \text{ iff } M \geq \max_j m_j; && \text{(A.5)} \\ \text{and that } \Theta(M) &\geq L_k\Theta(M) \text{ for any } M, && \text{(A.6)} \\ \text{and } \Theta(M) &= L_k\Theta(M) \text{ iff } m_k = \max_j m_j \text{ and } M \leq m_k. && \text{(A.7)} \end{aligned}$$

(A.5) 'if' & part of (A.4): Consider $M \geq \max_j m_j$.

In this case, $\phi_k(M) = M$, $P_k(M) = 1$; and $d_m P_k(m) = 0$ for $m \geq M$. So from (A.3):

$$\Theta(M) = M.$$

(A.5) 'only if' & rest of (A.4): Consider $M < \max_j m_j$.

Let $k = \arg\max_j m_j$. So $M < m_k$, and we have $\phi_k(M) > M$. When $M \leq m < m_k$, $\phi_k(m) > M$, and when $m \geq m_k$, $d_m P_k(m) = 0$. So from (A.3):

$$\begin{aligned}
\Theta(M) &> M\left(P_k(M) + \int_M^{m_k} d_m P_k(m)\right) \\
&= M P_k(m_k) \\
&= M, \text{ because } P_k(m_k) = 1.
\end{aligned}$$

So from (A.4) and (A.5):

$$\begin{aligned}
M < \max_j m_j &\Rightarrow \Theta(M) > M \\
M = \max_j m_j &\Rightarrow \Theta(M) = M \\
M > \max_j m_j &\Rightarrow \Theta(M) = M.
\end{aligned} \tag{A.8}$$

Now, define $\delta_k(m, x_k) = \phi_k(m, x_k) - L_k \phi_k(m, x_k)$.

Fixing x again, note that $\delta_k(m) \geq 0 \ \forall\ m$, and $\delta_k(m) = 0$ for $m \leq m_k$ and that

$$\Theta(M) - L_k \Theta(M) = \delta_k(M) P_k(M) + \int_M^\infty \delta_k(m)\, d_m P_k(m) \tag{A.9}$$

which follows from applying the one-step operator L_k to each side of (A.3), subtracting the result from (A.3), and applying the definition of δ_k to the RHS.

(A.6): $\delta_k(m) \geq 0$, and $P_k(m)$ is non-negative and non-decreasing, so from (A.9) we have

$$\Theta(M) \geq L_k \Theta(M).$$

(A.7): When $m_k \geq M$, $\delta_k(M) = 0$, so the first term on the RHS of (A.9) is 0.

When $m_k \geq M$ and $m_k = \max_j m_j$, the integral on the RHS of (A.9) is 0, because either $\delta_k(m) = 0$, or $d_m P_k(m) = 0$, or both. So we have

$$\Theta(M) = L_k\Theta(M).$$

(If either $m_k < M$ or $m_k < m_{\neg k}$, then at least one term on the RHS of (A.9) is positive.)

Now using (A.6) & (A.7) with the implications from (A.8), and with $k = \arg\max_j m_j$:

$$\begin{aligned} M < \max_j m_j = m_k &\Rightarrow \Theta(M) = L_k\Theta(M) \text{ and } \Theta(M) > L_{\neg k}\Theta(M) \\ &\Rightarrow \max_j L_j\Theta = L_k\Theta = \Theta > M, \text{ so } \{M \vee \max_j L_j\Theta\} = L_k\Theta; \end{aligned}$$

$$\begin{aligned} M = \max_j m_j = m_k &\Rightarrow \Theta(M) = L_k\Theta(M) \text{ and } \Theta(M) > L_{\neg k}\Theta(M) \\ &\Rightarrow \max_j L_j\Theta = L_k\Theta = \Theta = M, \text{ so } \{M \vee \max_j L_j\Theta\} = M \equiv L_k \end{aligned}$$

$$\begin{aligned} M > \max_j m_j = m_k &\Rightarrow \Theta(M) > L_k\Theta(M) \text{ and } \Theta(M) > L_{\neg k}\Theta(M) \\ &\Rightarrow \max_j L_j\Theta < \Theta = M, \text{ so } \{M \vee \max_j L_j\Theta\} = M. \end{aligned}$$

So Θ satisfies (A.1), that is $\Theta = M \vee \max_j L_j\Theta$, and the Gittins index policy is optimal. ■

Thus, $\Theta = \Phi$ and the following identity holds:

$$\Phi(M, x) = B - \int_M^B \prod_k \frac{\partial \phi_k(m, x_k)}{\partial m}\, dm. \tag{A.10}$$

Whittle [63, section 9] indicates that the proof can be modified to incorporate variable length project stages.

Assume that when one engages project k in state x_k then one is committed to it for a stage of length $s = s(k, x_k)$. We shall suppose that s and $x_k(t+s)$ are conditioned only by k and x_k, and not by t. The dynamic programming equations become recursions between discrete stages instead of between discrete periods, and we modify the definition of the one-step operator L_k:

$$L_kF(x) = R_k(x_k) + \mathrm{E}\left[\beta^s F(x(t+s)) \mid x(t) = x, k(t) = k\right]$$

where $R_k(x_k)$ is now the reward from the stage starting from state x_k.

