THEORIES OF SELF-REGULATION

Amelia Fletcher

Nuffield College

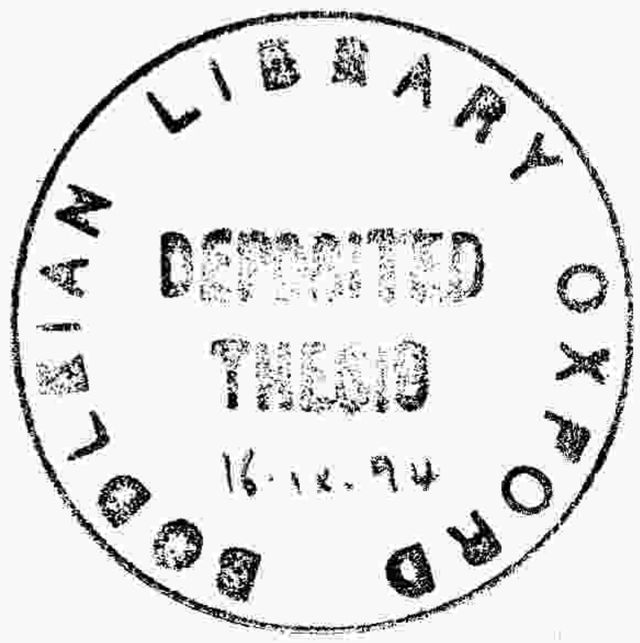
Oxford

December 1993

Thesis submitted

for the degree of

D. Phil. in Economics



THEORIES OF SELF-REGULATION - ABSTRACT

Self-regulation is becoming more and more prevalent in the UK, yet there has been very little theoretical examination of the issues involved. This thesis aims to fill that gap. The main intuition examined is that self-regulation will be a feasible form of regulation only if there exists some state, attainable via some form of regulation, which raises both industry profits and total social welfare, relative to the *laissez faire*.

Modelling self-regulation thus involves finding situations in which the *laissez faire* is not welfare-optimal; in which sellers can improve upon their profits, relative to those at the *laissez faire*; and in which this profit-improvement happens also to raise total social welfare. Such situations may exist when there is asymmetric information between sellers and consumers and/or externalities.

Chapter 3 examines quality choice under hidden irreversible action. It is shown that any comparison between the welfare under self-regulation and that under *laissez faire* will turn on the relative signs and values of (a) the net between-seller externality; (b) the difference in quality valuations between marginal and infra-marginal consumers; and (c) the effect of quality choice on total output. Under certain circumstances, however, self-regulation will certainly be the welfare-superior.

Chapter 4 examines quality choice under moral hazard. Self-regulation is shown to be feasible. It is also compared with external regulation: If qualities are observable (with a lag) but not verifiable, an SRO may have an advantage over an external regulator in that an SRO can condition its regulation on quality. Self-regulation in fact becomes more likely to be welfare-superior to external regulation (a) as fraud becomes more important relative to cheating on quality; (b) as sellers become more homogeneous in terms of capital; and (c) as the observation lag becomes greater.

Chapter 5 considers the factors that might influence the optimal number of SROs in an industry.

ACKNOWLEDGEMENTS

I would like to thank both Professor John Vickers and Meg Meyer for supervising this thesis and for offering many valuable suggestions.

Financial support was provided by the ESRC.

CONTENTS

1. Introduction

1.1 Empirical motivation 1

1.2 A rationale for self-regulation 4

1.3 Modelling self-regulation 7

1.3.1 Modelling quality 9

1.3.2 Players and their maximands 29

1.3.3 Move order 32

1.4 Thesis layout 33

1.5 Concluding remark 37

2. Literature Survey

2.1 Introduction 38

2.2 Leland (1979) 39

2.2.1 Leland's framework 41

2.2.2 The unregulated equilibrium 42

2.2.3 Introducing minimum quality levels 43

2.2.4 Self-regulation 44

2.2.5 The welfare-implications of self-regulation 45

2.2.6 Decreasing opportunity costs 47

2.2.7 Discussion of Leland's model 48

2.3 Shaked and Sutton (1981) 50

2.3.1 The framework 50

2.3.2 Equilibrium 52

2.3.3 A comparison with Leland's model 54

2.4 Shapiro (1983) 57

2.4.1 Shapiro's framework 59

2.4.2 The unregulated competitive equilibrium 61

2.4.3 Welfare analysis of the equilibrium 63

2.4.4 Implications for regulation 65

2.4.5 An extension of Shapiro's model 66
2.4.6 Implications for self-regulation 70
2.5 Mayer and Neven (1991) 73
2.5.1 The model 74
2.5.2 Regulatory contracts 74
2.5.3 Self-regulation 77
2.5.4 Problems with this model 79
2.6 Two other models 83
2.6.1 Shaked and Sutton (1982) 83
2.6.2 Shapiro (1986) 86

3. Self-Regulation Under Hidden Irreversible Action

3.1 Introduction 91
3.2 The model 92
3.3 The one-stage self-regulation game 97
3.4 Quality comparisons, for given output 98
3.5 Quality comparisons when output is allowed to vary 100
3.5.1 Comparing equilibria under quality collusion 103
3.5.2 Comparing self-regulation with *laissez faire* 107
3.6 The two-stage self-regulation game 109
3.7 Endogenizing entry 113
3.7.1 Perfect free entry 115
3.7.2 Compulsory SRO membership as an entry barrier 121
3.7.3 Threat-constrained entry 125
3.7.4 Endogenizing entry - conclusions 127
3.A Microfoundations for the inverse demand function $p(g_i, \bar{k}, z_i, X)$ 127
3.A.1 The negative externality 128
3.A.2 The positive externality 132
3.A.3 A simple 2-seller example 134

4. Self-Regulation Under Moral Hazard

4.1 Introduction 137

4.1.1 Comparing different forms of regulation 141

4.1.2 Layout and main conclusions of chapter 142

4.2 The quality choice game with moral hazard 143

4.2.1 The simple two-player quality choice game 143

4.2.2 The multi-player quality choice game 146

4.2.3 An information collation role for regulation 147

4.3 A model of regulation under moral hazard 151

4.3.1 Framework 152

4.3.2 The *laissez faire* equilibrium under perfect information transmission 155

4.3.3 The regulator/seller game 157

4.3.4 The buyer/regulator game 160

4.3.5 A characterization of the regulatory system 161

4.3.6 External regulation versus self-regulation 167

4.4 The seller/regulator game with tit-for-tat strategies 178

4.4.1 The model 178

4.4.2 External regulation versus self-regulation 183

4.A1 Tit-for-tat and trigger strategies 184

4.A2 The model with endogenous price 188

5. The Number Of SROS

5.1 Introduction 193

5.1.1 Changing the number of SROs 193

5.1.2 Layout and main conclusions of chapter 194

5.2 Shaked and Sutton (1981) - the two-SRO extension 198

5.2.1 The framework 199

5.2.2 The two-SRO equilibrium 199

5.2.3 Results 200

5.2.4 Discussion 203

5.3 The number of SROs - a 'Hotelling' model 205
5.3.1 The model 205
5.4 Case I - The industry with one SRO 208
5.4.1 Full collusion within each SRO (regime m) 209
5.4.2 Collusion on quality only (regime r) 210
5.4.3 A brief comparison of regimes (m) and (r) for the one-SRO case 211
5.5 Case II - The industry with two SROs 211
5.5.1 Full collusion within each SRO (regime m) 213
5.5.2 Collusion on quality only (regime r) 215
5.6 Results for regime (m): full collusion within SROs 217
5.6.1 The effect of actual entry (Comparing cases I and II) 218
5.6.2 Reaction curves 225
5.6.3 Entry deterrence 226
5.7 Results for regime (r): Collusion within SROs on quality only 228
5.8 The model with Leland-style supply-side 232
5.8.1 The model 233
5.8.2 Case I: One SRO and no non-SRO sector 235
5.8.3 Case II: One SRO (r) and a potential non-SRO sector (u) 237
5.8.4 Results of allowing free entry: a comparison of cases I and II 239
5.8.5 N SROs with a non-SRO sector 242

References 247

NOTATION - chapter 1

a	effort/quality
a^*	effort under perfect information
δ	discrete time discount rate
$E(\cdot)$	expectations operator
$f(\theta)$	density function of θ
ϕ	per seller fixed costs
k	training/quality
$\bar{k}$	average quality
J	initial payment to SRO
L	minimum quality level
λ_i	weights
$\mu(\cdot)$	(opportunity) cost of quality function
n	number of sellers regulated
$p(\cdot)$	price function
p^e	entrance price
Π	total industry profits
π	per seller profits
q	quality
$\bar{q}$	expected/average quality
$\sigma(n, \pi)$	SRO's maximand
θ	innate quality
$\hat{\theta}$	maximal quality
$\underline{\theta}$	minimal quality
$U(\cdot)$	utility function
$Var(q)$	variety of qualities available
w	observation lag
W	welfare
x	seller output
X	total industry output

1. INTRODUCTION

1.1 Empirical motivation

In the UK, a wide range of industries is subject to some degree of self-regulation, including such diverse examples as accountancy, chiropody, dry cleaning, advertising, travel agents, medicine, electrical contracting, law, soft drinks manufacturing, and financial services. This differs from most other countries where strict *de jure* regulation is more common (ref. Blair and Rubin (1980), OECD (1985)).

Self-regulatory arrangements take various forms, and are of varying degrees of importance, in different industries:

— Membership of a self-regulatory organization (SRO) is a prerequisite of functioning in the industry for some industries, while it is optional for others.

— Some SROs ensure good consumer protection (e.g. by enforcing obligatory minimum quality/qualification levels, by specifying procedures to be followed, and/or by running a consumer compensation scheme). Other SROs offer consumers no such assurances or, more seriously, give unsubstantiated assurances.

— Most of the main professional bodies and some other SROs have statutory backing, while many do not (which may limit the sanctions an SRO can impose on disobedient members).

— Different SROs lay different degrees of stress upon preventing malpractice/fraud and upon raising average quality/levels of competence.

— Some SROs have a very heterogeneous membership (in terms of size, quality and/or sector of industry), while the membership of others is a lot less segmented.

— In some industries there is only one SRO while, in others, there are a few. These may be similar; they may vary according to the quality or sector of industry of their sellers; or they may provide different self-regulatory functions (in which case sellers may belong to more than one SRO).

— Some industries face the threat of external statutory regulation if self-regulation fails to achieve sufficiently high consumer welfare.

Currently, there is a growing amount of self-regulation (particularly by trade associations). At the same time, though, there has been much criticism of watchdogs/SROs/professional bodies. For example:

— In the wake of the Maxwell pensions plundering affair, IMRO, the SRO for investment managers, was much criticized for being insufficiently alert and responsive to information available, and for being too uncritical of its members.

— The Calcutt report (1993) recommended replacing the self-regulatory Press Complaints Commission with a statutory external regulator. This proposal has not so far been taken up, but it has had a significant effect on the way the Commission is run.

— A survey carried out in March 1993 showed that 1 in 5 Conservative MPs (and 83% of Labour MPs) believed that self-regulation of the personal financial services industry should be scrapped in favour of statutory control (ref. Financial Times, 9.3.93, p.10).

— In June 1993, the Law Society was sued [by Miss Peggy Wood] for its failure to investigate her complaints against a particular firm of solicitors.

Much has been written about self-regulation in practice, and academic books and articles have been written on its costs and benefits for particular industries (ref. Goodhart and Seldon (1988), Franks and Mayer (1989)).

There has, however, been surprisingly little written about the ideas and theories behind self-regulation. Some potential advantages of self-regulation over external regulation/legislation have been suggested:

—— Self-regulatory codes of practice can be prepared/adjusted much more quickly and at less cost than can legislation (and this will ease technological change and innovation).

—— Monitoring will be less costly, since an SRO will tend to be better informed about what is going on in its industry than will an external regulator.

—— Self-regulatory codes can be more subjective and flexible than external regulation can be; and they can be applied in spirit, rather than just in the letter.

—— Redress can be quicker and less costly than legal proceedings (which is especially important given the comparatively minor nature of many consumer offences).

—— SROs can not only ban positively detrimental conduct, but they can also encourage high standards in excess of the basic minimum to which external regulation must restrict itself.

These potential advantages, however, will only be realized if the SRO has an *incentive* to prepare codes of practice, monitor members, punish wrongdoers, and encourage high standards. The usual incentive for self-regulation given is the threat of external regulation.[1] This cannot be the whole story, though, since much successful self-regulation takes place in industries which are (pretty much) free of this threat.

---

[1] For example, Sir Gordon Borrie (the recently retired Director General of Fair Trading) has written "I would argue that the shadow of legal regulation is frequently a powerful incentive to an industry to 'set its house in order'".

This thesis examines an alternative rationale/incentive for self-regulation. It also surveys some old, and develops some new, theories which help to explain some of the basic facts about self-regulation described above.

