UNIVERSITY OF SOUTHAMPTON

FACULTY OF LAW, ARTS, AND SOCIAL SCIENCES
School of Social Sciences

ESSAYS ON NEW INTERNET MARKETPLACES

GREG TAYLOR

Thesis for the degree of Doctor of Philosophy

September 2009

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF LAW, ARTS, AND SOCIAL SCIENCES

SCHOOL OF SOCIAL SCIENCES

Doctor of Philosophy

ESSAYS ON NEW INTERNET MARKETPLACES

by Greg Taylor

This thesis uses the techniques of economic theory to examine the behaviour of agents in new marketplaces that have formed on the Internet. It is divided into three chapters, each focusing upon a different aspect of market behaviour. I begin in chapter 2 by building a simultaneous, unit demand, heterogeneous good, ascending auction model, which I use to study the behaviour of bidders in Internet auctions. Using this model, I demonstrate that late bidding—often observed in eBay auctions—is not inconsistent with allocative efficiency. Moreover, I show that when all agents are fully rational, have private values, and face a perfect bid-transmission mechanism, there can still exit an incentive for bidders to systematically delay their final bid. In a manner consistent with earlier empirical observations, this incentive disappears when the hard-close ending rule is relaxed.

In chapter 3 I model strategic interaction amongst search engines that compete to serve consumer needs. Search engines generate revenue from advertisers, but also provide free organic search results. I demonstrate that, in an attempt to win market share, search engines compete not only against each other, but also against themselves: providing high-quality free links that compete for clicks with their own advertisements—thus cannibalising their advertising revenues. In particular, I find that in equilibrium consumers always (at least weakly) prefer to click on at least one non-paid-for link before clicking on a revenue generating advertisement so that some consumers never click an ad at all. That notwithstanding, revenue cannibalisation provides an incentive for quality degradation, even when the provision of quality is costless, and may engender low quality equilibria. When search engines show differentiated advertisements, the incentive to reduce quality is particularly strong.

Chapter 4 examines the relationship between the transmittability of information via advertisements and the fee structure used by the advertisement's publisher. For an advertiser, sending general advertisements with inflationary claims may attract additional consumers with whom it is poorly matched. This is costly for the firm when it must pay for the ads on a per-click basis (i.e. when it must pay for each consumer visit that it receives) since many of its visitors will not purchase. As a consequence, I find that perfect information transmission can always be sustained when adverts are priced per-click. By contrast, when firms pay for advertisements on a per-impression basis or on a per-sale basis, there is no disincentive to attracting poorly matched consumers, and maximum profits are obtained by attracting all consumers with some positive probability of purchase. This feature undermines the existence of fully informative equilibria under such fee structures, and may result in no information transmission being possible at all. Consumers benefit from increased informativeness, but distortions introduced by the market power given to advertisers imply that society may be better-off with no information transmission taking place.

Contents

1 Introduction **1**

2 Efficient Late Bidding in 'Hard-Close' Auctions **3**

I Introduction 3

I.1 Literature 4

II The Model 6

II.1 Base model: an iterative ascending auction. 6

II.2 Extendible end time auction 8

II.3 Hard close: a model of eBay proxy auctions 13

III Stochastic Entry and Arrival 19

III.1 Stochastic arrival in a hard close auction 19

III.2 Stochastic arrival in an extendible end time auction 21

IV Conclusion 22

A Omitted Proofs 23

3 Search Quality and Revenue Cannibalisation by Competing Search Engines **25**

I Introduction 25

I.1 Literature 27

II Simple Model 28

III Equilibrium Behaviour in the Simple Model 29

III.1 Optimal consumer behaviour 29

III.2 Equilibrium characterisation 32

IV Extension: Differentiated A-links 36

IV.1 Optimal consumer behaviour with differentiated A-links . . . 36

IV.2 Equilibrium analysis 37

IV.3 $p_g = p_y = q$ equilibria 39

IV.4 $p_g = p_y = p^{max}$ equilibria 40

IV.5 Non-existence of $p_g = p_y = \bar{\gamma}$ equilibria in the differentiated A-links game 42

IV.6 Comments 43

V Extension: market structure 44

V.1 n-search engine oligopoly 44

V.2 Monopoly search provider 47

V.3 Comment 48

VI Extension: Heterogeneous Visit Costs 48

VII Conclusion 55

A Omitted Proofs 56

A.1 Proofs from Section III 56

A.2 Proofs from Section IV 58

A.3 Proofs from Section V 58

A.4 Proofs from Section VI 59

4 The Informativeness of On-Line Advertisements 67

I Introduction 67

I.1 Literature 69

II Model 70

III Preliminaries 72

IV A Fully-Informative Equilibrium Always Exists Under Pay-Per-Click 73

V Pay-Per-Impression and Pay-Per-Sale 77

V.1 Inadmissibility of informative equilibria under PPI and PPS . 77

V.2 Partial pooling can sometimes be supported under PPI 78

V.3 Bidding and firm selection under PPS 83

V.4 Partial pooling under PPS 84

VI Comments on Welfare 88

VII Conclusion 93

A Omitted Equilibria 94

A.1 Partial pooling in a PPC environment 94

A.2 Pooling (babbling) in a PPC environment 101

A.3 Babbling in a PPI environment 103

A.4 Babbling in a PPS environment with sale-weighting 104

B Publisher Profits Under PPS Versus PPC and PPI 105

Bibliography **108**

List of Figures

2.1 Bid progression in the $dt \to 0$ case. 11

3.1 Organic (O) and Advertising (A) search results. 26

3.2 Minimum value of q consistent with a $p_y = p_g = q$ equilibrium for various combinations of λ and α. 34

3.3 Expected utility maximising consumer behaviour for the 'differentiated A-links' case. 38

3.4 Conditions for existence of a $p_g = p_y = q$ equilibrium in the differentiated A-links game. 41

3.5 Two classes of equilibria in the heterogeneous visit costs game. . . . 52

3.6 Reaction functions and equilibrium points with $\lambda = 1$, for various values of q. 54

3.7 The q in the range $0 \le q \le 1$ for which deviation yields profits equal to those from compliance are given by the roots of the cubic $8q^3-25q^2+14q-1$. 62

3.8 For $q > {}^1/_3\left(3-\sqrt{6}\right) \approx 0.1835$, (3.32) demands a value of $p_g < q$. 63

3.9 The q in the range $0 \le q \le 1$ for which $\pi_y(p_g^*(q), p_y^*(p_g^*(q)))-\pi_y(p_g^*(q), q)$ is negative are given by roots of a polynomial. 66

4.1 Consumer surplus under various equilibria for F uniform on $[0, 1/8]$, with $\alpha = 0.1$, and $n = 2$. The rank ordering of these equilibria is invariant to α and n. 90

4.2 Social welfare under various equilibria for F inverse quadratic on $[0, 1/8]$. Legend is as in Figure 4.1. 91

4.3 Publisher profit under various equilibria for F inverse quadratic on $[0, 1/8]$. Legend is as in Figure 4.1. 92

List of Tables

4.1 Welfare from PPC/PPI equilibria with large n, large α or small γ. . . . 88

4.2 Welfare from PPS equilibria with large α or small γ. 88

Declaration of Authorship

I, Greg Taylor, declare that the thesis entitled “Essays on New Internet Marketplaces” and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based upon work done by myself jointly with others, I have made clear eexactly what was done by others and what I have contributed myself;
- none of this work has been published before submission.

Signed: ....

Date: ....

For my mother, Rosemary, who set me on this path,
and for Andrea, who travels it with me.

Acknowledgements

A great many people have made contributions to the development of the project embodied by this thesis. Robin Mason has approached his task as my PhD supervisor with a characteristic degree of professionalism and commitment. The quality of Robin's input has been invaluable at every stage of the research process: question identification, model formulation, and execution. Moreover, Robin has provided guidance in my professional development that I am confident will stand me in good stead over the years to come.

Robin is not alone in having donated his academic wisdom. Juuso Välimäki's quickness of understanding and sound judgement have helped in guiding the overall direction of portions of this thesis. Along with Robin, the members of my PhD upgrade panel—Maximilian Kwiek and Jean-Yves Pitarakis—provided valuable feedback which has helped to refine the models that I present here as chapters 2 and 3. David Gill was also kind enough to offer his time to discuss chapter 3 of this thesis. Individuals from outside of the University of Southampton have also been most generous. In particular, I wish to highlight Justin P. Johnson, who carefully read and commented upon an earlier version of chapter 3.

I have presented the work contained herein at two conferences: the 63rd European Meeting of the Econometric Society, University Bocconi, Milan, Italy (chapter 2), and the Fifth bi-annual Conference on The Economics of the Software and Internet Industries, Institut D'Économie Industrielle, Toulouse, France (chapter 3). In the latter instance, I benefited from the kind input of my discussant, Markus Reisinger, as well as stimulating discussions with Alex White. In addition, I have given several internal seminars at the University of Southampton (based on chapters 2 and 3). I am grateful to participants in all of these events for their valuable feedback.

The PhD has been a personal journey too, and one that was made infinitely more tolerable by my fellow students and peers. In particular, Maryam Alriyami, Nayef Alsadoun, Reza Boostani, Jurica Brajkovic, Greg Bulczak, David Clifford, Juan Correa, Corrado Giulietti, Dimitrios Gkountanis, Jana Javornik-Skrbinsek, Priya Khambhaita, Tom King, Vicki Lambeth, Gloria Langat, Xinyi Li, Yu Liu, Alessandro Mennuni, Siti Mohd Daud, Chris Mukiza, Ronaldo Nazare, James Robards, Derya Tas, Xing Wang, Lorraine Waller, Clare Woodford, and Yun Zhu have been invaluable and inspirational companions. I am also acutely aware that a large team of support staff work behind the scenes to ensure that the research process runs smoothly. I am particularly indebted in this regard to Gill Crumplin, Chris Thorn, Marion Deicke, Fatima Fauzel, and Sandy McKinnon.

This thesis was examined by viva voce by Maximilian Kwiek and David Myatt. I am most grateful for this stimulating discussion, which has provided me with ideas that I will follow optimistically into my future life as an academic practitioner.

Financial support for this project was provided by the ESRC under their +3 post-graduate funding scheme. I am grateful that funding is available to ensure that enthusiastic scholars are able to pursue their interests, ideas, and aspirations regardless of financial background.

Chapter 1

Introduction

It seems unnecessary to discuss at length the social importance of the Internet—so absolutely pervasive is its influence on modern society. The Internet serves as a platform for the mediation of interactions and transactions between a multitude of agents of all kinds under a wide variety of institutional frameworks. In this thesis, I apply traditional methods of economic analysis to explore some of the novel behaviours emerging at the on-line frontier. Underlying this study is the postulate that, although new technologies have created new environments and enabled new possibilities for interaction, the basic behavioural forces that apply off-line should continue to govern behaviour in the networked age. Thus, behaviours that might seem surprising can be shown to be consistent with ordinary rational behaviour and, conversely, rationality applied in new contexts can give rise to novel behaviours. Amongst the wonders of the internet are often counted its effects on reducing transaction and search costs, and increasing overall efficiency. As in traditional markets, however, there often exist agents with an incentive to subvert or pervert the features of the on-line market environment to their own advantage, and this will be a recurring theme of the chapters to follow.

This study is divided across three chapters, each focusing on a different aspect of behaviour and institutional environment. I begin, in chapter 2 with the observed tendency for bidders in auctions with fixed ending times (such as those that have come to promenance on auction website eBay) to delay bidding until close to the auction's end. I show that this behaviour is consistent with rational bidders that (i) wish to coordinate in assigning winners to items and (ii) face the constant possibility of new bidders entering the market. Unlike previous work, the late bidding equilibrium that I identify results in an efficient allocation amongst the bidders, suggesting that delaying one's bid need not be socially harmful.

In chapters 3 and 4, I examine the effects that search costs can have on behaviour in on-line environments. I begin in chapter 3 by showing that search engines have an incentive to cannibalise their revenue from advertisement links by providing high quality free links that compete directly with them. This cannibalisation effect can provide an incentive for search engines to limit the quality of their search results, even when provision of that quality is costless. In chapter 4 I retain the theme of on-line advertising, and show that the fee structure used by an advertisement's publisher can have an important impact on the transmittability of information via those ads. Whilst consumers benefit from the informative advertisements made possible by a pay-per-click fee structure, an increase in informativeness can be harmful for society overall.

Chapter 2

Efficient Late Bidding in 'Hard-Close' Auctions

I Introduction

In this chapter I develop a heterogeneous good, unit demand, simultaneous auction model[1] with a view to comparing hard close and extendible end time auctions.[2] There has been considerable interest in the persistent tendency of bidders to submit bids close to the scheduled ending times of 'hard close' auctions such as those hosted by Internet auction site eBay. This chapter demonstrates that fully rational bidders facing similarly rational rivals and a perfect bid-transmission mechanism may have an incentive to systematically delay their final bid. Moreover, in contrast to previous work, I show that all bidders delaying their bidding in this fashion can be consistent with the attainment of allocative efficiency.

Bidders in a heterogeneous good auction have an incentive to coordinate in determining who should win each item so as to avoid unnecessary competition. This drives early, multiple, and cross-item bidding (all empirical features of eBay auctions) as agents attempt to signal their preferences to one another in what amounts to a matching game. I make use of a continuous time bidding algorithm (or strategy)—similar to that proposed by Demange, Gale, and Sotomayor (1986)—along with results from Leonard (1983) and Gul and Stacchetti (2000)

[1]Unit demand simultaneous auctions are sometimes referred to in the literature as 'assignment auctions'.

[2]A hard close rule has an auction end at some pre-specified time regardless of bidding activity, whilst an extendible end time auction is one whose duration is continuously extended until a period of pre-determined length elapses in which no new bids have arrived.

to establish an equilibrium in this setting when the end time is flexible. The imposition of a hard close rule changes the strategic dynamics of the model. A player can now outbid his rivals at the last possible minute and, in doing so, avoid a bidding war.[3] The rigidity of the auction's ending time ensures that those who have been outbid in this fashion have no opportunity to respond; instead, they seek to defend themselves against this type of strategic manipulation by submitting sincere valuation bids for the items assigned to them by the matching procedure.

When rival entry and arrival time is stochastic, there is a chance that the bidding activities of a new arrival will change each player's surplus-maximising item by changing the prices of the items for sale relative to one-another. Delaying valuation bidding until no further arrival is possible therefore has an option value to the players in a hard close auction, and stochastic arrival of bidders creates an incentive for systematic late bidding in such environments. Equally, since it is never possible to bid in the final period of an extendible end time auction, there exists a disincentive to heavy commitment in such environments. This is especially true close to the scheduled ending time since bids in this period trigger extension of the auction and thus provide an opportunity for new rivals to arrive. These findings are consistent with, for example, Roth and Ockenfels (2002) who find that late bidding is pervasive in eBay auctions, but not those extendible end time auctions taking place on Amazon.[4]

The remainder of this chapter is organised as follows: In the next subsection, I review some existing results from the literature. Section II.1 introduces the basic structure of the model, section II.2 provides a variant of the kind of bidding algorithm first proposed by Demange, Gale, and Sotomayor (1986), and interprets this in terms of bidding in an extendible end time auction. In section II.3, I impose a fixed end time and discuss the implications that this has for equilibrium behaviour and in section III I discuss the effects of permitting agents to arrive at the auction stochastically. Section IV concludes.

I.1 Literature

Wilcox (2000), Roth and Ockenfels (2002), Bajari and Hortaçsu (2003), and others provide empirical results from eBay.com and Amazon.com that have informed much of the theoretical effort on Internet auctions since. In particular, it has

[3]The desire to avoid a bidding war as a motive for sniping is supported by anecdotal evidence provided by, for example, Roth and Ockenfels (2002).

[4]In fact, Amazon no longer offer an auction marketplace. However, 'Amazon-style' extendible end time auctions remain of theoretical interest, and continue to provided a useful dichotomy with their hard-close counterparts.

been repeatedly demonstrated that there exists a strong propensity for late bidding—so-called 'sniping'—in eBay auctions. No such tendency was apparent in the extendible end time auctions formerly hosted by Amazon. To appreciate why this might have been a somewhat surprising result, it is necessary to know that in both auction forms bidding is conducted via a proxy bidding system: Each bidder is invited to enter the highest amount that he is willing to pay at a given auction. The system then compares all of the entered bids, declares that bidder having placed the highest 'reservation' bid to be the provisional winner, and enters a bid on that agent's behalf which is just high enough to outbid the now second-placed bidder. This process repeats each time a new reservation bid is provided by a bidder until, once the auction ends, the item is allocated to the current provisional winner who, in return, is obliged to pay the bid entered on his behalf by the proxy system.

If one imagines a proxy agent representing each actual agent in an incremental bidding process to determine the winner of an auction, then the proxy system is somewhat akin to a standard English auction. A casual inspection suggests that—from the buyer's perspective—the process of entering reservation prices into the system converts this mechanism into what is ostensibly a second price auction. However, it is well known that in such an auction with independent and private values, an agent can do no better than to report his true valuation and, when all agents behave in such a fashion, the timing of such a report ought to have little bearing on the actual outcome. The challenge, then, has been to provide a rational basis for the type of sniping behaviour observed at eBay auctions, whilst simultaneously explaining the absence of such behaviour at Amazon.

In their paper of 2002, Roth and Ockenfels provide an initial conjecture that the observed difference in behaviour is in large part due to differences in the ending rules in question. Auctions run at eBay have a pre-specified duration, which can range from 24 hours to ten days; the auction ends promptly once this time period has elapsed, with no further bidding possible under any circumstances. This is known as a 'hard close' rule. Amazon-style auctions also use a pre-specified ending time. However, an Amazon auction never ends if a new reservation bid has been entered within the preceding ten minutes; once the end time is reached, the auction is continuously extended until ten minutes have elapsed with no such bid updating taking place.

Ockenfels and Roth (2006) establish a theoretical basis to accompany their earlier empirical results. They postulate that bids submitted in the closing seconds of the auction are 'lost' with positive probability, giving rise to equilibria in which all players bid late and profit in expected terms from the possibility of rival bids being unsuccessfully transmitted. Also developed by Ockenfels and Roth (2006), as well as Bajari and Hortaçsu (2003), is an interdependent values model in which agents are aware of the fact that early bidding can reveal valuable private information to one's rival buyers.

These models provide a rationale for last-minute bidding, but not for the coexistence of sniping and early bidding in an eBay auction. This issue was rectified by Hossain (2008), who showed that agents may want to bid both early and late in order to learn about their value for an object. More recent work recognises the significance of concurrency in Internet auctions. In a primarily experimental paper, Ely and Hossain (2009) examine multi-unit, homogeneous good Internet auctions in a continuous time setting. The authors model agents as arriving sequentially and incrementally bidding in a uniform-price fashion until they are declared the high bidder for one of the available items. This again explains early and multiple bidding, although the assumption that some bidders are strategically naïve—failing to increase bids except when outbid by a rival—is required in order to induce last-minute bidding amongst rational agents. Evidence from Anwar, McMillan, and Zheng (2006) supports the prediction that agents bid across several competing auctions in the fashion predicted by Ely and Hossain (2009), whilst Wilcox (2000) and Chiang and Kung (2005) find evidence for multiple-bidding—also produced by the model, as well as by that of Hossain (2008).

II THE MODEL

II.1 Base model: an iterative ascending auction.

I begin by developing the basic model that will underlie the results to follow. There are a finite number of risk-neutral bidders $i \in \mathbb{I} = \{1, 2, \ldots, I\}$ who participate in a set of finitely many auctions, one for each element of a set of heterogeneous items $k \in \mathbb{K}$ to which the seller assigns a value of zero. Let $|\mathbb{K}| = K \geq 2$. The auctions run concurrently in continuous time. Where necessary, I shall refer to a specific point in time as t, with the starting time normalised to $t = 0$. Each $i \in \mathbb{I}$ is assigned a vector of (private) valuations, $\mathbf{v}_i$, with element k being his valuation for item k and denoted v_i^k. These valuations are drawn from the distribution F_i^k having support $\left[\underline{v}_i^k, \overline{v}_i^k\right]$ with $\underline{v}_i^k, \overline{v}_i^k \in [0, \infty)\ \forall\, k,\ \forall\, i$.

Denote by $\mathbf{b}_i\,(t)$ the vector of player i's bids at time t, with element $b_i^k\,(t) \geq 0$ being player i's bid on item k at time t, and $b_i^k\,(0) = 0\ \forall\, i,\ \forall\, k$. Players may place one bid per item per t for all t, with the restriction that previous bids cannot be withdrawn or reduced. At a given t, a player is announced to be the provisional winner for item k if $b_i^k(t) > \max_{j\in\mathbb{I}\setminus i}\left(b_j^k(t)\right)$, that is if his bid on k exceeds that of all of his rivals. Ties are broken immediately by some randomisation process. It is convenient to draw a distinction between the terms 'high bidder', which, at time t, refers to any member of the set $\mathbb{W}^k(t) = \left\{i : b_i^k(t) \geq \max_{j\in\mathbb{I}\setminus i}\left(b_j^k(t)\right)\right\}$ and 'provisional winner', the latter being the single member of $\mathbb{W}^k(t)$ selected by the tie-break rule and denoted by $w^k(t)$. At any point, players may observe the history of the game, $\mathcal{H}(t)$, comprised of the identities of those who have bid at all previous points in time and the precise timing of their bids, along with the second highest price for each k at each point in time—which I denote $\overline{b}^k(t)$ being the k^{th} element of the vector $\overline{\mathbf{b}}(t)$. Once the auction terminates, allocations are made according to standard second price auction rules: each $k \in \mathbb{K}$ is allocated to i if and only if $w^k(\bar{t}) = i$, where $\bar{t}$ denotes the last period in which a revised bid is received. Any k with no bids is left unallocated and has $w^k(t) = \varnothing$. Let $\mathbf{x}_i(t)$ be the vector whose k^{th} element is equal to 1 if $w^k(t) = i$, and equal to 0 otherwise. The transfer for each $i \in \mathbb{I}$ is then given by $-\mathbf{x}_i(\bar{t}) \cdot \overline{\mathbf{b}}(\bar{t})$. Sometimes, I shall refer to the price vector $\mathbf{p}$, with elements $p^k = \overline{b}^k(\bar{t})$. A strategy in this game is a vector-valued function mapping $\mathcal{H}$ and $\mathbf{v}_i$ into a bid vector $\mathbf{b}_i$ for each point in time.

As is standard in the literature, I assume that players have a utility function that is quasi-linear in a divisible unit which can be used to make any transfers demanded by the mechanism *viz.* money, and that the magnitude of the transfers that players may make is not restricted by a budget constraint. In addition, each bidder is assumed to have unit demand so that his final utility takes the form

$$U_i\left(\mathbf{x}_i(\bar{t})\right) = \max_{k\in\mathbb{K}}\left(x_i^k(\bar{t})\, v_i^k\right) - \mathbf{p}\cdot\mathbf{x}_i(\bar{t}).$$

The unit demand assumption is restrictive. However, Gul and Stacchetti (2000) demonstrate that the class of preferences for which a monotone (increasing) price auction is able to implement outcomes that are both efficient and strategy-proof does not extend to include more general gross substitutes preferences.[5] There is also evidence that, at least for some categories of items, unit demand bidders constitute a significant proportion of participants in Internet auctions. For example, Anwar, McMillan, and Zheng (2006) find that typically no more than ten or twenty percent of bidders exhibit multi-unit demand in eBay auctions for computer CPUs.

[5]Gul and Stacchetti (1999; 2000) do show that, if an economy with gross substitutes is replicated at least as many times as there are items for sale, a Bayesian equilibrium can be implemented in a monotone price auction. In this case the results that follow remain largely unchanged.

II.2 Extendible end time auction

In order to model auctions with an extendible end time, I suppose that time is initially defined on an interval thus:

$$t \in [0, T],$$

where $[0, T]$ is the scheduled time span of the auction, and T the scheduled end time. If a bid is received at some $\widetilde{t} \in (T - \tau, T]$ then the auction is extended to occupy $\left[0, \widetilde{t} + \tau\right]$, where τ is a parameter of extension time. A bid at some $\widetilde{t'} \in \left(\widetilde{t}, \widetilde{t} + \tau\right]$ triggers further extension so that the auction now runs until $\widetilde{t'} + \tau$, and so on until some $t^* \geq T$ is reached such that there has been no bidding in the interval $(t^*, t^*]$ at which point the auction is terminated.

I now propose an incremental strategy (or bidding algorithm) for the above-specified auction, similar to that originally proposed by Demange, Gale, and Sotomayor (1986), and then proceed to demonstrate that the proposed algorithm has a number of desirable properties. Define by $\gamma_i^k(t)$ the indicator function having value 1 for precisely one, randomly chosen, member of the set

$$\mathbb{S}_i(t) = \left\{ k : v_i^k > \overline{b}^k(t) \text{ and } k \in \arg\max_{m \in \mathbb{K}} \left(v_i^m - \overline{b}^m(t) \right) \right\},$$

and a value of 0 otherwise. If $\mathbb{S}$ is empty, then $\gamma_i^k(t) = 0 \; \forall \; k \in \mathbb{K}$. Thus, $\gamma_i(t)$ is the indicator function which selects one item from the set of “most appealing” items to i at price level $\overline{\mathbf{b}}(t)$, with $\sum_k \gamma_i^k(t) \leq 1$. The incremental strategy is then given by a starting bid vector, $\mathbf{b}_i(0) = (0, 0, \ldots, 0)$, and the following myopic rule for bid updating:

$$\text{(2.1)} \quad b_i^k(t + dt) = \begin{cases} \left(\overline{b}^k(t) + \alpha dt \right) \text{ if } \gamma_i^k(t) \prod_{k \in \mathbb{K}} \left(1 - x_i^k(t)\right) = 1 \\ b_i^k(t) \text{ otherwise,} \end{cases}$$

with α, a positive constant bounded away from both zero and infinity, being the parameter that controls the rate of bid updating. In words, (2.1) calls upon player i to increase his bid on item k if (i) he is not currently the provisional winner for any item, and (ii) at the current bid level he can extract a positive surplus from item k which is at least as great as that available to him from any other item. Specifying the strategy in this fashion allows me to link, explicitly, the rate of bid updating to that of time progression. Given that valuations must always be finite, this leads to the conclusion of lemma 1: that the auction must end after a finite time has elapsed when all i bid in accordance with (2.1).

Lemma 1 *The algorithm proposed in (2.1) causes the auction to terminate in finite time.*

Proof. Note that (2.1) calls for the rate of bid increase to be proportional to that of time progression. Thus, $\alpha t \leq \sum_{\mathbb{K}} \overline{b}^{k}(t)$ so long as updated bids are being submitted. Since (2.1) never calls upon a player to bid more than his valuation (see the definition of γ_i^k), it must always be true that $\overline{b}^{k}(t) \leq \max_{i\in\mathbb{I}}\left(\overline{v}_i^k\right)$ so that

$$\alpha\bar{t} \leq \sum_{k\in\mathbb{K}} \overline{b}^{k}\left(\bar{t}\right) \leq \sum_{k\in\mathbb{K}} \max_{i\in\mathbb{I}}\left(\overline{v}_i^k\right).$$

Observing that $\alpha \in (0,\infty)$, $\overline{v}_i^k$ is always finite by assumption, and that $\mathbb{K}$ is a finite set completes the proof. ■

It is helpful to spend some time considering the form of a *Walrasian* (or competitive) *equilibrium* within the present framework. I use the following definition:

Definition 2 (Walrasian Equilibrium) *An allocation* $\mathbf{X} = \mathbf{x}_1, \ldots, \mathbf{x}_I$, *and a price vector* $\mathbf{p}$ *constitute a Walrasian equilibrium if (i)* $\sum_{i\in\mathbb{I}} x_i^k \leq 1 \; \forall \; k$, *(ii)* $p^k = 0$ *whenever* $\sum_{i\in\mathbb{I}} x_i^k = 0$, *and (iii)* $(\mathbf{x}_i \cdot \mathbf{v}_i) - (\mathbf{x}_i \cdot \mathbf{p}) \geq (\mathbf{y}_i \cdot \mathbf{v}_i) - (\mathbf{y}_i \cdot \mathbf{p})$ *for all* i, *and any allocation* $\mathbf{y}_i$ *with* $y_i^k \leq 1 \forall k$.

These conditions respectively require that no item is allocated more than once, that unallocated items have a price of zero, and that—taking the Walrasian price as given—no buyer can do better than his Walrasian allocation.[6] Sellers play no part in this definition and, consequently, condition (ii) is required in order to rule out the sort of unattractive equilibria in which all prices are prohibitive and no player has (or wants to have) any item. With this definition in mind, it becomes possible to introduce lemma 3 as follows:

Lemma 3 *If* $dt \to 0$, *the algorithm proposed in (2.1) terminates in a Walrasian equilibrium allocation with corresponding supporting price vector.*

Proof. Consider the three conditions in definition 2: (i) is always satisfied by the auction mechanism since $w^k(t)$ is defined as a singleton; to see that (ii) must be true, note that $p^k > 0$ implies $b_i^k\left(\bar{t}\right) > 0$ for some subset of $\mathbb{I}$ and, in turn, that $w^k\left(\bar{t}\right) \in \mathbb{I}$. Thus, by definition, $\sum_{i\in\mathbb{I}} x_i^k \neq 0$. Finally, for condition (iii): if $w^k\left(\bar{t}\right) = i$

[6]Note that, when bidders have unit demand, condition (iii) requires that no bidder is allocated more than one item with $p^k > 0$.

then, from (2.1), when i last increased his bid on k, γ_i^k was equal to 1. Since the high bid for k has not since risen, and the high bid for every $m \in \mathbb{K}\backslash k$ cannot have fallen, it must remain the case that

$$\max_{m \in \mathbb{K}} \left(v_i^m - \overline{b}^m \left(\overline{t} \right) \right) = \left(v_i^k - \overline{b}^k \left(\overline{t} \right) \right) \geq 0.$$

Moreover, if $w^k \left(\overline{t} \right) \neq i \; \forall \; k \in \mathbb{K}$ then (2.1) implies that $\gamma_i^k = 0 \; \forall \; k$—no item can yield positive surplus to player i.

Thus, with unit demand, $\mathbf{x}_i$ can be no worse for i than any other allocation—including the null allocation—at prices $\mathbf{p} = \overline{\mathbf{b}} \left(\overline{t} \right)$, and condition (iii) is satisfied.
■

The following corollary is an immediate consequence of lemma 3 and the first fundamental welfare theorem.[7]

Corollary 4 *The outcome of the unit demand progressive auction when* $dt \rightarrow 0$*, and bidders use (2.1), is efficient.*

Aside from lending the desirable property of efficiency to the auction, attaining a Walrasian equilibrium is an important first step along the road to incentive compatibility. In the remainder of this chapter, I shall concern myself exclusively with the case of $dt \rightarrow 0$ so that bids are made in continuous time, with no minimum bid increment. Absent this assumption, agents will be unwilling to compete for their surplus-maximising bundle if the difference between their valuation and the standing high bid is less than the minimum bid increment; thus efficiency will not, in general, be obtained. Anticipation of such circumstances will distort a player's strategic incentives and, consequently, reliable convergence to an efficient allocation is key to producing an outcome that is also incentive compatible.

Example 1 *Figure 2.1 shows a simple example of bid progression in this continuous time case, with* $K = 2, I = 3$*, and valuations given by* $v_1 = (4, 8)$*,* $v_2 = (5, 2)$*, and* $v_3 = (0, 10)$ *. Initially, the high-bid level of item* 1*—depicted in the top panel—is constant at zero: there is no competition for this item, since player* 1 *(and player* 3*) would prefer to win item* 2*—depicted in the bottom panel—when* $\overline{b}^2 (\cdot) \leq 4 = v_1^2 - v_1^1$*. Once* $\overline{b}^2 (\cdot) > 4$ *player* 1 *competes against player* 2 *in the first auction and against player* 3 *in the second, updating his bids continuously to maintain equality of available surplus from each, so that both prices rise together*

[7] A proof of the first fundamental welfare theorem in the current context is provided in appendix A

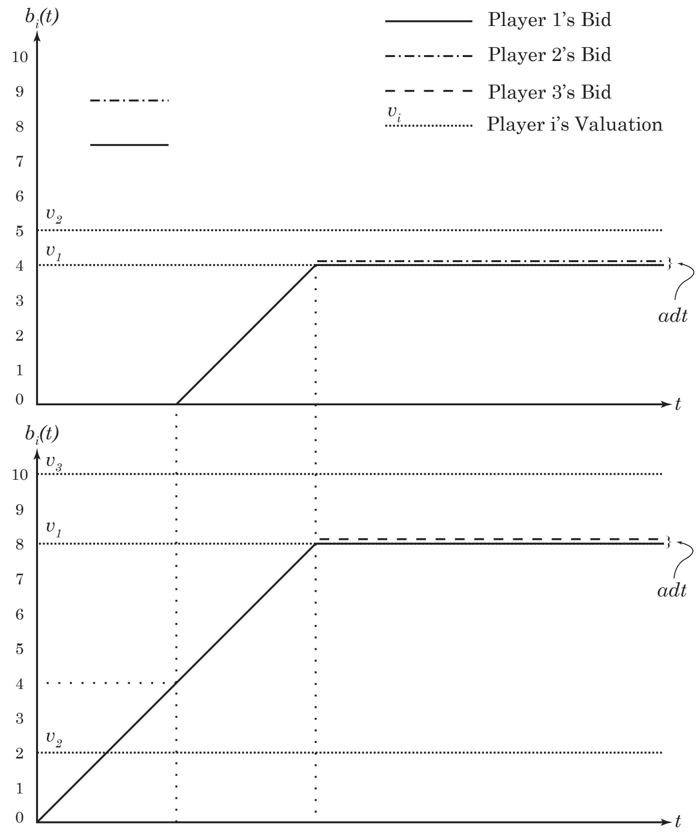

FIGURE 2.1 Bid progression in the $dt \to 0$ case.

at the same rate. Once player $1's$ value is reached, which happens simultaneously in both auctions, player 1 stops increasing his bid and his rivals outbid him by the minimum possible amount, αdt. The outcome reached is efficient, with player 2 winning item 1, and player 3 item 2. Observe that player 3 never submits a positive bid for item 1—since his value for this item is zero. Moreover, player two never submits a positive bid for item 2: player 2 would prefer to win item 1 whenever $\bar{b}^1(t) - \bar{b}^2(t) \leq v_2^1 - v_2^2 = 3$, a condition that is satisfied for all t in the given example.

Note that the total social surplus generated by an allocation is given by $\sum_{i \in \mathbb{I}} \mathbf{v}_i \cdot \mathbf{x}_i(\bar{t})$, and that $\mathbf{v}_i \cdot \mathbf{x}_i(\bar{t}) = U_i(\mathbf{x}_i) + \mathbf{p} \cdot \mathbf{x}_i(\bar{t})$. When $\left(U_i\left(\mathbf{x}_i\left(\bar{t}\right)\right), \mathbf{p} \cdot \mathbf{x}_i\left(\bar{t}\right)\right)$ is a Walrasian allocation, it is known as an imputation, with $U_i\left(\mathbf{x}_i\left(\bar{t}\right)\right)$ being the portion of the surplus generated from i's allocation that is retained by i and $\mathbf{p} \cdot \mathbf{x}_i(\bar{t})$ being the portion redistributed to the seller by means of a transfer payment. When i is optimally assigned item k, the latter of these terms is simply the equilibrium price, p^k. The astute reader will notice that there will generally be a multitude of price vectors that are capable of forming a Walrasian imputation. For example, if $I = K = 2$ and $\mathbf{v}_1 = (2, 1), \mathbf{v}_2 = (1, 2)$ then any $\mathbf{p}$ such that $p^1, p^2 \in [0,2]$ and $\left|p^1 - p^2\right| \leq 1$ supports the Walrasian equilibrium allocation with $w^1 = 1, w^2 = 2$. Given this multiplicity, one might ask: is there any logic to the means by which (2.1) chooses an equilibrium price vector? The answer is affirmative: the (Walrasian) equilibrium reached by (2.1) is supported by a price vector which is small in the strong sense that each item price, p^k, is no greater than the corresponding price from any other Walrasian price vector. Shapley and Shubik (1971) provide an existence proof for such a 'smallest Walrasian price vector'—which I denote by $\underline{\mathbf{p}}$—in the unit demand setting. Given this existence result, I am ready to demonstrate the convergence of (2.1) to $\underline{\mathbf{p}}$, and do so in lemma 5.