The single project return $\phi_k(m, x_k)$ defined in (A.2) is now in terms of the modified L_k, and the identity (A.10) after the end of the proof of the main result still holds between Φ and the ϕ_k; the Gittins index policy is optimal. ∎

A.1.2 Optimality of the Gittins index policy for bandit super-processes

We now show how Whittle's proof (outlined above) of the optimality of the Gittins index policy for simple processes (consisting of single-action projects) can be generalised to cover super-processes (consisting of multi-action projects).

Remember, a super-process is one in which, after a project has been chosen, there is a further decision to be made as to how to proceed, and this affects both the reward and the state transition of the chosen project. The proof of the optimality of the Gittins index policy for super-processes fails except in one special case, which is when the following condition holds: *the optimal subsidiary decision as to how to proceed with the chosen project is independent of the size of the fall-back.* (In other words, if a project is the only one available then your optimal action does not change when the fall-back varies over the range in which you prefer to continue with the project.) The proof below that this condition is sufficient elaborates on that in Whittle

[62]. That this condition is also necessary can be found in Glazebrook [30].

There are N projects, each project having possibly more than one available action when in a given state, and in each discrete period you can take only one action and thus work on only one project. The state of project k at time t is denoted by $x_k(t)$ and the state variable at time t is written as $x(t) = \langle x_1(t), x_2(t), \ldots, x_N(t)\rangle$. The set of available actions for project k in state x_k is denoted by $U_k(x_k)$, and the set of all available actions is the union over k of these, denoted by $U(x)$. Let $\kappa(\cdot)$ be the indicator function mapping available actions to projects, i.e. $\kappa(u) = k$ for $u \in U_k$. The action taken at time t is denoted by $u(t)$, and thus the project engaged at time t is $\kappa(u(t))$. If action u is taken at time t then you get an immediate expected reward of $R_{\kappa(u)}(x_{\kappa(u)}(t), u)$. Rewards are additive and discounted in time by a factor β. States of unengaged projects do not change and the state of the engaged project changes by a Markov transition rule: if $\kappa(u(t)) \neq k$ then $x_k(t+1) = x_k(t)$, and if $\kappa(u(t)) = k$ then the value of $x_k(t+1)$ is conditioned only by $u(t)$, k & $x_k(t)$.

Continue to assume that rewards are uniformly bounded:

$$-\infty < -B(1-\beta) \le R_k(x,u) \le B(1-\beta) < \infty.$$

When m is the available fall-back, $\phi_k(m, x_k)$ now solves

$$\phi_k = m \vee \sup_{u \in U_k} L_{k,u}\phi_k \tag{A.11}$$

where

$$L_{\kappa(u),u}\Phi(M,x) = R_{\kappa(u)}(x_{\kappa(u)}, u) + \beta\,\mathrm{E}\Big[\Phi(M, x(t+1)) \;\Big|\; M, x(t) = x, u(t) = u\Big]$$

As usual, the Gittins index of project k, denoted by $m_k(x_k)$, is the infimal value of m such that $m = \phi_k(m, x_k)$, namely the alternatives of stopping with

m_k and of embarking on project k (with the option of taking the fall-back staying open) are equitable, and so $m_k = \sup_{u \in U_k} L_{k,u}\phi_k$.

$\Theta(M, x)$ is defined as before, and we still have (A.3):

$$\Theta(M, x) = \phi_k(M, x_k)P_k(M, x) + \int_M^\infty \phi_k(m, x_k)\, d_m P_k(m, x)$$

so, having fixed x, the following ((A.4) & (A.5)) still hold:

$$\begin{aligned} \Theta(M) &\geq M \text{ for any } M, \\ \text{and } \Theta(M) &= M \text{ iff } M \geq \max_j m_j. \end{aligned}$$

The function $\delta(\cdot)$ is now action-specific not merely project-specific, so, for $u \in U_k$, define

$$\delta_{k,u}(m, x_k) = \phi_k(m, x_k) - L_{k,u}\phi_k(m, x_k).$$

Fixing x as before, to focus on m or M, note that

$$\Theta(M) - L_{k,u}\Theta(M) = \delta_{k,u}(M)P_k(M) + \int_M^\infty \delta_{k,u}(m)\, d_m P_k(m)$$

$$\begin{aligned} \text{so } \Theta(M) &- \sup_{u \in U_k} L_{k,u}\Theta(M) \\ &= \inf_{u \in U_k} \left(\delta_{k,u}(M)P_k(M) + \int_M^\infty \delta_{k,u}(m)\, d_m P_k(m)\right) \quad \text{(A.12)} \end{aligned}$$

We want to show that:

$$\begin{aligned} \Theta(M) &\geq \sup_{u \in U_k} L_k\Theta(M) \text{ for any } M, && \text{(A.13)} \\ \text{and } \Theta(M) &= \sup_{u \in U_k} L_k\Theta(M) \text{ iff } m_k = \max_j m_j \text{ and } M \leq m_k. && \text{(A.14)} \end{aligned}$$