1.2 A rationale for self-regulation

Equilibrium in a well-behaved perfectly competitive market, which has no informational problems and no externalities, will generally be pareto optimal. Thus there will be no need for regulation (although there may be a need for redistributive taxation). When sellers in an industry have some market power, however, or when there are informational problems, or when externalities exist, then the *laissez faire* equilibrium allocation may not maximize social welfare, and there may be a role for regulation (although, of course, the benefits of any proposed regulation must always be weighed up against its costs).

When there is a role for regulation, we need to consider what form this regulation should take. The literature on optimal external regulation is extensive (see Laffont and Tirole (1993)), and discusses (amongst other issues) the arguments for structural vs. conduct regulation (see Kay and Vickers (1988)); for price-cap vs. rate-of-return regulation; and for procurement vs. external regulation. This thesis tackles a different topic: the arguments for external regulation vs. self-regulation.

The essential point to be made about self-regulation, as opposed to external regulation, is that **we should never expect a rational self-regulated industry to act against its own interests**. Thus, whenever there is a direct conflict between what is best for social welfare and what is best for the industry, self-regulation cannot provide a solution. For example, the regulation of the pricing policies of monopoly utilities must necessarily be external, since the utility will always want higher prices (as long as price is below its monopoly value), while social welfare maximization will require lower prices (as long as price exceeds marginal cost).

Some *laissez faire* equilibria, however, may be such that certain movements away from the equilibrium will be in the interests of both the industry and its consumers. So an industry which acts purely in its own interest may also happen to improve social welfare. This is the rationale behind self-regulation, as I conceive it.

Throughout this thesis, then, I model self-regulation as a kind of cartel. A cartel involves sellers in an industry acting together to improve on the equilibrium allocation of profits which results when they all act non-cooperatively. When social welfare is increased as a by-product of this behaviour, it can be described as self-regulation.

To assess self-regulation properly, it needs to be compared not only with the *laissez faire* allocation, but also with external regulation and the welfare optimum. The following conditions are clearly necessary (though not sufficient) for self-regulation to be welfare-superior to external regulation:

(A) *Laissez faire* does not achieve welfare maximization;

(B) External regulation does not achieve welfare maximization;

(C) Equilibrium under self-regulation is different from, and welfare-superior to, that under *laissez faire*.

At some points in this thesis, I make direct comparisons between self-regulation and external regulation. More often, though, I model the equilibrium under self-regulation, and examine how it will compare with the *laissez faire* equilibrium and with the welfare optimum, but I do not go beyond this to make comparisons with external regulation. That is, for much of this thesis, I examine conditions (A) and (C) in some detail, but I hardly examine condition (B) at all.

The reason for omitting discussion of condition (B) (and also the direct comparison of self-regulation with external regulation) is that there are

essentially five reasons why external regulation may not achieve welfare-maximization:

(1) The regulator may not be omniscient. This is the main topic discussed in the regulation literature. In particular, the regulator may not know

(a) Industry costs (ref. Averch and Johnson (1962), Baron and Myerson (1982), Laffont and Tirole (1993) chap 2);

(b) Technological possibilities (ref. Baron and Besanko (1984), Laffont and Tirole (1988));

(c) Quality (ref. Spence (1975), Laffont and Tirole (1993) chap 4);

(d) Consumer demand functions (ref. Lewis and Sappington (1988), Spulber (1988)).

(2) The regulator may not be benevolent, especially if efficient regulation is costly or if regulation is carried out directly by government (ref. Demski and Sappington (1987), Spiller (1990), Laffont and Tirole (1990), Besanko and Spulber (1989)).

(3) External regulation may be very costly, in terms of gathering information, overcoming industry resistance, and formulating policies/ contracts (especially when optimal contracts need to depend on contingencies which are hard to foresee and formulate clearly);

(4) The regulator may be unable to pre-commit *ex ante* not to expropriate any sunk investments *ex post*, which may prevent efficient investment in the first place (ref. Gilbert and Newbery (1988), Salant and Woroch (1992));

(5) Even when observable, quality may not be verifiable, so that formal regulation, which needs to be conditioned upon quality, is impossible (ref. Lewis and Sappington (1988, 1991)).

Thus there are many different situations in which external regulation can diverge from the welfare optimum. For any particular industry, regulatory policy analysis needs to examine not only the expected self-regulation equilibrium, but also the external regulation equilibrium, which is likely to be different for each industry. Here, then, my aim is simply to examine self-regulation, and to leave comparisons with external regulation up to policy analysts.[2]

There is one main exception to this, in which a direct comparison between the two forms of regulation is fruitful. Consider a situation in which quality is observable but not verifiable (case (5) above). An external regulator may be unable to condition its regulation on sellers' qualities, since external regulation must usually be explicitly specified in a contract and enforceable by a third party (such as the courts or government). A self-regulatory organization (SRO), on the other hand, may be able to regulate in a looser and less explicit fashion, which means that it *may* be able to condition its regulation on sellers' qualities. In chapter 4, I consider whether self-regulation will tend to be welfare superior to external regulation in this situation, even when the external regulator is assumed to be totally benevolent and the SRO is assumed to be completely self-interested.

1.3 Modelling self-regulation

There are various ways in which a self-regulated industry may act differently from an atomistic unregulated industry (the *laissez faire*). These can be divided into two groups:

---

[2] Note that the crux of the comparison is whether an SRO's well-informed, less restricted pursuit of self-interest will be better or worse than an external regulator's ill-informed, more restricted, pursuit of welfare/ political objectives.

(1) An SRO may have better incentives to achieve and maintain socially beneficial levels of some variable z than individual sellers do because:

(a) It may be easier (or there may be stronger incentives) for consumers/media/government to observe a whole SRO's z (or average z) than the z of any individual seller;

(b) It may be easier (or there may be stronger incentives) for consumers/media/government to punish an SRO with undesirable z (or average z) than any individual seller;

(c) An SRO may enable the internalization of between-seller externalities, both positive and negative (collusion);

(d) An SRO may have proportionately less to gain and more to lose from milking its reputation (for desirable z) than individual sellers have;

(e) An SRO may care more about the future than (some) individual sellers do.

(2) An SRO may be better able to control what its members do (and who its members are) than the free market is because:

(a) An SRO may be better able both to observe and to punish sellers of undesirable z than consumers are, so that the set of credible z promises which can be made by sellers to consumers is enlarged;

(b) An SRO may have better access to channels of information dissemination (to consumers) than individual sellers do;

(c) An SRO may be able to impose minimum z levels on its members (or, at least, minimum levels of another variable y which will affect z in a desirable way);

(d) An SRO may be able to change how much its members care about the future (e.g. by imposing upwards sloping age-wage profiles on its members);

(e) An SRO may be able to prevent some sellers from entering the industry, even when these sellers could make positive profits.

The first step in building a model of self-regulation is to decide which of these differences between self-regulation and *laissez faire* to include.

The second step is to decide what z should denote. Decision variables for which self-regulation may be useful include vertical quality, horizontal quality, advertising, pollution control and R&D. The rest of this thesis looks at vertical quality only (which is hereafter just referred to as quality), but it should be remembered that some of the same arguments will apply to horizontal quality, advertising, pollution control and R&D. Indeed, research joint ventures could be seen as a kind of self-regulation of R&D, and some of the issues which arise in this thesis are also important in the literature on research joint ventures (ref. Ordover and Willig (1985), Grossman and Shapiro (1986), Katz (1986)).

1.3.1 Modelling quality

When consumers care about the quality (or variety of qualities) available in an industry, self-regulation will generally be socially beneficial only if the industry involves asymmetric information effects and/or externality effects. In this section, I examine these effects in some detail, and consider how they might be modelled. I divide the externality effects into:

— Between-seller externalities: These arise when one seller's payoffs are affected by the actions of other sellers (usually via the demand function).

—— Seller-consumer externalities: These arise when consumers are heterogeneous and sellers cannot price-discriminate (which means they cannot capture all of the consumer surplus).

Note that since my aim here is to give a general overview of theories of self-regulation, I make more modelling suggestions here than are actually followed up in the rest of this thesis.

A. Asymmetric information effects

Modelling asymmetric information involves making certain assumptions about how quality is actually determined and about its degree of reversibility. Essentially there are three possible components of quality:

(1) Intrinsic quality θ, chosen by nature and unchangeable by sellers;

(2) (Relatively) irreversible quality k, chosen by sellers and hard to change in the short run;

(3) Reversible quality a, chosen by sellers and easy to change.

It is easiest to think of θ as ability, k as training, and a as effort. For simplicity in what follows, all of these components are assumed to be deterministic, and quality q is taken to be equal to one, or the sum of more than one, of these three variables.

These three components of quality roughly correspond to three important potential quality problems:

(A1) Adverse selection;

(A2) Hidden irreversible action;[3] and

(A3) Moral hazard;

---

[3] Expressed loosely, this effect lies somewhere between adverse selection an moral hazard. The term is my own, but is relatively self-explanatory.

A fourth problem (A4), resulting from consumer risk aversion, may occur (for any form of quality) whenever consumers can neither observe nor deduce exactly what quality they are going to receive.

In addition to their different assumptions about quality reversibility, these four quality problems also involve different assumptions about what exactly consumers can observe.

(A1) Under adverse selection, it is assumed that consumers are not able to observe the qualities of individual sellers, but they are able either to observe or to deduce expected quality E(q).

(A2) Under hidden irreversible action, it is assumed that consumers cannot observe the qualities of individual sellers *perfectly*, but that they can either observe or deduce average quality $\bar{q}$. Seller i's inverse demand will thus be a function of both average quality $\bar{q}$ and seller i's actual quality q_i.

(A3) Under moral hazard, it is assumed that consumers cannot observe quality at all before purchase. Consumers do, though, know what quality is 'offered' by each seller, and (if a purchase is made) they can observe the actual quality which was purchased after w periods.

(A4) Suppose that consumers cannot observe the qualities of individual sellers, but that they can either observe or deduce both E(q) *and* something about the distribution of actual qualities. In this case, when consumers are risk averse, they will be worse off (for given E(q)) when the quality distribution has a high variance.

Let us now look at these 4 asymmetric information effects in more detail.

(A1) Adverse selection

The idea of adverse selection was first formalized by Akerlof (1970). Here, an example is described very informally. Assume that:

— Seller i's quality q_i simply equals an innate quality θ_i. That is, sellers cannot do anything to change their quality levels;

— θ is distributed on the range $[\underline{\theta}, \hat{\theta}]$ with a density function $f(\theta)$;

— Consumers cannot see individual sellers' θs, but either (i) they can somehow observe $E(q)$ or (ii) they know $f(\theta)$ and $[\underline{\theta}, \hat{\theta}]$ and are sophisticated enough to derive, from this information, their expected quality $E(q)$ which is given by

$$E(q) = \int_{\underline{\theta}}^{\hat{\theta}} \theta f(\theta)\, d\theta$$

— The price that consumers are willing to pay $p(E(q))$ for each unit is increasing in expected quality $E(q)$.

— Sellers' opportunity costs of production $\mu(\theta)$ are increasing in θ;

— A seller of quality θ will sell if and only if

$$p(E(q)) \geq \mu(\theta)$$

— $p(\theta) > \mu(\theta)$, for all θ. That is, sellers of every quality θ will want to sell under full information.

The adverse selection problem runs as follows:

For any quality range $[\underline{\theta}, \hat{\theta}]$ where $\hat{\theta} > \underline{\theta}$, it is clear that $E(q) < \hat{\theta}$. Thus, if $\mu'(\theta)$ is positive, there is a possibility that the price consumers are willing to pay for $E(q)$ will not be enough to cover the opportunity costs, $\mu(\hat{\theta})$, faced by the maximal quality seller. That is, $p(E(q)) < \mu(\hat{\theta})$. In this case, the maximal quality seller will drop out of the market. The effect of this, though, will be that $E(q)$ (and so $p(E(q))$) will fall still further. More sellers may now leave the market, which will further reduce

expected quality and thus price, which may induce yet more sellers to leave the market, ..., and so on.

In some cases, this process will carry on until it is impossible for the market to function at all (market closure). In others, the process will stop at some point, but the average quality available at this point will be socially suboptimal.

Minimum quality levels can be a useful tool in such situations. Raising the lower limit $\underline{\theta}$ of the available quality interval brings E(q) closer to $\hat{\theta}$, which raises price. Some sellers who were previously unable to cover their production costs will now be able to do so, and will enter the market, so that the average quality available at equilibrium rises.

Moreover, since this process will generally raise not only consumer surplus but also total industry profits, it is possible that self-regulation (of minimum quality levels) will be welfare-superior to the *laissez faire*. This idea is examined by Leland (1979) and by Shaked and Sutton (1981). Their models are discussed in chapter 2.