Lemma 5 *The competitive equilibrium reached by (2.1) is supported by* $\underline{\mathbf{p}}$*, the smallest price vector capable of supporting a Walrasian equilibrium in the market.*

Proof. Let $\underline{\mathbf{p}}$ denote the minimum Walrasian price vector and suppose that the lemma is false. Note that $\overline{b}^k(0) \leq \underline{p}^k$. Lemma 3 tells us that (2.1) always produces a competitive equilibrium outcome. Thus, if (2.1) does not lead to an allocation supported by $\underline{\mathbf{p}}$, there must be some $\widehat{t}$ such that (i) $\overline{\mathbf{b}}\left(\widehat{t}\right) \leq \underline{\mathbf{p}}$ and (ii) $\overline{b}^k\left(\widehat{t}+dt\right) > \underline{p}^k$ for some k. Denote by $\mathbb{A}$ the set of all such $k \in \mathbb{K}$. Since $\overline{b}^k\left(\widehat{t}+dt\right)$ is the second highest bid for k at time $\widehat{t}+dt$ this implies that, for each $k \in \mathbb{A}$, at least two members of $\mathbb{I}$ have valuations such that $v_i^k - \overline{b}^k\left(\widehat{t}+dt\right) \geq v_i^m - \overline{b}^m\left(\widehat{t}+dt\right)\ \forall\, m \in \mathbb{K}\backslash k$, and, more specifically, such that

(2.2) $v_i^k - \overline{b}^k\left(\widehat{t}+dt\right) \geq v_i^m - \overline{b}^m\left(\widehat{t}+dt\right)\ \forall\, m \in \mathbb{K}\backslash\mathbb{A}.$

Since (2.2) is true for at least two i for every $k \in \mathbb{A}$ and players never bid (or hold the provisional winner position) on more than one item simultaneously, it must be true for no fewer than $2\left|\mathbb{A}\right|$ members of $\mathbb{I}$.

Now compare $\overline{\mathbf{b}}\left(\widehat{t}+dt\right)$ to $\underline{\mathbf{p}}$. We know that $\underline{p}^k < \overline{b}^k\left(\widehat{t}+dt\right)\ \forall\, k \in \mathbb{A}$ and that $\underline{p}^k \geq \overline{b}^k\left(\widehat{t}+dt\right)\ \forall\, k \in \mathbb{K}\backslash\mathbb{A}$. Thus, there are at least $2\left|\mathbb{A}\right|$ bidders for whom $v_i^k - \underline{p}^k > v_i^m - \underline{p}^m\ \forall\, m \in \mathbb{K}\backslash\mathbb{A}$ for some $k \in \mathbb{A}$, but only $\left|\mathbb{A}\right|$ such items. There must, then, be at least one $k \in \mathbb{A}$ that is in excess demand at prices $\underline{\mathbf{p}}$, and hence $\underline{\mathbf{p}}$ is not an equilibrium. The necessary contradiction is thus produced. ■

Thus, the Walrasian equilibrium produced by (2.1) is, in fact, the minimum price equilibrium. A version of this convergence result was first demonstrated by Demange, Gale, and Sotomayor (1986). In figure 2.1, for example, a decrease in the final price for either item would result in 1's demand for that item being positive, and hence in excess demand. Thus the prices reached in this case are minimum Walrasian prices. In general, this minimum Walrasian price is exactly the same imputation that Leonard (1983) shows to be equivalent to the Vickrey Clarke Groves (VCG) outcome for the assignment game. Leonard's result was used as the basis for Gul and Stacchetti (2000) Theorem 5, which demonstrates that an algorithm that converges to $\underline{\text{p}}$ in a monotone fashion must also be incentive compatible. The core intuition for this result is that any deviation by a given player, i, results in an outcome that can be reproduced by a player with an artificially chosen valuation vector who bids in accordance with the algorithm. Thus, i's deviation amounts to an imitation of the artificially created player, and is therefore sub-optimal by the incentive compatibility of the VCG mechanism. This result implies that summarised in proposition 6.

Proposition 6 (Gul and Stacchetti, 2000) *The strategy specified in (2.1) constitutes a perfect Bayesian equilibrium of the extendible end time auction game when used by all players.*

II.3 Hard close: a model of eBay proxy auctions

Auctions on eBay operate a so-called 'hard close' rule, under which auctions run from a pre-specified start time until a similarly pre-specified end time. Once the end time is reached, the auction is concluded and there is no possibility of the auction's duration being extended beyond this point. An adjustment to the specification of time used in the 'base model' is therefore necessary, with the auction now occupying a fixed interval $t \in [0, T]$. I assume that there exists an α such that the algorithm in (2.1) always has sufficient time to complete before T is reached. This amounts to assuming that the interval $[0, T]$ has positive measure.

The important implication of this timing modification for the purposes of the current work is that agents now have the opportunity to place a final bid—at time T—to which no rival may respond. In seeking to understand the effects of a hard close rule, it is helpful to define some new compound strategies of the form (i's action during $[0, T)$; i's action at T); the three strategies of interest are (algorithm; abstain), (abstain; snipe) and (algorithm; snipe) where 'algorithm' denotes bidding in accordance with (2.1), 'abstain' abstinence from bidding all together and 'snipe' the placement of a bid equal to one's valuation on item $\arg\max_k \left(v_i^k - \bar{b}^k(T)\right)$.[8] The following lemma will also prove useful in subsequent proofs.

[8]Here, (algorithm; abstain) is analogous to the strategy shown to be an equilibrium of the extendible end time auction in proposition 6.

Lemma 7 *If* $\underline{\mathbf{p}}$ *is the minimum Walrasian price in a market with buyers* $\mathbb{I}$*, and* $\underline{\mathbf{p}}_{-i}$ *is similarly defined for a market with set of buyers* $\mathbb{I}\backslash i$ *then* $\underline{\mathbf{p}} \geq \underline{\mathbf{p}}_{-i}$*.*

Proof. Suppose that the lemma is false. If $\left\{k : \underline{p}^k < \underline{p}^k_{-i}\right\}$ is non-empty then, with set of players $\mathbb{I}\backslash i$, price vector $\underline{\mathbf{p}}$ must produce excess demand since $\underline{\mathbf{p}}_{-i}$ is a minimum Walrasian price vector. Since i's demand for each item is non-negative and players' demands depend only on prices and their own valuations, $\underline{\mathbf{p}}$ must also produce excess demand for set of players $\mathbb{I}$. But then $\underline{\mathbf{p}}$ is not an equilibrium price vector, and we have a contradiction. ■

The first ramification of the timing rule change can now be articulated as in lemma 8.

Lemma 8 *The strategy (algorithm; abstain) is weakly dominated the eBay (hard close) assignment auction.*

Proof. I demonstrate the existence of an expected surplus improving deviation for some i. Suppose that all players use the strategy (algorithm; abstain). The surplus extracted by i is then the greater of zero or $\max_{k\in\mathbb{K}}\left(v_i^k - \underline{p}^k\right)$. Now, suppose instead that all $j \in \mathbb{I}\backslash i$ use the strategy (algorithm; abstain) whilst i uses (abstain; snipe). During $[0, T)$, the algorithm converges to the lowest Walrasian prices for the reduced market which excludes player i, $\underline{\mathbf{p}}^k_{-i}$. The important point to note is that i is able to observe $\overline{\mathbf{b}}\,(\sup\,[0, T))$ and hence can still identify his surplus maximising item. Since $\underline{p}^k_{-i} \leq \underline{p}^k$ (by lemma 7), and players $j \in \mathbb{I}\backslash i$ have no opportunity to respond to i's time T bid, i's deviation yields to him a surplus of

$$\max\left\{0, \max_{k\in\mathbb{K}}\left(v_i^k - \underline{p}^k_{-i}\right)\right\} \geq \max\left\{0, \max_{k\in\mathbb{K}}\left(v_i^k - \underline{p}^k\right)\right\}.$$

Thus, when rivals are using (algorithm; abstain), the response (algorithm; abstain) is always (weakly) inferior to (abstain; snipe). ■

Thus, the equilibrium established for the extendible end time auction does not have a general analogue in the fixed end time case. Those players that use the (algorithm; abstain) strategy—increasing their bid only when a rival outbids them—in a fixed end time auction are typically described in the literature as naïve since they are behaving as one might expect a bidder to behave in a standard English auction, without acknowledging the additional strategic considerations introduced by the hard close rule. Such naïve behaviour is open to exploitation, as demonstrated in lemma 8; indeed, Ockenfels and Roth (2006) and Ely and Hossain (2009) have shown in a less general setting that (abstain; snipe) is a best response to such behaviour. In corollary 9, I am able to generalise this result to the concurrent, heterogeneous item case.

Corollary 9 *Pure sniping is a best response to naïve bidding of the form (algorithm; abstain).*

Proof. When using the strategy (abstain; snipe), against rivals who use (algorithm; abstain), i is able to buy any $k \in \mathbb{K}$ at a price of $\underline{p}^k_{-i}$ so that his pay-off is $\max\left\{0, \max_{k\in\mathbb{K}}\left(v^k_i - \underline{p}^k_{-i}\right)\right\}$. Suppose that i uses some other strategy. Clearly, if i abstains during $[0,T)$, he can do no better than $\max\left\{0, \max_{k\in\mathbb{K}}\left(v^k_i - \underline{p}^k_{-i}\right)\right\}$, regardless of his choice of time T bid. Suppose, instead, that i uses some strategy that has him bid during $[0,T)$ (possibly in addition to bidding at time T), whilst his rival bidders continue to use (algorithm; abstain). Denote by $\widetilde{\mathbf{X}}$ the allocation thus produced, and by $\widetilde{\mathbf{p}}$ the associated price vector. Thus i extracts a surplus that is, at most, equal to $\max\left\{0, \max_{k\in\mathbb{K}}\left(v^k_i - \widetilde{p}^k\right)\right\}$.[9]

Now, ignoring cases with $\sum_{k\in\mathbb{K}} \widetilde{x}^k_i > 1$ as clearly sub-optimal, consider the valuation functions given by

$$v^k_\psi = \begin{cases} 0 \text{ if } \widetilde{x}^k_i \neq 1 \\ \max_{j\in\mathbb{I}}\left(\overline{v}^k_j\right) + 1 \text{ if } \widetilde{x}^k_i = 1, \end{cases}$$

$$v^m_{I+k} = \begin{cases} \min\left\{b^m_i(T), \overline{b}^m(T)\right\} \text{ if } m = k \\ 0 \text{ if } m \neq k, \end{cases}$$

with $\{\psi, I+1, \ldots, I+K\} \cap \mathbb{I} = \emptyset$. If agents $\{\psi, I+1, \ldots, I+K\}$, having such utility functions, were to bid in i's place, and do so according to the strategy (algorithm; abstain) then the outcome obtained at time T would be similar to that produced by the arbitrary strategy used by i during $[0,T)$; having the same price vector, $\widetilde{\mathbf{p}}$, with i's items allocated instead to ψ. By lemma 5, since all members of $\{\psi, I+1, \ldots, I+K\} \cup \{j : j \in \mathbb{I}\backslash i\}$ are using (algorithm; abstain), $\widetilde{\mathbf{p}}$ must be the minimum Walrasian equilibrium price for the market with this set of bidders. It must, then, be true that $\widetilde{p}^k \geq \underline{p}^k_{-i}$ because adding players can never lower the minimum Walrasian price (lemma 7). It is thus shown that

$$\max\left\{0, \max_{k\in\mathbb{K}}\left(v^k_i - \widetilde{p}^k\right)\right\} \leq \max\left\{0, \max_{k\in\mathbb{K}}\left(v^k_i - \underline{p}^k_{-i}\right)\right\};$$

no response to (algorithm; abstain) can do better than can (abstain; snipe). ■

[9]The actual surplus may be somewhat less than this if, for example, i uses a strategy that calls upon him to make a large and early bid—leaving him committed to winning an item that is ultimately sub-optimal.

An important caveat attached to the behaviour described in proposition 9 is that the surplus maximising item for the sniper need not be the same item that he would otherwise have been allocated by the mechanism *viz.* his socially efficient allocation. Example 2 demonstrates how the incentive to snipe can induce this kind of inefficiency.

Example 2 *Consider valuation vectors given as follows*

$$\mathbf{v}_1 = (10, 10)\,;\ \mathbf{v}_2 = (5, 0)\,;\ \mathbf{v}_3 = (5, 0)\,;\ \mathbf{v}_4 = (0, 15)\,.$$

The efficient assignment is $w^1 = 1$ and $w^2 = 4$ with a total social surplus of 25 and prices $p^1 = 5,\ p^2 = 5$. However, if 1 uses the strategy (abstain; snipe) against rivals using (algorithm; abstain) then he is faced with a price of 5 for item 1, and prefers to win item 2 at a price of zero yielding a social surplus of 15.

The behaviour described in corollary 9 and example 2 does not, however, constitute a perfect Bayesian equilibrium of the auction game. In particular, as example 2 shows, those players that bid algorithmically face a positive probability of being outbid by the sniper, even when their valuation for their respective efficient items is higher than that of all of their rivals. The risk of surplus loss due to deviations can be partially mitigated if the algorithmic bidders also snipe—using the strategy (algorithm; snipe)—with the snipe bid being placed on the item to which they are allocated by the algorithm. Returning to example 2, if player 4 uses the strategy (algorithm; snipe) and bids $b_4^2(T) = 15$ then, in the face of sniping by player 1, he extracts a surplus of 5 from item 4—still worse than the efficient allocation at $\underline{\mathbf{p}}$, which left him with a surplus of 10, but non-the-less better than the outcome if he plays (algorithm; abstain), in which he receives zero surplus.

The adoption of the strategy (algorithm; snipe) by players $j \in \mathbb{I} \backslash i$ undermines the advantage of the (abstain; snipe) strategy for i, since the latter strategy exploits the general differential between the highest algorithmic bid placed by the provisional winner for each item and the valuation of that same bidder. The snipe portion of the (algorithm; snipe) strategy has the effect of eliminating this differential. In example 2, sniping against rivals using (algorithm; abstain) produces a positive increase in surplus of 5 for player 1, but using the same strategy when players switch to (algorithm; snipe) results in a fall in surplus of 5 for player 1 vis-à-vis the efficient outcome.

Once players adopt the (algorithm; snipe) strategy, i must take as given the fact that he will face a valuation bid on all items for which there is competition.

Since other players' valuations—and therefore the final price for any item won by i—are invariant to observed bidding behaviour, i's primary strategic concern is to ensure that he is allocated his surplus maximising item. There is thus hope that an efficient equilibrium with symmetric, algorithmic strategies might exist and indeed the existence of such an equilibrium is confirmed in proposition 10.

Proposition 10 *The strategy (algorithm; snipe) constitutes a symmetric perfect Bayesian equilibrium of the hard close auction game.*

Proof. The proof has three parts. I begin in parts (i) and (ii) by considering a deviation by i to some strategy that results in him winning, for concreteness, (only) item 1. I use a tilde to denote deviation variables, and a hat to denote the outcome that would prevail if i followed his deviation strategy during $[0, T)$, but abstained at time T.

(i) If $\widehat{w}^1(T) = i$, or if $\widetilde{p}^1 = 0$, then, since all $j \neq i$ are bidding incrementally, it must be true that $\widehat{p}^1 = \widetilde{p}^1$, and that at price $\widehat{\mathbf{p}}$, no $j \neq i$ would prefer to win item 1 over having their own $\widehat{\mathbf{x}}_j$. We can then introduce an artificial player, ψ having valuations as follows:

$$v_\psi^k = \begin{cases} 0 \text{ if } k \neq 1 \\ \max_{j \in \mathbb{I}} \left\{ \overline{v}_i^j \right\} + 1 \text{ if } k = 1, \end{cases}$$

Note that $\widehat{\mathbf{p}}$ is a Walrasian price in the market with $\psi \cup \mathbb{I} \backslash i$. Moreover, when ψ bids algorithmically in i's place, he wins item 1. Writing $\underline{\mathbf{p}}_\psi$ for the minimum Walrasian price thus reached, we have

$$\begin{aligned} \max\left\{0, \max_{k \in \mathbb{K}} \left(v_i^k - \underline{p}^k\right)\right\} &= i\text{'s non-deviation (VCG) surplus} \\ &\geq i\text{'s surplus when imitating } \psi \text{ in VCG} \\ &= \widetilde{\mathbf{x}}_i \cdot \mathbf{v}_i - \widetilde{\mathbf{x}}_i \cdot \underline{\mathbf{p}}_\psi \\ &\geq \widetilde{\mathbf{x}}_i \cdot \mathbf{v}_i - \widetilde{\mathbf{x}}_i \cdot \widehat{\mathbf{p}}, \\ &= \widetilde{\mathbf{x}}_i \cdot \mathbf{v}_i - \widetilde{\mathbf{x}}_i \cdot \widetilde{\mathbf{p}}, \end{aligned}$$

where the first inequality follows from the incentive compatibility of the VCG mechanism, and the second from the fact that $\underline{\mathbf{p}}_\psi$ is a minimum Walrasian price.

(ii) Suppose that $\widehat{w}^1(T) \neq i$ and $\widetilde{p}^1 > 0$ so that i wins item 1 by usurping player j with a snipe bid. It must, then, be the case that $\widetilde{p}^1 = v_j^1$. Replace i with K additional bidders, $\{I+1, \ldots, I+K\}$, having

$$(2.3) \quad v_{I+k}^m = \begin{cases} \widehat{b}_i^m(T) \text{ if } m = k \\ 0 \text{ if } m \neq k, \end{cases}$$

so that algorithmic bidding by all yields $w^1(T) = j$. If, in addition, we add the artificial player ψ so that the marketplace now consists of players $\psi \cup \{I + 1, \ldots, I + K\} \cup \mathbb{I}\backslash i$ then player ψ, bidding algorithmically, wins item 1 at the minimum Walrasian price for this enlarged market. This price is necessarily equal to the market's VCG price so that ψ's payment can be calculated as

$$\underbrace{\max_{\text{allocating } \mathbb{K}} \sum_{l \neq \psi} \mathbf{v}_l \mathbf{x}_l}_{A} - \underbrace{\max_{\text{allocating } \mathbb{K}\backslash 1} \sum_{l \neq \psi} \mathbf{v}_l \mathbf{x}_l}_{B} .$$

Since j was assigned item 1 by the algorithm under set of players $\{I + 1, \ldots, I + K\} \cup \mathbb{I}\backslash i$, term A can be re-written as

$$v_j^1 + \max_{\text{allocating } \mathbb{K}\backslash 1} \sum_{l \neq \psi, j} \mathbf{v}_l \mathbf{x}_l .$$

One possibility for calculating term B is to have $\sum_k x_j^k = 0$, and leave all other player's allocations unchanged from those in term A, in which case term B is equal to

$$\max_{\text{allocating } \mathbb{K}\backslash 1} \sum_{l \neq \psi, j} \mathbf{v}_l \mathbf{x}_l . \tag{2.4}$$

It may be possible to do better by allocating to j some $k \in \mathbb{K}\backslash 1$, so that (2.4) is a lower bound on term B. It follows that v_j^1 is an upper bound on ψ's payment. Now we have that

$$\begin{aligned}
\max\left\{0, \max_{k \in \mathbb{K}} \left(v_i^k - \underline{p}^k\right)\right\} &= \text{i's VCG surplus} \\
&\geq \text{i's surplus when imitating ψ in VCG against players $\mathbb{I}\backslash i$} \\
&\geq \text{i's surplus when imitating ψ in VCG against players} \\
& \ \{I + 1, \ldots, I + k\} \cup \mathbb{I}\backslash i \\
&\geq v_i^1 - v_j^1 \\
&= \text{i's deviation surplus,}
\end{aligned}$$

where the first inequality follows from the incentive compatibility of the VCG mechanism, the second from lemma 7, and the third from the argument presented above.

(iii) Ruling out deviations in which $\sum_k \widetilde{x}_i^k > 1$ is now a straightforward extension of parts (i) and (ii) of the proof: for concreteness, let item 1 be the k that i consumes. If $\widehat{w}^1(T) = i$, or if $\widetilde{p}^1 = 0$ then the proof follows the logic of part (i), but requires the addition of players similar to those defined in (2.3) for each $k \neq 1$ with $\widehat{w}^k(T) = i$. By lemma 7, such deviations must be worse for i than compliance—even if i only has to pay for the item he consumes. When $\widehat{w}^1(T) \neq i$

and $\widetilde{p}^1 > 0$, it must be true that $\widetilde{p}^1 = v_j^1$. Thus, i's deviation surplus is no greater than $v_i^1 - v_j^1$, which was shown in part(ii) to be less than i's compliance utility. ■

A perfect Bayesian equilibrium is thus provided for the eBay hard close auction. However, one should note that, whilst last minute bidding occurs in equilibrium here, this is a sniping strategy in name only: players are, in fact, indifferent between bidding their valuation at $t = T$, and doing so at any other time after the final algorithmic bid has been entered.[10] In order to break this indifference and make explicit the sniping incentive, I introduce in section III an additional feature into the model—namely stochastic entry—which creates an 'option value' for the delay of valuation bidding.

In addition, it should be noted that behaviour analogous to that of proposition 10 (namely, algorithmic bidding followed by an additional bid close to the auction's end) can also be supported in equilibrium in the extendible end time auction, so that systematic late bidding can arise in equilibrium at such auctions.[11] The introduction of stochastic entry creates a disincentive to this kind of behaviour in the extendible end time, and may thus explain the observed (relative) absence of late bidding at such auctions.

III STOCHASTIC ENTRY AND ARRIVAL

III.1 Stochastic arrival in a hard close auction

When an online (and often also an 'off-line') auction is taking place, it is typically difficult (if not impossible) to observe exactly the number and identities of those bidders that are participating. Moreover, Internet auctions are conducted over relatively lengthy periods, with new rival bidders liable to arrive at any instant. I model this behaviour as a Poisson process in which agents arrive randomly during $t \in [0, T - \epsilon)$, where ϵ is positive, but may be small.[12]

One would naturally like to know whether the results in lemmas 3 and 5, and propositions 6, and 10 continue to hold in the presence of such stochastic entry. I begin by observing that, by lemmas 5 and 7, the high bid, $\bar{b}$, for any item at the

[10] It is left for the reader to verify that proposition 10 remains valid with valuation bidding at any such t.

[11] For instance, it is a symmetric equilibrium of the extendible end time auction for all players to bid in accordance with (2.1) during $[0, t)$, bid one's valuation at time T on any item for which one is the provisional winner, and then to submit no further bids (during the period over which the auction is extended).

[12] The interval $[T - \epsilon, T]$, during which entry is prohibited, is required to have some positive measure in order to ensure that all players have sufficient time to complete the 'algorithm' portion of the (algorithm; snipe) strategy subsequent to the final arrival.

time of arrival of a new agent must be less than the minimum Walrasian price for the newly enlarged set of bidders. Moreover, the proof of lemma 3 remains unchanged. In this light, lemma 5 also continues to hold.[13] Given the veracity of lemmas 3 and 5, equivalence with the VCG outcome is maintained, and propositions 6 and 10 are verbatim transferable to the stochastic entry setting.

The point of interest emerging from all of this is that, in the presence of stochastic entry, agents in a hard close auction prefer to snipe strictly after the last possible arrival time, as well as after algorithmic bidding has concluded. Intuitively, this is because any bidder 'sniping' earlier is vulnerable to being heavily committed to winning an item which, subsequent to the arrival of a new rival, is not a member of his surplus-maximising set. The intuition is illustrated in example 3.

Example 3 *Consider valuations given by the following vectors:*

$$\mathbf{v}_1 = (10, 5), \mathbf{v}_2 = (3, 2).$$

The algorithm arrives at the allocation $w^1 = 1, w^2 = 2$ *supported by price vector* $\underline{\mathbf{p}} = (1, 0)$. *If, at time* $T - \epsilon$*, a third bidder arrives with a valuation vector of the form* $\mathbf{v}_3 = (9, 0)$ *then a bid of* $b_1^1(t) = 10$ *at some* $t < T - \epsilon$ *is clearly sub-optimal—leaving player* 1 *with a surplus of* 1 *as opposed to the surplus of* 3 *that he obtains by also using the algorithm for* $\mathbb{I} = \{1, 2, 3\}$ *and sniping (this time on item* 2*) only when algorithmic bidding is complete and no further arrival possible.*

It has been shown that the result of participating in the algorithm until the minimum Walrasian price for the set of all bidders is reached coincides with the VCG outcome, and is thus incentive compatible. Thus, deviations by a player having the potential to change the final outcome—of which premature valuation bidding is an example—can never be profitable since they amount to imitating a player of a different type and thus receiving that player's VCG outcome. This is formalised in proposition 11.

Proposition 11 *In the hard close auction with stochastic entry, it is a symmetric, perfect Bayesian equilibrium if player* i *behaves as follows: bid in accordance with (2.1) until some* t *such that* $t > T - \epsilon$ *and* $\overline{\mathbf{b}}(t) = \overline{\mathbf{b}}(t - dt)$*; bid* $b_i^k(t) = \max\left\{b_i^k(t - dt), x_i^k(t - dt) \cdot v_i^k\right\}$ *for all* t *thereafter.*

The proof here is identical to that of proposition 10, with the snipe portion of the (algorithm; snipe) strategy now strictly interpreted as meaning "snipe after $T - \varepsilon$,

[13] Note that this implies that 2.1 will converge to $\underline{\mathbf{p}}$ from any $\overline{\mathbf{b}} < \underline{\mathbf{p}}$.

and only once algorithmic bidding has finished". If one attempts to interpret the 'sniping' of proposition 10 as an earlier valuation bid then it becomes possible to construct an example profile similar to the case illustrated in example 3, so that one or more bidders prefers to delay their final bid.

In order to better understand the difference between the theoretical results contained within Ely and Hossain (2009), and the equilibrium of proposition 10, it is useful at this juncture to draw a comparison between the two. Since all items are identical in Ely and Hossain's paper, bidding one's valuation on the item having the lowest price as soon as that item has been identified is an equilibrium under stochastic (or sequential) arrival. The addition of naïve bidders is required to generate an incentive for the remaining rational bidders to snipe. However, the presence of bidders that do not bid optimally will typically induce inefficiency. By contrast, proposition 11 of the current chapter demonstrates that introducing item heterogeneity *and* stochastic entry creates an incentive to delay bidding—even in the absence of naïve bidders—by inducing bidders to care differentially about the future evolution of $\mathbf{p}$. Corollary 4 implies that the resulting allocation is efficient.

III.2 Stochastic arrival in an extendible end time auction

The effect that allowing stochastic entry has upon the extendible end time auction is quite different. Firstly, it has already been seen that the core results of section II.2 continue to hold under stochastic entry so that the equilibrium presented within that section remains valid.

Proposition 12 *Bidding in accordance with (2.1) remains a symmetric equilibrium of the extendible end time auction under stochastic entry.*

A bid at some $\widetilde{t}$ in the interval $(T-\tau,T]$ causes the auction to be extended and may create an opportunity for additional bidders to arrive at the auction during $\left(T,\widetilde{t}+\tau\right]$. It is precisely for this reason that agents in an extendible end time auction with stochastic entry will, in general, find it unprofitable to snipe: to do so potentially opens them to exactly the kind of problem that compels buyers in an eBay auction to bid late, and in any case increases the expected intensity of competition for all items. Thus, whenever i has valuations such that $v_i^k > \overline{b}_i^k\,(T-dt)$ for some $k \in \mathbb{K}$, he faces a strong disincentive to sniping. Indeed, when $\tau > \varepsilon$, algorithmic bidding followed by a larger 'snipe' bid during $[T-\varepsilon,T]$ will not generally be an equilibrium of the stochastic entry, extendible end time auction since any given bidder would be better off bidding algorithmically throughout the entire auction. The intuition for this point is illustrated in example 4.

Example 4 *Suppose that players bid in accordance with the equilibrium of section II.2 and have valuations given as follows:*

$$\mathbf{v}_1 = (10, 5), \mathbf{v}_2 = (6, 0).$$

If, by time $T - dt$*,* 1 *and* 2 *are still the only two bidders to have arrived then* $\overline{\mathbf{b}}(T - dt) = (5, 0)$ *with* $w^1 = 2, w^2 = 1$*. Now, consider a 'snipe' bid of* $b_1^2(T) = 5$ *so that the auction is extended until* $T + \tau$*. If some player with valuation* $\mathbf{v}_3 = (0, 4)$ *arrives during the interval* $(T, T + \tau]$ *then* 1 *is left with a surplus of* 1 *from item* 2*. By contrast, continuing to use the algorithm, which implies* $b_1^2(T) = b_1^2(T - \epsilon) = 0$ *will deny* 3 *the opportunity of participating and lead to the outcome with* $w^1 = 2, w^2 = 1, \underline{\mathbf{p}} = (5, 0)$*, and an improved surplus for* 1 *of* 5*.*

Thus, the same mechanism that promotes sniping in an eBay-style hard close auction may deter such behaviour in extendible end time environments. The lack of a clear equilibrium incentive to snipe, accompanied by the pitfalls of doing so illustrated in example 4 may go some way to explaining the relative paucity of observed Amazon sniping.

IV CONCLUSION

In this chapter I have addressed the question of what effects ending rule choice might have on equilibrium behaviour in multi-unit auctions. In particular, when bidders arrive stochastically, I have demonstrated the existence of an equilibrium of the eBay-style 'hard close' auction environment which is novel in the sense that the allocation is efficient despite the fact that many of the bidders systematically delay their final bid until the end of the auction. Moreover, this late bidding emerges with fully rational agents facing a perfect bid transmission technology. Thus, whilst I make no claim regarding the uniqueness of the equilibria presented here, I have been able to show that efficient outcomes are possible in equilibrium, even when auctions use a hard close rule, and thus that fixed ending times need not always be regarded with suspicion. Indeed, if a degree of certainty regarding the end time of the auction is of value to the participants, or if circumstances impose a natural deadline upon proceedings, then a hard close rule may be preferable. This notwithstanding, one remaining pitfall associated with the imposition of a fixed end time is the danger that bidders not have sufficient time to complete the matching process—a danger especially apparent when bidders may arrive very close to the end of the auction—and, in this respect, the extendible end time auctions retain their advantage.

I have also been able to demonstrate that early, late, and multiple bidding can coexist in a manner consistent with rational behaviour when auction markets are viewed as multi-unit auctions. Although the work here has been presented in a unit demand framework, more general preferences for which the minimum Walrasian price equilibrium is coincident with the VCG outcome (such as the K replica gross substitute economies of Gul and Stacchetti (1999)) should leave the above results largely unaltered.

Since both hard close and extendible end time auctions arrive at the same price vector, there is no revenue benefit to choosing one specification over the other in the equilibria presented above. Much like the single-unit English auction (and by the revenue equivalence theorem, other standard single-unit auctions) convergence to the minimum Walrasian price equilibrium implies that revenues are minimised over the set of efficient outcomes.

There is still more work to be done in attempting to understand the dynamics of Internet auctions. Not all eBay snipe bids are placed by the ultimate winner of an auction and, indeed, there are multiple snipe bids placed for many items. It remains to be seen whether these facts can be reconciled with rationality and efficiency. Other potential extensions to the present work include consideration of the effects of allowing the start and end times for the auctions to differ across items, and the strategic implications of allowing buyers to arrive at any point up to the end of the hard close auction.

A Omitted Proofs

Proof of First Fundamental Welfare Theorem. Let $\mathbf{X}$ be some competitive equilibrium allocation at price $\mathbf{p}$, and $\widehat{\mathbf{X}}$ be some allocation that Pareto dominates $\mathbf{X}$. Since all agents (weakly) prefer $\widehat{\mathbf{X}}$ to $\mathbf{X}$, from condition (iii) of definition 2, it must be true that

$$\sum_{i\in\mathbb{I}} \mathbf{p}\cdot\mathbf{x}_i < \sum_{i\in\mathbb{I}} \mathbf{p}\cdot\widehat{\mathbf{x}}_i. \tag{2.5}$$

Since, by condition (ii) of definition 2, $p^k = 0$ whenever $\sum_{i\in\mathbb{I}} x_i^k = 0$, (2.5) implies that there exists a k such that $\sum_{i\in\mathbb{I}} \widehat{x}_i^k > 1$ and hence $\widehat{\mathbf{X}}$ is not a feasible allocation. ■

Chapter 3

Search Quality and Revenue Cannibalisation by Competing Search Engines

I Introduction

In this chapter I build a simple model of the competitive environment within which search engines operate and examine the interplay between paid-for advertisements (or A-links) and free non-advertisement search results (so-called organic links, or O-links) when there are consumers that search optimally. Interesting questions include: why should search engines offer organic links at all?; given that organic links are offered, should we expect them to be useful?; and how might the presence of Organic links affect advertising revenues?

Intuitively, a search engine competes for A-link clicks, and thus has an incentive to provide a high quality of service in order to win market share. By providing high quality O-links a search engine attracts consumers to visit its site first. This is beneficial if the same consumers, in an attempt to minimise search costs, stay to also click on advertisements, rather than continuing their search elsewhere. However, there exists a countervailing effect since search engines face competition for A-link clicks not only from links at rival search engines, but also from their own organic links. Thus, the market is characterised by a kind of revenue cannibalisation that results in a delicate trade-off between the complementary effects of O-links (the incentive to compete for market share) and the desire to minimise the extent to which consumers substitute away from advertisements.

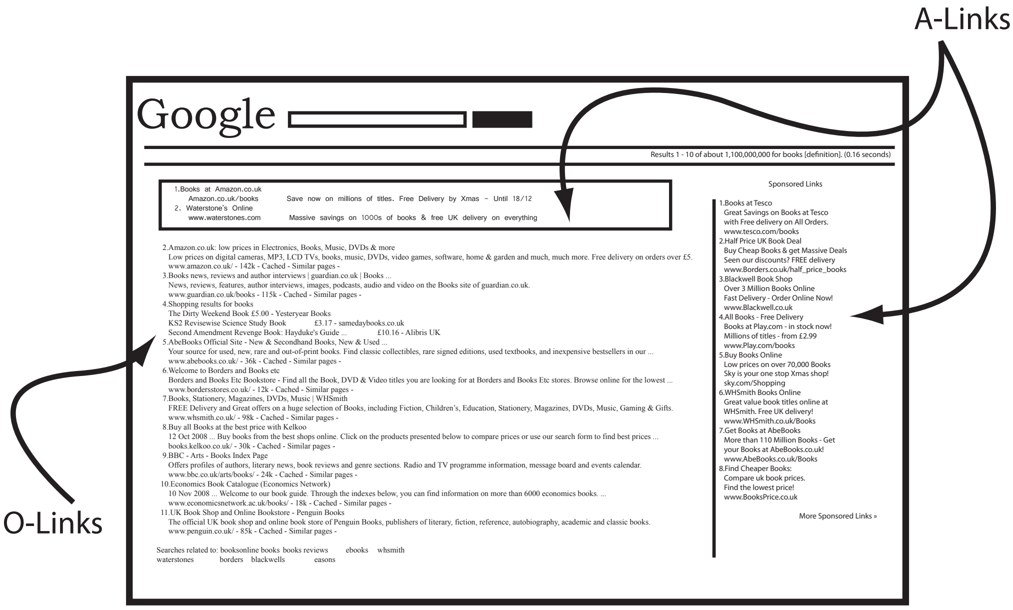

FIGURE 3.1 Organic (O) and Advertising (A) search results.