It is still the case that, for any $u \in U_k$, $\delta_{k,u}(m) \geq 0$ for all m, so inequality (A.13) still holds, and if we are able to assert that, for some $u \in U_k$, $\delta_{k,u}(m) = 0$ for $m \leq m_k$, then equality (A.14) also holds, by considering the RHS of (A.12). The assertion that such an action $u \in U_k$ exists is the same as saying that in the continuation region for the single project the optimal action is unique. ■

However, if there is not a unique optimal action $u \in U_k$ when $m \leq m_k$, then the RHS of (A.12) might be strictly positive for some $M \leq m_k$, in which case equality (A.14) would not hold, and the remainder of the proof would not go through.[2] To see this, suppose that a switch of actions occurs when the fall-back is $\hat{m}$, i.e. when m is such that $m \leq \hat{m}$ it is optimal to take action u', and when m is such that $\hat{m} \leq m \leq m_k$ it is optimal to take action u''. For action u' this implies that $\delta_{k,u'}(m) = 0$ when $m \leq \hat{m}$, & $\delta_{k,u'}(m) > 0$ when $\hat{m} < m \leq m_k$, and for action u'' this implies that $\delta_{k,u''}(m) > 0$ when $m < \hat{m}$, & $\delta_{k,u''}(m) = 0$ when $\hat{m} \leq m \leq m_k$. Consider $M < \hat{m}$, and suppose that the other projects under consideration are such that $P_k(M) > 0$ and $d_m P_k(m) > 0$ for $M \leq m < m_k$. Looking at the RHS of (A.12) for the two actions in turn we see that (a) the first term is zero because $\delta_{k,u'}(M) = 0$, but the integral is non-zero because neither $\delta_{k,u'}(m)$ nor $d_m P_k(m)$ is zero over $[\hat{m}, m_k]$, and (b) $\delta_{k,u''}(M) > 0$ and also the integral is non-zero (over $[M, \hat{m}]$). So the expression in parentheses on the RHS of (A.12) is strictly positive for either action, hence the infimum over the two actions is positive.

As in Appendix A.1.1, when the number of periods required to complete an action in a project is different for different actions and different projects,

[2] As an informal example of the second condition failing, consider a project with two actions: one leads to a state with a low mean value and a high variance; the other one leads to a state with a high mean value and a low variance. *Taking either action renders the other unavailable.* When the fall-back is high enough, it is optimal to take it. When the fall-back is lowered, it becomes optimal to take the more risky action, because if a poor outcome is realised there is always the fall-back. However, as the fall-back is lowered even further, it is no longer a good enough guarantee and so the optimal action switches to the less risky one.

the definition of the action of L_k can be suitably modified so that the above result remains valid.

Assume that when one takes action u in project $k = \kappa(u)$ in state x_k then one is committed to it for a stage of length $s = s(k, x_k, u)$. We shall suppose that s and $x_k(t+s)$ are conditioned only by k, x_k, and u, and not by t. The definition of the one-step operator L_k becomes:

$$L_{\kappa(u),u}F(x) = R_{\kappa(u)}(x_{\kappa(u)}, u) + \mathrm{E}[\beta^s F(x(t+s)) \mid x(t) = x, u(t) = u]$$

where $R_{\kappa(u)}(x_{\kappa(u)}, u)$ is now the reward from the stage starting from state x_k when action u is taken.

As before, the single project return $\phi_k(m, x_k)$ is now defined in terms of the modified L_k, and the identity (A.10) still holds between Φ and the ϕ_k; the Gittins index policy is optimal. ■

A.2 Incremental sunk costs

A.2.1 Proof of Result 9

Suppose first that the Result is true in all possible successor states to k^t. We show that it is then true for k^t. This implies, by backward induction from the state where all firms have completed the project, that the Result is true for all states.

The proof is broken down into a number of intermediate steps.

Step 1: The boundary in period $t+1$ will be either $x^b(k^t)$ or $x^b(k^t) - 1$, no matter what actions firms take in t.

The boundary in $t+1$ cannot be further from the end than $x^b(k^t)$ since the number of firms at $x^b(k^t)$ or closer can only be higher in $t+1$ than in t. Moreover the boundary in $t+1$ is at least as far from the end as $x^b(k^t) - 1$:

the number of firms which could profitably finish were all firms at $x^b(k^t)$ is more than the number of firms at $x^b(k^t) - 1$ or closer in period t, and so is more than the number of firms at $x^b(k^t) - 2$ or closer in period $t + 1$. This implies that the number of firms which could profitably finish were all firms at $x^b(k^t) - 1$ in period $t + 1$ is more than the number of firms at $x^b(k^t) - 2$ or closer, and so that the boundary in period $t + 1$ is at least as far as $x^b(k^t) - 1$.

Step 2: Assume the Result is true in all possible successor states to k^t. Then all firms at $x^b(k^t) - 1$ or closer in period t will invest in t, in any equilibrium.