Signalling can also be a useful tool in combatting adverse selection. Suppose that each seller's innate quality *can* be observed, but only by those consumers who try the seller's product. Since qualities cannot be changed, there is no asymmetric information problem between seller and consumer from this point on. Repeat purchases will then be made only if the initial purchase would have been made under full information (as long as zero switching costs are assumed).

The only problem that remains, therefore, is to get consumers to buy from a seller in the first place. If a high quality seller makes more from repeat purchases than a low quality seller does, then a high quality seller may be able to invest more in signalling (for example, via low introductory prices or wasteful expenditures (e.g. on advertising)) than a low quality seller can. In this case, we get a separating equilibrium, although pooling equilibria (where signalling is of no use) are also possible.

Signalling has received much attention (ref. Spence (1973), Nelson (1974) and, more recently, Milgrom and Roberts (1986)). There are three possible roles for self-regulation in a signalling framework: [4]

(a) If membership of an SRO is costly, then merely joining an SRO may signal high quality, even without the SRO itself doing anything. That is, joining an SRO can signal high quality simply by being a wasteful expenditure.

---

[4] Tirole (1989) provides a very simple model of signalling which can easily be manipulated to demonstrate these roles of self-regulation:

There are two sellers of qualities θ_1 and θ_2 ($\theta_2 < \theta_1$), and with costs μ_2 and μ_1 ($\mu_2 < \mu_1$). The full information price for θ_1 is p_1 and the full information price for θ_2 is zero. There are two periods and the price charged in the first period is p^*. Seller 1 thus receives

$$\pi_1 = (p^* - \mu_1) + \delta(p_1 - \mu_1)$$

where δ is the discount rate. π_1 must be non-negative at any equilibrium. For this to be a separating equilibrium, seller 2 must not be able to duplicate this initial price p^*. That is, it must be the case that $p^* - \mu_2 < 0$. Thus, a necessary condition for separating equilibrium is

$$(N^4.1) \qquad 0 \leq \pi_1 < \delta(p_1 - \mu_1) - (\mu_1 - \mu_2)$$

When $\delta(p_1 - \mu_1) < (\mu_1 - \mu_2)$, this cannot hold and there will only be pooling. But when $\delta(p_1 - \mu_1) > (\mu_1 - \mu_2)$, there exist p^* such that separating is possible.

Now, suppose that each seller must give an initial payment J to an SRO, which is only returnable if the seller is still in the industry at the end of the second period. In this case, seller 1's payoff is unchanged, but seller 2 will be unable to duplicate p^* if $p^* - \mu_2 - J < 0$. Thus, the necessary condition for separating equilibrium is now

$$(N^4.2) \qquad 0 \leq \pi_1 < \delta(p_1 - \mu_1) - (\mu_1 - \mu_2 - J)$$

This is a weaker condition than $(N^4.1)$, and so it is more likely to hold. Moreover, since J can be set such that $p^* = p_1$, an SRO can achieve separation without any loss in efficiency.

(b) An SRO might be able to achieve a separating equilibrium more efficiently (in terms of wasteful expenditures and low introductory prices) than under *laissez faire*;

(c) An SRO may be able to achieve a separating equilibrium where there is only pooling under *laissez faire*.

These possibilities are not considered in this thesis.

(A2) Hidden irreversible action

I want to distinguish between two very different aspects of (what is usually termed) moral hazard. Firstly, there is the problem of initial quality choice. This is what I want to call hidden irreversible action.

Let us consider an extreme case, in which sellers cannot be distinguished by consumers at all. For this case, assume that:

— Seller i's quality is equal to k_i. That is, quality is *ex ante* a matter of choice, but (once chosen) it is not easily reversible;

— Consumers cannot observe the quality levels of individual sellers, but they do somehow form an accurate expectation of quality E(k), which is equal to average industry quality $\bar{k}$;

— The price that consumers are willing to pay $p(\bar{k})$ is increasing in $\bar{k}$;

— Sellers' costs of quality $\mu(k)$ are convex and increasing in k.

— $p(\bar{k}) - \mu(k)$ is a convex single-peaked function, and is positive for at least some k.

The problem here is a free rider one. Under full information, each seller would receive $p(k_i)$, and would set quality such that $p_k = \mu'$. Under asymmetric information, all sellers must receive the same price $p(\bar{k})$, since

they are indistinguishable. Suppose that there are n sellers in the industry, each supplying one unit, so that the contribution of each seller's actual quality k_i to $p(\bar{k})$ will be $p_{\bar{k}}/n$. Taking the training choices of the other sellers as given, each seller will thus set $p_{\bar{k}}/n = \mu'$. This will result in average quality being suboptimal (and tending to zero as n gets large). Moreover, if equilibrium k is low enough, then $p(\bar{k}) - \mu(k)$ may be negative, in which case there is market closure.

As in free rider problems generally, the agents can all do better if they are able to collude when making their choice. The same is true in this case. If the sellers collude on quality, then they will set $p_{\bar{k}} = \mu'$ just as under full information.

Again, since this collusion process will generally raise both consumer surplus and total industry profits, it seems feasible that self-regulation (where this involves collusion on quality) will be welfare-superior to the *laissez faire*. This idea is examined in chapter 3 of this thesis.

(A3) Moral hazard

When each seller's quality is both unobservable *and* reversible, then not only is the initial quality choice problematic, as is described in (A2), but there may also be quality maintenance problems, because of sellers' incentives to suddenly reduce quality. Assume that:

— Seller i's quality is equal to a_i, where a_i is fully reversible;

— Consumers cannot observe any seller's actual quality at all before purchase, and because quality is perfectly reversible, past quality is no sure guide to future quality.

— Consumers *can* distinguish between sellers and, in particular, they know what quality is 'offered' by each seller.

— The price that consumers are willing to pay for quality a, p(a), is increasing in a;

— Sellers' costs of quality $\mu(a)$ are increasing in a.

— $p(a) - \mu(a)$ is a convex single-peaked function, and is positive for at least some a.

Under full information, sellers would set quality such that $p_{a^*} = \mu'$. Let this quality be a*. Under asymmetric information, because both prices and costs are increasing in quality, sellers have an incentive to offer high quality, but to actually provide (less costly) low quality. This is called 'cheating' on a quality offer.

Suppose that sellers offer a*, but actually supply zero quality. Rational consumers will realise that sellers have an incentive to act in this way, and so they will only pay a price p(0) in the first place. At equilibrium, then, sellers will provide zero quality and consumers will pay p(0) for it. If $p(0) \geq \mu(0)$ then this equilibrium will simply result in suboptimal quality, but if $p(0) < \mu(0)$ then there will be market closure.

Suppose now that, if a purchase is made, then the actual quality purchased *is* observed after w periods. In a repeated game, it may be possible to deter a seller from cheating by threatening some future punishment if actual quality is found to be less than 'offered' quality. Such incentive schemes are known as reputation mechanisms (ref. Klein and Leffler (1981), Shapiro (1983), Allen (1984)).

The maximum punishment that a single consumer can enforce is to refuse to buy from that seller again and also to warn as many people as possible not to buy from it. The threat of this punishment may be enough, when the incentive to cheat is not large. In many situations, however, stronger punishments may be required. Note, especially, that when the reputation mechanism is quite weak, it may work for one large seller but not for two or more smaller sellers. In this case, there may be a first-entrant

advantage, with the weakness of the reputation mechanism acting as a barrier to further entry (ref. Farrell (1986)).

This need for stronger punishments provides a rationale for regulation. A regulator may be able to reduce the incentive to cheat (e.g. by raising minimum possible quality levels, or by increasing the speed of detection of cheatings); or it may be able to impose additional penalties on cheating sellers (e.g. by improving information transmission between consumers or by enforcing further punishments). Moreover, since providing such incentives will generally raise both consumer surplus and total industry profits, self-regulation would seem to be a feasible form of regulation, and would certainly seem to be welfare-superior to the *laissez faire*.

These ideas are examined by Shapiro (1983) and Mayer and Neven (1991), whose models I discuss in chapter 2. Chapter 4 continues the discussion and modelling of these ideas.

(A4) Risk aversion

Risk aversion is not a separate problem from the three problems mentioned above. If any of these three problems exists in an industry, however, then risk aversion may exacerbate the situation, and give an extra rationale for self-regulation.

Consider an adverse selection set-up in which, at equilibrium, different sellers have different qualities, but consumers can only observe (or deduce) the distribution of these qualities, rather than how they are allocated between sellers. Suppose that consumers are risk averse. This means that the potential disutility from receiving low qualities outweighs the potential utility from getting high qualities, which lowers the price consumers are willing to pay for any E(q). Three implications of this are:

— The fall in p(E(q)) increases the likelihood that $p(E(q)) < \mu(\hat{\theta})$, and so it worsens the adverse selection problem;

— Minimum quality levels will again be a good method of dealing with the adverse selection problem, since these eliminate the possibility of consumers' receiving the highest disutility (lowest quality) products.

— Consumers will be willing to pay higher average prices when there are two distinguishable sets of sellers (e.g. two SROs) spread over quality intervals of (L', L'') and $(L'', \hat{\theta})$, [where $L' < L'' < \hat{\theta}$], than when there is just one set of sellers spread over the quality interval $(L', \hat{\theta})$. This is because the risk involved in buying from either of the smaller SROs is lower than is involved in buying from one big SRO. This is one argument for having multiple SROs (see chapter 5 of this thesis for further discussion of this topic).

Now consider a hidden irreversible action framework, in which the free rider problem has been solved by collusion on quality between sellers (self-regulation). If sellers face different costs of producing quality, it may be optimal (under risk neutrality) for them to choose different qualities at the collusive equilibrium, even though consumers only care about average quality. If, however, consumers are risk averse, any dispersion of available qualities will cause $p(\bar{k})$ to shift downwards. Thus, SRO members will have incentives to locate more closely together in quality space than they would under risk neutrality.

It is ambiguous whether the introduction of further different average quality SROs will be of help here. The benefit comes from the fact that sellers will be better able to relate the quality they provide to their costs of producing it. This, though, must be weighed against the fact that quality collusion is unlikely to occur as readily between SROs as within them, which means that the free rider problem is, to an extent, reintroduced.

Lastly, consider a situation with moral hazard. When consumers are risk averse, so that their potential utility losses from cheating are greater, the argument for regulation/self-regulation becomes even stronger.

I do not model risk aversion explicitly in any part of this thesis, but the remarks above suggest some of the effects that its introduction might have.

B. Between-seller externality effects

(B1) Positive spillovers

The distinction between asymmetric information effects and externality effects is not clear-cut. For example, the free rider problem deriving from hidden irreversible action (described in (A2)) is essentially an positive externality problem in which each seller, acting selfishly, doesn't take into account the effect of its own quality choice on the profits of other sellers. We can thus also label this effect (B1).

From an externality perspective, we can in fact describe this effect a little more generally than above. Assume that:

— Seller i's quality q_i is a choice variable;

— The price that consumers are willing to pay is increasing in both $\bar{q}$ and q_i. The rationale for this is that there may be a degree of asymmetric information between sellers and consumers, even when consumers *can* observe more about individual sellers than just $\bar{q}$;

— Sellers' costs of quality $\mu(q)$ are increasing in q.

In this set-up, an increase in q_i will increase $\bar{q}$, which will in turn increase the price received by other sellers. This is a positive between-seller externality which, if not internalized, will result in qualities being chosen too low for industry profit maximization.

As suggested above, an SRO can overcome this problem by allowing its members to collude on quality (i.e. by internalizing the externality). This may be beneficial if the qualities chosen under *laissez faire* are also too low for social welfare maximization.

(B2) Business-stealing effects

Suppose that the relative price received by each seller i depends on the level of q_i relative to the industry average $\bar{q}$, (i.e. $q_i - \bar{q}$).[5] In this case, quality choice involves a negative externality. Suppose there are n sellers, each supplying one unit. An increase in q_i will raise $\bar{q}$ (by 1/n times the change in q_i). This will raise seller i's relative quality (and so price), but will reduce the relative qualities of (and so the prices received by) all the other sellers in the industry. If sellers do not internalize this negative externality, chosen qualities will be too high for industry profit maximization.

Business-stealing effects are also important when considering the variety of qualities available, rather than their absolute level. In a Hotelling (1929) framework, as seller i locates closer and closer (in 1-dimensional product space) to another seller j, i will steal more and more of j's customers. That is, by moving towards j, i confers a negative externality on j. This effect means that sellers will tend to locate too close together (i.e. there will be too little product differentiation) for industry profit maximization. Note that Hotelling models also exhibit an *indirect* between-seller externality, which works via price (this is described in section D).

Again, an SRO can overcome either of these business-stealing problems by allowing its members to collude on quality (i.e. by internalizing the externality). This may be beneficial if the qualities chosen under *laissez faire* are also too high (or located too close together) for social welfare maximization. Comparisons between self-regulation and the *laissez faire*, under both (B1) and (B2), are examined in more detail in chapter 4 of this thesis.

---

[5] A formal justification for this is given in the appendix to chapter 3.