The extent of equilibrium cannibalisation is variable. Sometimes, global free-link quality can get 'too good' so that consumers find it profitable to switch search engines and continue clicking free links on another site, rather than clicking advertisements. In this case, there is a natural ceiling to how good the free links can get in equilibrium. In other cases, search engines set their quality to the maximum possible so that competition is, in some sense, maximal. The cannibalisation effect also engenders a class of equilibrium in which quality is comparatively low, even when the provision of quality is costless. What all oligopoly equilibria have in common is that the quality of non-paid-for links is at least as good as that of advertisements, so that consumers (at least weakly) prefer to click on a non-paid for link first. Cannibalisation is thus a pervasive feature of equilibrium. If the market exists as a monopoly then the 'competition for market share' effect disappears and revenue cannibalisation compels the monopolist search engine to set a low quality.

I also analyse competition when consumers have heterogeneous visit costs for each search engine. This gives each search engine a degree of monopoly power over its captive consumers, which weakens the competition effect relative to that of revenue cannibalisation so that the maximum equilibrium quality is lower. Nevertheless, I find that the result that organic links are at least as good as their advertising counterparts in equilibrium is robust to this relaxation of the model's assumptions.

I.1 Literature

Gandal (2001) conducts an empirical study of competition within the Internet search engine market. Two results are of particular interest for the current work: firstly, Gandal finds that (on average) consumers use more than one search engine, so that consumers are demonstrably willing to switch between engines in order to find what they are looking for. Secondly, it is shown that the relevance of search results is by far the most important determinant of search engine ranking (as measured by number of searches conducted). This suggests that search engines have a strong incentive to compete on result quality. Taken together with the obvious importance of search engines in modern society, these results motivate and justify the model developed below.

Most closely related to this work is the paper of White (2008), who obtains results similar in spirit to those presented here using a monopolist search engine that offers both advertisement and non-advertisement links. The non-advertising links reduce consumers' search costs and induce more consumers to search—increasing the demand faced by advertising merchants—but simultaneously provide competition for those consumers' business, thus reducing price in the final goods market. This competition amongst the sellers reduces the amount that they are willing to pay for advertisements. The search engine trades-off these two factors to maximise the profit that it is able to generate when charging a fixed fee to advertisers.

Telang, Rajan, and Mukhopadhyay (2004), and Pollock (2008) have theoretical models of the search industry with competing search engines. Telang, Rajan, and Mukhopadhyay (2004) model entry in the Internet search industry and attempt to explain the existence of low quality firms (*viz.* search engines) when consumers pay a price of zero. Pollock (2008) demonstrates a tendency for concentration in the internet search industry and then explores a number of welfare and regulation issues with a monopolist search provider. In both papers, search engine revenues are treated in a reduced form fashion. In this chapter, by contrast, I explicitly model consumer link choice, which makes the profitability of a consumer's visit endogenous to the search engine's chosen quality. This introduces a number of new equilibrium considerations which lead to the cannibalisation discussed above.

There is also a large literature on the relationship between advertising and other media content—for a summary see Bagwell (2007) section 10. Examples include Anderson and Coate (2005), who analyse the provision of programs and advertisements by radio and television broadcasters, and Gabszewicz, Laussel, and Sonnac (2001), who study advertising in the newspaper press. In both cases, much as for search engines, advertisements are provided alongside additional content that is attractive to consumers but does not generate revenue directly. The Internet search advertising market differs, however, in the fact that organic search results compete directly for clicks with the revenue generating advertisements on the same site.

This work also has links to other branches of the literature. In particular, Varian (2007); Edelman, Ostrovsky, and Schwarz (2007); and Athey and Ellison (2008) model the 'position' auction framework that search engines typically use to sell advertisements.

The remainder of this chapter is organised as follows: In section II, I establish a duopoly model of search engine competition in an environment with a mass of identical consumers and then, in section III, proceed to characterise equilibrium behaviour within this model. The subsequent sections consist of a series of extensions to this simple model. Firstly, in section IV, I relax the assumption that identical advertisements are displayed on all search engines. I generalise the duopoly model to the monopoly and oligopoly cases in section V and, In section VI, I relax the assumption of consumer homogeneity and analyse competition when consumers have heterogeneous costs for visiting each search engine. Section VII concludes.

II Simple Model

A mass 1 continuum of homogeneous risk-neutral consumers in this model have a particular need or desire that they seek to satisfy by searching the Internet. Each time a consumer visits (or re-visits) a search engine they must pay a visit cost, $S > 0$. When the search results are returned, the consumer may click on them as he or she pleases, but must pay a further cost, $s > 0$, for each link that is clicked. If a clicked link matches the consumer's need then the consumer receives a surplus, which I normalise to 1. I assume that the consumers exhibit unit demand so that a satisfied consumer can gain no further utility from searching. In the case of internet search, the switching cost, S, may be small. However, I demonstrate below that it is that size of this switching cost relative to the cost of clicking a further link at the current site that is of greatest importance.

Two search engines, g and y, provide search results to consumers at zero marginal cost. A search at a given search engine returns two results: one organic (non-advertising) search result (O-link), and one advertising search result (A-link). Let A_i and O_i respectively denote the A-link and O-link at site i. The search engines simultaneously choose a quality, denoted by

$p_i \in [0, p^{max}], p^{max} \leq 1, i = g, y$, for their O-links. This quality is the probability that the O-link at site i satisfies a given consumer's need, conditional on having been clicked. I assume that satisfaction is statistically independent across consumers and across links. In order to focus attention on the pure incentives for quality choice in search competition, rather than on the broader competition in R&D (which, although frantic, occurs over longer time scales), I assume that a commonly, and freely accessible search technology already exists, and permits qualities up to p^{max}, which can be thought of as the maximum technologically feasible algorithm quality. In particular, this allows clear identification of the extent to which equilibrium quality reduction is due to the cannibalisation of revenue effect.

The A-link also has a quality, denoted q, which is common to both g and y. For the purposes of the simple model, I assume that the same link appears in the A-link slot at both search engines (this assumption is relaxed in section IV).[1] The A-links and O-link point to websites each drawn from separate pools of firms, so that there are always three distinct links available to the consumer.

I assume that the consumers are able to observe the match probabilities p_g, p_y, and q.[2] When the A-link at a search engine is clicked, that search engine receives a per-click price, which I denote by b. The search engine receives nothing when its O-link is clicked.

To summarise: search engines move first and simultaneously select a quality p_i. Consumers observe p_g, p_y, q, S and s, and select whether, and in which order to click each link. The game ends when all consumers have either had their need met or do not wish to click any further links. Throughout the chapter, the solution concept that I use will be that of sub-game perfect Nash equilibrium.

III Equilibrium Behaviour in the Simple Model

III.1 Optimal consumer behaviour

The problem faced by each consumer in this model is to determine whether to click on each link and in which order to do so. A strategy for a consumer specifies

[1]To understand the rationale for this assumption, which is relaxed in section IV, suppose that the characteristics of consumers attracted by g and y are identical, and that each search engine sells its A-link slot by means of a second price auction. If advertiser-firm j makes profit π_j for each sale and converts proportion q_j of its clicks into sales, then it is a (weakly) dominant strategy for j to bid $\pi_j q_j$ per click for the advertising opportunity. Under such circumstances, the winner of each ad auction will be the same firm.

[2]For example, one can think of consumers who search regularly and rapidly learn the average result qualities, p and q.

these actions as a function of the choices of p_g and p_y, as well as the model parameters S, s, and q. Since the consumers are aware that both sites display the same A-link, they only ever click the A-link at one of the two sites. Once a consumer's need is met, that consumer always stops clicking on links. Note that since all consumers are assumed homogeneous they will agree on a preference ordering over the set of possible strategies.

If $q < s$ then no consumer ever clicks an A-link and search engines, which make zero profits, are indifferent across all choices of p. When $s \leq q < S + s$, consumers click on an A-link only if $max\{p_g, p_y\} \geq S + s$ (which will always be true in equilibrium), in which case the below analysis remains essentially unchanged. Throughout the remainder of this and the following two sections, I assume that $q \in (S + s, 1)$, so that consumers are always willing to click on A-links.[3]

Suppose, now, that $p_g, p_y > q$. Having paid the initial visit cost, S, for the first search engine, the consumer finds it optimal to click the O-link there first. If $\min\{p_g, p_y\}$ is particularly high, consumers may prefer to next switch from one search engine to the other—clicking the O-link at both—before finally clicking the A-link at the second site he visits. In contrast, if $\min\{p_g, p_y\}$ is only marginally larger than q, the consumer prefers to click both the O-link and A-link at the first site he visits before switching over to the second site. I shall refer to these behaviours respectively as 'switching' and 'sticking', and denote by γ the value of $\min\{p_g, p_y\}$ that makes the consumers indifferent between the two.

Calculation of the parameter γ is as follows: Suppose, for concreteness that $q < p_g \leq p_y$. Consumers find it optimal to click on O_y first. The total expected utility that each consumer gets from 'switching' to next click on O_g is

$$(3.1) \quad U = p_y(1 - S - s) + (1 - p_y)p_g(1 - 2S - 2s) + \\ (1 - p_y)(1 - p_g)q(1 - 2S - 3s) + (1 - p_y)(1 - p_g)(1 - q)(-2S - 3s).$$

Similarly, the utility from 'sticking' is

$$(3.2) \quad U = p_y(1 - S - s) + (1 - p_y)q(1 - S - 2s) + \\ (1 - p_y)(1 - q)p_g(1 - 2S - 3s) + (1 - p_y)(1 - q)(1 - p_g)(-2S - 3s).$$

Setting (3.1) and (3.2) equal to one another, and calculating γ as the p_g that

[3]In order to simplify the exposition, I am also ruling out the trivial case in which $q = 1$. Briefly, $q = 1$ gives rise to a continuum of uninteresting equilibria in which, search engines choose an arbitrary $p \in [0, 1)$, and consumers click an A-link at an arbitrarily chosen search engine resulting in immediate satisfaction.

yields indifference gives:

$$p_g = \frac{S+s}{s} q \equiv \gamma.$$

This simply says that the utility per unit of expenditure for clicking O_g and A_y should be equal. An increase in p_g to some $p'_g \in (\gamma, p_y)$ makes clicking $\{O_y, O_g, A_g\}$ strictly more attractive than $\{O_y, A_y, O_g\}$, and so consumers switch. Conversely, a reduction in p_g leaves consumers preferring to stick. A symmetric argument applies to the case of $p_y < p_g$.

It is now possible to characterise the utility maximising strategy for a consumer given $\{p_g, p_y, q\}$, and this is done in strategy 1. First, let

$$e \equiv \begin{cases} g \text{ if } p_g > p_y \text{ and } p_g \geq s \\ y \text{ if } p_g < p_y \text{ and } p_y \geq s \\ g \text{ w.p. } \alpha, y \text{ w.p. } 1-\alpha \text{ if } p_g = p_y \text{ or } \max\{p_g, p_y\} < s, \end{cases}$$

for some $\alpha \in (0,1)$, and

$$-e \equiv \{g, y\} - \{e\}.$$

Strategy 1 *Suppose that $q \geq S + s$. Any consumer best response strategy maps the link qualities $\{p_g, p_y, q\}$ into a click order $\{a_1, a_2, a_3\}$ in the following manner. Begin by clicking a_1 thus:*

$$a_1 = \begin{cases} O_e \textit{ if } p_e > q \\ A_e \textit{ if } p_e < q \\ A_e \textit{ w.p. } \lambda, O_e \textit{ w.p. } 1-\lambda \textit{ if } p_e = q. \end{cases}$$

If the consumer's need was met by a_1 then stop clicking (i.e. $a_2 = a_3 = \emptyset$), otherwise click a_2 as follows:

$$a_2 = \begin{cases} A_e \textit{ if } a_1 = O_e \textit{ and } p_{-e} < \gamma \\ O_{-e} \textit{ if } a_1 = O_e \textit{ and } p_{-e} > \gamma \\ A_e \textit{ w.p. } \phi, O_{-e} \textit{ w.p. } 1-\phi \textit{ if } a_1 = O_e \textit{ and } p_{-e} = \gamma \\ O_e \textit{ if } a_1 = A_e \textit{ and } p_e \geq s \\ \emptyset \textit{ if } p_e < s. \end{cases}$$

If the consumer's need was met by a_1 or a_2 then stop clicking (i.e. $a_3 = \emptyset$), otherwise click a_3 as follows:

$$a_3 = \begin{cases} O_{-e} \textit{ if } a_2 \neq O_e, \textit{ and } p_{-e} \geq S+s \\ A_{-e} \textit{ if } a_2 = O_{-e} \\ \emptyset \textit{ if } p_{-e} < S+s. \end{cases}$$

For the sake of notational (and, occasionally, algebraic) simplicity, I have assumed that the mixing parameters α, ϕ, and λ are constant and applied symmetrically across search engines. I also focus on the case with $\alpha \in (0,1)$, so that both search engines are visited first with some probability if they are of identical quality. On an intuitive level, it is fairly clear that there will exist Nash equilibria that are sustained by consumer threats—such as the threat not to click on any links at all unless organic link quality is set at p^{max}. However, such threats are not credible since any $\{a_1, a_2, a_3\}$ that maximises expected consumer utility must take the form given in strategy 1 and, in particular, unsatisfied consumers can always obtain positive utility by clicking on an A-link. Therefore, once one insists on sub-game perfection, any best response strategy for the consumer must take the form of strategy 1 with some choice of $\alpha, \phi, \lambda, \in [0,1]$. These probability parameters may therefore, without loss of generality, be interpreted as proportions of the population undertaking any given action.

III.2 Equilibrium characterisation

I now proceed to characterise the equilibria of this game, focusing on equilibria that are in pure strategies for the search engines. I refer to such equilibria as SE-pure. The first result establishes that, when search engines are constrained to use particularly poor algorithms, quality competition will generally be maximal.

Lemma 13 *If $p^{max} < q$ then $p_g = p_y = p^{max}$ is the unique sub-game perfect equilibrium search engine behaviour.*

The proof for this and other results from this chapter can be found in appendix A.

Note that cannibalisation is not in effect for values of p less than q since all consumers then prefer to click the A-link first. Thus, there is no offsetting force for the 'competition for market share' incentive, which is why the intensity of competition described in lemma 13 obtains. For the remainder of the chapter, I shall make the assumption that $p^{max} \geq q$. I next show that all SE-pure equilibria are necessarily symmetric in search engine strategies.

Lemma 14 *All sub-game perfect SE-pure equilibria have $p_g = p_y$.*

Now, in lemma 15, I foreshadow what is to come with a result demonstrating that search engines have a strong incentive to provide O-links that compete fiercely against their own A-links.

Lemma 15 *There is no sub-game perfect SE-pure equilibrium in which* $\min\{p_g, p_y\} < q$.

The intuition here is similar to that for lemma 13: There is no cannibalisation incentive for quality limitation so long as consumers prefer to click an A-link first, which is true whenever O-link qualities are less than q. Thus, so long as both search engines offer a $p < q$, there is no disincentive to increasing p in order to capture market share. When $p^{max} > q$, lemma 15 implies that the lowest $\{p_g, p_y\}$ pair that could ever be consistent with an SE-pure equilibrium is $p_g = p_y = q$. In equilibrium 1, I demonstrate that this is indeed an equilibrium, albeit under a restricted set of circumstances.

Equilibrium 1 *The search engine strategy* $p_g = p_y = q$*, and a consumer strategy of the form detailed in strategy 1, form a sub-game perfect equilibrium of the above game whenever*

$$(3.3)\quad q \geq \max\left\{\frac{1-\alpha}{1-\alpha+\alpha\lambda}, \frac{\alpha}{\lambda+\alpha-\alpha\lambda}\right\}.$$

Condition (3.3) is plotted in figure 3.2. The higher are λ and $\min\{\alpha, 1-\alpha\}$, the lower is the minimum value of q consistent with the above equilibrium. For $\lambda = 1, \alpha = 1/2$, the equilibrium in proposition 1 is sustained by any $q \geq 1/2$—if $q < 1/2$ then there is no $\{\lambda, \alpha\}$ such that neither g nor y wishes to deviate.

The intuition here is fairly straightforward: In equilibrium 1, search engines do not wish to reduce O-link quality since their competitor then captures the entire market. Each search engine can, in fact, capture the entire market for themselves with a (potentially small) increase in quality to some $q + \epsilon$. This, though, carries a cost: at the higher quality, consumers are induced to click the deviator's O-link first so that only those left unsatisfied by the O-link ever click on the A-link. In particular, when a search engine deviates to some $q + \epsilon$, the mass of consumers that eventually clicks its A-link is $(1 - q - \epsilon)$ so that the deviation becomes less attractive as q becomes large. When consumers play in such a way that 3.3 is satisfied, the loss in clicks that the cannibalisation effect generates outweighs the gain in market share so that the search engine prefers not to deviate. Thus, the striking result that, even with costless quality, search engines may wish to systematically limit quality provision in equilibrium obtains.

If q is low, so that 3.3 can not be satisfied, then any SE-pure equilibrium must have $p_g, p_y > q$. The result of the next equilibrium makes this statement more

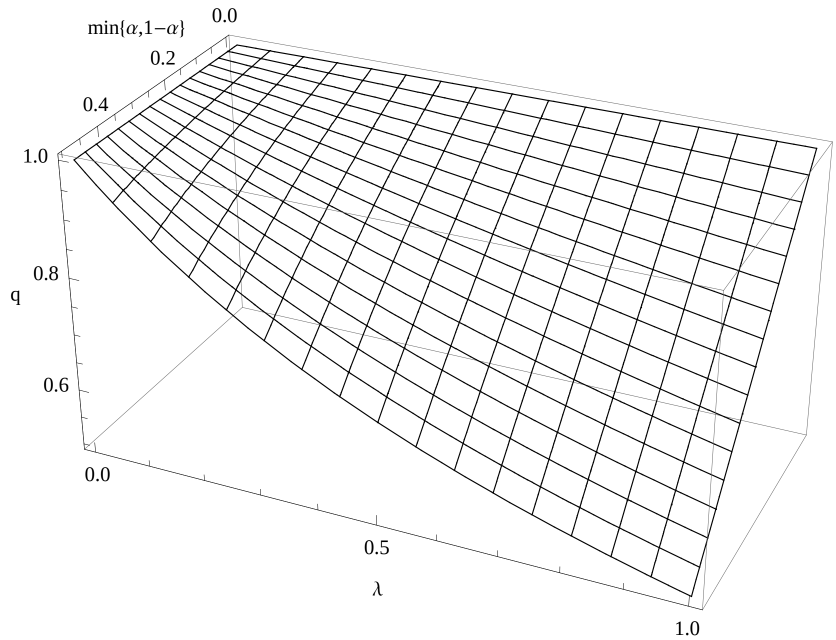

FIGURE 3.2 Minimum value of q consistent with a $p_y = p_g = q$ equilibrium for various combinations of λ and α.

stark. Under certain conditions on the parameters of the model, there is an equilibrium in which the intensity of competition is maximal in the sense that both p_g and p_y are set to p^{max}. This result is formalised in equilibrium 2.

Equilibrium 2 *If $\gamma > p^{max}$ then the search engine strategy $p_g = p_y = p^{max}$, and the consumer strategy detailed in strategy 1, form a sub-game perfect equilibrium of the above game.*

The condition for existence of equilibrium 2, *viz.* $\gamma > p^{max}$, is more likely satisfied if S is large relative to s: a search engine's power over its captive audience increases when sticking to click more links is much cheaper than continuing the search elsewhere—and hence so does the incentive to compete for visitors. More specifically, equilibrium 2 comes about because, when $\min\{p_g, p_y\} \in [0, \gamma)$, the consumer always prefers to play 'stick' rather than 'switch'. If both search engines choose some $p < p^{max}$ then there is an incentive for each engine to offer a slightly higher p than its rival in order to capture all of the A-link clicks. This result is clearly reminiscent of Bertrand price competition, and in this manner large portions of search engine profits may be dissipated. When $p^{max} = 1$, equilibrium 2 implies zero profits for search engines. In fact, it turns out that if $p^{max} = 1$, $p_g = p_y = 1$ is always an equilibrium.[4]

[4]By assumption, $q < 1$. Thus, with $p_g = p_y = 1$, the consumer clicks one of O_g or O_y, and is immediately satisfied. A-links are never clicked, and search engine profits are thus zero. A deviation

In equilibrium 2, search engines have an incentive to increase their p to the greatest extent possible. This is true whenever $p_g, p_y < \gamma$. In contrast, if $\min\{p_g, p_y\} > \gamma$, the consumer always switches and there is thus an incentive for each search engine to offer a p that is *lower* than that of its rival. It might seem natural, then, to look for an equilibrium in which $p_g = p_y = \gamma$. This is the business of equilibrium 3.

Equilibrium 3 *If $\gamma \leq p^{max}$ then the search engine strategy $p_g = p_y = \gamma$, and the consumer strategy detailed in strategy 1, with*

(3.4) $\phi \leq \dfrac{1-\gamma}{2-\gamma}$

form a sub-game perfect equilibrium of the above game.

Intuitively, equilibrium 3 works as follows: any reduction in O-link quality by a search engine is unprofitable because consumers do not have sufficient incentive to visit until its rival's links have been exhausted. Positive deviations are also unprofitable for the search engines since O-links are then 'too good'—when (3.4) is satisfied, consumers that are attracted by a high p are too likely to switch sites and continue clicking O-links, rather than stick around to try the A-link at the deviator's site. Search engines clearly have little incentive to compete for consumers who switch so readily.

In summary, the simple model has produced three classes of SE-pure equilibrium: those with $p_g = p_y = q$, those with $p_g = p_y = p^{max}$, and those with $p_g = p_y = \gamma$.[5] Equilibria from up to two of these classes can be supported at any one time. Firstly, there always exists precisely one $\{p_g, p_y\}$ with $p_g, p_y > q$ that is consistent with equilibrium. When the cost of switching search engines is high relative to that of clicking links (if $\gamma \geq p^{max}$) then there exists a family of equilibria of the form $p_g = p_y = p^{max}$, in which search engines compete fiercely for captive consumers. With lower switching costs (when $\gamma < p^{max}$), these equilibria can no longer be sustained, and are replaced by $p_g = p_y = \gamma$ equilibria, with appropriate choices of ϕ by the consumer.

In addition to the prevailing 'high quality' equilibria, a third class of 'low quality' equilibria can be supported whenever $q \geq 1/2$—these are equilibria in which search engines set $p_g = p_y = q$ and the consumer chooses α, λ to satisfy (3.3).

by i, namely setting $p_i < p_{-i} = 1$, results in the consumer clicking on O_{-i} first with probability one and thus i's profits remain at zero. Thus i has no profitable deviation.

[5] I write of classes of equilibria because there will typically by many combinations of α, λ, and ϕ that are consistent with equilibrium. However, the multitude of equilibria thus produced can be neatly divided into three classes—each having a uniqe $p_g = p_y$—along the lines specified in equilibria 1–3.

It is interesting to note that the equilibria from two of the identified classes—those with $p_g = p_y = q$ and $p_g = p_y = \gamma$—involve quality degradation, despite the fact that quality provision here has no pecuniary cost. This is a result of the trade-off between competition and cannibalisation faced by the search engines.

Since there exist two equilibrium quality levels for some parameter configurations, the question of equilibrium selection naturally arises. Search engines have an incentive to collude around the low quality equilibria—since this involves less revenue cannibalisation. The consumers, for their part, would prefer to steer search engines into the higher quality equilibria by setting mixing parameters α and λ that violate (3.3). Intuitively, it seems as though this type of coordination might be difficult when there are many 'small' consumers. However, the Internet search industry is characterised by highly asymmetric market shares, which appear to correspond to a large value for $|1/2 - \alpha|$.[6] This observation provides an empirical rationale for focussing on the 'high quality' class of equilibria. Intuition along these lines may also provide some hint as to why Google has been able to maintain such a dominant position in the search market: by behaving in this manner, consumers ensure that search engines are not able to form collusive equilibria—the incentive for low market share search engines to deviate would be too high.

IV Extension: Differentiated A-links

Although, empirically, there is often a good deal of overlap between the advertisement links shown by different search engines, it is not at all uncommon for these links to differ across sites. In recognition of this fact, I now relax the assumption that g and y share the same A-link (but maintain the assumption that A_g and A_y share a common quality, q). This implies that optimising consumers will potentially click both A_g and A_y. The model is otherwise as presented in section II.

IV.1 Optimal consumer behaviour with differentiated A-links

I begin, as before, with an examination of the optimal consumer strategy. The presence of a second A-link complicates the consumer's problem somewhat. It is, nonetheless, possible to identify a class of best responses to each possible $\{p_g, p_y\}$ pair. As before, when taken together, these responses form a class of strategies

[6] Whilst this model does not predict asymmetric market shares, it is not inconsistent with them.

that the consumers must play in any sub-game perfect equilibrium. When e and $-e$ are defined as in section III, the structure of these strategies can be outlined as in figure 3.3.

Here, $\bar{\gamma}$ (which is analogous to γ from the simple model) is the p_{-e} that makes consumers indifferent between the click orders $\{O_e, A_e, O_{-e}, A_{-e}\}$ and $\{O_e, O_{-e}, A_{-e}, A_e\}$, and its value is given by

$$\bar{\gamma} \equiv \frac{S + qs}{S + s - qs}.$$

For values of $p_{-e} > \bar{\gamma}$, consumers strictly prefer the latter ('switching') click order.

Now that consumers may click both search engines' A-link, another type of switching behaviour might arise: if O-link qualities are low enough, the consumer may find it profitable to use a click order of the form $\{A_{-e}, A_e, \cdot, \cdot\}$—preferring to pay the switching cost to click both A-links first, rather than sticking at the first site to click an O-link. More specifically, when $s \leq p_{-e} \leq p_e < q$,

$$\underline{\gamma}_1 \equiv \frac{qS + qs - S}{qS + s}$$

is the value of p_e that makes the consumer indifferent between the click orders $\{A_e, O_e, A_{-e}, O_{-e}\}$ and $\{A_{-e}, A_e, O_e, O_{-e}\}$. Similarly, if $p_{-e} < s < p_e < q$, then

$$\underline{\gamma}_2 \equiv \frac{qs}{S + s}$$

is the value of p_e that makes the consumer indifferent between the click orders $\{A_e, O_e, A_{-e}, \emptyset\}$ and $\{A_{-e}, A_e, O_e, \emptyset\}$. In both cases, an increase (reduction) in p_e away from these values would induce consumers to prefer the former (latter) click order.

These three cases each have an associated probability of sticking when consumers are indifferent between switching and sticking. These are respectively $\bar{\phi}$, $\underline{\phi}_1$, and $\underline{\phi}_2$. In keeping with the notation of section III, I again use λ as the mixing parameter when consumers are indifferent between clicking on an O-link or an A-link, and α to denote the probability that a consumer visits g first when indifferent.

IV.2 Equilibrium analysis

From figure 3.3, one can see that if $p^{max} < \max\left\{\underline{\gamma}_1, s\right\}$, any $p_g, p_y < s$ is consistent with equilibrium. As in section III, if $\max\left\{\underline{\gamma}_1, s\right\} \leq p^{max} < q$ then

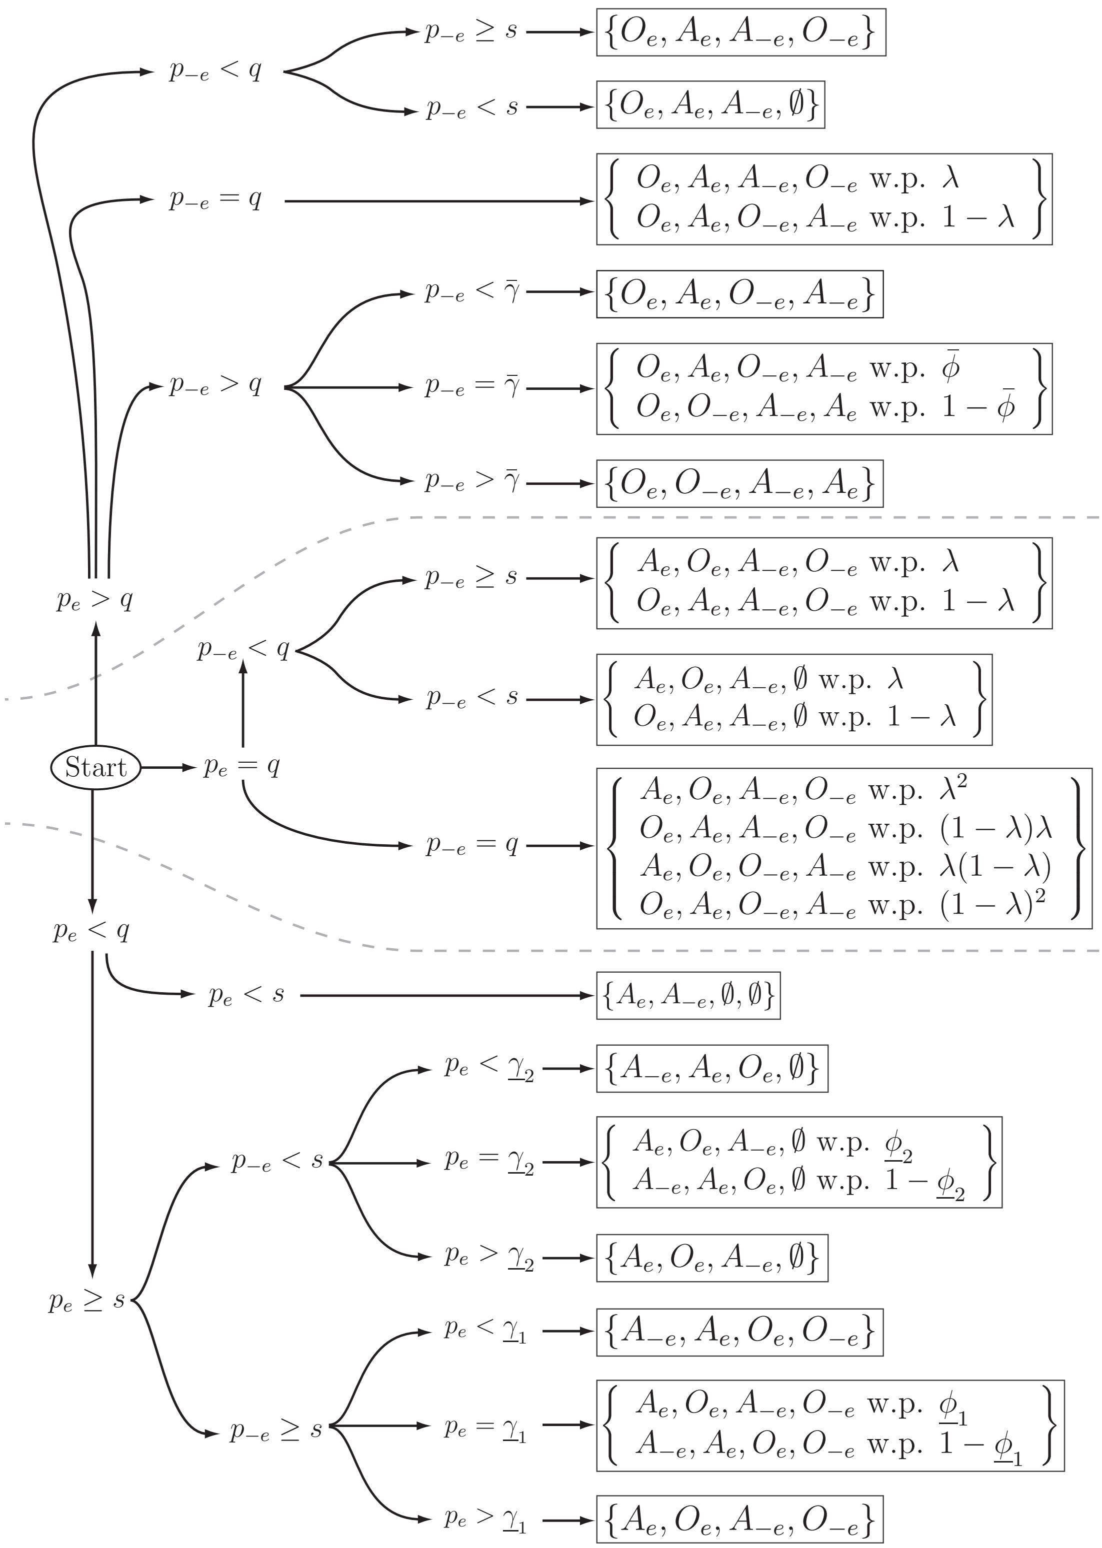

FIGURE 3.3 Expected utility maximising consumer behaviour for the 'differentiated A-links' case.

$p_g = p_y = p^{max}$ is the unique strategy pair consistent with sub-game perfect equilibrium. In the analysis that follows, I shall assume that $p^{max} \geq q$.

Under these circumstances, all sub-game perfect equilibria must have $p_g, p_y \geq q$.

Lemma 16 *Suppose that* $p^{max} \geq q$. *There is no sub-game perfect SE-pure equilibrium in which* $\min\{p_g, p_y\} < q$.

I now examine how the existence conditions for the three classes of equilibria identified in section III change in response to the introduction of differentiated A-links.

IV.3 $p_g = p_y = q$ *equilibria*

When $p_g = p_y = q$, the consumer is indifferent over which site he visits first, and which link he clicks first at each site but (for $S > 0$) strictly prefers to click both links at the first site before visiting the second. As before, I denote by λ the probability of clicking A_i before O_i. Profits for sites g and y when $p_g = p_y = q$ are then given by

$$\pi_g = \alpha\left[\lambda + (1-\lambda)(1-q)\right]b + (1-\alpha)\left[\lambda(1-q)^2 + (1-\lambda)(1-q)^3\right]b.$$

$$\pi_y = (1-\alpha)\left[\lambda + (1-\lambda)(1-q)\right]b + \alpha\left[\lambda(1-q)^2 + (1-\lambda)(1-q)^3\right]b.$$

When i deviates to some $p_i' > p_{-i} = q$, the rational consumer prefers to use a click order of the form $\{O_i, A_i, \cdot, \cdot\}$, so that i's profits become

(3.5) $\pi_i' = (1 - p_i')b.$

Since (3.5) is drecresing in p_i', it suffices to consider the limiting case with $p_i' = q$. The condition for any such deviation to be unprofitable for g is

(3.6) $$q \geq \frac{3-2\lambda}{2-2\lambda} - \frac{\sqrt{1+\alpha^2(1-2\lambda)^2 - \alpha(4\lambda^2 - 4\lambda + 2)}}{2(1-\alpha)(1-\lambda)},$$

and for y is

(3.7) $$q \geq \frac{3-2\lambda}{2-2\lambda} - \frac{\sqrt{1+(1-\alpha)^2(1-2\lambda)^2 - (1-\alpha)(4\lambda^2 - 4\lambda + 2)}}{2\alpha(1-\lambda)}.$$

Now, since the two sites' A-links are different, i may receive an A-link click even if it is not visited first. Specifically, any deviation to some $p_i < p_{-i} = q$ results in a best-response click order for the consumer of $\{\cdot, \cdot, A_i, \cdot\}$ with associated profits given by

$$\pi_i'' = (1-q)\,(1-p_{-i})\,b = (1-q)^2 b.$$

Since

$$\lim_{p_i' \to q} (1-p_i')b > (1-q)^2 b,$$

a necessary and sufficient condition for the existence of an equilibrium with $p_g = p_y = q$ is that both (3.6) and (3.7) are satisfied. Note that $\alpha = 1/2$, always satisfies these two conditions when λ is arbitrarily close to 1. Thus, the following proposition obtains:

Proposition 17 *When each site presents a unique A-link of common quality q, $p_g = p_y = q$ can always be sustained in a sub-game perfect equilibrium.*

Figure 3.4(a) shows the minimum q consistent with a $p_g = p_y = q$ equilibrium in the differentiated A-link game, as α and λ are varied. In figure 3.4(b), I compare the minimum q consistent with a $p_g = p_y = q$ equilibrium in the differentiated A-link (solid line) and identical A-link (dashed line) cases, when $\alpha = 1/2$. A $p_g = p_y = q$ equilibrium can be sustained in each case at any point above the respective surface or curve. It is immediately apparent that such an equilibrium is more easily sustained in the differentiated A-links case. In particular, when the two sites have different A-links, a $p_g = p_y = q$ equilibrium can always be sustained for $\alpha = 1/2$ and λ large enough. This is in contrast to the simple model, in which such an equilibrium can never be sustained for $q < 1/2$.