Recall that the Result is assumed true in all possible successor states to k^t, and that we have just shown that the boundary in period $t + 1$ will be either $x^b(k^t)$ or $x^b(k^t) - 1$. Together these imply that the largest possible number of firms which will finish subsequent to t is the number which could profitably finish were all firms at $x^b(k^t) - 1$. They also imply that all firms at $x^b(k^t) - 2$ or closer in period $t + 1$ will invest fully in each subsequent period until they have completed the project. Any firm at $x^b(k^t) - 1$ or closer in period t can thus guarantee that it will be among those completing the project, and that it will make a profit, by investing in period t so that it is at $x^b(k^t) - 2$ or closer in period $t + 1$. If any of these firms waited instead, they would not reduce the number of firms that finished in total, they might find a rival finished in their stead, and their payoff would be discounted, so that in any equilibrium, all firms at $x^b(k^t) - 1$ or closer will invest in period t.

Step 3:Assume the Result is true in all possible successor states to k^t and suppose that at t some firms are at the boundary $x^b(k^t)$. If the number at the boundary or closer is less than the number that could profitably finish were all firms at the boundary, then all

firms actually at the boundary invest in period t in any equilibrium, otherwise exactly $\max\left\{n \mid u(x^b(k^t), n) \geq 0\right\} - \#\{i \mid k_i^t < x^b(k^t)\}$ **of them do so.**

Suppose instead that fewer than this invest. In period $t+1$ the boundary will remain at $x^b(k^t)$, and since the Result is assumed true in all successor states, at least one firm at the boundary in $t+1$ will invest in each subsequent period and complete the project. If this firm had been at the boundary at t it would have done better to invest in t rather than wait for a period. If not, one of the firms which *was* at the boundary at t and which never invested would have done better to invest at t and so ensure that it was among the finishing firms. Suppose now that the number of firms at the boundary or closer is more than the number which could profitably finish were all firms at the boundary, and that more than $\max\left\{n \mid u(x^b(k^t), n) \geq 0\right\} - \#\{i \mid k_i^t < x^b(k^t)\}$ of the firms actually at the boundary at t invest. Subsequently either all of these firms, and all those which were closer than the boundary at t, will complete the project, and the firms which were at the boundary at t will make a loss. Alternatively at least one of the firms which was at the boundary at t and which invested does not invest further, in which case it cannot have been part of an equilibrium for that firm to invest at t.

Step 4. Assume the Result is true in all possible successor states to k^t. Then the Result is true in k^t.

This follows directly from steps 2 and 3 above.

Step 5. The Result is true in all states.

This follows by backward induction from the state in which all firms have completed the project. In this state, the Result is clearly true. Suppose all firms but one have finished the project, and the exception has just one stages to complete. The Result is true in all possible successor states, and so is true in this one by step 4. Suppose that all firms but one have finished the project, and the exception has two stages to complete. The Result is true in all possible successor states, and so is true in this one, and so on. Suppose all firms but two have finished the project, and the two exceptions have just one stage each to complete. The Result is true in all possible successor states and so is true in this one, and so on. In this way the Result is proved in all possible states. ■

A.2.2 Proof of Result 11

The proof that these strategies form the only symmetric perfect Nash equilibrium can be taken from the proof of Result 9 given in Appendix A.2.1 for all states except those where firms have the same number of stages to complete but if both invest fully to the end they would each make a loss, i.e. those states where $k_i = k_j$ and $u^d(k_i) < 0$ in Appendix A.2.1. We now show that in these states the only symmetric perfect Nash equilibrium has firms choosing mixed strategies.

Define $\overline{x}^d$ as the greatest number of stages a firm can complete and be sure of a positive payoff, no matter what the rival does: $\overline{x}^d = \max\left\{x \mid u^d(x) \geq 0\right\}$. Start with $k_i = k_j = \overline{x}^d + 1$. In the pure strategy equilibrium just one firm would invest but this cannot be so in a symmetric equilibrium. If one firm invests while its rival does not its value in the state next period will be $u^m(\overline{x}^d)$ and so its value this period is $u^d(\overline{x}^d + 1)$. If both invest both get $u^d(\overline{x}^d + 1)$. If a firm does not invest its value will be zero whether its rival invests or not (as strategies are time stationary). In a mixed strategy equilibrium the firm will be indifferent between investing and not. If its rival invests with probability $s_j^*(\overline{x}^d+1, \overline{x}^d+1)$ then:

$$0 = s_j^*(x,x)\, u^d(x) + \left(1 - s_j^*(x,x)\right)\, u^m(x), \quad x = \overline{x}^d + 1$$

which rearranges to give:

$$s_j^*(x,x) = \frac{1}{1 - u^d(x)/u^m(x)}, \quad x = \overline{x}^d + 1$$

as given in Result 11. The value of each firm in $(\overline{x}^d+1, \overline{x}^d+1)$ is clearly 0.