(B3) Excess entry effects

Suppose that the number of sellers in an industry is endogeneous. In this case, business-stealing effects become excess entry effects. Consider Hotelling's model. Since each potential entrant considers only its own costs of entry, and not the negative externality its entry will confer on existing sellers, we should expect to see too many sellers entering the industry (and thus too much variety) for industry profit maximization.

An essentially identical result can be seen in the monopolistic competition model studied by Dixit and Stiglitz (1977). By introducing a new product, a seller steals consumers from other sellers, which reduces the profits of these other sellers. If this effect cannot be internalized, sellers will tend to introduce too many products (i.e. there will be too much variety) for industry profit maximization.

The important issue for self-regulation, though, is not whether there will be more entry/variety than would maximize total industry profits, but whether there will be more entry/variety than would maximize total social welfare. When sellers can appropriate all consumer surplus, profits will correspond exactly to welfare, and so there *will* be too much entry/variety for welfare maximization. When consumer surplus is not fully appropriable, however, there is an seller-consumer externality ((C2) below) which may work in the opposite direction to (B3) and reverse the result. As a reference point, though, we do know that a simple homogeneous goods Cournot oligopoly will involve too much entry for welfare maximization.[6]

---

[6] This result is easily proved (see Mankiw and Whinston (1986)). Let X denote total industry output; p(X) denote price; μ denote unit costs; n denote the number of sellers; x denote individual outputs (= X/n); and ϕ denote per seller fixed costs. In this case welfare is given by

$$W = \int_0^X p(X')\,dX' - \mu X - n\phi$$

The derivative of welfare with respect to n is

Suppose that there *is* too much entry/variety for welfare maximization. A regulator which can impose extra entry/membership costs (on new sellers or new products) may be able to set these costs such that the equilibrium number of sellers/products is exactly right for welfare maximization. An SRO, though, which aims to maximize total industry profits will only set these entry/membership costs at an optimal level if sellers can appropriate all consumer surplus. Otherwise, the SRO may have incentives to restrict entry by more than is socially optimal. Moreover, since these extra costs may be justified by the SRO as required for checking that the new sellers/ products are of 'appropriate quality', it may be hard to gauge when SROs are behaving too restrictively.

The joint possibilities of excess entry under *laissez faire* and excessive entry restriction under self-regulation are considered by Leland (1979) (whose model is described in chapter 2 of this thesis), and also in the model in chapter 3 of this thesis (section 7).

These various between-seller externalities will also be important in determining the optimal number of SROs in an industry. As the number of SROs increases, the ability of SROs to restrict entry (internalize the (B3) externality) will be limited, but the degree to which externalities (B1) and (B2) can be internalized will also be reduced. The optimal number of SROs will thus tend to be increasing with the risk of excessive entry restriction, but decreasing with the overall magnitudes of (B1) and (B2).

---

[6] cont.

$$\frac{dW}{dn} = (p(X) - \mu)\frac{dX}{dn} - \phi$$

$$= (p(X) - \mu)\left[x + n\frac{dx}{dn}\right] - \phi$$

$$= \pi + n(p - \mu)\frac{dx}{dn}$$

where π are individual sellers' profits. Under free entry, $\pi = 0$. Thus, dW/dn has the same sign as dx/dn, which is negative. Thus there are too many sellers in the industry for welfare maximization. □

C. Seller-consumer externality effects

When demand is not completely elastic, and when perfect price-discrimination is not possible, an industry will never be able to appropriate all of the consumer surplus it provides. The most well-known effect of this is that, when sellers have market power, prices will tend to be too high (and quantities too low) for welfare maximization. Some less well-known effects are described below.

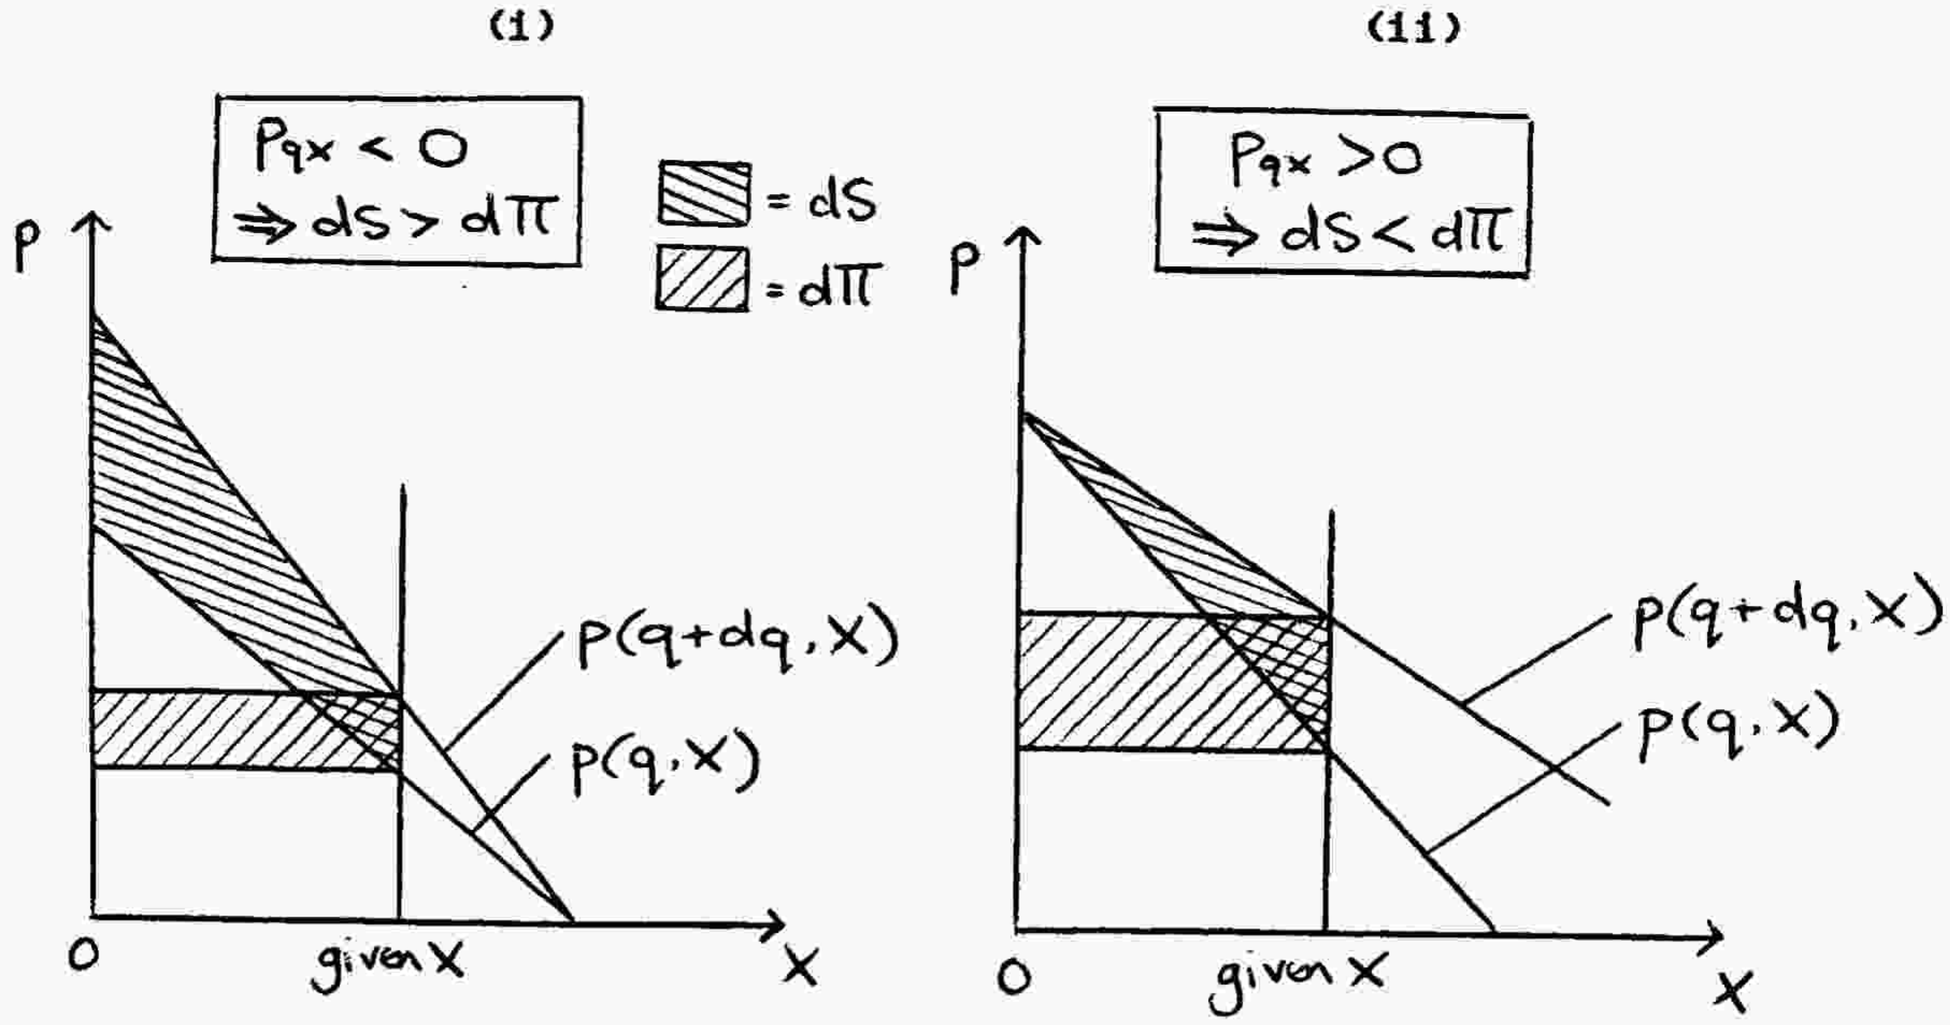
(C1) The Spence effect

Suppose that the seller(s) in an industry have a degree of market power. In this case, Spence (1975) shows that, for a given output level, each seller's profit-maximizing choice of quality will take into account only the quality valuation of the marginal consumer (i.e. the consumer who is just willing to pay the going price). Total social welfare, however, depends on the quality valuations of all the purchasing consumers.

For given output, then, a seller with market power will choose a socially optimal level of quality only if the marginal consumer is representative of all the consumers. Otherwise, each seller may choose too high or too low a quality level for social welfare maximization, depending on how the marginal consumer differs in his quality valuation from the infra-marginal consumers. These two possibilities are shown below, for the simple monopoly case.

Let $p(q, X)$ be the inverse demand function with respect to quality q and output X. In (i), where $p_{qX} < 0$, quality has a far smaller effect on the demand of marginal consumers than it does on infra-marginal consumers. Thus, for a given quantity, the monopolist will choose a socially suboptimal quality level. The converse is true in (ii), where $p_{qX} > 0$.

More formally, for given X, a seller with market power will undersupply quality relative to the welfare optimum when



$$(1.3.1) \quad \frac{1}{X} \int_0^X p_q(q, v)dv > p_q(q, X)$$

That is, for given X, the profit-maximizing quality will be welfare-suboptimal when the *average* marginal valuation of quality exceeds the *marginal* marginal valuation of quality. Assuming that the sign of p_{qx} stays the same for all q and X, this condition can be simplified to

$$(1.3.2) \quad p_{qx} < 0$$

The converse will be true for $p_{qx} > 0$.

To examine the effect of self-regulation on this situation, assume that an SRO will want to maximize total industry profits. As long as prices are below their monopoly levels, this will involve restricting industry supply, if this is at all possible. Consider the case where $p_{qx} < 0$, so that quality is welfare-suboptimal for given X. Any output restriction (lowering of X) will raise the return to quality p_q, which raises the profit-maximizing level of q (towards its welfare-optimal level). The converse is true when $p_{qx} > 0$. Note, though, that this welfare-improvement in terms of quality is at the expense of higher prices and lower output, so that the overall welfare effect is unclear.

The Spence effect plays an important role in Leland's model (see section 2.2) and in the model in chapter 3 of this thesis.

(C2) Variety effects of seller-consumer externalities

The monopolistic competition model analysed by Dixit and Stiglitz (1977) exhibits an effect which is essentially the same as the Spence effect. When the industry introduces a new product, the surpluses of both marginal and infra-marginal consumers will increase. If this is the only effect going on, and if the new product improves the surplus of the infra-marginal consumers more than the surplus of the marginal consumers, then there will be too few products available (too little variety) for welfare maximization. And conversely.

Note that self-regulation will only be useful in the situation where there are *too many* products. This is because self-regulation would work by internalizing the business-stealing externality (B2), which results in a reduction in the number of products available.