IV.4 $p_g = p_y = p^{max}$ *equilibria*

I now determine the conditions under which an equilibrium with $p_g = p_y = p^{max}$ can be sustained in the differentiated A-links game, assuming that $p^{max} > q$. As in the simple model, $\bar{\gamma} \geq p^{max}$ is a necessary condition for sustaining a $p_g = p_y = p^{max}$ equilibrium, since each firm has an incentive to reduce p_i when $p_g, p_y > \bar{\gamma}$. Given satisfaction of this condition, the profits generated when $p_g = p_y = p^{max}$ are

$$\pi_g = \left[\alpha(1-p^{max}) + (1-\alpha)(1-p^{max})^2(1-q)\right] b,$$

$$\pi_y = \left[(1-\alpha)(1-p^{max}) + \alpha(1-p^{max})^2(1-q)\right] b.$$

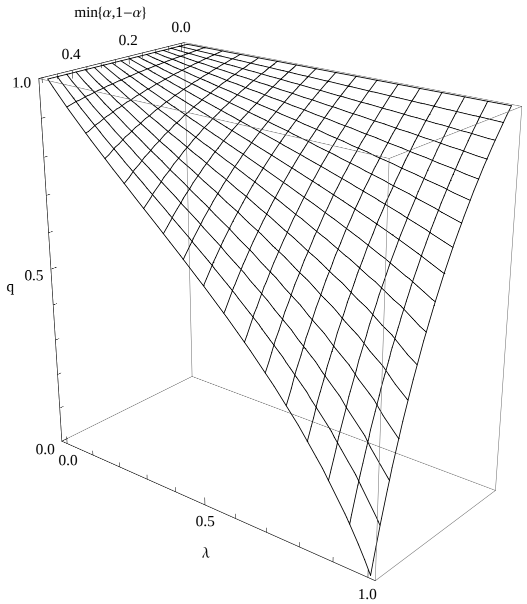

(a) Minimum q consistent with a $p_g = p_y = q$ equilibrium in the differentiated A-links game.

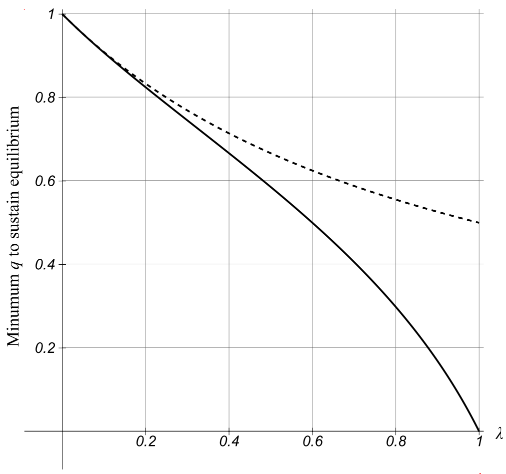

(b) Minimum q consistent with a $p_g = p_y = q$ equilibrium in the differentiated A-link (solid line) and identical A-link (dashed line) cases, with $\alpha = 1/2$.

FIGURE 3.4 Conditions for existence of a $p_g = p_y = q$ equilibrium in the differentiated A-links game.

Small reductions in p_i result in i's A-link being demoted to fourth in the click order, and are therefore less profitable than a deviation to some $p_i < q$, which compels consumers to use a click order of the form $\{O_{-i}, A_{-i}, A_i, \cdot\}$. Such a deviation is associated with profits of

$$\pi_i = (1 - p^{max})(1 - q)b.$$

Deviation is thus non-profitable for both firms when

$$q \geq \max\left\{\frac{(1-\alpha)p^{max}}{(1-\alpha)p^{max} + \alpha}, \frac{\alpha p^{max}}{\alpha p^{max} + (1-\alpha)}\right\}.$$

This condition is most easily satisfied when $\alpha = 1/2$. Making this substitution reveals the following:

Proposition 18 *When $\bar{\gamma} \geq p^{max} > q$, and each site presents a unique A-link of quality q, the strategies $p_g = p_y = p^{max}$ can be sustained in a sub-game perfect equilibrium if*

$$q \geq \frac{1}{1 + p^{max}} p^{max}.$$

Thus, this class of equilibrium is more easily sustained when q is close to p^{max}: deviation implies that a search engine's A-link is clicked only if its rival's A-link fails. When q is large, this happens so rarely that firms prefer not to deviate.

IV.5 Non-existence of $p_g = p_y = \bar{\gamma}$ equilibria in the differentiated A-links game

I next turn my attention to the question of whether it is possible to sustain $p_g = p_y = \bar{\gamma}$ in equilibrium. The profits for g from compliance with such an equilibrium are given by

$$\pi_g = \bar{\phi}\left[\alpha(1-\bar{\gamma}) + (1-\alpha)(1-\bar{\gamma})^2(1-q)\right]b + (1-\bar{\phi})\left[\alpha(1-\bar{\gamma})^2(1-q) + (1-\alpha)(1-\bar{\gamma})^2\right]b.$$

The analogous equation for y is obtained by replacing α with $(1-\alpha)$ in the above. A deviation to some $p_i' > \bar{\gamma}$ ensures that i is visited first, and yields profits of

$$\pi_i' = \bar{\phi}(1-p_i')b + (1-\bar{\phi})(1-p_i')(1-\bar{\gamma})(1-q)b.$$

As usual, this expression is decreasing in p_i', so I consider the limiting case of $p_i' = \bar{\gamma}$. Making this substitution and solving reveals that the deviation is non-profitable when

$$(3.8)\quad \bar{\phi} \leq \frac{(1-\bar{\gamma})q}{2q + \bar{\gamma} - 2q\bar{\gamma}}.$$

An alternative deviation for search engine i is to set $p_i < q$, thereby ensuring that A_i is the third link clicked. Profits in this case are

$$\pi_i'' = (1-\bar{\gamma})(1-q)b.$$

The condition for this deviation to be non-profitable for g is, then

$$(3.9)\quad \bar{\phi} \geq \frac{q - q\alpha - \bar{\gamma} + q\alpha\bar{\gamma}}{q - 2q\alpha - q\bar{\gamma} - \alpha\bar{\gamma} + 2q\alpha\bar{\gamma}},$$

whilst for y the condition is

$$(3.10)\quad \bar{\phi} \geq \frac{q - q(1-\alpha) - \bar{\gamma} + q(1-\alpha)\bar{\gamma}}{q - 2q(1-\alpha) - q\bar{\gamma} - (1-\alpha)\bar{\gamma} + 2q(1-\alpha)\bar{\gamma}}.$$

Simultaneous satisfaction of (3.9) and (3.10) is least demanding when $\alpha = 1/2$, in which case these two inequalities reduce to

$$(3.11)\quad \bar{\phi} \geq 2 - \frac{q(1+\bar{\gamma})}{\bar{\gamma}}.$$

Substituting $\bar{\gamma} = q$ into (3.8) yields

$$\bar{\phi} \leq \frac{1-q}{3-2q},$$

whilst making the same substitution into (3.11) yields

$$\bar{\phi} \geq 1 - q.$$

Clearly, both inequalities cannot hold simultaneously. We know that $\bar{\gamma} \geq q$ must be true, so the natural question is: are there any $\bar{\gamma} > q$ such that both (3.8) and (3.11) are satisfied? Differentiating the right hand sides of (3.8) and (3.11) respectively gives:

$$\frac{\partial}{\partial\bar{\gamma}} \frac{(1-\bar{\gamma})q}{2q+\bar{\gamma}-2q\bar{\gamma}} = -\frac{q}{(2q(1-\bar{\gamma})+\bar{\gamma})^2} < 0,$$

and

$$\frac{\partial}{\partial\bar{\gamma}} \left(2 - \frac{q(1+\bar{\gamma}}{\bar{\gamma}}\right) = \frac{q}{\bar{\gamma}^2} > 0.$$

It follows that, since either (3.8) or (3.11) must be violated by $\bar{\gamma} = q$, the same inequality must also be violated when $\bar{\gamma} > q$. The following proposition obtains:

Proposition 19 *It is impossible to support a $p_g = p_y = \bar{\gamma}$ equilibrium in the differentiated A-links game.*

IV.6 Comments

The introduction of differentiated A-links has weakened the 'high quality' equilibria and strengthened their 'low quality' counterparts. Overall, A-link differentiation increases the bite of cannibalisation, and therefore makes low values of p more compelling for the search engines. The intuition for this result is straightforward: In the identical A-links case, the benefit accruing to i from a deviation from $p_i = q$ to some $p_i' > q$ is that i is guaranteed to have its site visited first, which is a necessary condition for having A_i be clicked. The cost associated with this deviation is increased cannibalisation, owing to the fact that the consumer always clicks NA_i before clicking A_i. When A-links differ, the cost of any such deviation endures, but the benefit is now reduced since i's A-link is sometimes clicked, even when i is not visited first. Similarly, when $p_i = p^{max}$, deviation in the identical A-link case yields zero profit, whilst a search engine with differentiated A-links can still attract some clicks even when deviating to $p_i < p^{max}$.

V Extension: market structure

In this section I generalise the simple model to the oligopoly and monopoly search provider cases, and perform comparative statics analysis on the existence conditions for equilibria under these circumstances. For this purpose, I return to the case in which all search engines show identical A-links.

V.1 n-search engine oligopoly

With n search engines, optimal consumer behaviour is somewhat analogous to the duopoly case. In order to simplify the analysis, I assume that, whenever the consumer is indifferent between visiting two or more search engines, he visits each with equal probability, so that the consumer treats search engines symmetrically.[7] The best response of the consumer is then given in strategy 2 below.

Strategy 2 *Label the search engines so that $p_1 \geq p_2 \geq \ldots \geq p_n$. A consumer's best response strategy maps the link qualities $\{p_1, p_2, \ldots, p_n, q\}$ into a click order $\{a_1, a_2, \ldots, a_{n+1}\}$ as follows. Begin by clicking a_1 thus:*

$$a_1 = \begin{cases} O_1 \textit{ if } p_1 > q \\ A_1 \textit{ if } p_1 < q \\ A_1 \textit{ w.p. } \lambda, O_1 \textit{ w.p. } 1-\lambda \textit{ if } p_1 = q, \end{cases}$$

Now, for $1 \leq i \leq n$, if the consumer's need was met by a_i then stop clicking (i.e. $a_{i+1} = \ldots = a_{n+1} = \emptyset$), otherwise click a_{i+1} as follows:

[7] In the simple model, this corresponds to $\alpha = 1/2$.

$$a_{i+1} = \begin{cases} A_i \textit{ if } a_i = O_i \textit{ and } p_{i+1} < \gamma, \textit{ or if } a_i = O_i \textit{ and } i = N \\ O_{i+1} \textit{ if } a_i = O_i \textit{ and } p_{i+1} > \gamma \\ A_i \textit{ w.p. } \phi, O_{i+1} \textit{ w.p. } 1-\phi \textit{ if } a_i = O_i \textit{ and } p_{i+1} = \gamma \\ O_i \textit{ if } a_i \neq O_i \textit{ and } p_i \geq s \\ \emptyset \textit{ if } p_i < s. \end{cases}$$

Strategy 2 again makes reference to γ. Suppose that the consumer is weighing the click orders $\{\ldots, O_{i-1}, O_i, A_i, O_{i+1}, \ldots\}$ and $\{\ldots, O_{i-1}, A_{i-1}, O_i, O_{i+1}, \ldots\}$, *viz.* the consumer is comparing the utility from clicking the A-link at the i^{th} search engine he visits, to that from doing so at the $i-1^{th}$ site. The respective utilities are

$$(3.12) \quad \sum_{m=1}^{i-1}\left[p_m(1-mS-ms)\prod_{k=1}^{m-1}(1-p_k)\right] + \left[p_i(1-iS-is)\prod_{k=1}^{i-1}(1-p_k)\right] + \left[q(1-iS-(i+1)s)\prod_{k=1}^{i}(1-p_k)\right] + \sum_{m=i+1}^{n}\left[p_m[1-mS-(m+1)s](1-q)\prod_{k=1}^{m-1}(1-p_k)\right] - \left[[nS+(n+1)s](1-q)\prod_{k=1}^{n}(1-p_k)\right],$$

and

$$(3.13) \quad \sum_{m=1}^{i-1}\left[p_m(1-mS-ms)\prod_{k=1}^{m-1}(1-p_k)\right] + \left[q(1-(i-1)S-is)\prod_{k=1}^{i-1}(1-p_k)\right] + \left[p_i[1-iS-(i+1)s](1-q)\prod_{k=1}^{i-1}(1-p_k)\right] + \sum_{m=i+1}^{n}\left[p_m[1-mS-(m+1)s](1-q)\prod_{k=1}^{m-1}(1-p_k)\right] - \left[[nS+(n+1)s](1-q)\prod_{k=1}^{n}(1-p_k)\right].$$

The first, fourth and fifth terms in (3.12) and (3.13) are identical so that comparison of the two amounts to comparing the second and third terms. It transpires that the two expressions are equal precisely when $p_i = \gamma$. Thus, when $p_i = \gamma$, the consumer is indifferent between the two considered click orders.

Given this consumer strategy, I am ready to examine equilibrium search engine behaviour when each search engine faces $n-1 \geq 1$ competitors. When $q \leq p^{max} < 1$, it is fairly straightforward to transfer lemmas 14 and 15 to the n-search engine case:[8]

Lemma 20 *When $p^{max} < 1$, all sub-game perfect SE-pure equilibria in the n-search engine version of the simple model are symmetric in search engine strategies.*

Lemma 21 *There is no sub-game perfect SE-pure equilibrium with $p_i < q$ for some i in the n-search engine version of the simple model.*

It again must be the case, then, that any SE-pure equilibrium has $p_i \geq q \; \forall \; i$. The existence of equilibrium 1 rests upon the requirement that the consumer clicks A-links first with a probability sufficient to ensure that no search engine can obtain a big enough increase in (expected) A-link clicks from an increase in its p to make such a deviation profitable.

When there are $n \geq 2$ search engines, the probability with which each of the search engines is visited first is given by $\frac{1}{n}$. Recalculating (3.3) using this probability reveals the following:

Equilibrium 4 *The search engine strategy $p_i = q \forall i$, and the consumer strategy detailed in strategy 2 form a sub-game perfect equilibrium of the n-search engine oligopoly game if*

$$(3.14) \quad q \geq \frac{n-1}{n-1+\lambda}.$$

It is immediately apparent from (3.14) that the condition for existence of the 'low p' equilibrium becomes less demanding as the number of competing search engines in the industry is reduced so that less competitive industries are more succeptable to quality collusion.

In contrast, it turns out that the condition for existence of the analogue of equilibrium 2 remains unchanged for any $n > 1$. This is formalised in equilibrium 5.

Equilibrium 5 *When $\gamma \geq p^{max}$, the search engine strategies $p_1 = p_2 = \ldots = p_n = p^{max}$, and the consumer behaviour detailed in strategy 2 form a sub-game perfect equilibrium of the n-search engine oligopoly game.*

[8] When $p^{max} = 1$ and $n > 2$, it is possible to sustain trivial asymmetric equilibria of the form $\gamma > 1 = p^{max} = p_1 = p_2 = \ldots = p_k > q, p_{k+1}, \ldots, p_n$, in which all search engines make zero profits and there are no A-link clicks.

What, then, of the case with $\gamma < p^{max}$? It has already been demonstrated that, when $p_i = \gamma$, the consumer is indifferent between clicking the A-link at the i^{th} search engine he visits and doing so at the $i - 1^{th}$ site. Moreover, if it is also the case that $p_{i+1} = \gamma$ then the consumer is indifferent between the click orders $\{\ldots, O_{i-1}, O_i, A_i, O_{i+1}, \ldots\}$ and $\{\ldots, O_i, O_{i+1}, A_{i+1}, O_{i+2}, \ldots\}$. Thus, if $p_i = \gamma\ \forall i$, it can be established by transitivity that the consumer is indifferent between any two click orders which do not have him click on an A-link first.

Now, let us look for an equilibrium with $p_i = \gamma \forall i$. When i deviates by setting some $p_i' > \gamma$, he is visited first with probability 1. Thus, profits from deviation to some $p_i' > \gamma$ are given by

$$\text{(3.15)}\quad \pi_i' = (1 - p_i')\phi b,$$

which is maximised when p_i' is arbitrarily close to γ. Suppose, instead, that i complies with the suggested equilibrium. Profits then are given by[9]

$$\text{(3.16)}\quad \pi_i = \left[\frac{\phi}{n}\sum_{k=1}^{n-1}(1-\gamma)^k(1-\phi)^{k-1}\right]b + \left[\frac{1}{n}(1-\gamma)^n(1-\phi)^{n-1}\right]b.$$

Note that (3.15) approaches zero with ϕ, whilst the second term in (3.16) remains positive (since $\gamma < p^{max}$ implies $\gamma < 1$). Thus a $p = \gamma$ equilibrium can always be sustained by setting a low enough ϕ.

Proposition 22 *There exists a $\hat{\phi}$ such that, for $\phi < \hat{\phi}$, (3.15) is less than (3.16). Strategy 2 with any such ϕ sustains a sub-game perfect equilibrium in which $p_i = \gamma \forall i$.*

V.2 Monopoly search provider

I now extend the simple duopoly model of section II to examine the equilibrium considerations induced by a monopolist search engine. In the results above,

[9] The probability that site i is the k^{th} site to be visited is $1/n$. Conditional on being the k^{th} site (for $k < n$), i receives a profit of b if and only if all of the first $k - 1$ O-links, as well as i's own O-link fail to match the consumer's need (each link failing with probability $1 - \gamma$), and if the consumer chooses to switch at the first $k - 1$ sites and stick at i's site. This gives rise to the first term in (3.16). The second term comes from the fact that if i is the n^{th} site to be visited, and the consumer has switched at the first $n - 1$ sites, then the consumer sticks at i with probability 1 since there are no more search engines to switch to.

competition for visits prompts search engines to cannibalise their revenues from A-link clicks. When this competition is taken away so is the incentive to provide O-links of a high quality. Only the cannibalisation effect remains, and thus monopolists will generally set a low quality.[10]

Proposition 23 *With a monopoly search provider, any equilibrium must involve some $p \leq q$. If $\lambda < 1$ then $p < q$ must hold for equilibrium to be sustained.*

V.3 Comment

One of the issues in the regulation of Internet search has been to fully understand the effects of reduced competition in the industry. This is especially true since advertisers can easily substitute to other mediums, and the ability of search engines to exercise any market power over them is thus limited. However, the above results demonstrate that reduced competition may also spill-over into the quality of search services enjoyed by consumers: In the oligopoly case, search engines may have an incentive to consolidate or collude since this can create new equilibria with higher total industry profits, but lower O-link quality. This is even more true if the consolidated firm is a monopolist—in which case the industry profits are maximised and quality is particularly low. This may prove to be an important consideration if the consumer's search experience is part of the regulator's objective.

VI Extension: Heterogeneous Visit Costs

The results from section III are stark, but the Bertrand-like assumption of all-or-nothing profits is equally so. One question that naturally arises is whether or not the above results remain valid when there is some degree of continuity to the demand faced by each search engine. To this end, I invoke a standard Hotelling (1929) linear city type model in which consumers have heterogeneous visit costs distributed along a segment of the real line. As well as being of interest in its own right, this serves as a robustness check for the model of section II.

More concretely, suppose that a mass 1 of consumers are uniformly distributed along a line of unit length[11], with g located at point 0, and y at point 1. A consumer located at point x must pay a cost $S_g(x) = tx$ to visit g and $S_y(x) = t(1-x)$

[10] If $S + s > q$ for a large enough proportion of consumers then the monopolist may still wish to set a $p > q$ in order to induce those consumers to search.

[11] The main results of this section do not critically depend upon these distributional assumptions.

to visit y, where $t \in (0,1)$ is a parameter that scales costs. The model is otherwise as described in section II. One possible intuition for this model is the observation that some consumers have a search engine bookmarked, set as their browser's homepage, installed as a browser tool bar, or else use a particular search engine provider for other services such as email and calendaring—making the use of such a search engine relatively less costly vis-à-vis its competitors. Alternatively, consumers may just be more used to reading the results at one particular search engine, and hence find searching there easier.

Whereas in section III I maintained the condition that $1 > q > S + s$, I now now use the slightly stronger assumption that $1 > q > t + s$. Moreover, in order to simplify the analysis, I focus on the limiting case of s positive and tending to zero.[12] In the first instance, this rules out the potentially complicated switching behaviour (which, in general, will now be location dependent). The assumption of small s also admits the tractability of the analytic solutions that follow.

Given the structure of visit costs, consumers now determine which search engine to visit first in consideration not only of p_g and p_y, but also their own personal location, x. In particular, the utility from any click order involving clicking both links at g first is decreasing in x, whilst that from clicking both links at y first is increasing in x. Thus, for some x^*, the best response for a consumer is to visit g first whenever $x \leq x^*$ (otherwise begin at y), and once there to click on the O-link first if its quality exceeds that of the A-link, and *vice-versa*. As in section III, consumers stop clicking if their need is satisfied, or if the quality of the O-link at the second site is less than S. More formally, let

$$e(x) \equiv \begin{cases} g \text{ if } x \leq x^* \\ y \text{ if } x > x^* \end{cases}, -e(x) = \{g, y\} - \{e(x)\}.$$

Strategy 3 *A consumer's best response strategy, then, maps his position and the qualities into a click order* $\{a_1, a_2, a_3\}$ *thus:*

$$a_1 = \begin{cases} A_{e(x)} \textit{ if } q > p_{e(x)} \\ O_{e(x)} \textit{ if } q < p_{e(x)} \\ A_{e(x)} \textit{ with probability } \lambda, O_{e(x)} \textit{ with probability } 1-\lambda \textit{ if } q = p_{e(x)}, \end{cases}$$

$$a_2 = \begin{cases} A_{e(x)} \textit{ if } a_1 = O_{e(x)} \textit{ and consumer's need remains unmet} \\ O_{e(x)} \textit{ if } a_1 = A_{e(x)} \textit{ and consumer's need remains unmet} \\ \emptyset \textit{ if consumer's need was met by } a_1, \end{cases}$$

[12] This is equivalent to assuming that consumers prefer to click the higher quality of the two links at a site first, but pay no cost to do so.

and

$$a_3 = \begin{cases} O_{-e(x)} \text{ if } p_{-e(x)} \geq S_{-e(x)}(x) \text{ and consumer's need remains unmet} \\ \emptyset \text{ if } p_{-e(x)} < S_{-e(x)}(x) \text{ or if consumer's need was met by } a_1 \text{ or } a_2. \end{cases}$$

An immediate question is whether it is ever optimal for search engine i to choose some $p_i < t$—thus inducing some consumers to click only the two links at its rival. In fact, in a result analogous to lemma 15, I am able to show that, when consumers behave as defined in strategy 3, no $p_i < q$ is ever optimal for search engine i.

Proposition 24 *When consumers play a best response, choosing some $p_i < q$ with positive probability is never optimal for search engine $i \in \{g, y\}$.*

Intuitively, increasing p_i causes more consumers to visit i first. As long as p_i remains below q, the proportion of consumers visiting i who click A_i does not decrease (there is no cannibalisation effect). Thus, when $p_i < q$, i can always do better by finding some $p_i' \in (p_i, q)$, and the existence of such a p_i' is assured since $[0, q)$ has no largest element.

Although proposition 24 exists in the same spirit as does proposition 15, the former is a stronger result: In addition to ruling out equilibria with some $p_i < q$, proposition 24 demonstrates that—conditional on the consumer behaving rationally—any $p_i < q$ is strictly dominated for search engine i.

Given that consumers will always be induced to click all three links (so long as their need remains unsatisfied), one can use utility functions implied by $p_g, p_y \geq q$ to identify the location of the consumer indifferent between first visiting g and y as approaching

(3.17) $x_I^* = \dfrac{p_g + q - p_g q}{p_g + p_y + 2q - p_g q - p_y q}$

in the $s \to 0$ limit.

Thus, for $p_g, p_y > t$, given that consumers do not switch, and given that they click on links at the first site in declining order of quality, g's profits are given by

(3.18) $\pi_g = x_I^*(1 - p_g)b$

when $p_g > q$,

(3.19) $\pi_g = x_I^* b$

when $t \leq p_g < q$, and

(3.20) $\pi_g = x_I^* [(1-\lambda)(1-p_g)b + \lambda b]$

when $p_g = q$. The analogous profits for y are $\pi_y = (1 - x_I^*)(1 - p_y)b$, $\pi_y = (1 - x_I^*)b$, and $\pi_y = (1 - x_I^*) [(1-\lambda)(1-p_y)b + \lambda b]$ respectively.

From proposition 24, there can clearly be no equilibrium with some $p_i < q$. I now turn my attention to the possibility of equilibria with $p_g, p_y \geq q$, focusing again on SE-pure equilibria. It is immediately apparent that any equilibrium with some $p_i = q$ must have $\lambda = 1$, otherwise (3.19) is strictly greater than (3.20) for p_g less than, but sufficiently close to q (and likewise for y's profit functions). I am now able to establish the following equilibria:

Equilibrium 6 *There exists a* $\underline{q} = 0.0836$ *such that the search engine strategies* $p_g = p_y = q$*, and the consumer strategy detailed in strategy 3 with* $\lambda = 1$ *form a sub-game perfect equilibrium of the heterogeneous visit costs game whenever* $q \geq \underline{q}$*.*

Equilibrium 7 *There exists a* $\bar{q} = 0.1042$ *such that, whenever* $q \leq \bar{q}$*, the search engine strategies*

(3.21) $p_g = p_y = \dfrac{1-3q}{3-3q} \; (> q),$

and the consumer strategy detailed in strategy 3 form a sub-game perfect equilibrium of the heterogeneous visit costs game. Moreover, this is the only SE-pure equilibrium with $p_g, p_y > q$*.*

These two classes of equilibria are shown diagrammatically in figure 3.5. The discontinuity in g's profits at $p_g = q$ originates from the fact that, when p_g is increased from $p_g = q - \epsilon$ to $p_g' = q + \epsilon$ (ϵ small), the entire mass of consumers who click g's A-link switch from clicking A_g first to clicking O_g first so that g's profits jump discontinuously from (3.19) to (3.18). The profit function also has a kink at $p_g = \hat{p}$, where

$$\hat{p} = \frac{p_y q - p_y - 2q + t(1-q) + \sqrt{4(1-q)tq + (p_y q - p_y - 2q + t(1-q))^2}}{2(1-q)}.$$

For very low values of p_g ($< \hat{p}$), consumers close to g use the click order $\{A_g, O_g, O_y\}$, whilst all others click both links at y but never visit g. When p_g is increased slightly, the utility from the former click order increases, whilst that

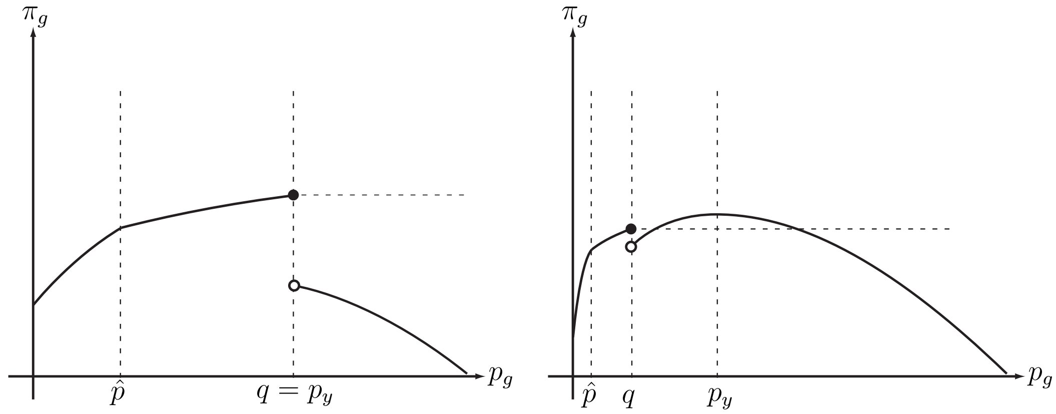

(a) $q \geq \underline{q}$, $p_g = p_y = q$, $\lambda = 1$ (b) $q \leq \bar{q}$, $p_g = p_y = (1-3q)(3-3q)^{-1} > q$

FIGURE 3.5 Two classes of equilibria in the heterogeneous visit costs game.

of the latter does not—consumers who never click O_g do not gain from an increase in its quality. It is thus fairly easy for g to attract the marginal consumer, and the market share of g thus increases rapidly in p_g below $\hat{p}$.

Once p_g exceeds $\hat{p}$, however, some consumers, having clicked both links at y, subsequently find it worth while to visit g and click O_g. The relevant indifference point for determining g's demand is now the point at which the consumer is indifferent between this behaviour and the click order $\{A_g, O_g, O_y\}$. Demand, then, becomes less responsive to changes in p_g above $\hat{p}$, since any increase in p_g not only increases the utility of visiting g first, but also that of the relevant alternative—namely starting the search at y.

Note that equilibrium 6 necessarily has $\lambda = 1$, which implies that the search engines do not incur any loss due to cannibalisation in this equilibrium. Cannibalisation, however, continues to play an important role in the form of this equilibrium by deterring search engines from setting a high quality—even though quality provision is costless.

Further insight into the mechanics of these equilibria may be obtained by examining the reaction functions that give rise to them. Taking a first order condition reveals the choice of p_g that maximises (3.18) to be

(3.22)
$$p_g = \frac{2q^2 - 2q - p_y\,(q-1)^2 + \sqrt{(q-1)\left(p_y^2(1-q)^2 + q(1-q) + p_y\,(1+q-2q^2)\right)}}{(q-1)^2}.$$

A symmetric function obtains for y. Provided this function evaluates to some

$p_g > q$, which is the case whenever

$$p_y > \frac{q - 3q^2}{(1-q)(2q-1)},$$

(3.22) represents the p_g that maximises (3.18).[13] Alternatively, g can choose $p_g = q$, in which case, profits are given by (3.20). Thus g's complete reaction function is formed by choosing between these two values of p_g to maximise profits. Such composite reaction functions are shown in figure 3.6 for various values of q, and for $\lambda = 1$ (recall that $\lambda = 1$ is necessary to sustain a $p_g = p_y = q$ equilibrium).

Firstly, there exists a $q_0 = 0.0573$ such that for any $q \leq q_0$, g's best response is to select some $p_g > q$, and play in accordance with (3.22)—irrespective of the value of p_y. In equilibrium 7, it was shown that (3.21), which lies in the range $[0, 1]$ whenever $q \leq 1/3$, is a fixed point of (3.22). Symmetry of the reaction functions ensures that this is a point of intersection between the two.

Increasing p_g above q has both a benefit and a cost for g. The benefit arises when the higher p attracts more consumers to visit g first, so that x^* increases. The manifestation of the cost is that consumers are induced to click O_g first, and thus g's A-link is clicked with probability $1 - p_g$, rather than with probability 1. When q is increased, so is the smallest p_g such that $p_g > q$. Thus, each extra consumer attracted by the higher p_g clicks on the A-link with an ever decreasing probability. As q becomes particularly high, the number of additional visitors necessary to compensate g for this fall soon exceeds the actual increase in x^* brought about by the original increase in p_g.

Thus, as q rises above q_0, it becomes optimal for g to respond to low values of p_y with $p_g = q$, rather than to use (3.22), and likewise for y to play $p_y = q$ when p_g is low. Moreover, as q is increased further, the size of the interval over which search engines wish to play in this manner increases, so that, by the time $q \geq \underline{q}$, $p_g = p_y = q$ is a mutual best response (provided that $\lambda = 1$) giving rise to a new equilibrium (see figure 3.6(c)). For yet higher values of q this effect is so strong that the $p_g, p_y > q$ equilibrium is undermined (see figure 3.6(e)) so that $p_g = p_y = q$ is the only remaining equilibrium behaviour for the search engines.

Reaction functions depend only on q, and are symmetric and (weakly) increasing *viz.* p_g and p_y are strategic complements. This implies that there are no asymmetric equilibria.

[13] If (3.22) yields a $p_g \leq q$ then, since (3.18) is concave, profits are decreasing in p_g everywhere above q.