When $k_i = k_j > \overline{x}^d + 1$ then if one invests while the other does not it gets $u^m(k_i)$ while if both invest each gets $-c$. This gives:

$$s_i^*(x,x) = \frac{1}{1 + c/u^m(x)}, \quad x = d, d-1, \ldots \overline{x}^d + 1$$

and the value is again 0. ■

A.2.3 Proof of Result 12

In order to find the probability that both firms complete the project, we first examine the conditions under which both finish. According to the equilibrium in Result 11 there is a bound $\overline{x}^d$ such that if both firms have $\overline{x}^d$ or fewer stages to complete both will certainly finish the project. If at least one has more than $\overline{x}^d$ stages to complete either both have the same number and choose mixed strategies, or one has fewer stages and it alone finishes. Therefore both firms complete the project if and only if they complete each of the symmetric states outside the boundary, (x,x), $x = d, d-1, \ldots, \overline{x}^d + 1$, at the same time as each other. As is clear from the equilibrium strategies the boundary $\overline{x}^d$ is defined by the condition that a firm completing $\overline{x}^d$ or fewer stages and then producing has a positive payoff even if its rival finishes first, but a negative payoff if it has to complete more than $\overline{x}^d$ stages and then share the market with its rival. In other words $\overline{x}^d$ is the largest x for which:

$$u^d(x) = \frac{\delta^x}{1-\delta}\pi(2) - \frac{1-\delta^x}{1-\delta}c \geq 0$$

or in terms of the parameters of the sunk cost model of concentration, it is the largest x for which:

$$u^d(x) = \hat{\delta}^{x/d}\Pi(2) - \frac{1-\hat{\delta}^{x/d}}{1-\hat{\delta}}\sigma \geq 0 \tag{A.15}$$

We have seen that both firms complete the project if and only if each completes every stage x, $x = d, d-1, \ldots, \overline{x}^d + 1$, at the same time as its rival. Denote the probability that both firms complete stage x together given that the market has arrived at state (x, x) by $p(x)$. Both choose probability $s_i^*(x, x)$. If both complete the stage at the same time this is either because both invested in the first period in which the state was (x, x), or because neither finished in this period and both finished together at a later time. Thus:

$$p(x) = (s_i^*(x,x)) + (1 - s_i^*(x,x))^2 p(x)$$

which rearranges to give:

$$p(x) = \frac{1}{2/s_i^*(x,x) - 1} \tag{A.16}$$

Inspecting this expression, and that in Result [ ], giving the optimal strategies s^*, confirms that $p(x)$ is non-increasing in x: a firm is less likely to complete a stage at the same time as its rival when it has more stages left to complete.

The probability that both firms finish the project is denoted ϕ and is the product of the probability that they complete each of the stages (x, x), $x = d, d-1, \ldots, \overline{x}^d + 1$ together.

$$\phi = \prod_{x=d}^{\overline{x}^d+1} p(x)$$

We can place a lower bound on this product by considering just the smallest possible value of $p(x)$, i.e. $p(d)$. Denote this lower bound by:

$$\underline{\phi} = [p(d)]^{\overline{x}^d+2-d} \leq \phi$$

Substituting for $p(x)$ using Equation A.16 and the strategies given in Result 11 gives:

$$\underline{\phi} = \left[\frac{1}{1 + \frac{2c}{u^m(d)}}\right]^{\overline{x}^d + 2 - d}$$

and we will consider the limit of this expression in the limit as $d \to \infty$.

Two parts of the expression for $\underline{\phi}$ above depend on d: the power of the expression, $\overline{x}^d + 2 - d$, and the cost c in the bracket (from the definition of the parameters in Equation 2.5 we have $u^m(d) = \hat{\delta}\Pi(1) - \sigma$, no matter d). We consider each in turn.

Consider first the term $\overline{x}^d + 2 - d$. The boundary $\overline{x}^d$ is the furthest *integer* distance from completion, x, satisfying Equation A.15. Define ξ as the *real* value satisfying:

$$u^d(\xi d) = \hat{\delta}^{\xi}\Pi(2) - \frac{1 - \hat{\delta}^{\xi}}{1 - \hat{\delta}}\sigma = 0$$

Note that ξ is independent of d. In the limit as $d \to \infty$ there is always a distance x such that x/d is arbitrarily close to any real value, so that:

$$\lim_{d \to \infty} \bar{x}^d = d\xi$$

Turning to the cost c we have, from Equation 2.5:

$$c = \frac{1 - \hat{\delta}^{1/d}}{1 - \hat{\delta}}\sigma$$

We can find the limit of this as d becomes large by taking a binomial expansion of $\hat{\delta}^{1/d} = \left(1 - (1 - \hat{\delta})\right)^{1/d}$ and ignoring all high powered terms in $1/d$. This gives:

$$\begin{aligned}
\lim_{d \to \infty} \left(1 - (1 - \hat{\delta})\right)^{1/d} &= 1 - \frac{1}{d}(1 - \hat{\delta}) + \frac{(1/d)\,(1/d - 1)}{2!}(1 - \hat{\delta})^2 \\
&\quad - \frac{(1/d)\,(1/d - 1)\,(1/d - 2)}{3!}(1 - \hat{\delta})^3 - \dots \\
&= 1 - \frac{1}{d}(1 - \hat{\delta}) + \frac{1}{d}k
\end{aligned}$$

where:

$$k = -\frac{(1-\hat{\delta})^2}{2} + \frac{(1-\hat{\delta})^3}{3} - \cdots$$

which converges. Thus:

$$\lim_{d\to\infty} c = \frac{\sigma}{d}\left(1 - \frac{k}{1-\hat{\delta}}\right) \tag{A.17}$$