A similar seller-consumer externality exists in Hotelling models, due to the fact that sellers do not take account of consumers' transport costs when making their location decisions (except insofar as transport costs affect their demand curves). This will have two implications:

— When there is a given number of sellers in the industry, the failure to take account of transport costs will result in quality locations which are not sufficiently well dispersed. Instead, there will either be too much clumping or too much quality differentiation for welfare maximization.

— When the number of sellers is endogenous in the model, a potential entrant will only take its own entry costs into account when deciding whether or not to enter, and not the positive externality (in the form of reduced transport costs) which its entry would confer on consumers. Thus there is a tendency for too few sellers to enter the industry.

Self-regulation is unlikely to be able to overcome these latter problems. Their importance lies in the fact that they may act in the opposite direction to the between-seller externalities mentioned above. So they might eliminate the need for self-regulation in these cases.

D. Indirect effects of quality

Lastly, we should consider possible indirect effects of quality. Suppose that sellers' quality choices in a game also affect the equilibrium outcome of the game, in terms of other variables. In this case, quality choice may be strategic, and this role of quality is likely to have important welfare implications, both in terms of quality and in terms of the other variable.

(D1) Indirect effects of quality on output

Consider a simple adverse selection model in which sellers' innate qualities θ are distributed uniformly over an interval $[0, 1]$, but in which sellers can only enter the industry if they have $\theta \in [L, \hat{\theta}]$, where L is the minimum quality level, and $\hat{\theta}$ is the maximal quality available (which may be less than or equal to 1). This is essentially Leland's (1979) model.

In this situation, average quality $\bar{q} = (\hat{\theta} + L)/2$ and total supply $X = \hat{\theta} - L$. Thus, average quality and total supply are interlinked. An SRO facing a demand function with $p_x < 0$ will generally want to reduce X from its welfare optimal level. Because $\bar{q}$ and X are interlinked in this model, an SRO can only do this by raising average quality. Thus, there is a positive *output restriction effect* on $\bar{q}$ at the self-regulation equilibrium. The model in chapter 3 exhibits a similar output restriction effect.

(D2) Strategic effects of quality on competition

Quality may also have an effect on the degree of price competition in an industry. Consider a two-stage quality model (either a Hotelling type or a monopolistic competition type), in which sellers choose qualities at stage one and outputs at stage two. Assume that all sellers face an identical

cost of quality function which is increasing and convex. Because price/output competition intensifies as the disparity between sellers' qualities shrinks, sellers will try to differentiate themselves as much as they can, enabling high prices to be charged.

In a vertical (rather than horizontal) quality framework, we can consider the introduction of a minimum quality level. Ronnen (1991) shows that, because this effectively limits the range over which sellers can differentiate qualities (due to the convexity of the cost of quality function), price/output competition will intensify. Prices will thus fall and (generally) outputs will rise. This tends to improve consumer surplus, although the change in qualities on offer might act in the opposite direction (especially if some consumers prefer low quality goods).

The effect of the minimum quality level on sellers' profits is not straightforward. The higher-quality sellers will be worse off, since differentiating themselves from lower-quality sellers will now be more expensive and less effective. The profits of the lower-quality sellers, however, may well improve.

The reasoning behind this increase in profits is easily seen in a two seller example. It is easy to show that the lower-quality seller's revenue will be increasing in the higher-quality seller's quality. Thus, when the lower-quality seller has a first-mover advantage it will choose a higher quality level than it does at the Nash equilibrium, in order to force the higher-quality seller up to an even higher quality. The effect of a minimum quality level will be (at least at first) to push the Nash equilibrium towards this Stackelberg equilibrium. Thus, the lower-quality seller's profits will improve.

For his model, Ronnen shows that the introduction of a minimum quality level will (at least at first) increase total welfare. This only provides an argument for self-regulation, though, if it is possible for one part of the industry (the lower-quality part) to impose minimum quality standards without requiring the agreement of all sellers in the industry.

The ability of SROs to use quality choices to manipulate outputs is also an important factor in determining the optimal number of SROs. As the number of SROs increases, the ability of SROs to use qualities strategically in this way is reduced, so qualities move back towards their non-strategic levels. If there are no other effects going on, this will tend to be good for social welfare. This intuition is examined further in chapter 5.

1.3.2 Players and their maximands

In the above discussion, and in most of this thesis, it is assumed that SROs aim to maximize the total profits of their membership. This is a simple assumption to make, but it is also a very strong one. In this section, then, we consider some alternative assumptions.

The question of what exactly an SRO maximizes is very similar to the much discussed problem of defining an objective function for trade unions (ref. Oswald (1985)). The most general expression for an SRO's maximand is $\sigma(n, \boldsymbol{\pi})$, where n is its number of members, and $\boldsymbol{\pi}$ is the vector of members' profits. Various special cases can be defined:

(a) The most common assumption is that an SRO is simply a set of sellers colluding to maximize their joint profits. In this case, then

$$\sigma(n, \boldsymbol{\pi}) = \sum_{i=1}^{n} \pi_i$$

(which equals $n\pi$ at a symmetric equilibrium).

(b) Another possibility, considered by Shaked and Sutton (1981) (whose model is discussed in chapter 2), is that an SRO aims to maximize *per seller* profits, so that $\sigma(n, \boldsymbol{\pi}) = \pi$. Note that this is the same as (a) when entry is not an issue (i.e. n is fixed).

(c) If an SRO is separate from (although organised by) the industry, then it may prefer to maximize its own total revenue. If membership

charges are set externally, then such an SRO will simply have $\sigma(n, \pi) = n$. This possibility is considered in chapter 3. If charges are endogenous, however, they will generally depend on sellers' profits, so that $\sigma(n, \pi)$ cannot be simplified.

(d) Control of power in an SRO may not be spread evenly between member sellers. Asymmetric information problems and seller-seller externalities are likely to be less important for larger, more profitable sellers than for smaller, less profitable sellers. If this means that the former have less to gain from self-regulation from the latter, then the former may either have a greater degree of power in the SRO (because they have greater bargaining strength) or less power (because they expend fewer resources on influencing self-regulation decisions). In either case, $\sigma(n, \pi)$ may be a very complicated function, and it will certainly be no simpler than

$$\sigma(n, \pi) = \sum_{i=1}^{n} \lambda_i \pi_i \quad \text{(where } \lambda_i \text{ are weights).}$$

(e) Lastly, it has frequently been argued that managerial motives in capitalist firms result in a desire for large size or revenue (ref. Marris (1964)). In so far as this desire can be separated from a desire for market share, it seems reasonable that an SRO (which is made up of such firms) might include size and/or revenue as arguments of its maximand.

Turning now to consumers, each consumer is assumed to maximize his own expected utility (subject to various constraints), and to proceed with his optimal purchase as long as this gives him non-negative expected utility. For simplicity, the following strong assumption is sometimes made: When quality is an imperfectly observable random variable $\mathbf{q}$, each consumer's expected utility from quality $EU(\mathbf{q})$ can be rewritten as the utility he gets from expected quality $U(E(\mathbf{q}))$.

Given this, there are (at least) three quality-related arguments which may occur in a consumer's utility function:

(a) Expected quality E(q): This may equal a seller's actual quality, average industry quality, a convex combination of the two (as described above), or the minimum possible quality. Utility may be increasing in E(q) (maybe at different rates for different consumers), or it may be increasing only until the consumer reaches his 'most preferred' quality level (where 'most preferred' quality levels may also differ between consumers). In much of this thesis, the former assumption is made, but in chapter 5 (which looks at the optimal number of SROs), I allow 'most preferred' quality levels.

(b) The difference between expected quality and average industry quality, $E(q_i) - \bar{q}$. A utility function will be increasing in this difference when consumers care about relative (as well as absolute) quality. This is assumption provides one (but not the only) justification of the price function used in chapter 3, which includes $(q_i - \bar{q})$ as an argument.

(c) The variety of qualities available, denoted Var(**q**). Utility will be increasing in the variety of qualities available, if consumers value choice and if sellers' qualities are either (i) at least partially observable or (ii) credible/deducible. The assumption that consumers value choice is usual in monopolistic competition models (e.g. Dixit and Stiglitz (1977)). Utility will be decreasing in the variety of qualities available when consumers are risk averse and cannot distinguish between the qualities of different sellers.

Lastly, each seller is assumed to want to maximize her own profits, given the actions of all the other players, where her profits are simply equal to her ((price × output) - costs).

1.3.3 Move order

A self-regulation game may begin from one of two initial states:

(a) Pre-industry-formation, where sellers' decisions (on entry, quantity, quality, etc) all have to be made from scratch;

(b) Post-industry-formation, where sellers are initially at the *laissez faire* equilibrium (N.B. a variation on this initial state involves sellers' having expectations of future self-regulation).

The difference between these two initial states is particularly important in Shapiro (1983), which is discussed in chapter 2 of this thesis.

At stage 0 of the game, an external body decides on the institutional structure of the industry with, say, the aim of maximizing social welfare (although, in the real world, political objectives may also be important). Regulatory choices at this stage will be structural (conduct regulation comes at a later stage). Possibilities might be:

— Promoting competition within the industry, by breaking large firms up into smaller ones;

— Imposing external (conduct) regulation for certain decision variables;

— Promoting/imposing self-regulation for certain decision variables;

— Allowing self-regulation for certain decision variables;

— Banning self-regulation for certain decision variables;

— Determining the number of SROs that should exist, and how they should differ (if at all);

— Determining whether membership of an SRO should be compulsory for practising in the industry;

— Determining conditions under which an SRO should be allowed to refuse membership to an applicant;

— Determining how the work of the SRO should be funded, and how its decisions should be made.

The models in this thesis generally miss out stage 0. Instead, assumptions are made about what decisions have been made at stage 0, and the implications are examined. However, these implications should (and do) provide us with some arguments for what should happen at stage 0.

The final stage in every model involves consumers deciding whether to buy, and how much to buy of, the good, given p, $E(q)$, $E(q) - \bar{q}$ and $V(q)$. An inverse demand function can be derived from this stage, which is used for calculating outcomes in earlier stages.

The stages in between stage 0 and the final stage differ between models of self-regulation. Generally, though, entry by SROs and memberhip decisions occur first, quality and minimum quality level choices occur next, and output/price decisions follow.

1.4 Thesis layout

Chapter 2 is a critical survey of some of the literature which is relevant to this thesis. I focus on four papers: Leland (1979); Shaked and Sutton (1981); Shapiro (1983); and Mayer and Neven (1990). Each is discussed in some detail, and I also explain how they relate to the rest of the thesis. Two other papers are examined in less detail.

The model in chapter 3 looks at quality choice when there is a hidden irreversible action problem. In some ways, it can be seen as an extension of Leland (1979), in that it captures all of the important charateristics

of Leland's model (plus some more). The modelling approach, however, is very different:

Quality is assumed to be a (relatively) irreversible choice variable (denoted k), and self-regulation is modelled as explicit quality collusion between sellers. Quality collusion allows sellers to overcome any between-seller externalities, and it also affects equilibrium output. Whether self-regulation is likely to be welfare-superior to *laissez faire* is shown to depend on the relative signs and values of

— the *net* between-seller externality (the difference between the positive (B1) and the negative (B2) between-seller externalities);

— $p_{\bar{k}x}$: this captures the Spence effect, as described above (C1);

— $\partial(p - \mu_x)/\partial\bar{k}$: this determines how total industry output changes with average industry quality, and thus captures the output restriction effect (D1).

It also depends on whether quality is chosen before, or at the same time as, outputs.

The intuition for the importance of the first two of these effects is straightforward. When there is a net positive between-seller externality, self-regulation will cause quality to rise. If marginal consumers value quality less than infra-marginal consumers ($p_{\bar{k}x} < 0$), quality will be socially sub-optimal (for given X) under *laissez faire*. Thus, an increase in quality will raise both industry profits and consumer surplus. The converse holds when there is a net negative between-seller externality and $p_{\bar{k}x} > 0$. For further and more detailed results and intuitions, see chapter 3.

I also consider entry in this model, and examine whether welfare will ever be higher when SROs are allowed to choose the number of sellers in the industry. It is shown that self-regulation under free entry results in too

much entry for welfare-maximization (an excess entry effect (B3)). Self-regulation when n is a choice variable, however, results in too little entry for welfare-maximization. An alternative policy is allow the SRO both to restrict its membership and to threaten to abolish itself in the face of any entry into the industry by any non-members. I call this threat-constrained entry, and show that it will result in too much entry for welfare-maximization, but less entry than under free entry.

Chapter 4 looks at the moral hazard problem (A3), and how regulatory contracts can be designed to solve it. The model in this chapter can be seen partly as an extension of Shapiro (1983), and partly as a response to Mayer and Neven (1991), whose model is unattractive in several ways. The model also applies some of the ideas in Milgrom, North and Weingast (1990) to the issue of regulation.