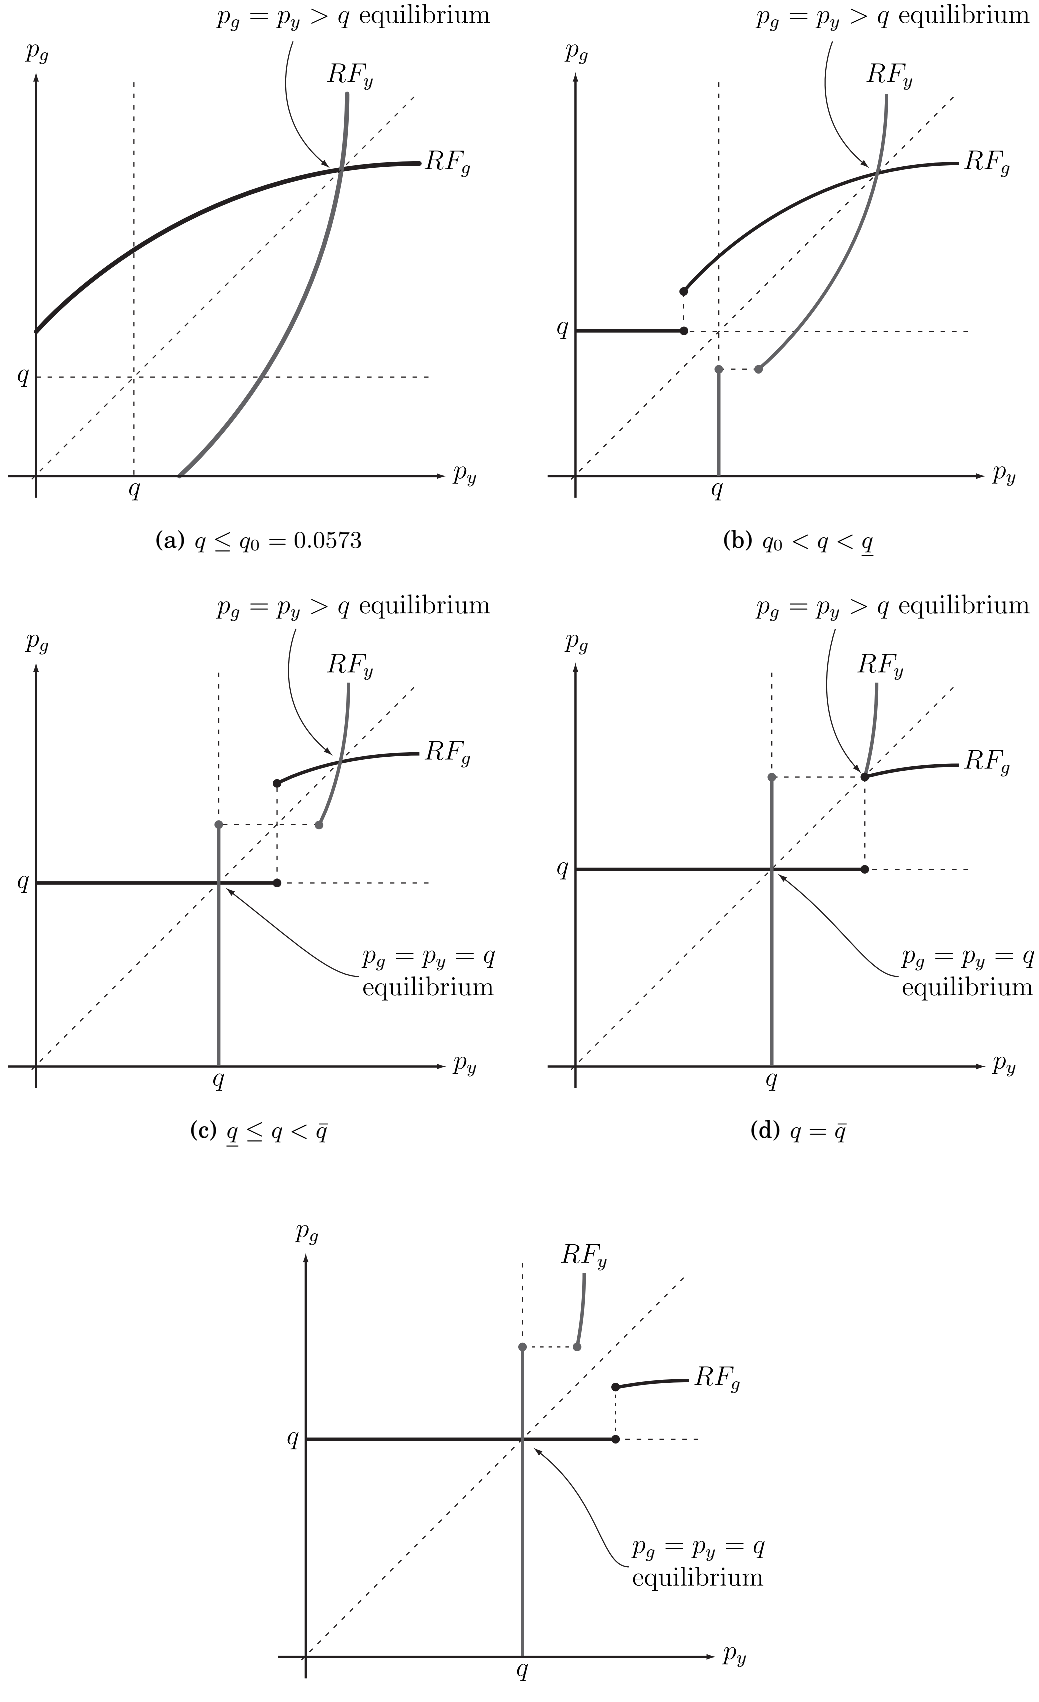

(a) $q \leq q_0 = 0.0573$

(b) $q_0 < q < \underline{q}$

(c) $\underline{q} \leq q < \bar{q}$

(d) $q = \bar{q}$

(e) $q > \bar{q}$

FIGURE 3.6 Reaction functions and equilibrium points with $\lambda = 1$, for various values of q.

Proposition 25 *There are no SE-pure equilibria with asymmetric search engine strategies in the heterogeneous visit costs game.*

Taken together, proposition 24, proposition 25, and equilibrium 7 imply that equilibria 6 and 7 are the only sub-game perfect SE-pure equilibria of the heterogeneous visit costs game.

VII Conclusion

In this chapter I have examined equilibrium behaviour in a simple model of the Internet search market. If search engines compete on result quality and consumers select search engines according to link relevance—as evidenced by Gandal (2001), equilibrium quality competition is strong in the sense that organic search results are at least as good as profit-yielding advertising links. Search engines provide such high quality links in an attempt to attract consumers who may stay to click on advertisements. It is in this fashion that search engines 'cannibalise' their own revenue streams, since the organic links that a search engine provides compete for clicks with its own advertisement links. In fact, unless the quality of advertising links is high (in the sense of equation (3.3)), and the number of competing firms small, the only equilibria in pure strategies for the search engines involve non-advertising link qualities that strictly exceed the quality of the advertising links.

The core results of the model are robust to a number of extensions. When search engines show differentiated advertisements, the incentive to compete for consumers is weakened, and cannibalisation has more 'bite' in the sense that low quality equilibria are more easily sustained. I have found that reductions in the strength of competition in the search industry may create new equilibria with lower link quality, which may be an important consideration for regulators. Introducing heterogeneity in visit costs also moderates the incentive to compete for consumer visits since search engines have a degree of monopoly power over consumers located nearby. These results notwithstanding, I again find that non-advertising links must be at least as good as (and, for low advertising link quality, strictly better than) their advertising counterparts in equilibrium so long as a search engine faces at least one competitor.

A Omitted Proofs

A.1 Proofs from Section III

Proof of Lemma 13. In equilibrium, the consumer uses a strategy of the form given in strategy 1. For concreteness, suppose that $p_i \leq p_{-i}$. If $p_i, p_{-i} < p^{max}$ then at least one search engine is visited with probability less than 1, and can profit by setting $p = p^{max}$, which results in that search engine being visited and receiving an A-link click with probability 1. Thus, equilibrium requires that at least one search engine sets $p = p^{max}$ with probability 1.

If $p_i < p_{-i} = p^{max}$ then i makes zero profits, and (since $\alpha \in (0,1)$) can make positive expected profits by deviating to $p_i = p^{max}$. It follows that playing any $p_i < p^{max}$ with positive probability is not consistent with equilibrium. ■

Proof of Lemma 14. In equilibrium, the consumers' strategies must have the form of strategy 1. I now proceed by establishing a proof by contradiction. For concreteness, suppose that $p_i > p_{-i}$. If $p_i \leq q$ then $-i$ makes zero profits, but can make positive profits by setting $p'_{-i} \geq \max\{s, p_i\}$. When $p_i > q$ and $p_{-i} < \gamma$, i receives a profit of $(1-p_i)b$—which is decreasing in p_i—and therefore prefers to reduce p_i slightly. If $p_{-i} = \gamma$ then i's profits are $\phi(1-p_i)b$, so that i again prefers to decrease p_i. Finally, if $p_i > p_{-i} > \gamma$ then the consumer uses click order $\{O_i, O_{-i}, A_{-i}\}$ and i's profits are zero. There is thus a profitable deviation for i which has it set $p'_i \in (\gamma, p_{-i})$. ■

Proof of Lemma 15. By lemma 14, all SE-pure equilibria are symmetric in search engine strategies. Thus, $\min\{p_g, p_y\} < q$ implies that $p_g = p_y = p < q$. Both $i \in \{g, y\}$ must, then have their A-link clicked with probability less than 1. It follows that both g and y have a profitable deviation, namely to set some $p'_i \in (p, q)$. ■

Proof of Equilibrium 1. When $p_g = p_y = q$, the consumers are indifferent about which site they visits first, and also about which link they click first. Suppose that the consumers visit search engine g first with probability α and, having visited site i first, click on A_i first with probability λ. Search engine expected profits are then

$$(3.23)\quad \pi_i = \begin{cases} \alpha\left(\lambda + (1-\lambda)(1-p_g)\right) b \text{ if } i = g \\ (1-\alpha)\left(\lambda + (1-\lambda)(1-p_y)\right) b \text{ if } i = y. \end{cases}$$

Consider a deviation by search engine i to $p_i < p_{-i} = q$. The optimal behaviour for the consumer now involves clicking on both A_{-i} and O_{-i} before visiting i to click on O_i. Search engine i's profits are zero and the deviation is not profitable.

Suppose, instead, that i sets $p'_i > p_{-i} = q$. The optimal click-order for the consumer is now $\{O_i, A_i, O_{-i}\}$, which gives an expected profit for i of

(3.24) $\pi'_i = \left(1 - p'_i\right) b.$

Since (3.24) is decreasing in p'_i, it suffices to consider the limiting case with p'_i arbitrarily close to q. For the deviation to be non-profitable, (3.23) must be greater than (3.24) for both g and y, which gives (3.3) when q has been substituted in place of p_i and p'_i.

Since the consumer is indifferent about click order when $p_g = p_y = q$, any λ constitute a best response so that the proposed strategies form an equilibrium. Moreover, since strategy 1 details a best response for any search engine actions, the equilibrium is sub-game perfect. ■

Proof of Equilibrium 2. Values of p_i greater than p^{max} are not possible. Consider a deviation in which i sets $p_i < p_{-i} = p^{max}$. From strategy 1, $p_i < \min\{p_{-i}, \gamma\}$ implies that the consumer never clicks A_i, and i's profits are thus zero.

Noting that strategy 1 details a best response for any search engine actions completes the proof. ■

Proof of Equilibrium 3. With $p_g = p_y = \gamma > q$, the expected profit for i is

$$\pi_i = \begin{cases} [\alpha\phi(1-p_g) + (1-\alpha)(1-\phi)(1-p_y)(1-p_g)]\, b \text{ if } i = g \\ [(1-\alpha)\phi(1-p_y) + \alpha(1-\phi)(1-p_g)(1-p_y)]\, b \text{ if } i = y. \end{cases}$$

A deviation by i that has it set $p_i < p_{-i} = \gamma$ leaves it with a profit of zero since the consumer never clicks on A_i if $p_i < \min\{\gamma, p_{-i}\}$. Suppose instead that i deviates with $p'_i > p_{-i} = \gamma$. The expected pay-off for i becomes

$$\pi'_i = \phi(1 - p'_i)b.$$

Since this pay-off is decreasing in p'_i, it suffices to consider the limiting case of $p'_i = p_i = \gamma$. The deviation is not profitable so long as $\pi_i \geq \pi'_i$. Substituting γ for p'_i, p_i, and p_{-i} in this expression and rearranging yields (3.4). Thus neither search engine has a profitable deviation so long as (3.4) holds. Given that the consumer is, by definition, indifferent over all $\phi \in [0, 1]$ when $p_g = p_y = \gamma$, satisfaction of (3.4) is consistent with equilibrium. Moreover, since strategy 1 details a best response for any search engine actions, the proposal describes a sub-game perfect equilibrium. ■

A.2 Proofs from Section IV

Proof of Lemma 16. By contradiction: suppose for concreteness that g sets $p_g < q$.

If $\lambda < 1$ then it is a best response for y to set some $p_y \in \left(\max\{p_g, s, \underline{\gamma}_1\}, q\right)$, but then g would wish to deviate to some $p_g \in (p_y, q)$ to be sure of having its A-link clicked first.

If $\lambda = 1$ then $p_y = q$ is also a best response for y. With $p_g < p_y = q$, the consumer uses click order of the form $\{\cdot, \cdot, A_g, O_g\}$, whereas $p_g = p_y = q, \lambda = 1$ implies that the consumer uses click order $\{A_y, O_y, A_g, O_g\}$ with probability $(1 - \alpha)$, and $\{A_g, O_g, A_y, O_y\}$ with probability α. Thus, for any $\alpha \in (0, 1)$, $p_g < p_y = q$ is not a best response for g. ■

A.3 Proofs from Section V

Proof of Lemma 20. Define $\mathbb{I}_<$ to be the set of all i with $p_i < \gamma$, $\mathbb{I}_=$ all those i with $p_i = \gamma$, and $\mathbb{I}_>$ to be the set of all i with $p_i > \gamma$. Now, if $\mathbb{I}_= = \mathbb{I}_> = \emptyset$, there can clearly be no asymmetric equilibrium since $\arg\min_i\{p_i\}$ makes zero profit, and can make positive profit by setting $p_i' \in (\max\{\max_j\{p_j\}, s\}, \gamma)$. If $\mathbb{I}_<$ and at least one of $\mathbb{I}_=$ or $\mathbb{I}_>$ are non-empty then all $i \in \mathbb{I}_<$ make zero profit, but can make positive expected profit from $p_i' = \gamma$. Thus, there can be no asymmetric equilibria in which some i sets $p_i < \gamma$.

If $\mathbb{I}_>$ is non-empty then $\arg\max_{i\in\mathbb{I}_>}\{p_i\}$ has a profitable deviation: Firstly, if $\phi > 0$ or $\mathbb{I}_= = \emptyset$, choosing some $p_i' \in (\gamma, \min_{j\in\mathbb{I}_>}\{p_j\})$ decreases the extent of revenue cannibalisation, and guarantees a non-zero probability of an A-link click that is no less than prior to the deviation. Conversely, if $\mathbb{I}_=$ is non-empty and $\phi = 0$, then i can profit by deviation to $p_i' = \gamma$, which is the only p that can yield positive profits.

Since there can never exist an asymmetric equilibrium in which either $\mathbb{I}_<$ or $\mathbb{I}_>$ (or both) is non empty, it is clear that an SE-pure asymmetric equilibrium can never exist. ■

Proof of Lemma 21. By lemma 20, any equilibrium with some $p_i < q$ must have $p_1 = p_2 = \ldots = p_n = p < q$, the corresponding profits for each i being b/n. If some i deviates to $p_i' \in (\max\{p, s\}, q)$, it is visited first with probability 1 and makes profit b. The proposed deviation is therefore profitable. ■

Proof of Equilibrium 4. When $p_1 = p_2 = \ldots = p_n = q$, the consumers are indifferent about which site they visits first, and also about which link they click first. Suppose that the consumers visit search engine i first with probability $1/n$ and, having visited site i first, click on A_i first with probability λ. Search engine expected profits are then

$$\text{(3.25)} \quad \pi_i = \frac{1}{n}\left(\lambda + (1-\lambda)(1-p_i)\right)b$$

Consider a deviation by search engine i to $p_i < q$. The optimal behaviour for the consumer now involves clicking the A-link at a site other than i before visiting i to click on O_i. Search engine i's profits are zero and the deviation is not profitable.

Suppose, instead, that i sets $p_i' > q$. The optimal click-order for the consumer is now $\{O_i, A_i, \ldots\}$, which gives an expected profit for i of (3.24). Since this is decreasing in p_i', it suffices to consider the limiting case with p_i' arbitrarily close to q. For the deviation to be non-profitable, (3.25) must be greater than (3.24), which gives (3.14) when q has been substituted in place of p_i and p_i'.

Since the consumer is indifferent about click order when $p_i = q \forall i$, any λ constitute a best response so that the proposed strategies form an equilibrium. Moreover, since strategy 2 details a best response for any search engine actions, the equilibrium is sub-game perfect. ■

Proof of Equilibrium 5. Expected profits from compliance with the equilibrium are given by $\pi_i = (1/n)(1-p^{max})b \geq 0$. Values of p greater than p^{max} are not possible. Consider a deviation in which i sets $p_i < p^{max}$. From strategy 2, $p_i < \min\{p_1, \ldots, p_{i-1}, p_{i+1}, \ldots, p_n, \gamma\}$ implies that the consumer never clicks A_i, and i's profits are zero. Thus, i has no profitable deviation.

Noting that strategy 2 details a best response for any search engine actions completes the proof. ■

Proof of Proposition 23. The best response for a consumer is now to click the O-link first if $p > q$, to click the A-link first with probability 1 if $p < q$, and to click the A-link first with some probability λ if $p = q$. Given this behaviour, the monopoly search engine's best response is to select a p that ensures the A-link is clicked first—this implies $p < q$ (or $p \leq q$ if $\lambda = 1$). ■

A.4 Proofs from Section VI

Proof of Proposition 24. Suppose that $p_g < q$. Since s is small, consumers click on A_g if and only if they visit g first. Denote by x^* the mass of all such consumers. If p_g is a best response, then it must be the case that $x^* > 0$, since g can always induce nearby consumers to visit it first by setting a p_g close enough to p_y. Similarly, $x^* = 1$ implies that $p_y < p_g$, and in this case a symmetric argument for y establishes that p_y is not a best response.

Consider, then, the interior case with $0 < x^* < 1$. Since $p_g < q$, rational consumers always click on A_g before they do O_g. Thus g's profits are given by the mass of consumers that visit g first, x^*, multiplied by b. A consumer that clicks A_g first must either use the click order $\{A_g, O_g, O_y\}$ or else use $\{A_g, O_g, \emptyset\}$, which implies utility functions

$$U(A_g, O_g, O_y) = q(1 - tx - s) + (1 - q)p_g(1 - tx - 2s)+ \\ (1 - q)(1 - p_g)p_y(1 - t - 3s) + (1 - q)(1 - p_g)(1 - p_y)(-t - 3s), \tag{3.26}$$

and

$$U(A_g, O_g, \emptyset) = q(1 - tx - s)+ \\ (1 - q)p_g(1 - tx - 2s) + (1 - q)(1 - p_g)(-t - 2s). \tag{3.27}$$

Any rational consumer who does not find it optimal to use either of the above two click orders must visit y first. The possible click orders in use by such consumers are $\{A_y, O_y, O_g\}$; $\{A_y, O_y, \emptyset\}$; $\{O_y, A_y, O_g\}$; and $\{O_y, A_y, \emptyset\}$. These are respectively associated with the following utility functions.

$$U(A_y, O_y, O_g) = q(1 - t(1 - x) - s) + (1 - q)p_y(1 - t(1 - x) - 2s)+ \\ (1 - q)(1 - p_y)p_g(1 - t - 3s) + (1 - q)(1 - p_y)(1 - p_g)(-t - 3s), \tag{3.28}$$

$$U(A_y, O_y, \emptyset) = q(1 - t(1 - x) - s)+ \\ (1 - q)p_y(1 - t(1 - x) - 2s) + (1 - q)(1 - p_y)(-t(1 - x) - 2s). \tag{3.29}$$

$$U(O_y, A_y, O_g) = p_y(1 - t(1 - x) - s) + (1 - p_y)q(1 - t(1 - x) - 2s)+ \\ (1 - p_y)(1 - q)p_g(1 - t - 3s) + (1 - p_y)(1 - q)(1 - p_g)(-t - 3s), \tag{3.30}$$

and

$$U(O_y, A_y, \emptyset) = p_y(1 - t(1 - x) - s)+ \\ (1 - p_y)q(1 - t(1 - x) - 2s) + (1 - p_y)(1 - q)(1 - t(1 - x) - 2s). \tag{3.31}$$

Now, given that the consumer's cost for visiting g (y) is continuously increasing (decreasing) in x, the set of x for which consumers click A_g (A_y) first must be a connected interval, and must include point $x = 0$ ($x = 1$). Thus, every consumer with an $x < x^*$ must be visiting g first, and every consumer for whom $x > x^*$ must be using some click order that has him visit y first. By the continuity of x (and since utility varies continuously with x), there must exist a marginal consumer at x^* who is just indifferent between the click orders in use by those consumers at $x^* + \epsilon$ and $x^* - \epsilon$, with ϵ small. That is to say, the maximum of (3.26) and (3.27) must be equal to the maximum of (3.28), (3.29), (3.30), and (3.31) at $x = x^*$.

Consider a small increase in p_g to $p'_g \in (p_g, q)$. The derivatives of (3.26) and (3.27) with respect to p_g are positive, and are in every case both greater than those of (3.28), (3.29), (3.30), and (3.31). Thus, the increase in p_g causes the marginal consumer to strictly prefer some click order that has him click A_g first to all others: the mass of consumers clicking A_g, (and hence g's profits) is thus increased. By the continuity of p_g, there exists such a p'_g for all $p_g < q$. Choosing a $p_g < q$ with positive probability is thus not optimal.

A symmetric argument holds for y. ■

Proof of equilibrium 6. The consumer's strategy specifies a best response to any combination of p_g and p_y. It remains to show that the search engine strategies are optimal responses to one another, given this subsequent consumer behaviour. I show that there is no profitable deviation for g, and appeal to symmetry to complete the argument for y.

With $\lambda = 1$, (3.20) collapses to (3.19). Since, by proposition 24, profit from any $p_g < q$ is increasing in p_g it suffices to consider deviations to some $p_g > q$. Such a deviation yields profits for g given by (3.18). Substituting $p_y = q$ and calculating the first order condition gives g's reaction function:

$$(3.32)\quad p_g = \frac{4q^2 - 3q - q^3 + Z}{(q-1)^2},$$

where

$$Z = \sqrt{(q-1)^2 q(2 + 3q - 4q^2 + q^3)}.$$

Now, substituting (3.32) into (3.18) yields g's deviation profits thus:

$$(3.33)\quad \pi_g = \frac{\left(2q - 3q^2 + q^3 - Z\right)\left(3q^2 - q^3 - q - 1 + Z\right)}{Z(q-1)^2} b.$$

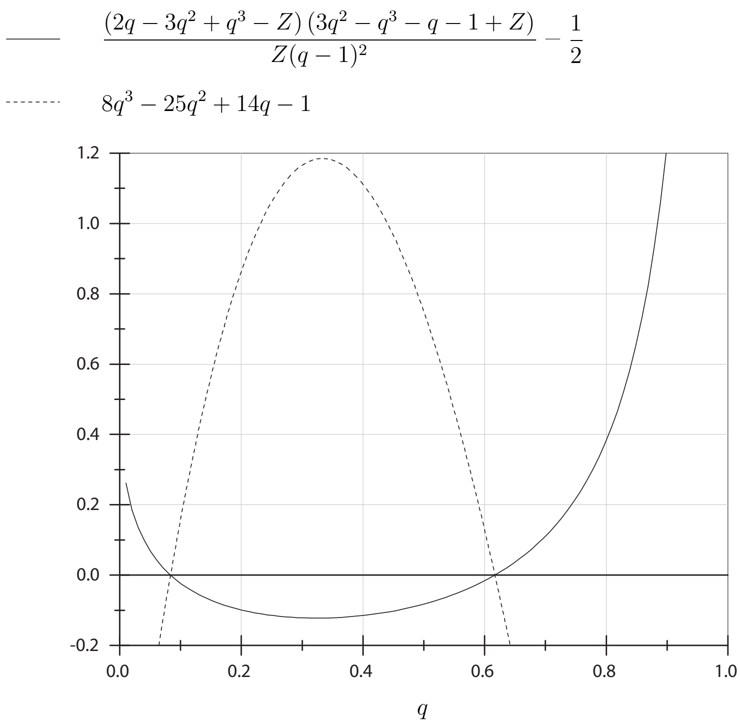

FIGURE 3.7 The q in the range $0 \leq q \leq 1$ for which deviation yields profits equal to those from compliance are given by the roots of the cubic $8q^3 - 25q^2 + 14q - 1$.

These are to be compared with the profits from compliance with the candidate equilibrium, which are given by (3.20) with $p_g = p_y = q$, and $\lambda = 1$, which gives

$$(3.34)\quad \pi_g = \frac{b}{2}.$$

Finding the values of $q \in [0,1]$ that equate (3.33) and (3.34) is equivalent to finding the roots of the following cubic (see figure 3.7):

$$8q^3 - 25q^2 + 14q - 1,$$

which can be achieved using, for example, the method of Nickalls (1993). The two solutions to this equation which lie within the $[0,1]$ interval are given by

$$(3.35)\quad q = \frac{25}{24} - \frac{17}{12}\sin\left[\frac{\pi}{6} \pm \frac{1}{3}\cos^{-1}\left(\frac{3889}{4913}\right)\right] \approx \{0.08356, 0.616981\},$$

However, for any $q > (1/3)\left(3 - \sqrt{6}\right) \approx 0.1835$, (3.32) demands a value of $p_g < q$. Moreover, twice differentiating (3.18) with respect to p_g yields the following

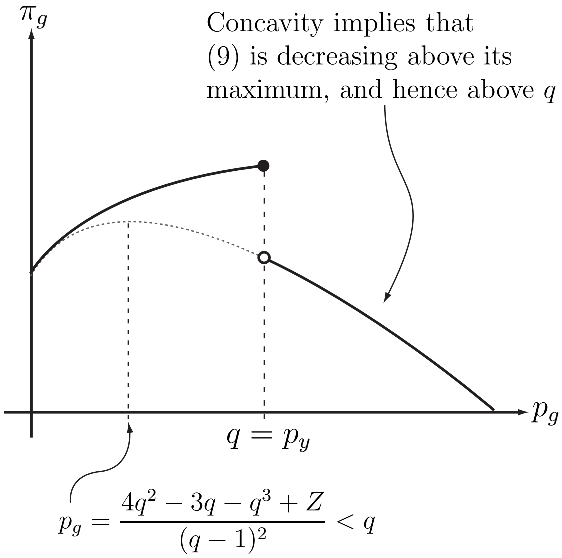

FIGURE 3.8 For $q > 1/3\left(3 - \sqrt{6}\right) \approx 0.1835$, (3.32) demands a value of $p_g < q$.

(having substituted $p_y = q$):

$$\frac{\partial^2 \pi_g}{(\partial p_g)^2} = \frac{2q\left(2 + q - 7q^2 + 5q^3 - q^4\right)}{\left(p_g(q-1) + (q-3)q\right)^3} b,$$

which is negative for $q \in (0,1)$ so that profits are concave—this implies that profits are decreasing above q when the optimal p_g is less than q. The final piece of the puzzle is to note that, for any p_y, (3.18) is less than (3.19) when both are evaluated at $p_g = q$. Taken together, these facts imply that the second root in (3.35) can be ignored, and that, for $q \geq 0.08356$, deviation to some $p_g > q$ is not profitable. This argument is is summarised in figure 3.8. ■

Proof of Equilibrium 7. The consumer's strategy specifies a best response to any combination of p_g and p_y. It remains to show that the search engine strategies are optimal responses to one another, given this subsequent consumer behaviour. I show that there is no profitable deviation for g, and appeal to symmetry to complete the argument for y.

Suppose that the proposed equilibrium is valid. $(1-3q)/(3-3q)$ is greater than q for $q < 1/3\left(3 - \sqrt{6}\right) \approx 0.1835$, so that the appropriate profit function for g is (3.18). Taking the derivative of (3.18) with respect to p_g, yields the quasi-reaction function

$$p_g^* = \frac{2q^2 - 2q - p_y\,(q-1)^2 + \sqrt{(q-1)\left(p_y^2(1-q)^2 + q(1-q) + p_y\left(1 + q - 2q^2\right)\right)}}{(q-1)^2},$$

which is valid for $p_g > q$. The corresponding function for y is symmetric. Differentiating g's quasi-reaction function with respect to p_y, and substituting $p_y = 1$ yields

$$\left.\frac{\partial p_g^*}{\partial p_y}\right|_{p_y=1} = 0.06066.$$

Now, taking the second derivative of p_g^* gives

$$\frac{\partial^2 p_g^*}{(\partial p_y)^2} = -\frac{(1-q)^4}{4\left[(1-q)^2\left((p_y)^2(1-q)^2 + q(1+q) + py(1+q-2q^2)\right)\right]^{3/2}} < 0.$$

Since the second derivative is negative, and the first derivative is positive at $p_y = 1$, the first derivative of p_g^* must be positive for all $p_y \in [0,1]$. That the quasi-reaction functions are symmetric, increasing and concave implies that there can be at most two points of intersection, and that both of these must have $p_y = p_g$.

Imposing $p_y = p_g$ for symmetry and solving the quasi-reaction function gives

$$(3.36)\quad p_g = p_y = \frac{1-3q}{3-3q},.$$

and

$$p_g = p_y = \frac{q}{q-1}.$$

The second solution is non-positive for all $0 < q < 1$. Since these are the only two points of intersection of the two $p > q$ quasi-reaction functions, the only possible equilibrium behaviour in which $p_g, p_y > q$ is given by (3.36).

By construction, when p_y plays according to (3.36), no $p_g > q$ can yield a higher profit for g than will compliance with (3.36). Moreover, by proposition 24, any deviation to $p'_g < q$ is less profitable than some $p''_g \in (p'_g, q)$. It suffices, then, to show that the limit of (3.19) as $p_g \to q$ can not be higher than the profit from compliance with the proposed equilibrium.

Substituting (3.36) into (3.18) gives profits for compliance with the candidate equilibrium thus:

$$(3.37)\quad \pi_g = \pi_y = \frac{1}{3-3q}b.$$

Substituting $p_g = q$, and $p_y = (1-3q)/(3-3q)$ into (3.19) gives an expression for the maximal deviation profits:

$$(3.38)\quad \pi_g = \frac{3q^2-6q}{3q^2-6q-1}b.$$

Equating (3.37) and (3.38) yields the cubic

$$-9q^3 + 24q^2 - 12q + 1 = 0,$$

which can, again, be solved using the method of Nickalls (1993). There are two roots that lie in the interval $[0, 1]$, namely

$$\frac{8}{9} - \frac{4}{9}\sqrt{7}\sin\left[\frac{\pi}{6} \pm \frac{1}{3}\cos^{-1}\left(\frac{241}{112\sqrt{7}}\right)\right] \approx \{0.1042, 0.5228\}.$$

For $0 < q \leq 0.1042$ profits from compliance exceed those from deviation; for $0.1042 < q < 0.5228$ deviation appears strictly profitable, and for $0.5228 \leq q$ compliance is again optimal. However, for $q \geq 1/3\left(3 - \sqrt{6}\right) \approx 0.1835$, (3.36) demands a $p_g \leq q$. Thus, (3.36) constitutes a valid equilibrium strategy only when $0 < q \leq 0.1042$. ■

Proof of Proposition 25. By proposition 24, $p_i < q$ is never consistent with equilibrium. By equilibrium 7, there are no asymmetric equilibria with $p_g, p_y > q$. It follows that, if there exists an asymmetric equilibrium, it must have $p_e = q, p_{-e} > q$. Since a condition for any equilibrium with $p_e = q$ existing is $\lambda = 1$, I suppose that this is the behaviour used by consumers.

Now, inverting the logic of equilibrium 6, a condition for g wanting to choose some $p_g > q$, when $p_y = q$ is that $q \leq 0.0836$. Demonstrating that $p_y = q$ being a best response to this optimal p_g requires some $q > 0.0836$ will thus suffice to complete the proof. Denote by $p_e^*(p_{-e})$ e's best response in the range $(q, p^{max}]$ to the specified p_{-e}, and define $\pi_e(p_g, p_y)$ as e's profits when g plays p_g and y plays p_y. Thus, solving

$$\frac{\partial \pi_g(p_g > q, q)}{\partial p_g} = 0$$

for p_g yields

$$p_g^*(q) = \frac{4q^2 - 3q - q^3 + Z}{(q-1)^2},$$

where

$$Z = \sqrt{(q-1)^2 q(2 + 3q - 4q^2 + q^3)}.$$

Substituting $p_g^*(q)$ into $p_y^*(p_g)$ gives

$$p_y^*(p_g^*(q)) = \frac{q - 2q^2 + q^3 + 8q^2 - Z + \sqrt{(1-q)\left(10q^4 - 2q^5 - 16q^3 + Z - 4qZ + 2q^2 Z\right)}}{(1-q)^2}.$$

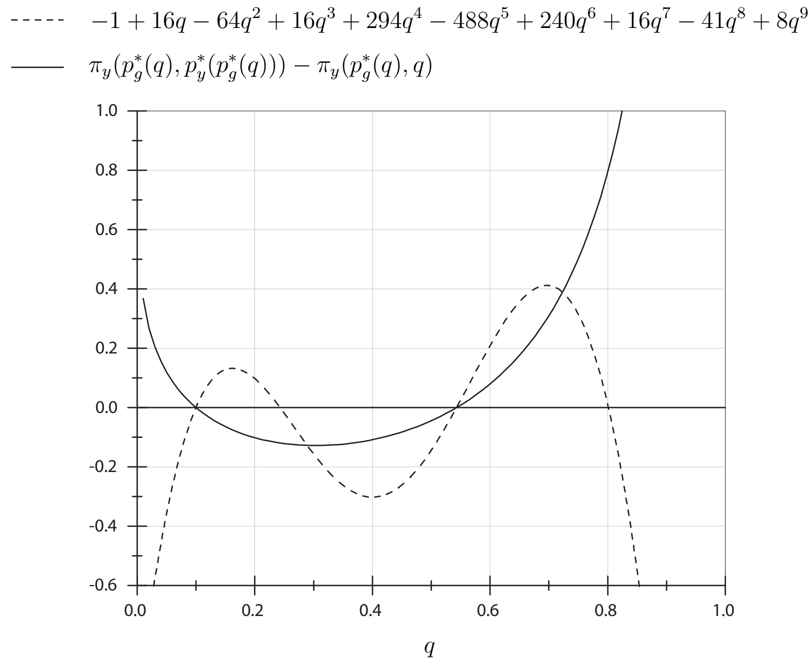

FIGURE 3.9 The q in the range $0 \leq q \leq 1$ for which $\pi_y(p_g^*(q), p_y^*(p_g^*(q))) - \pi_y(p_g^*(q), q)$ is negative are given by roots of a polynomial.

Substituting $p_y = p_y^*(p_g^*(q)), p_g^*(q)$ into $\pi_y = (1 - x_I^*)(1 - p_y)b$, and $p_y = q, p_g = p_g^*(q)$ into $\pi_y = (1 - x_I^*)b$ gives expressions for $\pi_y(p_g^*(q), p_y^*(p_g^*(q)))$ and $\pi_y(p_g^*(q), q)$. Finding the values of $q \in [0, 1]$ for which the former is less than the latter can be achieved by identifying roots of the following polynomial (see figure 3.9)

$$-1 + 16q - 64q^2 + 16q^3 + 294q^4 - 488q^5 + 240q^6 + 16q^7 - 41q^8 + 8q^9 = 0.$$

In particular, there are two roots of $\pi_y(p_g^*(q), p_y^*(p_g^*(q))) - \pi_y(p_g^*(q), q)$ for q in the interval $[0, 1]$, namely $\{0.0998, 0.5432\}$. Since $\pi_y(p_g^*(q), p_y^*(p_g^*(q))) - \pi_y(p_g^*(q), q)$ is positive when $q < 0.0998$, y strictly prefers to deviate to some $p_y > q$ in this range.

As required, it is thus proven that $p_y = q, p_g > q$ cannot simultaneously be best response strategies. A symmetric argument applies to the case of $p_g = q, p_y > q$. ■

Chapter 4

The Informativeness of On-Line Advertisements

I Introduction

This chapter investigates the role played by the prevailing fee structure in determining the extent to which information about the products that a firm sells may be transmitted via advertisements. The main result is that a fully informative equilibrium always exists when advertisers pay per-click (per-visit), and never exists when advertisers pay per-impression or per-sale. At the core of this result is the following simple idea: a per-click advertising fee can serve as a disincentive to sending uninformative advertisements of general appeal by making it costly for an advertiser to attract visits from consumers with whom it is poorly matched. By contrast, when an advertiser pays on a per-impression or per-sale basis, its incentive is to attract a visit from any consumer that will purchase with a positive probability. This incentive undermines the existence of informative equilibria, reduces the usefulness of advertisements, and damages consumer welfare, but, for some specifications of search costs, can be shown to be beneficial for publishers and society as a whole.

The advent of the mainstream Internet has had a dramatic impact on the advertising industry. Advertising expenditure is swiftly migrating towards the new medium and with this change in focus has come a similarly significant change in the terms of business for advertisers and advertisement publishers. The unique capacity of Internet publishers to monitor and track users' activity and (in some cases) even their identity has given rise to a number of novel structures for the pricing of advertisement (ad) facilities. The three classes of fee structure that I shall consider here are as follows:

1. **Pay-per-click (PPC):** Advertisers pay each time a consumer clicks on their ad. PPC pricing has been successfully used by search engines such as Google and Yahoo to sell sponsored search results.