We are now able to consider the limit of the lower bound to the probability that both firms complete the project as the number of stages becomes large.

$$\begin{aligned}\lim_{d\to\infty} \underline{\phi} &= \lim_{d\to\infty}\left[\frac{1}{1+\frac{2}{u^m(d)}\left(1-\frac{k}{1-\hat{\delta}}\right)\frac{\sigma(\xi-1)}{d(\xi-1)}}\right]^{d(\xi-1)} \\ &= \exp\left(-\frac{2\sigma(\xi-1)}{u^m(d)}\left(1-\frac{k}{1-\hat{\delta}}\right)\right) > 0\end{aligned}$$

■

A.2.4 Proof of Result 13

The state space is sufficiently small that we can construct equilibrium backwards state by state starting at $(0,0,0)$ in which no firms can take any action. In $(0,0,1)$ a single firm can choose to invest and get $u(1,3)$. Since this is positive investment is optimal. In $(0,1,1)$ two firms can choose to invest. Both will finish since if either finishes it gets a positive payoff no matter what its rival does. There is no incentive to wait and both invest straightaway. In $(0,1,2)$ the firm with 1 stage to complete has a dominant strategy: since $u(1,3) \geq 0$ it will invest no matter what the other firm does. Given this the firm with 2 stages to complete does not invest since the best it can hope for is a payoff of $u(2,3)$ which is negative. In $(0,2,2)$ there is no symmetric pure strategy equilibrium: it is profitable for one firm to finish but not both. If one invests while its rival does not the state next period will be $(0,1,2)$, it will invest again, and so its payoff is $u(2,2)$. If both invest both get $u(2,3)$. If a firm does not invest its value will be zero: it will never arrive at a state in which it invests. Each firm will be indifferent between investing and not when its rival invests with probability q_1 and if it invests it also gets an expected payoff of zero:

$$0 = q_1\, u(2,3) + (1 - q_1)\, u(2,2)$$

which rearranges to give:

$$q_1 = \frac{u(2,2)}{u(2,2) - u(2,3)}$$

Both will invest with this probability in any symmetric equilibrium. Strategies in the other states follow in a similar way, except $k = (2,2,2)$.

It is straightforward to show that in $(2,2,2)$ in any symmetric equilibrium all invest with a probability q_2 which is strictly less than 1. They must choose the same probability since strategies are symmetric by assumption. Moreover they cannot all invest with certainty since if they do the state next period will be $(1,1,1)$, all will invest again, giving an initial value $u(2,3)$ which is negative. The difficulty is in showing that firms invest with a strictly positive probability. As we have seen, when two firms invest the third gets a negative payoff if it invests as well. It is a sufficient condition for $q_2 > 0$ that when two firms do not invest the third has a positive payoff if it does. This sufficiency follows simply from the facts that (a) any interior probability is optimal if and only if the payoff from investing is zero, (b) a firm's payoff if it invests is a continuous function of its rivals' investment probabilities.

Next we confirm that if two firms do not invest the third does indeed get a positive payoff from investing. If one of the other two invests the third firm will certainly get a positive payoff from investing itself: the state next period will be $(1,1,2)$, the two lead firms alone will invest and so its payoff is $u(2,2) \geq 0$. However, a firm with one stage to go may prefer that one rival also has just one stage to go that both still had two stages to complete. This can arise because if two still have two stages to go they will choose mixed strategies and both may end up completing the project, reducing the profits earned by the third. We need an explicit expression for the value of a firm when it has one stage to go and the other two have two, i.e. for $V_1(1,2,2)$. Suppose that initially the state is $(1,2,2)$. The first firm invests with certainty and the two rivals invest with probability q_1. If neither rival completes the first stage before t and at t just one completes the stage, the first firms payoff, evaluated initially, is:

$$u(1,1) - \delta^t \Pr(1 \text{ firm completes first stage at } t)\,(u(1,1) - u(1,2))$$

and similarly if two firms complete the first stage at t. The first firm's initial value is thus:

$$\begin{aligned} V_1(1,2,2) &= u(1,1) \\ &\quad -\sum_{t=0}^{\infty} \delta^t \Pr(1 \text{ firm completes the first stage at } t)\,(u(1,1) - u(1,2)) \\ &\quad -\sum_{t=0}^{\infty} \delta^t \Pr(2 \text{ firms complete the first stage at } t)\,(u(1,1) - u(1,3)) \end{aligned}$$

The probability that neither rival completes the first stage before t and at t just one completes the stage is:

$$(1-q_1)^{2t-1}\,q_1$$

and similarly for the probability that both rivals complete the first stage at t. Substituting these and rearranging gives:

$$V_1(1,2,2) = u(1,1) - \frac{q_1}{1-\delta(1-q_1)^2}\left((u(1,1) - u(1,2)) + q_1\,(u(1,2) - u(1,3))\right)$$