Sellers are assumed to have different innate quality levels, but they may cheat on them, and provide zero quality. They may also commit direct fraud. Qualities are assumed to be observable (w periods after purchase), but not verifiable. The role of a regulator is two-fold. A regulator acts as a collator of information about cheatings, which enables the information transmission needed for the reputation mechanism to function. Plus, a regulator adds to the punishment inflicted upon cheating sellers with a fine J, which sellers may or may not have to pay in advance of the actual cheating (depending on how easily they can be forced to pay it *ex post*).

Because qualities are never verifiable, self-regulation may have an advantage over external regulation, in that an SRO can condition its regulation on quality. On the other hand, an SRO will have to be given incentives not to cheat on its regulation promises (although this will also be true of a non-benevolent external regulator). Moreover, the regulation which an SRO *does* provide may not achieve welfare maximization, if this conflicts with the maximization of industry profits.

There is thus a trade-off between being able to condition regulation on quality and having (a restricted form of) regulation set in a welfare

maximizing way. This model examines conditions under which the former effect is likely to outweigh the latter (so that self-regulation is welfare superior to external regulation). It is shown that self-regulation becomes more likely to be welfare-superior to external regulation as fraud becomes more important, as sellers become more homogeneous in terms of capital, and as the observation lag becomes greater.

Chapter 5 considers the factors that might influence the optimal number of SROs in an industry. I first examine a two-SRO extension by Shaked and Sutton (1981) of their one-SRO model. In their model, membership of an SRO is a necessary prerequisite of functioning in the industry. They show that entry (or even the *threat* of entry) by a small second SRO will generally be welfare improving, because it will restrict the incumbent SRO's ability to control total industry output. This idea is intuitively reasonable, but the model is unrealistic in some important ways.

This leads me to develop a new model, which is a variant of the Hotelling (1929) model. As in Shaked and Sutton, membership of an SRO is a necessary prerequisite of functioning in the industry, and the entry of a second SRO into a one-SRO industry is considered.

When entry is purely member-stealing (that is, when the second SRO steals all of its members from the incumbent), welfare will definitely be higher after entry, as will consumer surplus. When there is output competition between members of an SRO, the incumbent's profits will also be higher under entry, so the incumbent will actually welcome entry. When there is output collusion between members of an SRO, however, the incumbent's profits will be lower under entry. It may thus try to deter entry by lowering its quality level (with ambiguous welfare consequences).

When entry is purely member-augmenting (that is, the second SRO recruits completely new sellers into the industry), the results again depend on whether or not there is output competition between members of an SRO. If there is, then welfare, consumer surplus and profits will actually be lower under entry. When there is output collusion between members of an SRO,

however, welfare and consumer surplus will definitely be higher under entry and profits (even total industry profits) will definitely be lower. Again, then, there will be a risk of entry-deterring behaviour.

This model can be extended to look at entry into a one-SRO industry by an unregulated sector. In this case, welfare and consumer surplus will be higher under entry, but the SRO's profits will be lower. This result seems intuitive, but a little too simplistic to be really satisfactory.

1.5 Concluding remark

The purpose of this chapter has been to provide an introduction to the issues involved in self-regulation; to discuss some ways in which self-regulation might be modelled; and to provide a taxonomy for the various quality effects, many of which reappear later in this thesis.

NOTATION - chapter 2

a	effort/quality
α	probability of hearing outcome of a unit
AS	aggregate supply
AD	aggregate demand
C	capital ϵ $[0, \bar{C}(\theta)]$
CR	capital requirement
δ	discrete time discount rate
$E(\cdot)$	expectations operator
ϕ	additional costs per seller
$g(\sigma, \tau)$	probability distribution over expectation parameters σ and τ
$H(\phi)$	distribution function for ϕ
h	labourer's services consumed per buyer
J	punishment for cheating
K	minimum training level
k	training
L	minimum quality level
L^r	L which maximizes Π
L*	L which maximizes W
$M(h, \bar{q})$	quantity of labourers' services a consumer is willing to forego for one unit of quality $\bar{q}$ lawyers' services
m	$(\delta^{-\omega} - 1)$
$\mu(\cdot)$	opportunity cost of quality function
μ_L	opportunity cost of L
N	maximum number of sellers
n	number of sellers regulated
ν	taste for quality
ω	willingness to pay for product
$p(\cdot)$	price function
p^e	entrance price
Π	total industry profits
$\pi(.)$	profits per seller (for honest sellers)

2. LITERATURE SURVEY

2.1 Introduction

There is relatively little existing literature on theories of self-regulation, and that literature which does exist is not very widely known. So, instead of giving brief résumés of a number of papers, this chapter examines four relevant papers in some detail: discussing not only the results, but also how they are derived, how realistic the assumptions are, and how the results might change under different (perhaps more realistic) assumptions. These four models are chosen partly because of their importance and partly because they relate in some way to the models set out in the rest of this thesis.

Briefly, these are the four papers.

— Leland (1979) examines how minimum quality levels may be used to improve welfare in an industry facing adverse selection, and derives conditions under which an SRO will choose too high/too low a minimum quality level for welfare-maximization. The two main effects driving this comparison reappear in the irreversible quality choice model in chapter 3

— Shaked and Sutton (1981) look at a similar situation, but in a different framework. Their model has firmer microeconomic foundations than Leland's, but their results turn out to be less general. Their model has the benefit that it can be easily extended to a two-SRO regime, as is shown in chapter 5.

— Shapiro (1983) provides a model of reputation formation in a competitive industry with moral hazard, and examines the welfare effects of varying the minimum quality level and/or the detection lag (the time taken for consumers to detect cheatings). Shapiro himself draws no conclusions about self-regulation from his analysis, but such conclusions are easily made. I also manipulate the model, to see how

these results are affected by a change in the model's assumptions. Reputation formation is revisited in chapter 4.

— Mayer and Neven (1990) also look at moral hazard. Their emphasis, though, is on how moral hazard can be eased through the use of capital requirements and penalties for cheating. I argue that the assumptions made in this model are problematic, but some of Mayer and Neven's intuitions reappear in the main model in chapter 4.

Section 2.6 examines, in less detail, two other relevant papers:[1]

— Shaked and Sutton (1982), which provides a simple communications argument for why bigger SROs may be more satisfactory than small ones;

— Shapiro (1986), which examines the use of minimum qualification levels (input regulation), rather than minimum quality levels (output regulation).

Before proceeding, I should make two terminological points. First, the papers by Leland and by Shaked and Sutton actually discuss 'professions' rather than 'SROs', but I shall frame them here in terms of SROs. Second, I shall, throughout this thesis, speak of consumers as male, sellers as female, and any regulator as neutral ('it').

2.2 Leland (1979)

Leland analyses a model in which the sellers and potential sellers of a good or service are of differing intrinsic qualities, which they cannot improve upon through training or effort and which consumers cannot observe in advance of making a purchase. All that consumers are assumed to be able to observe or infer is the average quality of those sellers participating in the industry. This means that there is a classic adverse selection problem as described in chapter 1.

As usual, the main result of this adverse selection problem is that average quality $\bar{q}$ in an unregulated market will be welfare suboptimal. This symptom may be eased by the introduction of a regulator which can set a minimum quality level L. By keeping some low quality sellers out of the industry completely (if their innate qualities are less than L), the introduction of minimum quality levels will cause average quality to rise, which will improve welfare. However, the introduction of minimum quality levels will also raise total industry costs, and may cause total output X to fall, both of which will be bad for welfare.

Leland examines this trade-off between improved average quality, increased total costs and reduced total output. In particular, he considers the welfare optimal trade-off, and how this will compare with the trade-off chosen by an SRO, which aims to maximize total industry profits. Two separate effects are important in this welfare-comparison:

— If the quality valuations of marginal consumers are different from those of infra-marginal sellers, then the Spence effect will be important (this is effect (C1) in chapter 1). An SRO will only care about the quality valuations of marginal consumers, whereas social welfare maximization must take all consumers' quality valuations into account. Only if the quality valuations of all consumers are identical, should we expect the Spence effect to be zero.

— An increase in minimum quality levels will generally have an effect on output. Any reduction in output will be bad for social welfare, but may be good for total industry profits. Unless $dX/dL = 0$, therefore, an SRO will tend to manipulate L so as to achieve lower X, whereas welfare-maximization requires the reverse.

Leland's conclusion is that, under certain realistic assumptions, while some regulation is generally better than none, an SRO will tend to set too high a minimum quality level for social welfare maximization. As is shown below, the realism of these assumptions can be questioned. Changing the

assumptions does not affect the substance of Leland's model, though. It just changes the implications drawn.

Leland's model does not allow for quality to be a choice variable for sellers. He simply looks at the straight adverse selection case discussed in chapter 1 (denoted (A1)). However, the two effects he focusses upon may also be important in situations with hidden irreversible action (denoted (A2) above). In chapter 3, I provide a model of this situation, which gives a role to both the Spence effect and the output restriction effect.

2.2.1 Leland's framework

The quality level q_i of each seller i's product is equal to that seller's innate quality level θ_i. That is, sellers are unable to change their quality level. For simplicity, and without any loss of generality, let the range of θ which are potentially available be $[0, 1]$, let supply be uniformly distributed on this interval, and let the total supply over the whole interval be normalized to one.

Consumers can observe the average quality $\bar{q}$ available in the industry, but they cannot observe each seller i's quality θ_i. Consumers have a marginal willingness to pay (inverse demand curve) which depends on industry output X and on expected quality E(q). Expected quality is equal to the average quality available, $\bar{q}$. Thus, the inverse demand curve is

$$p(\bar{q}, X) \qquad [p_{\bar{q}} > 0, \; p_X < 0]$$

Sellers face an opportunity cost $\mu(\theta)$ of supplying a unit of quality θ. $\mu(\theta)$ is initially assumed to be increasing in θ, although this assumption is relaxed in section 2.2.6.

A regulator can introduce a minimum quality level L into the industry. An SRO will choose that L which maximizes total industry profits. In what follows, this profit-maximizing L is compared with the welfare optimal choice of L.

2.2.2 The unregulated equilibrium

A seller of quality θ will sell if and only if the price she receives is greater than or equal to her opportunity cost. Let $\hat{\theta}$ be the maximal quality service or good which is being supplied in the market. If there is no minimum quality level in the industry, it is clear that market supply X will be equal to $\hat{\theta}$, and the average quality $\bar{q}$ available will be equal to $\hat{\theta}/2$. At equilibrium, price will just cover the opportunity costs of the maximal quality seller. Thus, the equilibrium maximal quality level $\hat{\theta}_e$ is given by

$$\mu(\hat{\theta}_e) = p(\bar{q}_e, X_e) = p(\hat{\theta}_e/2, \hat{\theta}_e)$$

To compare this equilibrium $\hat{\theta}_e$ to the quality level which is 'socially optimal', let net benefits be denoted by

$$W = \int_0^X p(\bar{q}, X')dX' - \int_0^{\hat{\theta}} \mu(\theta)d\theta \qquad (2.2.1)$$

Substituting $\hat{\theta}/2$ and $\hat{\theta}$ for $\bar{q}$ and X and differentiating with respect to $\hat{\theta}$ gives:

$$\frac{dW}{d\hat{\theta}} = \int_0^{\hat{\theta}} \tfrac{1}{2}p_{\bar{q}}(\hat{\theta}/2, X')dX' + p(\hat{\theta}/2, \hat{\theta}) - \mu(\hat{\theta})$$

At $\hat{\theta} = \hat{\theta}_e$, we know that $p(\hat{\theta}/2, \hat{\theta}) = \mu(\hat{\theta})$, so

$$\left.\frac{dW}{d\hat{\theta}}\right|_{\hat{\theta}_e} = \int_0^{\hat{\theta}_e} \tfrac{1}{2}p_{\bar{q}}(\hat{\theta}_e/2, X')dX'$$

This term must be positive, since $p_{\bar{q}}$ is positive. Thus an unregulated open market (that is, a market where L = 0 and $\hat{\theta}_e > 0$) will underprovide quality relative to that which is socially optimal. The intuition here is simple. At equilibrium, the marginal seller has an opportunity cost equal to the

social value of an extra unit of *average* quality, but she provides a unit of *above-average* quality. The social value of her unit thus exceeds its opportunity cost, and the margin should be pushed to higher quality levels.

2.2.3 Introducing minimum quality levels

Consider the imposition of a minimum quality level L. Market supply X will now be equal to $\hat{\theta} - L$, and average quality $\bar{q}$ will be $(\hat{\theta} + L)/2$. Thus, at equilibrium

$$p[(\hat{\theta}_{=} + L)/2, (\hat{\theta}_{=} - L)] = \mu(\hat{\theta}_{=}) \qquad (2.2.2)$$

Differentiating W, as given by equation (2.2.1), at $\hat{\theta}_{=}$, with respect to L gives

$$\frac{dW}{dL} = \left[\frac{d\hat{\theta}_{=}}{dL} + 1\right] \cdot \int_{0}^{\hat{\theta}_{=}-L} \tfrac{1}{2} p_{\bar{q}}(\bar{q}_{=}, X')dX' - (\mu(\hat{\theta}_{=}) - \mu(L)) \qquad (2.2.3)$$

Clearly, if this expression is positive at L = 0, then some minimum quality level L > 0 is desirable.