2. **Pay-per-impression (PPI):** Advertisers are charged each time their ad is shown to a consumer, regardless of whether that consumer takes any further action. This pricing method was extensively used for the display "banner" ads that dominated early on-line advertising. Its analogue is also common when selling advertisements in more traditional media.

3. **Pay-per-sale (PPS):** An advertiser must pay for each consumer that clicks on its advertisement and subsequently makes a purchase from it.

Although the focus of this chapter is on-line advertising, the work is of broader interest insofar as the above fee structures (or their analogues) are used in traditional media advertising.

To better understand the mechanics of this chapter, consider the simple example of a consumer who sees a digital camera advertisement and may be interested in buying either a high-end digital SLR camera, or a consumer-level compact camera. A specialist professional photography supplier that sells only the former may truthfully advertise itself as such, in which case consumers seeking a compact camera are unlikely to visit. Alternatively, the firm could simply advertise "great deals on digital cameras", and attract consumers of both types. Most of those (visiting) consumers that are looking for a point-and-shoot camera will leave without purchasing since they and the firm are a poor match. However, it also seems plausible that a small (or even a tiny) fraction of them will learn something new about the SLR cameras on offer and purchase one from the firm irrespective of their original intentions. Under a PPC regime, the firm would be heavily penalised for this transgression—having to pay for the visit of each of those consumers that does not purchase. PPI and PPS, on the other hand, create no explicit disincentive to such attempts at deception.

Often, the law too provides limited disincentives to such behaviour. Whilst outright deception is typically prohibited, a claim such as "great deals on digital cameras" offers no obvious grounds for regulatory intervention. Moreover, enforcement becomes considerably more complicated on the Internet, where a huge proliferation of small firms broadcast advertisements that transcend international boundaries and jurisdictions. Publishers may intervene to prevent obvious abuse of their advertising resource, but again the large and shifting volumes of advertisements handled by major publishers makes perfect enforcement impractical.

Arguably, the most effective check on this behaviour are the consumers themselves. Consumers that view many advertisements each day and are 'media savvy' are likely to be intrinsically aware of the appropriate degree of trust associated with the messages contained therein. That is to say, rational consumers cannot be systematically deceived: when the environment is one that is not conducive to honest reporting, consumers are likely to place little stock in the claims that they see advertised. The result is that meaningful communication between advertiser and consumer becomes difficult in such environments. It is this intuition that rests at the core of the present chapter.

I.1 Literature

This chapter is related to several strands of literature. Nelson (1974) pioneered the study of the informational role of advertising—most notably with the idea that the very existence of an advertisement may be informative about a product's quality. A simple example of how this might be true is the fact that high-quality producers that expect repeat purchases may be able to profitably transmit advertisements, whilst low quality firms that sell only once cannot—advertisement expenditure can thus serve as a signal and therefore a dimension for separation in equilibrium. Key ideas first introduced in Nelson's work were later formalised by Schmalensee (1978) and Milgrom and Roberts (1986). Nelson also acknowledged that advertisements may not always be honest, and that firms may have an incentive to lie when consumers face a transport cost.

The literature has also treated the case in which advertisements are explicitly informative. In a highly influential paper, Grossman and Shapiro (1984) examined the problem of oligopolists with (horizontally) differentiated products whose advertisements convey information about the sender's location in a circular city. They found that advertising will typically be excessively provided from society's standpoint. More recently, on informative advertising, Meurer and Stahl (1994) study a model in which consumers, only some of whom receive an advertisement, are unsure which product offers them the best match. Advertising in Meurer and Stahl's model can be either under- or over-provided. Anderson and Renault (2006) show that informative advertising can be used by a monopolist to mitigate the 'hold-up' problem when an advertiser is able to transmit threshold information about consumers' values. All of these papers share the assumption that any informational content of an ad must be truthful.

There is a small but growing literature on the efficacy of various advertising fee structures in on-line advertising—with the primary focus having been on vertically differentiated products. Dellarocas and Viswanathan (2008) show that pay for performance fee structures tend to favour low quality firms and yield lower surpluses for all vis-à-vis flat-rate pay-per-impression regimes. Meanwhile, Sundararajan (2003) shows that whilst performance-based pricing of digital marketing is generally profit increasing, it can not screen out low quality advertisers. Agarwal, Athey, and Yang (2009) develop a model to explore the problem of a publisher that must aggregate bids across multiple actions in a per-per-action environment. Differences in the estimates of action probabilities used by the publisher and advertiser create an incentive for the advertiser to skew its bidding in favour of those actions whose probability has been overestimated. This incentive can lead to reversals in the ranking of firm bids and lower publisher profit.

More closely related to this chapter, as an extension of their model of position auctions, Athey and Ellison (2008) find that pay-per-click auctions are less admissive to the obfuscation of advertisements than pay-per-action in a model where each firm has some idiosyncratic probability of satisfying a given consumer—so that advertisers are essentially vertically differentiated. This supports the core results of the model presented below, and also demonstrate that the intuition is robust to environments with ordered lists of advertisers—a matter of practical importance. In contrast to Athey and Ellison's model, I relax the assumption of (horizontal) firm symmetry and fully develop the role of consumer beliefs in equilibrium formation when advertisers are both horizontally and (weakly) vertically differentiated. I also consider the role of price choice by advertisers, which exposes important welfare differences between fee structures and has a role in firm selection under pay-per-sale. Finally, in a related literature, Ellison and Ellison (2009) and Ellison and Wolitzky (2008) discuss firms' incentives to obfuscate consumer search. Intuitively, if firms artificially increase the marginal cost to consumers of an additional firm visit, consumers are induced to search less, and price competition is thus weakened.

II Model

The idea here is to build as simple a model as possible whilst capturing the key stages of on-line product search: (i) submitting a query to a search provider or visiting a publisher, (ii) finding an advertisement—typically consisting of a short description—from which an initial assessment of relevance must be formulated, (iii) visiting a link that appears to be relevant in order to obtain more detailed information, and (iv) (potentially) making a purchase.

There are n *firms* that produce (at zero cost) a collection of goods from the set of goods, $\{A, B\}$. In particular, let $\theta_j \in \{\{A\}, \{B\}, \{A, B\}\}$ denote the (privately known) set of goods sold by firm j. For notational convenience, I shall omit the set notation and write $\theta_j \in \{A, B, AB\}$. I assume that each firm is of type A with probability $\gamma/2$, of type B with probability $\gamma/2$, and of type AB with probability $1-\gamma$, and that these probabilities are publicly known. In order to sell to the consumers, a firm must occupy an advertising opportunity that is controlled by a monopolist *publisher*. There is one such advertising opportunity, and firms must therefore bid in a second price, sealed-bid auction for the right to become the unique *advertiser*. The price reached by the auction determines the advertising fee that the advertiser must pay, with the fee being paid per-impression, per-click or per-sale—depending on the fee structure, $\phi \in \{I, C, S\}$, that is in use. I assume that ϕ is public knowledge.

There is a unit mass of *consumers* identified by their type $t_i \in \{A, B\}$, with each type having equal frequency. Each consumer's type represents their ideal product from a set of products, $\{A, B\}$. A consumer is perfectly matched with an advertiser when $t_i \in \theta_j$, and 'imperfectly matched' if $t_i \notin \theta_j$. More specifically, this is a search good model in which the consumer learns about his value by visiting firms. Upon arriving at firm j and finding $t_i \in \theta_j$, the consumer discovers himself to be matched with j and learns his value, v_i, which is uniformly distributed on $[0, 1]$, as well as j's price, p_j. If i finds that $t_i \notin \theta_j$ then i and j are unmatched; however, there is a positive probability $\alpha \in (0, 1)$ that i discovers A and B to be perfect substitutes, in which case i again learns his value, v_i, and the firm's price. With probability $(1-\alpha)$, the consumer finds that an unmatched firm offers him zero value. When consumer i visits the advertiser, denote by $\hat{v}_i \in \{0, v_i\}$ his realised value there.

I assume that the consumer's search or browsing activity publicly reveals an interest in $\{A, B\}$, but not the specific value of t_i. When a consumer is using a search engine, this may be because the consumer's search term is too general.[1] Otherwise, $\{A, B\}$ may be indirectly revealed through browsing activity—say because the consumer visits a website about $\{A, B\}$.

The advertiser must choose a price, p. It can also costlessly transmit a publicly-observable signal (advertisement), $m(\theta_j) \in \{A, B, AB\}$, to the consumers. Visiting the advertiser in order to obtain full product specifications and pricing information is costly—specifically, a cost of $s_i \geq 0$ is incurred by consumer i upon

[1] I am agnostic on why it might be that consumers do not fully specify their preferences in their search phrase. In fact, general phrases are often more highly valued than specific ones in on-line ad auctions.

visiting. Denote by $F(\cdot)$ the CDF of s_i, with density $F' = f > 0$ everywhere in its support, $[0, 1/8]$.[2] The main question addressed in this chapter is: to what extent can we expect m to be informative about θ_j?

In summary, the game proceeds as follows: firstly, firms submit bids for the ad auction—the winning firm being designated as the advertiser. The advertiser then chooses a p and an m, with the latter being transmitted to the consumers. Consumers decide whether to visit the advertiser or not, and those that do visit decide whether or not to purchase. Lastly, consumer surplus, and firm and publisher profits are realised. I look for perfect Bayesian equilibria that are in pure strategies.

III Preliminaries

In a perfect Bayesian equilibrium, the consumers use Bayes' rule to form a posterior belief about the true type of the advertiser given m and the equilibrium strategy profile being played. In particular, let $\mu(\theta|\phi, m)$ be the posterior probability distribution over θ when the consumer observes ϕ, m. Given these beliefs, the consumer's optimal strategy is to visit if and only if

$$E\left(\hat{v}_i - p|\mu(\cdot), t_i\right) \geq s_i,$$

and (conditional on visiting) to purchase when $\hat{v}_i \geq p$.

Nesting the three fee structures together, the advertiser's profit then has the general form

$$\pi^a\left(\theta, m, p, \phi, \mu(\cdot)\right) = \lambda^1\left(\theta, \lambda^2(\cdot)\right)\left[(1-p)(p-b_S)\right] - \lambda^2\left(m, \phi, \mu(\cdot)\right)b_C - b_I, \tag{4.1}$$

where b_S, b_C, and b_I denote the fee per-sale, per-click and per-impression respectively. I shall refer to $\lambda^1(\cdot)$ as the advertiser's potential demand. For any given profile of player strategies, this is the mass of consumers that visit and find $\hat{v}_i > 0$. This may not be the same as the number of visitors to the advertiser—which I denote by $\lambda^2(\cdot)$. Note that so long as prices are not advertised, they do not enter as arguments of $\lambda^2(\cdot)$. The first derivative of (4.1) with respect to p is

$$\frac{\partial \pi^a}{\partial p} = \lambda^1(\cdot)\left(1 - 2p + b_S\right).$$

This implies that the optimal price for the advertiser is the standard monopoly

[2]It will turn out that consumers with search costs greater that $1/8$ are never willing to search, and I therefore normalise the population so that such consumers are excluded.

price:

(4.2) $p^* \equiv \frac{1+b_S}{2},$

and consumers will rationally anticipate this as being the price charged. Given this price, the gross surplus accruing to a unit mass of matched consumers can be calculated in the usual fashion by integrating under the demand curve implied by the distribution of v_i:

$$GCS = \int_0^{\frac{1-b_S}{2}} \left(1 - q - \frac{1+b_S}{2}\right) dq = \frac{1}{8}(1-b_S)^2.$$

IV A FULLY-INFORMATIVE EQUILIBRIUM ALWAYS EXISTS UNDER PAY-PER-CLICK

In this section I constrain the publisher to use $\phi = C$—that is $b_I = b_S = 0$ and $b_C = b^{(n-1)}$, where $b^{(n-1)}$ denotes the second highest bid from the ad auction—and look for equilibria in the resulting game involving the firms and consumers. I begin with the matter of whether or not some information transmission can take place in a PPC regime. As it turns out, *full* information transmission is always possible. That is to say, the PPC environment can always support a fully-separating equilibrium in which each firm type sends a unique message.

To see that this is the case, start by combining (4.1) and (4.2), and imposing $b_I = b_S = 0$ to give PPC advertiser profits:

(4.3) $\pi^a(\theta, \cdot, p^*, C, \cdot) = \frac{\lambda^1(\cdot)}{4} - \lambda^2(\cdot) b_C.$

Setting this expression greater than or equal to zero and solving for b_C reveals the maximum that each firm is prepared to pay in order to occupy the advertising opportunity. Standard second-price auction arguments dictate that it is weakly dominant for firms to bid up to this level during the auction phase of the game. Thus, the optimal bid is given by

(4.4) $b(\theta_j, C) = \frac{\lambda^1(\cdot)}{4\lambda^2(\cdot)}.$

Equilibria in which each firm type sends a distinct signal, $m^*(\theta)$ are essentially equivalent to one another—the most intuitive case is 'truth-telling', in which $m^*(\theta) = \theta$. Consumers with $s_i \leq \alpha/8$ always find it optimal to search. Since a truth-telling equilibrium fully reveals a firm's type, consumers with $s_i > \alpha/8$

click if and only if $t_i \in (m^*)^{-1}(m(\theta))$, where $(m^*)^{-1}$ is the inverse of $m^*(\cdot)$. Thus, when consumers expect honest reporting, sending a single-product-firm message generates

(4.5) $$\lambda^2\left(m^*(A), C, \mu^*\left(m^*(A)\right)\right) = \lambda^2\left(m^*(B), C, \mu^*\left(m^*(B)\right)\right) = \frac{1}{2}\left[1 + F\left(\frac{\alpha}{8}\right)\right]$$

clicks, where $\mu^*(m)$ is the posterior belief that places probability one on the advertiser being of type $(m^*)^{-1}(m)$. A proportion α of the unmatched consumers that click realise a positive $\hat{v}_i$. Thus, single-product advertisers that report honestly have

(4.6) $$\lambda^1\left(\theta, \lambda^2(\cdot)\right) = \frac{1}{2}\left[1 + \alpha F\left(\frac{\alpha}{8}\right)\right],$$

for $\theta \in \{A, B\}$. This in-turn implies optimal single-product firm bids of

$$b(A, C) = b(B, C) = \frac{1 + \alpha F\left(\frac{\alpha}{8}\right)}{4\left[1 + F\left(\frac{\alpha}{8}\right)\right]}.$$

Note that this bid is decreasing in $F(\alpha/8)$. When $F(\alpha/8)$ is large there are many consumers that can search so cheaply that they will visit even when they know the firm to be a poor match. The greater is the number of unmatched consumers that will visit a firm, the lower is that firm's conversion rate and hence its value per-click.

Under μ^*, an advertiser sending $m = m^*(AB)$ receives clicks from all consumers so that

$$\lambda^2(m^*(AB), C, \mu^*\left(m^*(AB)\right)) = 1.$$

If the advertiser is indeed multi-product then it matches with all of its visitors, implying $\lambda^1(AB, \lambda^2(\cdot)) = 1$, and thus

$$b(AB, C) = \frac{1}{4}.$$

Note that the optimal bids exhibit the (socially) desirable property that a multi-product firm wins the auction whenever one or more such firms participate, which happens with probability $1 - \gamma^n$.

Now, it is immediately evident that a type B advertiser never wishes to imitate a type A firm, and vice-versa—this would entail giving up a mass of matched consumers to gain an equal mass of unmatched consumers. It is also clear that a multi-product advertiser never wishes to send a single-product signal. It remains to be demonstrated that single-product advertisers have no incentive to transmit

$m = m^*(AB)$. Honest single-product advertisers attract all of their matched consumers, and continue to do so when sending $m = m^*(AB)$, so that the effect of imitating a multi-product firm is to attract an additional mass equal to

$$\frac{1}{2}\left[1 - F\left(\frac{\alpha}{8}\right)\right]$$

of *unmatched* consumers. Thus, the resulting net addition to advertiser profit is

$$\Delta\pi^a = \frac{1}{2}\left[1 - F\left(\frac{\alpha}{8}\right)\right]\left(\frac{\alpha}{4} - b^{(n-1)}\right).$$

It follows that a necessary condition for the deviation to be profitable is

$$\text{(4.7)} \quad \frac{\alpha}{4} > \min_{\theta \in \{A,B,AB\}} \{b(\theta, C)\}.$$

Now, this inequality is clearly violated for all α by the multi-product firm's bid of $1/4$. A single-product firm's bid is decreasing in $F(\alpha/8)$, so that (4.7) is most easily satisfied when $F(\alpha/8) = 1$. Making this substitution reveals that there is no α such that (4.7) is not violated. Thus, conditional on having won the auction, no firm wishes to deviate from the informative reporting strategy and, given this, the bidding strategy is optimal. This leads us to the first main result of the chapter in proposition 26 and equilibrium 8.

Proposition 26 *A fully informative equilibrium can always be sustained under the pay-per-click regime.*

Equilibrium 8 *When the publisher is constrained to use a PPC fee structure there exists a fully informative equilibrium in which firms bid up to (4.4). The successful firm transmits a fully informative message, $m^*(\theta)$, and sets $p = p^*$. Consumers update their beliefs to μ^*, visit if (i) $t_i \in (m^*)^{-1}(m)$, or (ii) $s_i \leq \alpha/8$, and purchase (conditional on having visited) if $\hat{v}_i > p^*$.*

It is appropriate to briefly discuss the mechanism responsible for the above result. A firm's willingness to pay for a PPC advertisement is determined by the rate at which it can convert clicks into sales. Firms that wish to transmit more general advertisements in order to attract a broader range of potential customers therefore become less competitive in the ad auction, and are thus unlikely to ever have the opportunity to perpetrate such obfuscation. Since winning the ad auction is a precondition for profitability, maintaining competitiveness in the bidding process is the overriding consideration here. Once a firm has bid the per-click fee up to the level at which it makes zero profits when reporting honestly, any deviation in messaging strategy decreases the average value of a click for the firm and thus results in negative profits.

Consumer surplus under this equilibrium can be written as

$$CS = \left(1 - \frac{\gamma^n}{2}\right)\left(\frac{1}{8} - \int_0^{1/8} sf(s)\,ds\right) + \frac{\gamma^n}{2}\left(F\left(\frac{\alpha}{8}\right)\frac{\alpha}{8} - \int_0^{\alpha/8} sf(s)\,ds\right),$$

where the first term is the surplus accruing to consumers that find $t_i \in \theta_j$, and the second to those with $t_i \notin \theta_j$.

This equilibrium yields three possible outcomes for the publisher, whose profit is given by $E(b^{(n-1)}\lambda^2)$. Firstly, with probability γ^n, there are only single-product firms. Secondly, with probability $n(1-\gamma)\gamma^{n-1}$ there is precisely one multi-product firm—in which case the per-click cost is set at the single-product firms' bid level, but the number of clicks is that for a multi-product firm. Finally, with the complementary probability there are two or more multi-product firms. Thus,

$$\pi^P = \gamma^n \frac{1+\alpha F\left(\frac{\alpha}{8}\right)}{4\left[1+F\left(\frac{\alpha}{8}\right)\right]}\left[\frac{1}{2} + \frac{1}{2}F\left(\frac{\alpha}{8}\right)\right] +$$
$$\left[n(1-\gamma)\gamma^{n-1}\right]\frac{1+\alpha F\left(\frac{\alpha}{8}\right)}{4\left[1+F\left(\frac{\alpha}{8}\right)\right]} + \left[1 - \left(\gamma^n + n(1-\gamma)\gamma^{n-1}\right)\right]\frac{1}{4}.$$

All single-product firms bid up to their value and therefore make zero profit. A multi-product firm makes a positive profit if and only if it is the unique such firm, in which case it pays the single-product firm bid. Thus,

$$\pi\left(AB, m^*(AB), p^*, C, \mu^*(AB)\right) = \gamma^{n-1}\left[\frac{1}{4} - \frac{1+\alpha F\left(\frac{\alpha}{8}\right)}{4\left[1+F\left(\frac{\alpha}{8}\right)\right]}\right].$$

Lastly, total social welfare can be calculated as the sum of consumer and publisher surplus from above along with the expected profits of the advertiser. A more direct method is to note that each visiting consumer with $\hat{v}_i > 0$ generates

$$\int_0^{q^*} 1 - q\,dq = \int_0^{p^*} 1 - q\,dq = \frac{3}{8}$$

units of surplus, and expends s_i units on search. Thus, total welfare is

$$W = \left(1 - \frac{\gamma^n}{2}\right)\left(\frac{3}{8} - \int_0^{1/8} sf(s)\,ds\right) + \frac{\gamma^n}{2}\left(F\left(\frac{\alpha}{8}\right)\frac{3\alpha}{8} - \int_0^{\alpha/8} sf(s)\,ds\right).$$

Both consumer surplus and social surplus are created whenever a consumer realises $\hat{v}_i > 0$, and lost whenever a consumer pays the search cost; the two expressions are therefore similar.

The advertiser's message to the consumers is costless and unverifiable, making this a game of 'cheap talk'. Multiplicity of equilibria is pervasive in such games (see Crawford and Sobel, 1982), but a common feature is the existence of a pooling (or babbling) equilibrium,[3] and such an equilibrium does indeed exist in the pay-per-click environment. In addition, it can be shown that the PPC regime supports a series of partial pooling equilibria.

Proposition 27 *The PPC environment can always support a symmetric partial pooling equilibrium in which A- and B-type firms pool together, as well as a symmetric pair of partial pooling equilibria in which AB-type firms pool with one of the two types of single-product firm.*

Proposition 28 *The PPC environment can always support a babbling equilibrium in which all firms use the same messaging strategy.*

The equilibria described in propositions 27 and 28 are derived in appendix A.1 and appendix A.2.

V Pay-Per-Impression and Pay-Per-Sale

V.1 Inadmissibility of informative equilibria under PPI and PPS

In the previous section it was shown that a fully informative equilibrium always exists in the PPC environment. In stark contrast to this result, I show in this section that when the publisher is constrained to use either pay-per-impression or pay-per-sale, a non-trivial fully informative equilibrium can never exist.[4]

Proposition 29 *In pay-per-impression and pay-per-sale environments, there exists no non-trivial equilibrium with fully informative advertising messages.*

[3]The logic is thus: if the receiver of a cheap-talk message (the consumers in our case) chooses to ignore the message's contents then it is a best response for the sender (the advertiser) to send an uninformative message. Given this signalling strategy, ignoring the message is indeed optimal for the receiver.

[4]There will generally exist trivial equilibria of the PPS environment in which firms submit very high bids and the advertiser sets a very high price so that no transactions take place and the surplus for all parties is zero regardless of m.

Proof. Suppose that the proposition is false. If a single-product advertiser deviates from a truth-telling equilibrium by sending $m = m^*(AB)$, then the resulting change in its profits is

$$(4.8)\quad \Delta\pi^a = \Delta\lambda^1(\cdot)\left[(1-p)(p-b_S)\right].$$

Consumers that are matched with the advertiser expect the same surplus from an AB-type firm as from a firm having the advertiser's true type, and the deviation therefore has no effect on the visit decision of such consumers. Consumers that are unmatched with the advertiser obtain expected surplus

$$\frac{1}{8}(1-b_S)^2 - s_i$$

from visiting an AB-type firm, and

$$\frac{\alpha}{8}(1-b_S)^2 - s_i$$

from visiting a firm of the advertiser's true type. The deviation therefore yields an increase in $\lambda^1(\cdot)$ of

$$\Delta\lambda^1(\cdot) = \alpha\frac{1}{2}\left[F\left(\frac{1}{8}(1-b_S)^2\right) - F\left(\frac{\alpha}{8}(1-b_S)^2\right)\right] > 0.$$

Thus, (4.8) is positive for any $1 > p > b_s$, which is always true for non-trivial equilibria. ■

V.2 Partial pooling can sometimes be supported under PPI

Given proposition 29, a natural question that remains is: how informative can equilibrium advertising messages be in a PPI environment? The issue here is whether or not there exists one or more partially pooling equilibria—in which the typespace is partitioned into sub-sets, with firms from within each sending the same message—when the advertiser must pay on a per-impression basis. There are three candidates for a (pure-strategy) partially pooling equilibrium. Firstly, the two single-product firm types may pool together, sending a different message to multi-product firms. It is fairly straightforward to verify that no such equilibrium can be supported: after observing $m(A) = m(B)$ consumers update, placing probability $1/2$ on the firm being each of A and B.[5] A consumer's

[5] Symmetry of the game implies that both A- and B-type firms bid the same and win the auction with equal probability.

expected utility is then

$$\frac{1}{2}\left(\frac{1+\alpha}{8}\right) - s_i,$$

implying

$$\lambda^2(m(A), \cdot, \mu(m(A))) = \lambda^2(m(B), \cdot, \mu(m(B))) = F\left(\frac{1+\alpha}{16}\right) < 1.$$

By contrast, any firm transmitting $m(AB)$ is identified as being multi-product, so that

$$\lambda^2(m(AB), \cdot, \mu(m(AB))) = \lambda^2(m(AB), \cdot, \mu^*(m(AB))) = 1.$$

Thus, single-product advertisers face an incentive to imitate their multi-product counterparts.

The two alternative partial pooling configurations have one of the two single-product firms pooling with the multi-product firm. Since these two cases are completely symmetric, I shall focus on the case in which $m(A) = m(AB) \equiv m(A \cup AB) \neq m(B)$. PPC and PPI share the same p^*; it follows that upon observing $m(B)$ the consumers face an identical problem to that when they observe $m^*(B)$ in the fully-informative PPC case. Therefore, $\lambda^2\left(m(B), I, \mu\left(m(B)\right)\right)$ and $\lambda^1\left(B, \lambda^2\left(\cdot\right)\right)$ are equal to the λ's given in equations (4.5) and (4.6) respectively.

Now, the equilibrium bid for each firm type induces a probability of that firm type being the winner of the auction. Let $z(AB, \gamma, n)$ be the probability that the winner is type AB, and $z(A, \gamma, n)$ the probability that the winner is type A. After observing $m(A \cup AB)$, the consumers use Bayes' rule to update their beliefs and consider the advertiser to be of type AB with probability

$$\frac{z(AB, \gamma, n)}{z(AB, \gamma, n) + z(A, \gamma, n)}.$$

A given B-type consumer then expects utility of

$$\left(\frac{z(AB, \gamma, n)}{z(AB, \gamma, n) + z(A, \gamma, n)}\right)\frac{1}{8} + \left(\frac{z(A, \gamma, n)}{z(AB, \gamma, n) + z(A, \gamma, n)}\right)\frac{\alpha}{8} - s_i \equiv s^{PPIpp} - s_i$$

from visiting, whilst A types expect $1/8 - s_i$. This implies that

$$\lambda^2\left(m(A \cup AB), \cdot, \mu(m(A \cup AB))\right) = \lambda^1\left(AB, \lambda^2(\cdot)\right) = \frac{1}{2}\left[1 + F\left(s^{PPIpp}\right)\right],$$

and

$$\lambda^1\left(A, \lambda^2(\cdot)\right) = \frac{1}{2}\left[1 + \alpha F\left(s^{PPIpp}\right)\right].$$

In order to determine the values of $z(AB, \cdot)$ and $z(A, \cdot)$, it is necessary to derive the firms' bids. Advertiser profits at the monopoly price are given by

$$\pi^a = \lambda^1(\cdot)(1 - p^*)p^* - b_I,$$

which implies that firms will bid up to

(4.9) $b\left(\theta, I\right) = \lambda^1(\cdot)(1 - p^*)p^*.$

Thus, $b(AB, I) > b(A, I) > b(B, I)$, so that a multi-product firm is again selected by the auction whenever one exists, and it immediately follows that

$$z(AB, \gamma, n) = 1 - \gamma^n,\ z(A, \gamma, n) = \left(1 - \frac{1}{2^n}\right)\gamma^n.$$

Since λ^1 is always equal to λ^2 for multi-product firms, and sending $m(B)$ induces a lower λ^2 than does $m(A \cup AB)$, AB-type firms have no incentive to deviate from their messaging strategies. Similarly, type A firms have no incentive to deviate: imitating a B-type advertiser results in

$$\lambda^1(A, \cdot) = \frac{1}{2}\left(\alpha + F\left(\frac{\alpha}{8}\right)\right),$$

which implies that A-types do not wish to transmit $m(B)$ so long as

$$1 - F\left(\frac{\alpha}{8}\right) \geq \alpha\left[1 - F\left(s^{PPIpp}\right)\right].$$

This inequality is always satisfied. A type B firm that sends $m(A \cup AB)$ realises

$$\lambda^1(\cdot) = \frac{1}{2}\left[\alpha + F\left(s^{PPIpp}\right)\right].$$

Combining this with (4.6) reveals that B types prefer not to imitate firms from the $A \cup AB$ pool so long as

(4.10) $1 - F\left(s^{PPIpp}\right) \geq \alpha\left[1 - F\left(\frac{\alpha}{8}\right)\right].$

Thus, if (4.10) is satisfied then, conditional on having won the auction, no firm wishes to deviate from the equilibrium reporting strategy and, given this, the bids computed above are also optimal.

Lemma 30 *If the publisher is constrained to use a PPI fee structure then there exists a partial pooling equilibrium if and only only if (4.10) is satisfied.*

Equilibrium 9 *When the publisher is constrained to use a PPI fee structure and (4.10) is satisfied there exists a partially pooling equilibrium in which advertiser types A and AB transmit one signal, and advertisers of type B transmit a second, distinct message. In such an equilibrium, firms bid in accordance with (4.9). B-type consumers visit if they observe $m(B)$ or have $s_i \leq s^{PPIpp}$, A-types visit if they observe $m(A \cup AB)$ or have $s_i \leq \alpha/8$. Conditional on having visited, consumers purchase if $\hat{v}_i > p^*$.*

Symmetry of the game ensures that, when (4.10) is satisfied, a similar equilibrium exists with B and AB-type firms pooling together. The left-hand side of (4.10) is increasing in γ so that a higher γ is more conducive to the existence of partial pooling equilibria: as γ becomes low, consumers that observe $m(A \cup AB)$ assign a high-likelihood to the advertiser being multi-product and sending such a signal becomes tempting for type B advertisers.

When $m(A) = m(AB) \neq m(B)$, consumer surplus is the sum of three terms, reflecting the fact that the type of the advertiser depends upon the number of firms of each type:

$$CS = \underbrace{(1-\gamma^n)\frac{1}{2}\left(\frac{1}{8} - \int_0^{1/8} sf(s)\,ds + \frac{1}{8}F\left(s^{PPIpp}\right) - \int_0^{s^{PPIpp}} sf(s)\,ds\right)}_{\text{At least one multi-product firm.}} +$$

$$\underbrace{\gamma^n\left(1-\frac{1}{2^n}\right)\frac{1}{2}\left(\frac{1}{8} - \int_0^{1/8} sf(s)\,ds + \frac{\alpha}{8}F\left(s^{PPIpp}\right) - \int_0^{s^{PPIpp}} sf(s)\,ds\right)}_{\text{No multi-product firms, but at least one } A\text{-type firm.}} +$$

$$\underbrace{\frac{1}{2}\left(\frac{\gamma}{2}\right)^n\left(\frac{1}{8} - \int_0^{1/8} sf(s)\,ds + \frac{\alpha}{8}F\left(\frac{\alpha}{8}\right) - \int_0^{\alpha/8} sf(s)\,ds\right)}_{\text{Only } B\text{-type firms.}}.$$

The publisher makes $E\left(b^{(n-1)}\right)$ under a PPI regime. In order to calculate this expectation, it is useful to know that there are $n-1$ or more B-type firms with probability

$$\psi^1 = \left[\left(\frac{\gamma}{2}\right)^n + n\left(\left(\frac{\gamma}{2}\right)^{n-1}(1-\gamma) + \left(\frac{\gamma}{2}\right)^n\right)\right],$$

and two or more AB-type firms with probability

$$\psi^2 = \left[1 - \left(\gamma^n + n(1-\gamma)\gamma^{n-1}\right)\right].$$

Publisher profits can, then, be written as

$$\pi^p = \frac{\psi^2}{8}\left[1 + F\left(s^{PPIpp}\right)\right] + \frac{\psi^1}{8}\left[1 + \alpha F\left(\frac{\alpha}{8}\right)\right] + \frac{1-\psi^1-\psi^2}{8}\left[1 + \alpha F\left(s^{PPIpp}\right)\right].$$

A multi-product firm makes positive profits only when there are no other AB-type firms (which occurs with probability γ^{n-1} and implies that the sole multi-product firm wins the auction with probability 1). If, in addition, there are no A-type firms, the successful multi-product firm pays a lower price. Profits are then

$$\pi\left(AB, m(AB), p^*, I, \mu(AB)\right) = \gamma^{n-1}\left[\underbrace{\frac{1}{8}\left(1 + F\left(s^{PPIpp}\right)\right)}_{\text{Revenue}} - \frac{1}{2^{n-1}}\underbrace{\frac{1}{8}\left(1 + \alpha F\left(\frac{\alpha}{8}\right)\right)}_{B\text{-type bid}} - \left(1 - \frac{1}{2^{n-1}}\right)\underbrace{\frac{1}{8}\left(1 + \alpha F\left(s^{PPIpp}\right)\right)}_{A\text{-type bid}}\right].$$

B-type firms make zero profit in equilibrium 9. A-type firms, though, make positive profits provided that all rival firms are of type B. One can then write an A-type firm's expected profits as

$$\pi\left(A, m(A), p^*, I, \mu(A)\right) = \left(\frac{\gamma}{2}\right)^{n-1}\left[\frac{1}{8}\left(1 + \alpha F\left(s^{PPIpp}\right)\right) - \frac{1}{8}\left(1 + \alpha F\left(\frac{\alpha}{8}\right)\right)\right].$$

Lastly, total social welfare is again calculated analogously to consumer surplus, noting that—as in the PPC case—each match generates $3/8$ units of social surplus:

$$W = (1-\gamma^n)\frac{1}{2}\left(\frac{3}{8} - \int_0^{1/8} sf(s)\,ds + \frac{1}{8}F\left(s^{PPIpp}\right) - \int_0^{s^{PPIpp}} sf(s)\,ds\right) + \frac{1}{2}\gamma^n\left(1 - \frac{1}{2^n}\right)\left(\frac{3}{8} - \int_0^{1/8} sf(s)\,ds + \frac{3\alpha}{8}F\left(s^{PPIpp}\right) - \int_0^{s^{PPIpp}} sf(s)\,ds\right) + \frac{1}{2}\left(\frac{\gamma}{2}\right)^n\left(\frac{3}{8} - \int_0^{1/8} sf(s)\,ds + \frac{3\alpha}{8}F\left(\frac{\alpha}{8}\right) - \int_0^{\alpha/8} sf(s)\,ds\right).$$

Thus it is established that there can sometimes be partial information transmission in a PPI environment. When (4.10) is not satisfied, no information transmission can ever take place in a (pure strategy) equilibrium of the pay-per-impression environment, and any equilibrium must involve advertiser babbling.

Corollary 31 *When (4.10) is not satisfied, the only pure strategy equilibrium of the PPI game is the babbling equilibrium.*

The form that such a babbling equilibrium must take is detailed in appendix A.3. In fact, when one calculates the surplus accruing to the various agents under a PPI babbling equilibrium, the following result obtains:

Proposition 32 *The pay-per-click babbling equilibrium (equilibrium 13), and the pay-per-impression babbling equilibrium (equilibrium 14) are equivalent in the sense that they yield equal surplus for all agents.*

This result is also established in appendix A.3, and stems from the fact that consumers face an equivalent problem (i.e. the same price and information), and therefore act in the same manner in both cases. Given this, the firms' values for the advertising opportunity in the two equilibria will also be equal so that firms make the same expected payment under both.