Note first that this can be less than $V_1(1,1,2) = u(1,2)$. The second term is always positive and so $V_1(1,2,2)$ is less than $u(1,1)$. If the first firm to finish does not mind when a second firm finishes because $u(x,1) = u(x,2)$ then it will also be less than $V_1(1,1,2)$. For our purposes, this means that we need to show explicitly that the third firm will invest in $(2,2,2)$ if the other two do not as we cannot rely on the fact that $u(2,2) = \delta V_1(1,1,2) - c \geq 0$ to argue that $\delta V_1(1,2,2) - c$ is likewise positive. Writing out the payoff from investing when the other two do not gives:

$$\begin{aligned}
\delta V_1(1,2,2) - c &= \delta u(1,1) - c \\
&\quad - \frac{\delta q_1}{1 - \delta(1-q_1)^2}\left((u(1,1) - u(1,2)) + q_1\left(u(1,2) - u(1,3)\right)\right) \\
&= u(2,1) \\
&\quad - \frac{q_1}{1 - \delta(1-q_1)^2}\left((u(2,1) - u(2,2)) + q_1\left(u(2,2) - u(2,3)\right)\right)
\end{aligned}$$

Note that the expression for q_1 found earlier gives $q_1\left(u(2,2) - u(2,3)\right) = u(2,1)$ so that:

$$\delta V_1(1,2,2) - c = u(2,1)\left(1 - \frac{q_1}{1 - \delta(1-q_1)^2}\right)$$

This is positive when $q_1 < 1 - \delta(1-q_1)^2$, i.e. when $\delta(1-q_1) < 1$, which is always the case.

This concludes the proof. ∎

Bibliography

[1] P Aghion, P Bolton, C Harris, and Jullien B. Optimal learning by experimentation. *Review of Economic Studies*, (58):621–654, 1991.

[2] T Barker and M Robbins. *A History of London Transport*, volume 1. 1974.

[3] A J Barton and J T Everest. Express coaches in the three years following the 1980 Transport Act. Technical Report 1127, Transport and Road Research Laboratory, 1984.

[4] J Beath, Y Katsoulacos, and D Ulph. Sequential product innovation and industry evolution. *Economic Journal*, 97:32–43, 1987. Supplement.

[5] D Berry and B Fristedt. *Bandit Problems: Sequential Allocation of Experiments.* Chapman & Hall, London, 1985. Chapter 6.

[6] C Budd, C Harris, and J Vickers. A model of the evolution of duopoly: Does the asymmetry between firms tend to increase or decrease? *Review of Economic Studies*, 60:543–574, 1993.

[7] D N Chester. *Public Control of Road Passenger Transport.* Manchester, 1936.

[8] Commons Select Committee on Transport. *Local Bus Deregulation.* HMSO, 1995.

[9] P Dasgupta and J.E. Stiglitz. Industrial structure and the nature of innovative activity. *Economic Journal*, pages 266–293, 1980.

[10] C D'Aspremement, J J Gabszewicz, and J F Thisse. On Hotelling's 'stability in competition'. *Econometrica*, 47:1145–1170, 1979.

[11] E Davis. Express coaching since 1980: Liberalisation in practice. *Fiscal Studies*, 5(1):76–86, 1984.

[12] S Davis and J Haltiwanger. Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics*, 107:819–864, 1992.

[13] A Dixit. Entry and exit decisions under uncertainty. *Journal of Political Economy*, 97(3):620–638, 1989.

[14] J Dodgson and Y Katsoulacos. Quality competition in bus services: Some welfare implications of bus deregulation. *Journal of Transport Economics and Policy*, 22:263–281, September 1988.

[15] J Dodgson, Y Katsoulacos, and C Newton. An application of the economic modelling approach to the investigation of predation. *Journal of Transport Economics and Policy*, 27:153–170, 1993.

[16] J Dodgson, C Newton, and Y Katsoulacos. A modelling framework for the empirical analysis of predatory behaviour in the bus services industry. *Regional Science and Urban Economics*, 22:51–70, 1992.

[17] T Dunne, M Roberts, and L Samuelson. Dynamic patterns of firm entry, exit, and growth. *Rand Journal of Economics*, 19:495–515, 1988.

[18] R Ericson and A Pakes. Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, pages 53–82, 1995.

[19] A Evans. A theoretical comparison of competition and other economic regimes for bus services. *Journal of Transport Economics and Policy*, 21(1):7–36, January 1987.

[20] D Evans. The relationship between firm size, growth and age: Estimates for 100 manufacturing industries. *Journal of Industrial Economics*, 35:567–581, 1987.

[21] D Evans. Tests of alternative theories of firm growth. *Journal of Political Economy*, 94:657–674, 1987.

[22] C Foster and J Golay. Some curious old practices and their relevance to equilibrium in bus competition. *Journal of Transport Economics and Policy*, 20:191–216, 1986.

[23] D Fudenberg, Gilbert, J Stiglitz, and J Tirole. Preemption, leapfrogging, and competition in patent races. *European Economic Review*, 22:3–31, 1983.

[24] D Fudenberg and J Tirole. *Game Theory.* MIT Press, Cambridge MA., 1991.

[25] T Furusawa. Trial and error in R&D. mimeo University of Wisconsin-Madison, 1994.

[26] P Geroski. Some data-driven reflections on the entry process. In P Geroski and J Schwalbach, editors, *Entry and Market Contestability.* Basil Blackwell, 1991.