In fact, the sign of this expression is ambiguous. However, some general tendencies can be revealed through the use of numerical examples. These are summarised in the following informally stated result:

Result 2.2.1: Minimum quality standards will tend to be more advantageous in markets with:

(a) greater sensitivity to quality variations [$p_{\bar{q}} >> 0$];

(b) low elasticity of demand [$p_x << 0$];

(c) low marginal cost of providing quality [μ' small]; and

(d) low value placed on low-quality service [$p(0,0)$ small].

To see the intuition behind these conditions, remember that raising L increases welfare via its quality effects, and decreases welfare via its cost and output effects. Clearly, the more that L affects quality and the less that it affects costs and supply, the more likely it is that an increase in L will benefit welfare.

In many situations, then, minimum quality standards will be socially beneficial.

2.2.4 Self-regulation

As well as improving welfare, the introduction of minimum quality standards will tend to increase industry profits. Thus, an SRO will raise L upwards from zero, which will also (at least at first) be beneficial to welfare. Is it possible, then, that self-regulation will be socially optimal?

Unfortunately, the answer to this question will generally be no. Rather, an SRO will tend to choose too high or too low a level of L for welfare maximization. This is explained below.

The SRO will seek to maximize total industry profits Π, which are given by

$$(2.2.4) \qquad \Pi = p[\bar{q}, X].X - \int_{L}^{\hat{\theta}} \mu(\theta)\, d\theta$$

$$= p[(\hat{\theta} + L)/2, (\hat{\theta} - L)].(\hat{\theta} - L) - \int_{L}^{\hat{\theta}} \mu(\theta)\, d\theta$$

The first derivative of this with respect to L is

$$\frac{d\Pi}{dL} = [p_x(\bar{q}, X).X + p(\bar{q}, X)].\left[\frac{d\hat{\theta}}{dL} - 1\right]$$

$$+ \frac{1}{2}.p_{\bar{q}}(\bar{q}, X).X.\left[\frac{d\hat{\theta}}{dL} + 1\right] - \mu(\hat{\theta}).\frac{d\hat{\theta}}{dL} + \mu(L)$$

Denote the profit-maximizing minimum quality level L^r. At L^r, $d\Pi/dL = 0$. Thus, at the profit-maximizing equilibrium (where $p(\bar{q}_m, X_m) = \mu(\hat{\theta}_m)$), we have

$$(2.2.5) \quad \frac{d\Pi}{dL} = \left[\frac{1}{2}. p_q(\bar{q}_m, X_m). X_m + p_x(\bar{q}_m, X_m). X_m\right]. \left[\frac{d\hat{\theta}_m}{dL} - 1\right] + p_q(\bar{q}_m, X_m). X_m - [\mu(\hat{\theta}_m) - \mu(L^r)] = 0$$

2.2.5 The welfare-implications of self-regulation

Solving equation (2.2.5) for $[\mu(\hat{\theta}_m) - \mu(L^r)]$ (remembering that $dX_m/dL = d\hat{\theta}_m/dL - 1$), and substituting this into expression (2.2.3), gives

$$(2.2.6) \quad \left.\frac{dW}{dL}\right|_{L^r} = \frac{1}{2}. \left[\int_0^{\hat{\theta}_m - L^r} p_q(\bar{q}_m, X')dX' - p_q(\bar{q}_m, X_m). X_m\right]. \left[\frac{d\hat{\theta}_m}{dL} + 1\right] - p_x(\bar{q}_m, X_m). X_m. \frac{dX_m}{dL}$$

This equation is made up of two terms (one on each line). The first term results from the positive effect of L on $\bar{q}$. Indeed, if it were somehow the case that X was independent of L, then this first term would be the only term in equation (2.2.6).

As discussed above, the positive effect of L on $\bar{q}$ causes both W and Π to increase as L rises from zero. However, it need not be the case that the L which maximizes this positive effect on W should also be that L which maximizes the effect on Π. Indeed, this will only hold true if the first term in equation (2.2.6) is zero. Otherwise, since $d\bar{q}_m/dL$ is positive, this term will have the same sign as

$$(2.2.7) \quad \frac{1}{X_m}. \int_0^{X_m} p_q(\bar{q}_m, X')\, dX' - p_q(\bar{q}_m, X_m)$$

Thus, the first term in (2.2.6) will be non-positive only if the average marginal willingness to pay for quality, over levels of consumption from 0 to X_m (the first term in (2.2.7)) is less than the marginal willingness to pay for quality at the current level of consumption (the second term in (2.2.7)).

This condition is the Spence effect, described above. Welfare depends on consumers' *average* marginal willingness to pay for quality. An SRO, however, only cares about the *marginal* consumer's marginal willingness to pay for quality. If the latter exceeds the former, then an SRO will have an incentive to set L higher than is socially optimal. Note that, as long as p_{qx} always has the same sign, condition (2.2.7) will have the opposite sign to p_{qx}, and will be satisfied with equality if and only if $p_{qx} = 0$.

The second term in equation (2.2.6) results from the effect of L on X. If raising L reduces supply, then an SRO will want to raise L beyond its socially optimal level in order to increase Π. Conversely, if raising L increases supply, then an SRO will want to reduce L to below its socially optimal level in order to increase Π. By total differentiation of the equilibrium condition (2.2.2), we get

$$(2.2.8) \qquad \frac{dX_m}{dL} = \frac{p_q - \mu'}{\mu' - p_x - \tfrac{1}{2}p_q}$$

whose denominator must be positive for equilbrium to be stable. This means that dX_m/dL has the same sign as $(p_q - \mu')$, which can be shown to be negative (positive), as L rises from L^r, if μ is a strictly convex (concave) function. Thus the second term in equation (2.2.6) will be negative (positive) if μ is a strictly convex (concave) function.

From this discussion, we can derive the following. Denote the socially optimal minimum quality level as L*.

Result 2.2.2: If $p_{max} \geq 0$ and $\mu'' \geq 0$, then $L^r \geq L^*$;

If $p_{max} \leq 0$ and $\mu'' \leq 0$, then $L^r \leq L^*$; and if

If $p_{max} = 0$ and $\mu'' = 0$, then $L^r = L^*$.

Otherwise, the two effects in condition (2.2.6) will act against each other, and it will be ambiguous whether or not $L^r > L^*$.

Leland suggests that, in a typical industry, μ will be a convex function and p_{max} will be approximately zero. This means that, although welfare can be improved by the introduction of a minimum quality levels in such an industry, an SRO will set L at a higher level than is welfare-optimal.

2.2.6 Decreasing opportunity costs

This result extends to cases in which opportunity costs decrease with quality. A good example of this is the market for blood. At zero price, supply comes from voluntary donors only, who typically have disease-free blood. As the price paid increases, supply increases, but the average quality tends to drop since paid donors tend to have a higher rate of serum hepatitis.

In fact, for markets where $\mu' < 0$, it will *always* be the case that minimum quality standards will be desirable. This is because there is a different trade-off between quality, costs and supply. At the unregulated equilibrium with $\mu' > 0$, the available quality range will be $[\underline{\theta}_m, 1]$. A slight rise in L from $\underline{\theta}_m$ will still increase average quality, but it will also reduce total costs and definitely decrease supply (since $X_m = 1 - L$).

The decreased supply remains deleterious to welfare. However, the increase in $\bar{q}$ and reduction in total costs are both socially beneficial, and it can be shown that the change in sign of the cost effect is enough to make the overall effect of minimum quality standards on welfare unambiguously positive.

Turning to self-regulation, the fact that dX_-/dL is negative means that the output restriction effect will involve an increase in L by SROs. Thus, unless the Spence effect is strongly in the opposite direction ($p_{qx} << 0$), self-regulation will result in minimum quality levels being set too high for welfare-maximization.

2.2.7 Discussion of Leland's model

The above conclusions seem intuitively reasonable, but there are caveats:

First, we might want to question Leland's belief that a typical industry (which fits the above model in most ways) will (a) have opportunity costs which are increasing and convex, and will (b) face a demand function whose p_{qx} is roughly zero.

(a) On the cost side, it seems reasonable to expect a seller which exhibits high quality in one industry to exhibit similarly high quality in another. Since the opportunity cost of selling in the first industry is essentially just the price the seller would receive in the second industry, it seems reasonable that opportunity costs should be increasing in quality. There seems little reason to expect costs to be convex, however. Indeed, if prices in the second industry happen to be concave in quality, as is common, then we should expect opportunity costs in the first industry also to be concave.

(b) On the demand side, we know that marginal consumers are willing to pay less for the good than are inframarginal consumers, by definition. One common underlying explanation for this demand pattern relies on consumers' income. Suppose that the good provides all consumers with the same utility. A rich consumer will have a lower marginal utility of money than will a poor consumer, and so the former will be willing to pay more than the latter for a good of given utility. On the same grounds, even if both kinds of consumer have the same utility valuation of quality, the richer consumer will have the higher valuation in price terms. Thus, we might expect inframarginal

consumers to value quality more highly (in price terms) than marginal consumers, which means that $p_{max} < 0$.

If an industry exhibits concavity of opportunity costs and negative p_{max}, however, then from result 2.2.2 we know that an SRO will tend to set L too *low* for welfare-maximization, not too high. That is, self-regulation will certainly be a welfare-improvement on the *laissez faire*, but it will not raise L by as much as is needed for welfare-maximization.

A second caveat is that, even where an SRO does have incentives to set L at too high a level for social welfare maximization, it may be illegal for it to do so. In many countries there exists some degree of legislation on restrictive trade practices, according to which, if an SRO were to try to set L too restrictively, then it would tend to be severely disciplined, or even disbanded.

It is fair to say that Leland's and Shaked and Sutton's models were in fact motivated by the situation in the professions, which have traditionally been exempt from much of this legislation. However, as mentioned earlier, if and when the 1989 white paper on restrictive trade practices becomes law, the professions will be subject to the same rules as all other SROs and trade associations. Under these conditions, Leland's result may become much less important as an argument against self-regulation

Finally, even when opportunity costs are convex and p_{max} is zero, it can be shown that allowing a small amount of quality choice by sellers may ameliorate or even reverse Leland's result (see Fletcher (1991), chapter II, section 4). The idea behind this is that sellers which find themselves beneath the minimum quality level L can do some extra training to push themselves above it again. This reduces the trade-off between higher quality and lower output, and may even reverse its sign. Thus, an SRO which would otherwise set L too high for welfare-maximization is less able (or unable) to use increased L as a method of restricting output, and will tend to choose a lower L.

2.3 Shaked and Sutton (1981)

Shaked and Sutton also look at a situation in which qualities are innate. Their model is similar to Leland's in some ways, but it exhibits important differences. In particular, an implication of Shaked and Sutton's model is that an SRO will definitely set minimum quality levels higher than is socially optimal. This is a stronger result than is obtained from Leland's model, where this result only holds true if certain assumptions about opportunity costs and demand are satisfied.

In this section, I describe the Shaked and Sutton model, and explain how they manage to obtain such a strong result.

2.3.1 The framework

There are N workers who can be employed either as "lawyers" or as "labourers". While workers are identical when employed as "labourers", they differ from one another in the quality of service they offer when employed as lawyers, and they are ranked along the interval [0, N] on this basis, in ascending order of potential quality. It is assumed that the "lawyers" can form themselves into an SRO (or profession) and can impose a minimum quality level L upon themselves. The SRO will then supply X = N - L units of legal services, while labourers provide L units.

Now, let the expected quality of a profession with a minimum quality level of α and a maximum quality level of β be given by $\bar{q}(\alpha, \beta)$, where this function represents a preference ordering which is complete, reflexive, transitive and continuous, and where if $(\alpha_1, \beta_1) \preccurlyeq (\alpha_2, \beta_2)$, then $\bar{q}(\alpha_1, \beta_1) \leq \bar{q}(\alpha_2, \beta_2)$. Clearly, β will always equal N, and when there is a minimum quality level L imposed on the SRO then α will equal L. This definition of $\bar{q}$ as $\bar{q}(L, N)$ is more general than Leland's assumption that $\bar{q} = (L + \hat{\theta})/2$.

There are v consumers, each of which can purchase either one unit of legal services, or none. The services of labourers are a residual category of

expenditure. The utility level achieved by a representative consumer may be described by two utility functions:

— $V(h)$ gives the utility of an individual who purchases h units of labourers' services and none of lawyers'; and

— $U(h, \bar{q})$ gives the utility of and individual who purchases h units of labourers' services and *one* unit of service from a lawyer whose SRO is of quality $\bar{q}$.

Assume that $V(h)$ and $U(h, \bar{q})$ are increasing in all arguments; that $U(h, \bar{q}) > V(h)$, and that $U_h > V_h$, for all $h > 0$; and that $U(0, \bar{q}) = V(0) = 0$.