V.3 Bidding and firm selection under PPS

Substituting (4.2) into (4.1), and imposing $b_C = b_I = 0$ yields

$$\pi\left(\theta, m, p^*, S, \mu(\cdot)\right) = \lambda^1\left(\theta, \lambda^2(\cdot)\right)\left[\left(1 - \frac{1+b_S}{2}\right)\left(\frac{1+b_S}{2} - b_S\right)\right].$$

This is positive for all $b_S < 1$ so that the optimal bid for all firms is 1. Such behaviour engenders a trivial equilibrium in which $p = 1$, consumers do not search, and consumer surplus, publisher profit, and firm profit are all equal to zero. In order to address this issue, the publisher can compute the price implied by each firm's bid and (4.2), and impose a kind of 'sale-weighting'—weighting each firm's bid by some factor that ensure the optimal (i.e. profit maximising) outcome for the auction. When, as seems reasonable, the consumers do not observe the auction's outcome, this is achieved by selecting the bid, b, that maximises

$$(4.11)\quad \left(1 - \frac{1+b}{2}\right)b.$$

This is simply the demand implied by (4.2) (the term in brackets) multiplied by the fee per sale, b.[6]

Sale-weighting induces the bid that maximises the payment to the publisher, taking the type of the advertiser as given; however, the publisher is still faced with the problem that the auction may not select the optimal *type* of firm as the advertiser. Indeed, under sale-weighting it is weakly dominant for all firm types to bid $b(\theta, S) = 1/2$. Thus, the sale-weighted PPS auction is equivalent to a take-it or leave-it offer with some tie-breaking rule. With random tie-breaking (which I shall assume below), the successful advertiser will be of type A and B each with probability $\gamma/2$, and of type AB with probability $1 - \gamma$.

V.4 Partial pooling under PPS

In light of the results of section V.1, the most informative equilibrium possible under pay-per-sale is, at best, a partially pooling equilibrium. Equilibria with $m(A) = m(B) \equiv m(A \cup B) \neq m(AB)$ can be ruled out in a similar manner to that used in section V.2: upon observing $m(A) = m(B)$, consumers expect

$$U_i = \frac{1}{2}\left(\frac{1+\alpha}{32}\right) - s_i,$$

so that

$$\lambda^2(m(A \cup B), S, \mu(m(A \cup B))) = F\left(\frac{1+\alpha}{64}\right) < F\left(\frac{1}{32}\right),$$

whilst

$$\lambda^2(m(AB), S, \mu(m(AB))) = \lambda^2(m(AB), S, \mu^*(m(AB))) = F\left(\frac{1}{32}\right).$$

In both cases, $\lambda^1(\cdot) = \lambda^2(\cdot)(1+\alpha)/2$: the deviation increases $\lambda^1(\cdot)$—and hence profits are also increased.

It remains to be seen whether $m(A) = m(AB) \equiv m(A \cup AB) \neq m(B)$ (and its symmetric counterpart) can be supported under pay-per-sale. Suppose that such an equilibrium does exist. Upon observing $m(B)$ consumers update their beliefs to $\mu^*(m(B))$. It has already been shown that $b(\theta, S) = 1/2$ implies a gross

[6] If consumers are able to observe the auction outcome so that the publisher can credibly commit to a different weighting factor then its choice of weight becomes an equilibrium variable. This is because the publisher can influence consumers' beliefs about p indirectly by inducing a given b_S, and can therefore influence the mass of consumers that visits in a given equilibrium. In such circumstances, given equilibrium messaging strategies, the publisher's incentive is to maximise the expected product of (4.11) and λ^1. Allowing this behaviour would not qualitatively influence the results below, but would significantly add to the model's complexity.

consumer surplus of $1/32$, and it follows that an advertiser transmitting $m(B)$ will receive

$$\lambda^2(m(B), S, \mu^*(m(B))) = \frac{1}{2}\left[F\left(\frac{1}{32}\right) + F\left(\frac{\alpha}{32}\right)\right].$$

Thus, a truthfully reporting B-type advertiser faces a potential demand of

$$\lambda^1\left(B, \lambda^2(m(B), S, \mu^*(m(B)))\right) = \frac{1}{2}\left[F\left(\frac{1}{32}\right) + \alpha F\left(\frac{\alpha}{32}\right)\right]. \tag{4.12}$$

Since (conditional on being present) an A-type and AB-type firm are equally likely to win the auction, a B-type consumer that observes $m(A \cup AB)$ expects to receive

$$U_i = \left(\frac{1-\gamma}{1-\gamma+\frac{\gamma}{2}}\right)\frac{1}{32} + \left(\frac{\frac{\gamma}{2}}{1-\gamma+\frac{\gamma}{2}}\right)\frac{\alpha}{32} - s_i \equiv s^{PPSpp} - s_i,$$

surplus from visiting, whilst an A-type expects $(1/32) - s_i$. The implied consumer behaviour is

$$\lambda^2\left(m(A \cup AB), S, \mu(m(A \cup AB))\right) = \frac{1}{2}\left[F\left(\frac{1}{32}\right) + F\left(s^{PPSpp}\right)\right],$$
$$\lambda^1\left(A, \lambda^2(\cdot)\right) = \frac{1}{2}\left[F\left(\frac{1}{32}\right) + \alpha F\left(s^{PPSpp}\right)\right],$$
$$\lambda^1\left(AB, \lambda^2(\cdot)\right) = \frac{1}{2}\left[F\left(\frac{1}{32}\right) + F\left(s^{PPSpp}\right)\right].$$

Advertiser profit is given by$\lambda^1(\cdot)(1-p^*)(p^* - b_S) = \lambda^1(\cdot)/16$, and firms therefore wish to deviate in messaging strategy if and only if there is an alternative signal resulting in a higher $\lambda^1(\cdot)$. Since $\lambda^2(m(A \cup AB), \cdot) > \lambda^2(m(B), \cdot)$, and $\lambda^1(AB, \lambda^2(\cdot)) = \lambda^2(\cdot)$, AB-type firms do not wish to deviate. By sending $m(B)$, an A-type advertiser would realise

$$\lambda^1\left(A, \lambda^2(m(B), \cdot)\right) = \frac{1}{2}\left[\alpha F\left(\frac{1}{32}\right) + F\left(\frac{\alpha}{32}\right)\right],$$

which is less than $\lambda^1\left(A, \lambda^2(m(A \cup AB), \cdot)\right)$; it follows that A-type firms have no profitable messaging strategy deviation. Lastly, B-type advertisers that deviate and send $m(A \cup AB)$ have

$$\lambda^1\left(B, \lambda^2(m(A \cup AB), \cdot)\right) = \frac{1}{2}\left[\alpha F\left(\frac{1}{32}\right) + F\left(s^{PPSpp}\right)\right],$$

which is not greater than (4.12) whenever

$$\frac{F\left(\frac{1}{32}\right) - F\left(s^{PPSpp}\right)}{F\left(\frac{1}{32}\right) - F\left(\frac{\alpha}{32}\right)} \geq \alpha. \tag{4.13}$$

Much as in equilibrium 9, this condition is more easily satisfied when γ is large.

Lemma 33 *If the publisher is constrained to use a PPS fee structure then there can exist a partially pooling equilibrium only if (4.13) is satisfied.*

Thus, conditional on having won the auction and (4.13) being satisfied, no firm wishes to deviate from the prescribed reporting strategy and, given this, the specified bidding strategy is optimal. One can, then, characterise a partial pooling equilibrium of the pay-per sale environment thus:[7]

Equilibrium 10 *When the publisher is constrained to use a PPS fee structure with sale-weighting and (4.13) is satisfied, there exists a partial pooling equilibrium in which advertiser types A and AB transmit signal $m(A \cup AB)$, and advertisers of type B transmit $m(B)$. In such an equilibrium, firms bid $b = 1/2$. B-type consumers visit if they have $s_i < 1/32$ and observe $m(B)$, or have $s_i \leq s^{PPSpp}$; A-types visit if they have $s_i < 1/32$ and observe $m(A \cup AB)$, or have $s_i \leq \alpha/32$. Conditional on having visited, consumers purchase if $\hat{v}_i > p^*$.*

As in the PPI partial-pooling case, consumer surplus in this equilibrium is the sum of three terms, one for each realisation of advertiser type:

$$CS = \underbrace{\frac{1-\gamma}{2}\left[F\left(\frac{1}{32}\right)\frac{1}{32} - \int_0^{1/32} sf(s)\,ds + F\left(s^{PPSpp}\right)\frac{1}{32} - \int_0^{s^{PPSpp}} sf(s)\,ds\right]}_{AB\text{-type advertiser}} +$$

$$\underbrace{\frac{\gamma}{4}\left[F\left(\frac{1}{32}\right)\frac{1}{32} - \int_0^{1/32} sf(s)\,ds + F\left(s^{PPSpp}\right)\frac{\alpha}{32} - \int_0^{s^{PPSpp}} sf(s)\,ds\right]}_{A\text{-type advertiser}} +$$

$$\underbrace{\frac{\gamma}{4}\left[F\left(\frac{1}{32}\right)\frac{1}{32} - \int_0^{1/32} sf(s)\,ds + F\left(\frac{\alpha}{32}\right)\frac{\alpha}{32} - \int_0^{\alpha/32} sf(s)\,ds\right]}_{B\text{-type advertiser}}.$$

The profits for the publisher are given by the product of three terms. Firstly, the fee per sale is equal to $1/2$. Secondly, the rate at which matches are converted to

[7] As in the previous case, there is a similar, symmetric equilibrium with $m(B) = m(AB) = m(B \cup AB) \neq m(A)$.

sales is $1 - p^* = 1/4$. Lastly is the (expected) value of λ^1. Thus,

$$\pi^p = \frac{1}{2} \times \frac{1}{4} \times \left[\frac{1-\gamma}{2} \left(F\left(\frac{1}{32}\right) + F\left(s^{PPSpp}\right) \right) + \frac{\gamma}{4} \left(F\left(\frac{1}{32}\right) + \alpha F\left(s^{PPSpp}\right) \right) + \frac{\gamma}{4} \left(F\left(\frac{1}{32}\right) + \alpha F\left(\frac{\alpha}{32}\right) \right) \right]. \tag{4.14}$$

Each firm wins the auction with probability $1/n$. With $p = p^* = 3/4$, firms sell to $1/4$ of the consumers that realise $\hat{v}_i > 0$, and make $p^* - 1/2 = 1/4$ per sale. Thus, multi-product firm profits are

$$\pi\left(AB, m(A \cup AB), p^*, I, \mu(m(A \cup AB))\right) = \frac{1}{16n} \left[\frac{1}{2} \left(F\left(\frac{1}{32}\right) + F\left(s^{PPSpp}\right) \right) \right],$$

A-type firm profits are

$$\pi\left(A, m(A \cup AB), p^*, I, \mu(m(A \cup AB))\right) = \frac{1}{16n} \left[\frac{1}{2} \left(F\left(\frac{1}{32}\right) + \alpha F\left(s^{PPSpp}\right) \right) \right],$$

and B-type firm profits are

$$\pi\left(B, m(B), p^*, I, \mu(m(B))\right) = \frac{1}{16n} \left[\frac{1}{2} \left(F\left(\frac{1}{32}\right) + \alpha F\left(\frac{\alpha}{32}\right) \right) \right].$$

Lastly, as usual, social welfare has a similar form to consumer surplus:

$$W = \frac{1-\gamma}{2} \left[F\left(\frac{1}{32}\right)\frac{7}{32} - \int_0^{1/32} sf(s)\,ds + F\left(s^{PPSpp}\right)\frac{7}{32} - \int_0^{s^{PPSpp}} sf(s)\,ds \right] + \frac{\gamma}{4} \left[F\left(\frac{1}{32}\right)\frac{7}{32} - \int_0^{1/32} sf(s)\,ds + F\left(s^{PPSpp}\right)\frac{7\alpha}{32} - \int_0^{s^{PPSpp}} sf(s)\,ds \right] + \frac{\gamma}{4} \left[F\left(\frac{1}{32}\right)\frac{7}{32} - \int_0^{1/32} sf(s)\,ds + F\left(\frac{\alpha}{32}\right)\frac{7\alpha}{32} - \int_0^{\alpha/32} sf(s)\,ds \right].$$

As was the case for PPI, when (4.13) is not satisfied no information transmission can take place in a (pure strategy) equilibrium of the pay-per-sale environment. Under these circumstances any equilibrium must involve advertiser babbling.

Corollary 34 *When (4.13) is not satisfied, the only (non-trivial) pure strategy equilibrium of the PPS game is the babbling equilibrium.*

The form that such a babbling equilibrium must take is detailed in appendix A.4.

VI COMMENTS ON WELFARE

Although the model presented above does not admit a generally tractable comparative static analysis, in this section I offer some brief observations on the matter of welfare, along with some numerical analysis.

As n becomes large the PPC and PPI equilibria realise an AB-type advertiser with probability close to 1. As γ becomes small the realised advertiser type is again almost always AB. Moreover, as α approaches 1, all firms serve a similar function to AB-types. In all three cases, the resulting asymptotic welfare values are the same for all PPC and PPI equilibria, and are as presented in table 4.1.

CS	π^p	π	W
$1/8 - E(s_i)$	$1/4$	0	$3/8 - E(s_i)$

TABLE 4.1 Welfare from PPC/PPI equilibria with large n, large α or small γ.

PPS firm selection is essentially random so that welfare depends upon n only insofar as each individual firm's probability of selection approaches zero as n is increased—consumer and social welfare, along with realised advertiser and publisher profits are all invariant to n. With small γ or large α all firms are multi-product or behave as though they were multi-product. This gives rise to the asymptotic welfare values in table 4.2.

CS	π^p	π	W
$F(1/32)/32 - \int_0^{1/32} sf(s)\,ds$	$F(1/32)/8$	$F(1/32)/16n$	$7F(1/32)/32 - \int_0^{1/32} sf(s)\,ds$

TABLE 4.2 Welfare from PPS equilibria with large α or small γ.

Given equilibrium pricing and messaging strategies, consumers determine whether or not to visit based upon a comparison of their (expected) surplus from doing so and their idiosyncratic search cost, s_i. As consumers become better informed they are able to more accurately forecast the former, and therefore make better visit decisions. Indeed, in a fully informative equilibrium, $E(\hat{v}_i|\mu^*(m^*(\theta)), t_i) = E(\hat{v}_i|\theta, t_i)$, so that the consumers visit precisely when it is (privately) optimal for them to do so. Thus, up to the choice of p^*, the fully-informative equilibrium maximises consumer surplus. More generally, equilibria having the same price exhibit a perfect rank correlation between informativeness and ex ante expected consumer surplus.[8]

[8]From the perspective of A-type consumers, a partial pooling equilibrium with $m(A) = m(AB) \neq m(B)$ reveals all relevant information. Moreover, whenever the realised advertiser's type is B, this partial pooling equilibrium is also fully informative. Thus less than perfect information

But what of pay-per-sale, which exhibits a different price to PPC and PPI? It can be shown generally that the best pay-per-sale equilibrium (the partial pooling equilibrium) always yields lower consumer surplus than does the worst PPC/PPI equilibrium (namely, the babbling equilibrium of appendix A.2). It follows that PPS always yields lower consumer surplus than do PPC and PPI. This can easily been seen by noting that each consumer expects to receive the greater of his expected surplus from visiting and his reservation value of zero. In the PPC/PPI babbling case, a lower bound on the former is $1/16 - s_i$, which occurs when $\gamma = 1$, and $\alpha = 0$. In the PPS partial pooling equilibrium the gross consumer surplus generated by a match is $1/32$. It follows that an upper bound on a given visiting consumer's expected surplus in this case is $1/32 - s_i$. Since consumers visit whenever their expected gain from doing so is non-negative, and each individual consumer expects strictly more from visiting under PPC/PPI babbling than under PPS, it follows that the consumer surplus from PPS can be no higher than either alternative fee structure. This is a basic artefact of the double marginalisation problem: both the publisher and advertiser are acting as monopolists, so that prices are effectively distorted twice. As a result, consumer surplus per match is only $1/4$ of that in the PPC/PPI cases.

Figure 4.1 shows consumer surplus as a function of γ for F uniform on $[0, 1/8]$, with $\alpha = 0.1$, and $n = 2$.[9] Note that the equilibria of the PPC and PPI environments (which share $p^* = 1/2$) are ordered according to the amount of information contained within equilibrium messages. Asymmetric partial pooling equilibria in the pay-per-click environment provide higher surplus than those under pay-per-impression. This is because the PPC equilibria systematically select single-product firms of the type that sends a unique message over those that pool with the multi-product type (by penalising the latter for attracting many unmatched consumers)—making such equilibria more informative. The symmetric PPC partial pooling equilibrium of equilibrium 11 (see appendix A.1) is not shown, but lies between the asymmetric partial pooling and babbling PPC equilibria. The pay-per-sale equilibria are also ranked according to informativeness, but collectively provide significantly lower surplus than do the PPS/PPI equilibria.[10]

occurs in such an equilibrium only when a B-type consumer observes $m(A \cup AB)$. Conditional on $\theta \in \{A, AB\}$, the partial pooling equilibrium induces privately optimal visit decisions by B-types, whilst the corresponding babbling equilibrium induces excessive visiting by B-type consumers with s_i close to, but greater than $\alpha/8$. Similarly, conditional on $\theta \in \{A, B\}$, a symmetric partial pooling equilibrium with $m(A) = m(B) \neq m(AB)$ promotes optimal visit decisions, whereas the babbling equilibrium induces excess visiting. When $\theta = AB$, the symmetric partial pooling equilibrium induces the optimal visit decision for all consumers.

[9]Surplus from the three equilibria derived above, along with that from the babbling and asymmetric partial pooling equilibria derived in appendix A are shown.

[10]For low values of γ (to the left of the black points), the PPS and PPI partial pooling equilibria cannot be supported.

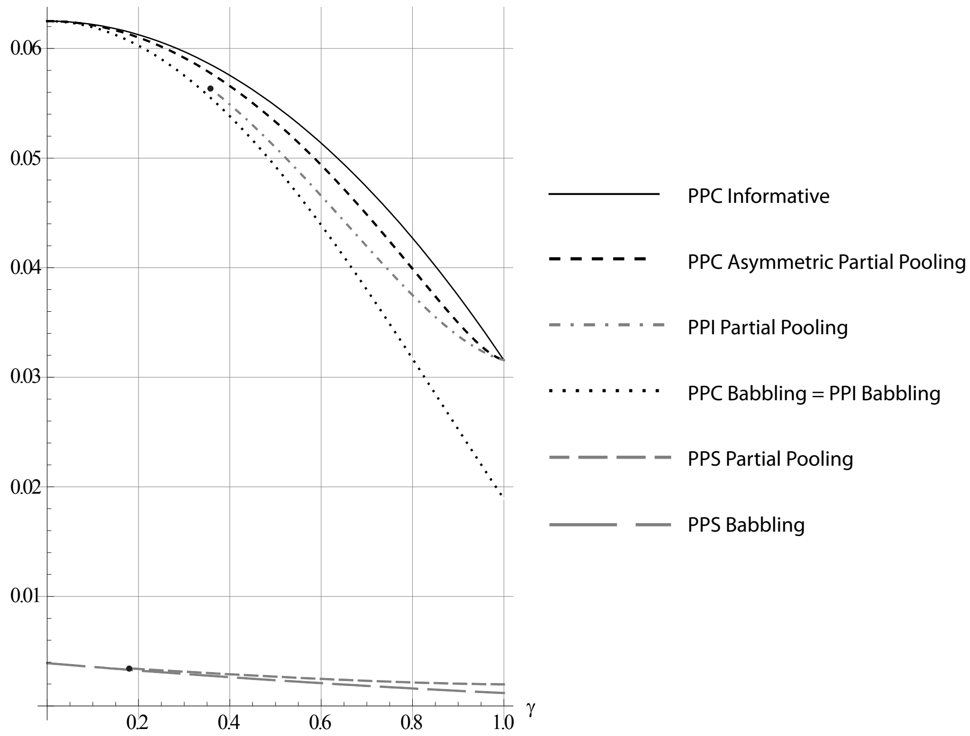

FIGURE 4.1 Consumer surplus under various equilibria for F uniform on $[0, 1/8]$, with $\alpha = 0.1$, and $n = 2$. The rank ordering of these equilibria is invariant to α and n.

For high values of γ, the asymmetric partial pooling PPC and PPI equilibria reveal all relevant information (by allowing consumers to differentiate perfectly between A- and B-type firms), and these equilibria therefore provide the same consumer surplus as does the fully informative equilibrium when $\gamma = 1$. By contrast, the symmetric partial pooling equilibrium reveals no useful information at all when γ is close to 1. Consumer surplus from this equilibrium is therefore equal to that from the PPC/PPI babbling equilibria under such circumstances.

Social welfare for F uniform on $[0, 1/8]$ behaves much as does consumer surplus: increasing with equilibrium informativeness. However, this is not necessarily the case more generally. Figure 4.2, for example, shows social welfare when F is given by the inverse quadratic distribution $F(s) = 192s^2 - 1024s^3$. For low α (figure 4.2(a)), this looks much like the uniform case. However, for intermediate values of α the welfare ranking of equilibria undergoes a dramatic reversal. Figure 4.2(b) shows welfare for $\alpha = 0.6$, in which case the PPC/PPI babbling equilibria provide the highest level of welfare.

To understand why it should be the case that improved information transmission might lower social welfare for some parametric configurations it is important

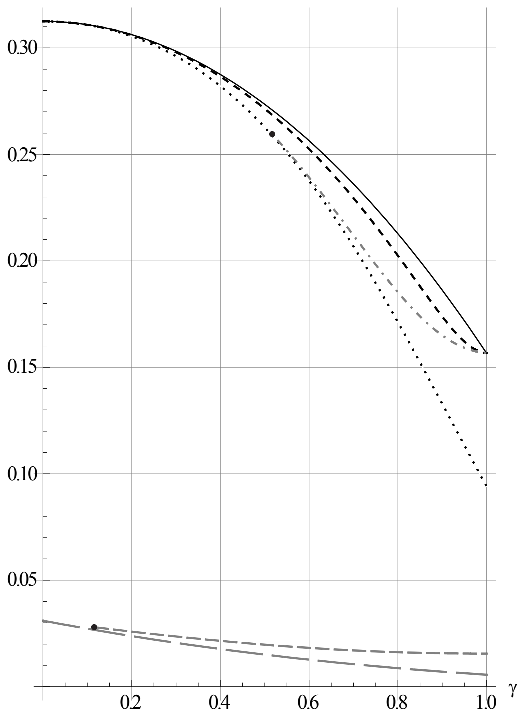

(a) $\alpha = 0.1, N = 2$

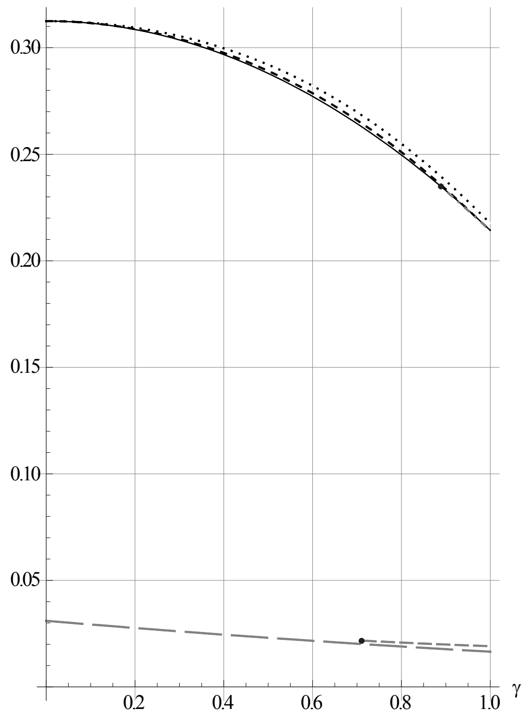

(b) $\alpha = 0.6, N = 2$

FIGURE 4.2 Social welfare under various equilibria for F inverse quadratic on $[0, 1/8]$. Legend is as in Figure 4.1.

to note that, in order to extract profits from their advertisement opportunities, publishers must be able to endow their advertisers with some degree of market power. Thus this environment is inherently second best. That advertisers and the publisher retain some of the surplus generated by a sale drives a wedge between consumer surplus and social welfare so that consumer visits, which depend on the private benefit of visiting, are inefficiently low. When the distribution of search costs is such that a reduction in information transmission induces additional visits from those consumers for whom s_i lies between the private and social benefit of visiting, an increase in overall surplus can result.

Much like social welfare, publisher profit can not generally be ranked according to equilibrium informativeness. Figure 4.3(b), for example, shows that for intermediate values of α, PPC/PPI babbling equilibria may yield higher profits than the fully informative PPC equilibrium. Profits from the informative equilibrium depend on F in two respects. Firstly, single-product firms that receive fewer unmatched consumers will bid more per-click, so that $E(b^{(n-1)})$ is decreasing in $F(\alpha/8)$. Secondly, the expected number of clicks is increasing in $F(\alpha/8)$. Thus, publisher profit will not, generally, be monotone in $F(\alpha/8)$. The PPC babbling

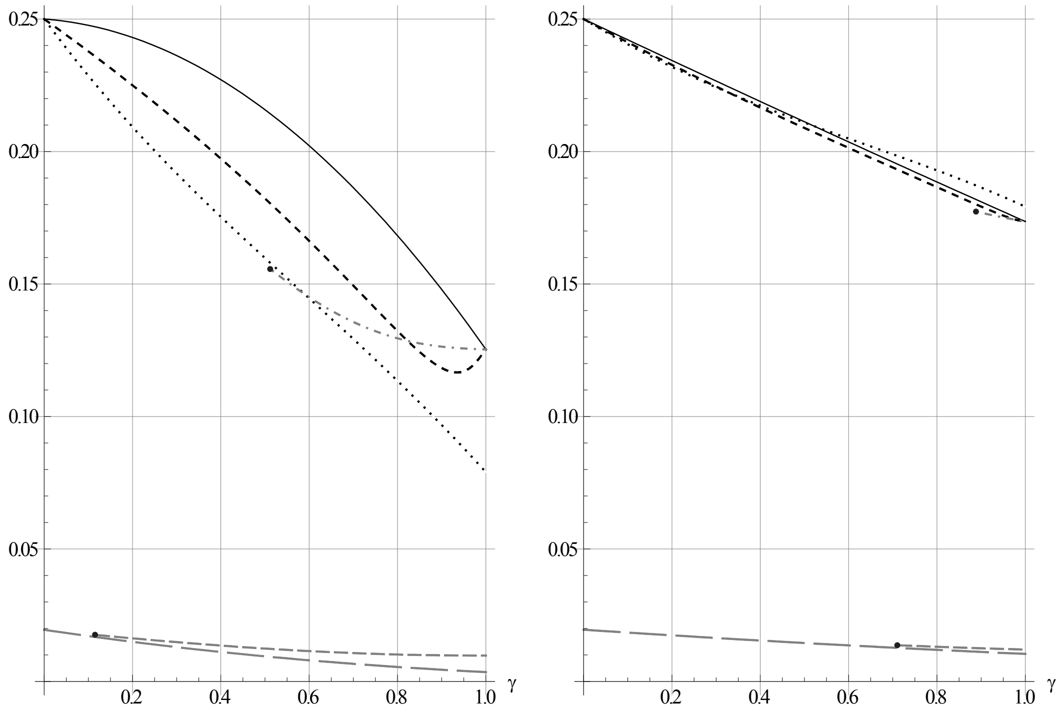

(a) $\alpha = 0.1, N = 2$ (b) $\alpha = 0.6, N = 2$

FIGURE 4.3 Publisher profit under various equilibria for F inverse quadratic on $[0, 1/8]$. Legend is as in Figure 4.1.

equilibrium,[11] by contrast, has bids invariant to F, and therefore has π^P increasing in the number of visits $F(s^{PPCb})$. The relative profitability of these two equilibria is therefore sensitive to the relative size of $F(\alpha/8)$, and $F(s^{PPCb})$.

What can be said unambiguously is that both pay-per-sale equilibria will always yield lower profits than either the informative equilibrium or the PPC/PPI babbling equilibria. Details of both results are presented in appendix A. Further, whilst an analytic comparison of PPC partial pooling profits with those from the PPS equilibria is intractable, numerical analysis suggests that PPS is also inferior to the PPC partial pooling equilibrium of appendix A.1 in this regard. Together, these results imply that a publisher should generally prefer to implement a PPC regime rather than use a PPS fee structure.

From the firms' perspective: multi-product firms find informative equilibria to be particularly appealing when $F(\alpha/8)$ is high—under these circumstances, single-product firms attract many unmatched consumers in spite of their honesty and therefore submit small per-click bids. Pay-per-sale equilibria are attractive to any firm facing a low γ—this increases the likelihood of there being two or

[11]recall, that the babbling PPC and PPI equilibria are profit-equivalent

more multi-product firms, in which case all firms make zero profit under PPC or PPI and positive profit under PPS (see tables 4.1 and 4.2).

VII Conclusion

This chapter has been concerned with the effect of fee structure choice on the equilibrium transmittability of information in on-line advertising. Charging advertisers on a per-click basis provides a disincentive to attracting poorly matched consumers, and can therefore encourage the transmission of informative advertisements. Advertisers who pay for advertisements on a per-impression basis, or pay only in the event of a sale, find attracting visits from poorly matched consumers to be profitable provided that those consumers will buy with some positive probability—however small that probability is. Such fee structures are therefore much less conducive to the equilibrium transmission of information. A variety of more complicated specifications for the type space and distribution of values appear to maintain a similar incentive structure with regard to creating informative advertisements so that similar results to those presented here ought to hold in a range of environments. Aside from being less informative, a pay-per-sale environment distorts prices by effectively increasing an advertiser's marginal cost, and offers no obvious basis for the systematic selection of multi-product firms, which are able to satisfy a broader range of consumers.

Informative equilibria are good for consumers since they allow proper optimisation of the visit decision, and therefore reduce time wasted visiting unattractive firms. The role of information transmission in determining the welfare of other parties is less straightforward. Since consumers do not retain the full surplus created by a purchase their search decisions are not always socially optimal. Meanwhile, if there are many consumers with moderately high search costs, publishers may benefit from an increase in consumer participation when moving from informative to uninformative equilibria. The sensitivity of publisher and social welfare to the parameters of the environment in general, and to the distribution of search costs in particular, suggests that special attention should be paid to these matters when forming policy. Pay-per-sale introduces large price distortions and consequently yields lower profits for the publisher than the worst PPC equilibrium.

There are several possible extensions that are worthy of note. The game that I have considered is 'one-shot'. In reality, consumers that experience poor search results are less likely to return, and this is likely to create an additional incentive for the publisher to encourage coordination around an informative equilibrium.[12] The publisher's preference ordering over equilibria is also likely to be influenced by competition with other publishers, in which case publisher must compete for both consumers and advertisers by satisfying both. Lastly, I have focused on the case in which there is a single advertisement slot available. When there are several competing advertisers, the incentive to 'cheat' under honest reporting in pay-per-sale and pay-per-impression remains so long as there is a positive probability (however small) that some consumers will buy upon visiting a firm with which they believe themselves (ex ante) to be poorly matched. This might be the case if, for example, some consumers searching for a PC learn upon visiting the Apple website that a Mac better suits their needs.

A OMITTED EQUILIBRIA

A.1 Partial pooling in a PPC environment

There are three candidates for a (pure-strategy) partially pooling equilibrium. Firstly, the two single-product firm types may pool together, sending a different message to multi-product firms. After observing $m(A) = m(B)$ consumers update, placing probability $1/2$ on the firm being each of A and B. A consumer's expected utility is then

$$\frac{1}{2}\left(\frac{1+\alpha}{8}\right) - s_i,$$

implying

$$\lambda^2(m(A), \cdot, \mu(m(A))) = \lambda^2(m(B), \cdot, \mu(m(B))) = F\left(\frac{1+\alpha}{16}\right).$$

Since the visiting consumers represent an equal mix of A and B types, single product firms have

$$\lambda^1(\cdot) = \frac{1+\alpha}{2}\lambda^2(\cdot).$$

This implies the optimal bid is

$$b(\theta, C) = \frac{1+\alpha}{8},$$

for $\theta \in \{A, B\}$.

Multi product firms match all of their visitors (since $m(AB)$ identifies the multi-

[12] Google, for instance, actively monitors the correspondence between ad text and website content in order to maintain consumer confidence in the 'sponsored search results'.

product firm as being such, it attracts a visit from all consumers). It follows that, as usual, $b(AB, C) = 1/4$. Thus, an AB type advertiser wins the auction whenever such a firm is present.

AB firms clearly do not wish to deviate from their prescribed messaging strategy. Single-product firms win the auction if and only if there are no multi-product firms, in which case they make zero profits. By sending $m(AB)$ such firms attract visits from all consumers. However, the average value of a click to the firm remains $(1 + \alpha)/8$, so that the firm's profits remain at zero and deviation is not profitable. Thus we have

Equilibrium 11 *When the publisher is constrained to use a PPC fee structure, there exists a (symmetric) partial pooling equilibrium in which advertiser types A and B transmit one signal, and advertisers of type AB transmit a second, distinct message. In such an equilibrium, firms bid in accordance with (4.4). Consumers visit if they observe $m(AB)$ or have $s_i \leq (1 + \alpha)/16$. Conditional on having visited, consumers purchase if $\hat{v}_i > p^*$.*

Consumer surplus in this equilibrium is given by

$$CS = (1-\gamma^n)\left(\frac{1}{8} - \int_0^{1/8} sf(s)\,ds\right) + \gamma^n \left[\left(\frac{1+\alpha}{16}\right) F\left(\frac{1+\alpha}{16}\right) - \int_0^{(1+\alpha)/16} sf(s)\,ds\right].$$

This equilibrium yields three possible outcomes for the publisher, whose profit is given by $E(b^{(n-1)}\lambda^2)$. Firstly, with probability γ^n, there are only single-product firms. Secondly, with probability $n(1-\gamma)\gamma^{n-1}$ there is precisely one multi-product firm—in which case the per-click cost is set at the single-product firms' bid level, but the number of clicks is that for a multi-product firm. Finally, with the complementary probability there are two or more multi-product firms. Thus,

$$\pi^P = \gamma^n \frac{1+\alpha}{8} F\left(\frac{1+\alpha}{16}\right) + n(1-\gamma)\gamma^{n-1}\frac{1+\alpha}{8} + \left[1 - \left(\gamma^n + n(1-\gamma)\gamma^{n-1}\right)\right]\frac{1}{4}.$$

All single-product firms bid up to their value and therefore make zero profit. A multi-product firm makes a positive profit if and only if it is the unique such firm, in which case it pays the single-product firm bid. Thus,

$$\pi\left(AB, m^*(AB), p^*, C, \mu^*(AB)\right) = \gamma^{n-1}\left[\frac{1}{4} - \frac{1+\alpha}{8}\right].$$

Lastly, total social welfare can be calculated as

$$W = (1-\gamma^n)\left(\frac{3}{8} - \int_0^{1/8} sf(s)\,ds\right) + \gamma^n\left[3\left(\frac{1+\alpha}{16}\right)F\left(\frac{1+\alpha}{16}\right) - \int_0^{(1+\alpha)/16} sf(s)\,ds\right].$$

The two alternative partial pooling configurations have one of the two single-product firms pooling with the multi-product firm. Since these two cases are completely symmetric, I shall focus on the case in which $m(A) = m(AB) = m(A \cup AB) \neq m(B)$. Suppose that such an equilibrium exists. Upon observing $m(B)$ the consumers face an identical problem to the fully-informative case, and $\lambda^2(m(B), I, \mu(m(B)))$ and $\lambda^1(B, \lambda^2(\cdot))$ are equal to the λ's given in equations (4.5) and (4.6) respectively.