[27] J Gittins. Bandit processes and dynamic allocation indices. *Journal of the Royal Statistical Society B*, (41):148–164, 1979.

[28] J Gittins. *Multi-armed Bandit Allocation Indices.* Wiley, Chichester, 1989. Chapter 3.

[29] J Gittins and D Jones. A dynamic allocation index for the sequential design of experiments. In J Gani, editor, *Progress in Statistics*, chapter A Dynamic Allocation Index for the Sequential Design of Experiments, pages 241–266. North-Holland, Amsterdam, 1974. (J Gani ed.).

[30] K Glazebrook. On a sufficient condition for superprocesses due to whittle. *Journal of Applied Probability*, (19):99–110, 1982.

[31] M Gort and S Klepper. Time paths in the diffusion of product innovations. *Economic Journal*, 92:630–653, 1982.

[32] G. R. Grimmett and D.R. Stirzaker. *Probability and Random Processes.* Oxford University Press, Oxford, 1982.

[33] C Harris and J Vickers. Perfect equilibrium in a model of a race. *Review of Economic Studies*, 52:193–209, 1985.

[34] O Hart and B Holmstrom. The theory of contracts. In T Bewley, editor, *Advances in Economic Theory: Fifth World Congress.* Cambridge University Press, Cambridge, 1987.

[35] H Hopenhayn. Entry, exit and firm dynamics in long run equilibrium. *Econometrica*, 60(5):1127–1150, 1992.

[36] S M Jaffer and D J Thompson. Deregulating express coaches: A reassessment. *Fiscal Studies*, 7(4):45–68, November 1986.

[37] S Klepper. Entry, exit, growth and innovation over the product life cycle. *American Economic Review*, 86(3):562–583, 1996.

[38] S Klepper and E Grady. The evolution of new industries and the determinants of market structure. *Rand Journal of Economics*, 21:27–44, 1990.

[39] R Koopmans. *Assymmetric Industry Structures: Multiple Technologies, Firm Dynamics and Profitability.* PhD thesis, London School of Economics, 1995.

[40] V Lambson. Industry evolution with sunk costs and uncertain market conditions. *International Journal of Industrial Organisation*, 9(2):171–196, 1991.

[41] P Mackie and J Preston. *The Local Bus Market: A Case Study of Regulatory Change.* Avebury, 1997.

[42] A McLennan. Price dispersion and incomplete learning in the long run. *Journal of Economic Dynamics and Control*, pages 331–347, 1984.

[43] P Nash. *Optimal Allocation of Resources between Research Projects.* PhD thesis, Cambridge University, 1974.

[44] W Novshek. Equilibrium in simple spatial (or differentiated) product models. *Journal of Theoretical Economics*, 22, 1980.

[45] A Pakes and R Ericson. Empirical implications of alternative models of firm dynamics. mimeo, Yale University, 1990.

[46] A Pakes, G Gowrisankaran, and P McGuire. Code for implementing the pakes-mcguire algorithm for computing markov perfect equilibrium. Mimeo, Yale University, 1993.

[47] A Pakes and P McGuire. Computing Markov-perfect Nash equilbiria: Numerical implications of a dynamic differentiated product market. *Rand Journal of Economics*, 25(4):555–589, 1994.

[48] D Quah. Empirics for economic growth and convergence. Discussion Paper EM/95/281, STICERD, London School of Economics, February 1995.

[49] D Robbins and P White. The experience of express coach deregulation in Great Britain. *Transportation*, 13(4), 1986.

[50] K Roberts and M Weitzman. Funding criteria for research, development, and exploration projects. *Econometrica*, (49):1261–1288, 1981.

[51] M Rothschild. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, (9):185–202, 1974.

[52] I Savage. *The Deregulation of Bus Services.* Gower, Aldershot, 1985.

[53] A Shaked and J Sutton. Relaxing price competition through product differentiation. *Review of Economic Studies*, 49:3–13, 1982.

[54] J Sutton. *Sunk Costs and Market Structure.* MIT Press, Cambridge, MA, 1991.

[55] J Sutton. Gibrat's legacy. *Journal of Economic Literature*, 1996. forthcoming.

[56] J Sutton. *Technology and Market Structure.* MIT, 1997. forthcoming.

[57] D Thompson and A Whitfield. Express coaching: Privatisation, incumbent advantage and the competitive process. Mimeo, 1990.

[58] J Tirole. *The Theory of Industrial Organisation.* MIT Press, 1988.

[59] F Vega-Redondo. Industrial dynamics, path dependence, and technological change. mimeo Universidad de Alicante, 1993.

[60] J Vickers. The evolution of market structure when there is a sequence of innovations. *Journal of Industrial Economics*, 35:1–12, 1986.

[61] M Weitzman. Optimal search for the best alternative. *Econometrica*, (47):641–654, 1979.

[62] P Whittle. Multi-armed bandits and the gittins index. *Journal of the Royal Statistical Society B*, (42):143–149, 1980.

[63] P Whittle. *Optimization over Time: Dynamic Programming and Stochastic Control Vol. I.* Wiley, New York, 1982. Chapter 14.