In a pay-per-click environment, the weakly dominant bidding strategy is given by (4.4). Since $\lambda^1(\cdot) = \lambda^2(\cdot)$ for multi-product firms, $b(AB, C) = 1/4$. Single-product firms have $\lambda^1(\cdot) < \lambda^2(\cdot)$, and therefore bid less than their multi-product counterparts. With $p = p^* = 1/2$, A- and B-type firms are both visited by all of their respective matched consumers. Since A-type firms pool with AB-types, B-type consumers expect a greater surplus from visiting when they observe $m(A \cup AB)$ than do A-type consumers having observed $m(B)$. It follows that an A-type advertiser will face a greater proportion of unmatched consumers in equilibrium than will B-types, and hence

$$\frac{\lambda^1(A,\cdot)}{\lambda^2(A,\cdot)} < \frac{\lambda^1(B,\cdot)}{\lambda^2(B,\cdot)}.$$

Thus, equilibrium bids must be ordered as follows:

$$b(AB, C) > b(B, C) > b(A, C).$$

This implies that the winning firm is of type AB with probability $1 - \gamma^n$, and type A with probability $\gamma^n 2^{-n}$. Using Bayesian updating, a B-type consumer expects a surplus of

$$U_i = \left(\frac{1-\gamma^n}{1-\gamma^n+\gamma^n 2^{-n}}\right)\frac{1}{8} + \left(\frac{\gamma^n 2^{-n}}{1-\gamma^n+\gamma^n 2^{-n}}\right)\frac{\alpha}{8} - s_i \equiv s^{PPCpp} - s_i$$

from visiting when $m(A \cup AB)$ has been observed. The same message leaves A-types with an expected surplus of $1/8 - s_i$. These surpluses imply that the number of visitors prompted by $m(A \cup AB)$ will be

$$\lambda^2(m(A \cup AB), C, \mu(m(A \cup AB))) = \frac{1}{2}\left[1 + F\left(s^{PPCpp}\right)\right],$$

and that

$$\lambda^1\left(A, \lambda^2(\cdot)\right) = \frac{1}{2}\left[1 + \alpha F\left(s^{PPCpp}\right)\right].$$

As usual, $\lambda^1\left(AB, \lambda^2(\cdot)\right) = \lambda^2(\cdot)$.

Since λ^1 is always equal to λ^2 for multi-product firms and sending $m(B)$ induces a lower λ^2 than does $m(A \cup AB)$, AB-type firms have no incentive to deviate from their messaging strategies. Similarly, type A firms have no incentive to deviate: imitating a B-type advertiser results in

$$\lambda^1(A, \cdot) = \frac{1}{2}\left(\alpha + F\left(\frac{\alpha}{8}\right)\right),$$

which implies that A-types do not wish to transmit $m(B)$ so long as

$$1 - F\left(\frac{\alpha}{8}\right) \geq \alpha\left[1 - F\left(s^{PPCpp}\right)\right].$$

This inequality is always satisfied. A type B advertiser that sends $m(A \cup AB)$ realises

$$\lambda^1(B, \cdot) = \frac{1}{2}\left[\alpha + F\left(s^{PPCpp}\right)\right].$$

The resulting profits for the b-type advertiser are then

$$\frac{1}{8}\left[\alpha + F\left(s^{PPCpp}\right)\right] - \frac{1}{2}\left[1 + F\left(s^{PPCpp}\right)\right] b^{(n-1)}. \tag{4.15}$$

Substituting in the equilibrium values of $b^{(n-1)}$ for each type of second-placed firm reveals that (4.15) is always negative for any F. In particular, (4.15) is clearly negative for $b^{(n-1)} = 1/4$; if the second highest bid is placed by a single-product firm then $b^{(n-1)}$ is of the form

$$\frac{1 + \alpha F(x)}{4\left(1 + F(x)\right)}.$$

Substituting this into (4.15), rearranging and expanding reveals that (4.15) is positive when

$$\alpha - 1 + F\left(s^{PPCpp}\right) F(x) - \alpha F\left(s^{PPCpp}\right) F(x) > 0,$$

which is never true for $\alpha < 1$. Thus, conditional on having won the auction, no firm wishes to deviate from the equilibrium reporting strategy and, given this, the bids computed above are also optimal.

Equilibrium 12 *When the publisher is constrained to use a PPC fee structure,* *there exists an (asymmetric) partial pooling equilibrium in which advertiser types A and AB transmit one signal, and advertisers of type B transmit a second, distinct message. In such an equilibrium, firms bid in accordance with (4.4). B-type consumers visit if they observe $m(B)$ or have $s_i \leq s^{PPCpp}$, A-types visit if they observe $m(A \cup AB)$ or have $s_i \leq \alpha/8$. Conditional on having visited, consumers purchase if $\hat{v}_i > p^*$.*

By symmetry, the analogous equilibrium with $m(AB) = m(B) = m(B \cup AB) \neq m(A)$ can also by supported under pay-per-click.

Consumer surplus from equilibrium 12 is calculated is a similar manner to the equilibria above:

$$CS = \underbrace{\frac{1-\gamma^n}{2}\left[\frac{1}{8} - \int_0^{1/8} sf(s)\,ds + \frac{1}{8}F\left(s^{PPCpp}\right) - \int_0^{s^{PPCpp}} sf(s)\,ds\right]}_{\text{Advertiser is type } AB} +$$
$$\underbrace{\frac{1}{2}\left(\frac{\gamma}{2}\right)^n\left[\frac{1}{8} - \int_0^{1/8} sf(s)\,ds + \frac{\alpha}{8}F\left(s^{PPCpp}\right) - \int_0^{s^{PPCpp}} sf(s)\,ds\right]}_{\text{Advertiser is type } A} +$$
$$\underbrace{\frac{\gamma^n}{2}\left(1 - \frac{1}{2^n}\right)\left[\frac{1}{8} - \int_0^{1/8} sf(s)\,ds + \frac{\alpha}{8}F\left(\frac{\alpha}{8}\right) - \int_0^{\alpha/8} sf(s)\,ds\right]}_{\text{Advertiser is type } B}.$$

Calculation of producer profit—$E\left[b^{(n-1)}\lambda^2(\cdot)\right]$—is more laborious since it depends on the type of both the first and second placed bidders. There are six relevant alternative cases. Firstly, the probability of there being two or more AB type firms is calculated as the probability of success in at least two out of n Bernoulli trials, with success probability $1-\gamma$:

$$\Pr(1^{\text{st}}\ AB, 2^{\text{nd}}\ AB) = \sum_{k=2}^{n}\binom{n}{k}(1-\gamma)^k\gamma^{n-k} = 1 - \left(\gamma^n + n(1-\gamma)\gamma^{n-1}\right).$$

The probability of the winner being type AB, and the per-click price being set by a B-type firm is calculated at the product of the probability of there being precisely one AB-type firm and the (conditional) probability of there being at least one B type firm in the remaining $n-1$ firms:

$$\Pr(1^{\text{st}}\ AB, 2^{\text{nd}}\ B) = n(1-\gamma)\gamma^{n-1} \times \left(1 - \frac{1}{2^{n-1}}\right).$$

The case with one AB firm and $n-1$ A-type firms has a probability that can be calculated analogously:

$$\Pr(1^{\text{st}}\ AB, 2^{\text{nd}}\ A) = n(1-\gamma)\gamma^{n-1} \times \left(\frac{1}{2^{n-1}}\right).$$

The probability that both the first and second placed firms are of type B is the product of the probability of there being no AB-type firm, and the probability of there being at least two B-type firms conditional on all firms being single-product:

$$\Pr(1^{\text{st}}\ B, 2^{\text{nd}}\ B) = \gamma^n \sum_{k=2}^{n} \binom{n}{k} \left(\frac{1}{2}\right)^n = \frac{\gamma^n(2^n - n - 1)}{2^n}.$$

The first and second placed firms are respectively of type B and A when there are no AB-type firms and there is precisely one B-type firm:

$$\Pr(1^{\text{st}}\ B, 2^{\text{nd}}\ A) = \gamma^n \binom{n}{1} \left(\frac{1}{2}\right)^n = n\left(\frac{\gamma}{2}\right)^n.$$

Last is the case in which all firms are of type A:

$$\Pr(1^{\text{st}}\ A, 2^{\text{nd}}\ A) = \left(\frac{\gamma}{2}\right)^n.$$

Piecing these together, along with the bids and λ^2's calculated above yields

$$\text{(4.16)}\quad \pi^p = \underbrace{\left[1-\left(\gamma^n + n(1-\gamma)\gamma^{n-1}\right)\right]\left[\frac{1}{4}\right]\left[\frac{1}{2}\left(1+F\left(s^{PPCpp}\right)\right)\right]+}_{1^{\text{st}}\ AB, 2^{\text{nd}}\ AB}$$
$$\underbrace{n(1-\gamma)\gamma^{n-1}\left(1-\frac{1}{2^{n-1}}\right)\left[\frac{1+\alpha F\left(\frac{\alpha}{8}\right)}{4\left(1+F\left(\frac{\alpha}{8}\right)\right)}\right]\left[\frac{1}{2}\left(1+F\left(s^{PPCpp}\right)\right)\right]+}_{1^{\text{st}}\ AB, 2^{\text{nd}}\ B}$$
$$\underbrace{n(1-\gamma)\gamma^{n-1}\left(\frac{1}{2^{n-1}}\right)\left[\frac{1+\alpha F\left(s^{PPCpp}\right)}{4\left(1+F\left(s^{PPCpp}\right)\right)}\right]\left[\frac{1}{2}\left(1+F\left(s^{PPCpp}\right)\right)\right]+}_{1^{\text{st}}\ AB, 2^{\text{nd}}\ A}$$
$$\underbrace{\frac{\gamma^n(2^n-n-1)}{2^n}\left[\frac{1+\alpha F\left(\frac{\alpha}{8}\right)}{4\left(1+F\left(\frac{\alpha}{8}\right)\right)}\right]\left[\frac{1}{2}\left(1+F\left(\frac{\alpha}{8}\right)\right)\right]+}_{1^{\text{st}}\ B, 2^{\text{nd}}\ B}$$
$$\underbrace{n\left(\frac{\gamma}{2}\right)^n\left[\frac{1+\alpha F\left(s^{PPCpp}\right)}{4\left(1+F\left(s^{PPCpp}\right)\right)}\right]\left[\frac{1}{2}\left(1+F\left(\frac{\alpha}{8}\right)\right)\right]+}_{1^{\text{st}}\ B, 2^{\text{nd}}\ A}$$
$$\underbrace{\left(\frac{\gamma}{2}\right)^n\left[\frac{1+\alpha F\left(s^{PPCpp}\right)}{4\left(1+F\left(s^{PPCpp}\right)\right)}\right]\left[\frac{1}{2}\left(1+F\left(s^{PPCpp}\right)\right)\right]}_{1^{\text{st}}\ A, 2^{\text{nd}}\ A}.$$

A multi-product firm makes a profit if and only if it faces no other such firm, and faces a lower price per click if all of its rivals are of type A. Thus, profits for a multi-product firm are

$$\pi = \gamma^{n-1}\left(1-\frac{1}{2^{n-1}}\right)\left[\frac{1}{8}\left(1+F\left(s^{PPCpp}\right)\right)-\frac{1}{2}\left(1+F\left(s^{PPCpp}\right)\right)\frac{1+\alpha F\left(\frac{\alpha}{8}\right)}{4\left(1+F\left(\frac{\alpha}{8}\right)\right)}\right]+$$
$$\left(\frac{\gamma}{2}\right)^{n-1}\left[\frac{1}{8}\left(1+F\left(s^{PPCpp}\right)\right)-\frac{1}{8}\left(1+\alpha F\left(s^{PPCpp}\right)\right)\right].$$

A-type firms make zero profit. B-type firms, in contrast, are able to make a profit when facing only A-type rivals. Expected profits, then, amount to

$$\pi = \left(\frac{\gamma}{2}\right)^{n-1}\left[\frac{1}{8}\left(1+\alpha F\left(\frac{\alpha}{8}\right)\right)-\frac{1}{2}\left(1+F\left(\frac{\alpha}{8}\right)\right)\frac{1+\alpha F\left(s^{PPCpp}\right)}{4\left(1+F\left(s^{PPCpp}\right)\right)}\right].$$

As usual, social welfare can be expressed by taking the equation for consumer surplus and replacing the $1/8$ units of consumer surplus (generated by a realisation of $\hat{v}_i > 0$) with the $3/8$ units of social surplus (generated by the same event).

$$W = \frac{1-\gamma^n}{2}\left[\frac{3}{8} - \int_0^{1/8} sf(s)\,ds + \frac{3}{8}F\left(s^{PPCpp}\right) - \int_0^{s^{PPCpp}} sf(s)\,ds\right] +$$
$$\frac{1}{2}\left(\frac{\gamma}{2}\right)^n\left[\frac{3}{8} - \int_0^{1/8} sf(s)\,ds + \frac{3\alpha}{8}F\left(s^{PPCpp}\right) - \int_0^{s^{PPCpp}} sf(s)\,ds\right] +$$
$$\frac{\gamma^n}{2}\left(1-\frac{1}{2^n}\right)\left[\frac{3}{8} - \int_0^{1/8} sf(s)\,ds + \frac{3\alpha}{8}F\left(\frac{\alpha}{8}\right) - \int_0^{\alpha/8} sf(s)\,ds\right].$$

A.2 Pooling (babbling) in a PPC environment

I look for a 'babbling' equilibrium of the PPC regime in which all θ transmit $m(\theta) = \widetilde{m}$. PPC advertiser profits are given in (4.3), which implies that the optimal (i.e. weakly dominant) bid is of the form given in (4.4).

Now, given that all advertiser types use the same messaging strategy in a babbling equilibrium, the set of consumers that visit (and hence λ^2) must be invariant to the type of the advertiser and symmetric across consumer types. A single-product firm therefore matches with half of its visitors, and also provides a positive $\hat{v}_i$ for a fraction α of the remaining half, implying

$$\lambda^1\left(A, \lambda^2(\cdot)\right) = \lambda^1\left(B, \lambda^2(\cdot)\right) = \left(\frac{1}{2} + \frac{\alpha}{2}\right)\lambda^2\left(\widetilde{m}, C, \widetilde{\mu}\right),$$

where $\widetilde{\mu}$ denotes the consumer's belief over firm types when $m(\theta) = \widetilde{m}\forall\theta$. From (4.4), we can therefore compute

$$b(A,C) = b(B,C) = \frac{1+\alpha}{8}.$$

Multi-product firms match with all consumers that visit them so that

$$\lambda^1\left(AB, \lambda^2(\cdot)\right) = \lambda^2\left(\widetilde{m}, C, \widetilde{\mu}\right),$$

and

$$b(AB, C) = \frac{1}{4}.$$

It follows that the winning bidder will be an AB-type, provided that such a firm exists. The posterior probability, $\mu(AB|C,\widetilde{m})$, that the winning advertiser is a multi-product firm is therefore $1-\gamma^n$. Since A- and B-type firms both submit equal bids, the posterior probability assigned to each type having won the auction is $\gamma^n/2$. A given consumer thus makes (in expectation)

$$U_i = (1-\gamma^n)\frac{1}{8} + \frac{\gamma^n}{2}(1+\alpha)\frac{1}{8} - s_i \equiv s^{PPCb} - s_i. \tag{4.17}$$

Consumers will click on the advertiser's link whenever this is positive, implying

$$\lambda^2(\widetilde{m}, C, \widetilde{\mu}) = F\left(s^{PPCb}\right).$$

Since firm's messages are ignored in a babbling equilibrium, there is no profitable deviation in m. Given this, the bids calculated above are optimal.

Equilibrium 13 *When the publisher is constrained to use a PPC fee structure, there exists a babbling equilibrium in which firms bid up to (4.4), with the successful firm transmitting $m(\theta) = \widetilde{m}$ and setting $p = p^*$. Consumers update their beliefs to $\widetilde{\mu}$, visit if and only if $s_i \leq s^{PPCb}$, and purchase (conditional on having visited) if $\hat{v}_i \geq p^*$.*

Each visiting consumer receives an expected payoff of s^{PPCb}, and must pay s_i. Overall consumer surplus is therefore

$$CS = F\left(s^{PPCb}\right)s^{PPCb} - \int_0^{s^{PPCb}} sf(s)\,ds.$$

Profit for the publisher is calculated as $E(b^{(n-1)})\lambda^2(\cdot)$, where $\lambda^2(\cdot)$ is outside of the expectations operator since it is constant in a babbling equilibrium. If there are two or more multi-product firms then $b^{(n-1)} = 1/4$, otherwise $b^{(n-1)} = (1+\alpha)/8$. The probability of there being no AB firm is γ^n, and the probability of there being exactly one is $n\gamma^{n-1}(1-\gamma)$. Expected publisher profit can thus be written as

$$\pi^P = \left[\left(\gamma^n + n\gamma^{n-1}(1-\gamma)\right)\frac{1+\alpha}{8} + \left[1-\left(\gamma^n + n\gamma^{n-1}(1-\gamma)\right)\right]\frac{1}{4}\right]F\left(s^{PPCb}\right).$$

Single-product firms make zero profits when all firms play as described above. A multi-product advertiser makes positive profits if and only if it is the only such firm, in which case it also wins the auction with certainty. Expected multi-product firm profits are therefore

$$\pi(AB, \widetilde{m}, p^*, C, \widetilde{\mu}) = \gamma^{n-1}\left(\frac{1}{4} - \frac{1+\alpha}{8}\right)F(s^{PPCb}),$$

where the term within the parentheses is the difference between gross producer surplus per click and the bid of a single-product firm.

Lastly, total social welfare can be calculated as above, noting that each visiting consumer with $\hat{v}_i > 0$ generates

$$\int_0^{p^*} 1 - q\, dq = \frac{3}{8}$$

units of surplus, and expends s_i units on search. Thus, total welfare is

$$\begin{aligned} W &= \left[(1-\gamma^n) F(s^{PPCb}) + \gamma^n F(s^{PPCb}) \frac{1+\alpha}{2} \right] \frac{3}{8} - \int_0^{s^{PPCb}} s f(s)\, ds \\ &= 3F\left(s^{PPCb}\right) s^{PPCb} - \int_0^{s^{PPCb}} s f(s)\, ds. \end{aligned}$$

A.3 Babbling in a PPI environment

When (4.10) is not satisfied, no information transmission can take place in a (pure strategy) equilibrium of the pay-per-impression environment, and any equilibrium must involve advertiser babbling.

The optimal bid under PPI is given in (4.9). If all firms send $m(\theta) = \widetilde{m}$, and price at p^* then the consumer's problem in a PPI babbling equilibrium is identical to that in the PPC case of section A.2. It follows that the λ's, and hence the ranking of firm's bids ($b(AB, \cdot) > b(A, \cdot) = b(B, \cdot)$) will also be the same in the two cases.

Equilibrium 14 *When the publisher is constrained to use a PPI fee structure, there exists a babbling equilibrium in which firms bid up to (4.9), with the successful firm transmitting $m(\theta) = \widetilde{m}$ and setting $p = p^*$. Consumers do not update their beliefs, visit if and only if $s_i \leq s^{PPCb}$, and purchase (conditional on having visited) if $\hat{v}_i \geq p^*$.*

That consumers face the same price and information set, and that each type of firm wins the auction with the same probability as in the PPC case implies that consumer surplus in equilibria 13 and 14 are identical. In equilibrium 13, the publisher makes

$$\pi^p = E(b^{(n-1)})\lambda^2(\cdot),$$

whereas in equilibrium 14 publisher profit is

$$\pi^p = E(b^{(n-1)}).$$

Noting that $b(\theta, I) = \lambda^2(\cdot)b(\theta, C)\forall\theta$ reveals that the two equilibria are also equivalent from the publisher's perspective. A similar exercise obtains for firm profits: in both equilibria, multi-product advertisers make revenues of $\lambda^1(AB, \lambda^2(\cdot))/4$. The payment to be made in a pay-per-click equilibrium is $\lambda^2(\cdot)b^{(n-1)}$, and the payment under PPI is $b^{(n-1)}$. Lastly, since consumers, the publisher, and firms enjoy the same surplus under equilibria 13 and 14, it follows that total social surplus is also equal in both.

A.4 Babbling in a PPS environment with sale-weighting

In this section I characterise a babbling equilibrium under PPS. Given the expression for gross consumer surplus derived in section III, and that $b(\theta, S) = 1/2$, an individual consumer expects to make

$$U_i = (1-\gamma)\frac{1}{32} + \gamma\left(\frac{1+\alpha}{2}\right)\frac{1}{32} - s_i = s^{PPSb} - s_i.$$

Equilibrium 15 *When the publisher is constrained to use a PPS fee structure and institutes sale-weighting, there exists a babbling equilibrium in which all advertiser types transmit message $\widetilde{m}$. In such an equilibrium, firms bid $1/2$ and the advertiser sets $p = p^* = 3/4$. Consumers visit if they have $s_i \leq s^{PPSb}$ and, conditional on having visited, purchase if $\hat{v}_i > p^*$.*

It follows immediately that

$$\text{(4.18)}\quad \lambda^2\left(\widetilde{m}, S, \widetilde{\mu}\right) = F\left(s^{PPSb}\right) < F\left(s^{PPCb}\right) = \lambda^2\left(\widetilde{m}, C, \widetilde{\mu}\right) = \lambda^2\left(\widetilde{m}, I, \widetilde{\mu}\right),$$

so that the number of visitors is lower under PPS babbling than the corresponding PPC/PPI equilibria. The reasoning for this result is twofold: (i) as noted by Dellarocas and Viswanathan (2008), a per-sale fee functions as a marginal cost for the advertiser: distorting prices upwards and reducing the consumers' expected surplus, and (ii) PPS auctions fail to systematically select multi-product firms and therefore reduce each consumer's probability of finding a match. Overall consumer surplus is given by

$$CS = F\left(s^{PPSb}\right)s^{PPSb} - \int_0^{s^{PPSb}} sf(s)\,ds.$$

From (4.2), the monopolist price is $p^* = 3/4$, so that a visiting consumer finding $\hat{v}_i > 0$ buys with probability $1/4$. Expected profits for the publisher are calculated as the product of the per-sale fee $(1/2)$, the rate of sales to visitors with $\hat{v}_i > 0$, and the expected value of λ^1:

$$\pi^p = \frac{1}{2} \times \frac{1}{4} \times \left[(1-\gamma) + \gamma\left(\frac{1+\alpha}{2}\right)\right] F\left(s^{PPSb}\right) = 4F\left(s^{PPSb}\right) s^{PPSb}. \tag{4.19}$$

In equilibrium, each firm wins the sale-weighted auction with probability $1/n$. Firms sell to $1/4$ of the consumers that realise $\hat{v}_i > 0$, and make $p^* - (1/2) = 1/4$ per sale. Thus, multi-product firm profits are

$$\pi\left(AB, \widetilde{m}, p^*, I, \widetilde{\mu}\right) = \frac{F\left(s^{PPSb}\right)}{16n},$$

and single-product firm profits are

$$\pi\left(\theta, \widetilde{m}, p^*, I, \widetilde{\mu}\right) = \left(\frac{1+\alpha}{2}\right) \frac{F\left(s^{PPSb}\right)}{16n}$$

for $\theta \in \{A, B\}$.

To calculate social welfare, I begin by integrating under the demand curve implied by the distribution of valuations in order to determine the expected surplus generated by each match:

$$\int_0^{1/4} 1 - q\,dq = \frac{7}{32}.$$

Welfare itself is calculated analogously to consumer surplus as follows:

$$W = 7F\left(s^{PPSb}\right) s^{PPSb} - \int_0^{s^{PPSb}} sf(s)\,ds$$

B Publisher Profits Under PPS Versus PPC and PPI

In this section I prove two propositions regarding the profitability for the publisher of equilibria under PPS:

Proposition 35 *Publisher profits are always at least as high under the informative PPC equilibrium as under either equilibrium 15 or 10.*

Proof. Publisher profit under pay per sale is given by the product of the per-sale price, the mass of visitors realising $\hat{v}_i > 0$, and the proportion of these that will ultimately buy thus:

$$\pi^p(S) = \frac{1}{2} \times E(\lambda^1) \times \frac{1}{4} = \frac{E(\lambda^1)}{8}.$$

It follows that

$$\pi^p(S) \leq \frac{1}{8}.$$

In the informative PPC equilibrium, both firm bids and the number of visitors are independent of γ and n. Increasing n or decreasing γ increases the probability of realising a multi-product advertiser, and therefore increases profits. Thus, informative equilibrium profits are at a minimum when $n = 2$ and $\gamma = 1$, in which case

$$\pi^p = \left[\frac{1 + \alpha F\left(\frac{\alpha}{8}\right)}{4\left(1 + F\left(\frac{\alpha}{8}\right)\right)}\right] \frac{1}{2}\left[1 + F\left(\frac{\alpha}{8}\right)\right] = \frac{1 + \alpha F\left(\frac{\alpha}{8}\right)}{8}.$$

Given the assumption placed upon α and F (i.e. that $\alpha > 0$ and F has positive density in $[0, 1/8]$), this expression is greater than $1/8$, and the proposition is verified as true. ■

Proposition 36 *Publisher profits are always at least as high under the PPC/PPI babbling equilibria as under either equilibrium 15 or 10.*

Proof. I prove the result for the PPC case. The PPI result then immediately follows from proposition 32. Publisher profit under pay per sale is given by the product of the per-sale price, the mass of visitors realising $\hat{v}_i > 0$, and the proportion of these that will ultimately buy thus:

$$\pi^p(S) = \frac{1}{2} \times E(\lambda^1) \times \frac{1}{4} = \frac{E(\lambda^1)}{8}.$$

Profit under pay-per-click is given by the product of the second highest bid and the number of clicks

$$\pi^p = E(b^{(n-1)}\lambda^2).$$

Since $b > 1/8$ for all θ in a babbling equilibrium, and since it is necessarily the case in any given equilibrium that $\lambda^2 \geq \lambda^1$, it suffices to show that the number of visitors under a pay-per-click babbling equilibrium, $F(s^{PPCb})$, is never less than than under either of the PPS equilibria. The result has already been shown in (4.18) for the babbling PPS case. The number of visitors under PPS partial pooling is at most given by

$$(4.20)\quad \frac{1}{2}\left[F\left(\frac{1}{32}\right)+F\left(s^{PPSpp}\right)\right].$$

Since s^{PPSpp} is a weighted average of $1/32$ and $\alpha/32$, its value is no greater than $1/32$, and thus (4.20) can be no greater than F(1/32). From (4.17), $F(s^{PPCb})$ is never less than $F(1/16)$, and the proof is thus complete. ■

Bibliography

AGARWAL, N., S. ATHEY, AND D. YANG (2009): "Skewed Bidding in Pay Per Action Auctions for Online Advertising," *American Economic Review Papers and Proceedings*, 99(2), 441–447.

ANDERSON, S. P., AND S. COATE (2005): "Market Provision of Broadcasting: A Welfare Analysis," *Review of Economic Studies*, 72(4), 947–972.

ANDERSON, S. P., AND R. RENAULT (2006): "Advertising Content," *American Economic Review*, 96(1), 93–113.

ANWAR, S., R. MCMILLAN, AND M. ZHENG (2006): "Bidding Behavior in Competing Auctions: Evidence from eBay," *European Economic Review*, 50(2), 307–322.

ATHEY, S., AND G. ELLISON (2008): "Position Auctions with Consumer Search," Working Paper.

BAGWELL, K. (2007): "The Economic Analysis of Advertising," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. Porter, vol. 3, pp. 1701–1844. Elsevier.

BAJARI, P., AND A. HORTAÇSU (2003): "The Winner's Curse, Reserve Prices and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34(2), 329–355.

BARBARO, S., AND B. BRACHT (2006): "Shilling, Squeezing, Sniping: Explaining late bidding in online second-price auctions," Working Paper, January 31, 2006.

BAYE, M. R., AND J. MORGAN (2001): "Information Gatekeepers on the Internet and the Competitiveness of Homogeneous Product Markets," *American Economic Review*, 91(3), 454–474.

BAYE, M. R., J. MORGAN, AND P. SCHOLTEN (2006): "Information, Search, and Price Dispersion," in *Handbook on Economics and Information Systems*, ed. by T. Hendershott, pp. 323–371. Elsevier, Amsterdam, The Netherlands.

BEATH, J., AND Y. KATSOULACOS (1991): *The Economic Theory of Product Differentiation*. Cambridge University Press, Cambridge, UK.

CAILLAUD, B., AND B. JULLIEN (2001): "Competing Cybermediaries," *European Economic Review*, 45(4–6), 797–808.

——— (2003): "Chicken & Egg: Competition among Intermediation Service Providers," *RAND Journal of Economics*, 34(2), 309–328.

CHIANG, K. C. H., AND A. W. P. KUNG (2005): "Bidding dynamics in multi-unit auctions: empirical evidence from online auctions of certificates of deposit," *Journal of Financial Intermediation*, 14(2), 239–252.

CRAWFORD, V. P., AND J. SOBEL (1982): "Strategic Information Transmission," *Econometrica*, 50(6), 1431–1451.

DELLAROCAS, C., AND S. VISWANATHAN (2008): "The Holy Grail of Advertising? Allocative Efficiency and Revenue Implications of "Pay-per-Action" Advertising in Environments with Quality Uncertainty," Working Paper, February, 2008.

DEMANGE, G., D. GALE, AND M. SOTOMAYOR (1986): "Multi-Item Auctions," *Journal of Political Economy*, 94(4), 869–872.

DIAMOND, P. (1971): "A Model of Price Adjustment," *Journal of Economic Theory*, 3(2), 156–168.

EDELMAN, B., M. OSTROVSKY, AND M. SCHWARZ (2007): "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords," *American Economic Review*, 97(1), 242–259.

ELLISON, G., AND S. F. ELLISON (2009): "Search, Obfuscation, and Price Elasticities on the Internet," *Econometrica*, 77(2), 427–452.

ELLISON, G., AND A. WOLITZKY (2008): "A Search Cost Model of Obfuscation," Working Paper, September, 2008.

ELY, J. C., AND T. HOSSAIN (2009): "Sniping and Squatting in Auction Markets," *American Economic Journal: Microeconomics*, 1(2), 68–94.

FUDENBERG, D., AND J. TIROLE (1991): *Game Theory*. MIT Press, Cambridge, Massachusetts.

GABSZEWICZ, J. J., P. G. GARELLA, AND N. SONNAC (2005): "Newspapers' Market Shares and the Theory of the Circulation Spiral," Core Discussion Paper 2005-59.

GABSZEWICZ, J. J., D. LAUSSEL, AND N. SONNAC (2001): “Press Advertising and the Ascent of the Pensée Unique,” *European Economic Review*, 45(4–6), 641–651.

GANDAL, N. (2001): “The Dynamics of Competition in the Internet Search Engine Market,” *International Journal of Industrial Organization*, 19(7), 1103–1117.

GEHRIG, T., AND R. STENBACKA (2005): “Two at the Top: Quality Differentiation in Markets with Switching Costs,” University of Copenhagen Centre for Industrial Economics Discussion Paper 2005-09.

GIBBONS, R. (1992): *A Primer in Game Theory*. Pearson Education, Harlow, England.

GROSSMAN, G. M., AND C. SHAPIRO (1984): “Informative Advertising with Differentiated Products,” *Review of Economic Studies*, 51(1), 63–81.

GUL, F., AND E. STACCHETTI (1999): “Walrasian Equilibrium with Gross Substitutes,” *Journal of Economic Theory*, 87(1), 95–124.

——— (2000): “The English Auction with Differentiated Commodities,” *Journal of Economic Theory*, 92(1), 66–95.

HOSSAIN, T. (2008): “Learning by Bidding,” *RAND Journal of Economics*, 39(2), 509–529.

HOTELLING, H. (1929): “Stability in Competition,” *The Economic Journal*, 39(153), 41–57.

KLEMPERER, P. (2004): *Auctions: Theory and Practice*. Princeton University Press, Princeton, New Jersey.

KRISHNA, V. (2002): *Auction Theory*. Academic Press, Sand Diego, California.

LEONARD, H. B. (1983): “Elicitation of Honest Preferences for the Assignment of Individuals to Positions,” *Journal of Political Economy*, 91(3), 461–479.

MEURER, M., AND D. O. STAHL (1994): “Informative Advertising and Product Match,” *International Journal of Industrial Organization*, 12(1), 1–19.

MILGROM, P. (2004): *Putting Auction Theory to Work*. Cambridge University Press, Cambridge, UK.

MILGROM, P., AND J. ROBERTS (1986): “Price and Advertising Signals of Product Quality,” *Journal of Political Economy*, 94(4), 796–821.

NELSON, P. (1974): "Advertising as Information," *Journal of Political Economy*, 82(4), 729–754.

NICKALLS, R. W. D. (1993): "A New Approach to Solving the Cubic: Cardan's Solution Revealed," *The Mathematical Gazette*, 77, 354–359.

OCKENFELS, A., AND A. E. ROTH (2006): "Late and Multiple Bidding in Second Price Internet Auctions: Theory and Evidence Concerning Different Rules for Ending an Auction," *Games and Economic Behavior*, 55(2), 297–320.

PARKES, D. C. (2006): "Iterative Combinatorial Auctions," in *Combinatorial Auctions*, ed. by P. Cramton, Y. Shoham, and R. Steinberg, pp. 41–77. MIT Press, Cambridge, Massachusetts.

PETERS, M., AND S. SEVERINOV (1997): "Competition among Sellers Who Offer Auctions Instead of Prices," *Journal of Economic Theory*, 75(1), 141–179.

PHLIPS, L. (1989): *The Economics of Imperfect Information*. Cambridge University Press, Cambridge, UK.

POLLOCK, R. (2008): "Is Google the next Microsoft? Competition, Welfare and Regulation in Internet Search," Working Paper, September 23, 2008.

ROCHET, J.-C., AND J. TIROLE (2003): "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1(4), 990–1029.

ROTH, A. E., AND A. OCKENFELS (2002): "Last-minute bidding and the rules for ending second-price auctions: Evidence from eBay and Amazon auctions on the Internet," *American Economic Review*, 92(4), 1093–1103.

SCHMALENSEE, R. (1978): "A Model of Advertising and Product Quality," *Journal of Political Economy*, 86(3), 485–503.

SHAKED, A., AND J. SUTTON (1982): "Relaxing Price Competition Through Product Differentiation," *Review of Economic Studies*, 49(1), 3–13.

SHAPLEY, L. S., AND M. SHUBIK (1971): "The Assignment Game I: The Core," *International Journal of Game Theory*, 1(1), 111–130.

SUNDARARAJAN, A. (2003): "Pricing Digital Marketing: Information, Risk Sharing and Performance," Working Paper, 2003.

TELANG, R., U. RAJAN, AND T. MUKHOPADHYAY (2004): "The Market Structure for Internet Search Engines," *Journal of Management Information Systems*, 21(2), 137–160.

TIROLE, J. (1988): *The Theory of Industrial Organization*. MIT Press, Cambridge, Massachusetts.

VARIAN, H. R. (2007): "Position auctions," *International Journal of Industrial Organization*, 25(6), 1163–1178.

WHITE, A. (2008): "Search Engines: Left Side Quality versus Right Side Profits," Preliminary Draft, September 30, 2008.

WILCOX, R. T. (2000): "Experts and Amateurs: The Role of Experience in Internet Auctions," *Marketing Letters*, 11(4), 363–374.