Given h and $\bar{q}$, let $M(h, \bar{q})$ be defined to be the quantity of labourers' services the consumer is willing to forego in return for 1 unit of lawyers' services of quality $\bar{q}$. That is,

$$U(h - M(h, \bar{q}), \bar{q}) = V(h)$$

which gives, on differentiating,

$$M_h = 1 - [V_h/U_h]$$

Since $\bar{q} = \bar{q}(L, N)$ and N is a constant, we know that $M_{\bar{q}}\bar{q}_\alpha$ is the quantity of labourers' services which an individual consuming legal services is willing to forego per unit increase in L. It is straightforward to show that $M_{\bar{q}}\bar{q}_\alpha$ is positive. Let it further be assumed that

$$\text{for all } h, \bar{q}, \beta: \quad (M_{\bar{q}}\bar{q}_\alpha)'_\alpha < 0.$$

This assumption is required to ensure that welfare is concave in L.

Lastly, it is assumed that an SRO will aim to maximize *per seller* earnings for its members, either relative earnings (relative to labourers' earnings) or absolute earnings.

2.3.2 Equilibrium

At equilibrium, the v consumers are divided into two groups:

—— those who enjoy h units of labourers' services only; and

—— those who buy one unit of legal services and enjoy h - M units of labourers' services.

There must be X consumers in the latter group, since X units of legal services are used, and there must be v - X consumers in the former group. Total demand for labourers' services is thus $X(h - M) + (v - X)h$, while total supply is L. From this, equilibrium in the market for labourers occurs when

$$M = M(h, \bar{q}) \qquad \text{and}$$

$$X(h - M) + (v - X)h = L$$

the latter of which is better written as

$$(2.3.1) \qquad vh = XM + L$$

The assumption that $U_{h} > V_{h}$, for all $h > 0$, implies that $1 > M_{h} > 0$, and thus that $h - M(h, \bar{q})$ is strictly increasing in h. From this and the fact that $M(0, \bar{q}) = 0$, we see that there must exist a unique equilibrium solution for any parameter values $L, \bar{q}$.

Writing $\bar{q}$ and X as functions of L, viz $\bar{q} = \bar{q}(L, N)$ and $X = N - L$, and differentiating (2.3.1) with respect to L, we get

$$(2.3.2) \qquad vh_{L} = XM_{h}h_{L} + XM_{\bar{q}}\bar{q}_{\alpha} - M + 1$$

which, on rearranging, gives

$$(2.3.3) \qquad h_L = \frac{XM_q\bar{q}_x - (M - 1)}{v - XM_h}$$

Meanwhile, differentiation of $M = M(h, \bar{q})$ gives

$$(2.3.4) \qquad M_L = M_h h_L + M_q \bar{q}_x$$

For an SRO to be viable, it must be the case that the per seller earnings of lawyers p are greater than or equal to those of labourers μ. That is, vaiblity requires that $p/\mu \geq 1$. It is clear that p/μ represents, at equilibrium, the number of units of labourers' services a consumer is willing to forego in return for one unit of lawyers', i.e. M. Thus, we can derive the following result.

Result 2.3.1: There exists some quality $\bar{q}'$ such that all single SROs with quality $\bar{q} \geq \bar{q}'$ are viable, while no single SRO of quality $\bar{q} < \bar{q}'$ is viable.

Proof: Viability requires that $M \geq 1$, so the lowest viable quality of SRO will be associated with an equilibrium value of M equal to 1. But when M = 1, each consumer enjoys the same total number of units of labour services h, so that $h = N/v$. Thus $\bar{q}'$ is given by

$$M(N/v, \bar{q}') = 1$$

From equation (2.3.3), when M = 1, h_L is given by

$$h_L = \frac{XM_q\bar{q}_x}{v - XM_h}$$

Since $M_h < 1$ and $X < v$, h_L will be positive. From (2.3.4), positive h_L implies that M_L is also positive at M = 1. This means that once M exceeds unity at some L, it will never fall below unity as L increases. □

This model can now be used to show that an SRO which is allowed to choose its own L will choose too high a level of L for welfare maximization.

Consider that the level of utility V(h) represents, at equilibrium, the level of utility achieved by all consumers. Thus V(h) can be taken as a measure of welfare. Now, maximizing V(h) as a function of L is equivalent to maximizing h as a function of L, and so welfare can also be identified with h. Suppose that lawyers care about their per seller income relative to that of labourers, M, so that the SRO chooses L to maximize M. It is clear from (2.3.4) that when $h_L = 0$, M_L is positive. So, when h is maximized ($h_L = 0$), the SRO can get a higher relative income per seller by increasing L beyond its socially optimal level. That is,

Result 2.3.2: An SRO which maximizes relative income per seller will set L higher than is socially optimal.

This result corresponds to result 2.2.2 in Leland's model.

Note that, if the SRO wants to maximize the absolute income per seller, rather than its relative income per seller, the same result holds. Let the income of each consumer be Y. A consumer who does not purchase legal services spends his entire income on h units of labourers' services. Thus $Y = \mu h$, which means that $M = p/\mu = ph/Y$, and so $p = MY/h$. Thus,

$$p_L = \frac{Y(M_L h - M h_L)}{h^2}$$

When $h_L = 0$, p_L has the same sign as M_L, and so an SRO which wants to maximize p will also set L higher than is socially optimal.

2.3.3 A comparison with Leland's model

Shaked and Sutton's main result (result 2.3.2) seems stronger than the equivalent result in Leland's model in that it states *unconditionally* that an SRO will set L higher than is socially optimal. If the assumptions of the two models were equally weak, this difference in results would be important. However, it turns out that this additional strength is easily explained by examining the different assumptions of the two models.

The most important differences in assumptions are:

(1) Shaked and Sutton assume that all sellers have the same expected income outside of the industry (or opportunity cost of being in the industry). That is, $\mu'(\theta) = 0$. Leland allows $\mu' \neq 0$.

(2) Shaked and Sutton model the 'world outside the industry' explicitly, so that μ, while the same for all sellers, is now an endogenous variable. $\mu(\theta)$ is exogenous in Leland's model.

(3) Shaked and Sutton do not assume that SROs maximize industry profits, as in Leland's model. Rather, SROs are assumed to maximize the *per seller* 'price', where this 'price' is either the absolute price each seller receives, p, or the ratio of this price to opportunity costs, p/μ.

(4) Shaked and Sutton provide better microeconomic foundations for consumer demand than Leland does. They assume, however, that all consumers have identical utility functions, so the Spence effect has no role in this model.

To see what is driving the Shaked and Sutton result, let $p_{max} = 0$ (since there is no Spence effect) and consider the effect of assumptions (1) and (3) on the Leland model.

The assumption that all sellers face the same opportunity cost μ seems unrealistic. In reality, the kind of people who would make good "lawyers" would also often make good doctors/professors/etc, so that opportunity costs might be expected to rise with quality, as is allowed for in Leland's model. When $\mu' = 0$, however, there is no problem of 'the bad driving out the good'. Instead, in terms of Leland's model, the maximal quality available in the industry $\hat{\theta}$ will simply equal the highest potential quality available, so that $d\hat{\theta}_m/dL = 0$ and $dX_m/dL = -1$. This means that the output restriction effect is strong and will easily dominate the average quality effect, resulting in the SRO setting L too high for welfare maximization.

More formally, under the assumption that $\mu' = 0$, it is clear that the second term in condition (2.2.6) in Leland's model will be negative, since $p_x < 0$ and $dX_m/dL < 0$. Thus, in Leland's model, constant μ and $p_{\neg x}$ ensures that an SRO will set L too high for welfare-maximization.

The assumption that an SRO will maximize *per seller* earnings is also questionable. Even a monopolist in an industry doesn't usually aim to maximize the price it can receive. Rather, it tries to maximize its profits, which are an increasing function not only of price but also of quantity. The question of what an SRO will want to maximize was mentioned in section 1.3.2 and I don't want to review it here. I just want to make the simple point that, if an SRO is concerned only with prices and not with quantities, it is going to have higher incentives to raise L than it would have if it did care about supply.

More formally, under the assumption that an SRO only cares about incomes per seller, the SRO's maximand in Leland's model becomes

$$\Pi' = p(\bar{q}_m, X_m) - [1/X_m] \cdot \int_L^{\hat{\theta}_m} \mu(\theta) d\theta$$

Maximizing this and deriving the corresponding version of equation (2.2.6) (for $p_{\neg x} = 0$) gives

$$\left.\frac{dW}{dL}\right|_{L_p'} = -\left[p_x(\bar{q}_m, X_m) \cdot X_m - \mu(\hat{\theta}) + \frac{1}{X_m} \cdot \int_L^{\hat{\theta}_m} \mu(\theta) d\theta\right] \frac{dX_m}{dL}$$

Under Leland's assumption that $\mu'(\theta) > 0$, the bracketted part of this expression is more negative than in the total profit-maximization case. Thus, in Leland's model with negative dX_m/dL, an SRO which cares only about *per seller* earnings, will raise L even higher relative to its socially optimal level than will an SRO who cares about supply as well as price.

Clearly, then, it is these two assumptions, combined with the assumption of identical consumers which ensures that $p_{sx} = 0$, that drive Shaked and Sutton's result. These assumptions may or may not be realistic for any given industry. However, because Leland's model can be used to explain Shaked and Sutton's results, it seems likely that Leland's model will be the more useful model, in general. An extra benefit of Shaked and Sutton's model, though, is that it can be easily extended to a two-SRO situtation, as is shown in chapter 5.

2.4 Shapiro (1983)

Shapiro's paper is concerned with examining the characteristics of competitive equilibrium in which reputation formation is used to solve the moral hazard problem (effect (A3) in chapter 1).[1]

Consider an industry in which each seller chooses the quality of her product, and then buyers - who observe each seller's quality 'offer' but cannot observe actual quality prior to purchase - decide whether to purchase. The lack of observability of actual quality creates an incentive problem: Any seller with a reputation for good quality will want to

---

[1] This model of reputation formation is chosen because it is intuitive and easy to work with. It should be noted, however, that other models do exist, for example Klein and Leffler (1981) and Allen (1984). These models do not differ from each other in terms of reputation mechanism, though, but rather in terms of how the necessary premium of price over marginal cost can be made consistent with the zero profit assumption of perfect competition/free entry. Klein and Leffler suggest that competition will occur in costly non-price dimensions until all positive profits are dissipated. Allen's model includes fixed, as well as unit, costs of quality, and sellers' outputs then move to equate price with *average* costs which gives zero profits. This seems realistic for manufacturing industries, but less so for industries in which output is not easily variable. Shapiro, on the other hand, allows sellers to produce just one unit (or none) per period, but suggests that there is an initial period during which sellers' qualities are not known and in which they all receive an identical (low) entry price. The difference between this entry price and the costs of high quality constitutes an initial investment upon which the later positive profits for high quality sellers are a return.

unexpectedly reduce her quality to the minimum possible quality level, so as to make short-run gains before buyers catch on.

In order to prevent this 'milking' of reputations, Shapiro suggests that equilibrium in a repeated game will involve a price-quality schedule such that high quality items sell at a premium above their cost. As well as being an incentive payment to induce quality maintenance, this premium can also be interpreted as a return to the initial investment required to achieve a good reputation in the first place.

Shapiro examines this price-quality schedule. He looks at how changes in the detection lag and in the minimum possible quality level will affect the position of the equilibrium price-quality schedule, and then how its position will affect welfare. He himself does not use his model to examine regulation. Indeed, one potential role for regulation is simply assumed away: When information transmission between consumers is not perfect, there is a role for a regulator to act as a collator of information. This possibility is examined in chapter 4. In Shapiro's model, however, perfect information transmission is assumed.

This section examines a different potential role for regulation, which derives from its ability to

(i) reduce the time taken to detect cheatings (which will always improve welfare); and/or

(ii) raise minimum quality levels from zero (which will, at least at first, improve welfare).

The role for external regulation is clear here. However, self-regulation will only be a feasible form of regulation if profits are also increasing as detection times fall and minimum quality levels rise. In Shapiro's model as it stands, however, this turns out not to be true. Indeed, faster detection and higher minimum quality levels *lower* the price-cost premiums for quality given to existing sellers, and so lower profits.

In Shapiro's model, then, an SRO (made up of existing sellers) will always want the opposite - slower detection times and lower minimum quality levels. Since this is directly opposed to consumer surplus maximization, self-regulation is not a feasible form of regulation, and external regulation is clearly superior. The aim in sections 2.4.5 and 2.4.6, therefore, is to extend Shapiro's model, and to show that self-regulation may be a feasible form of regulation after all.

2.4.1 Shapiro's framework

There are N identical potential sellers in the industry (where N is greater than the number of potential consumers), each of which is assumed to produce (in each period) either one unit of the good or nothing.

A seller's quality is equal to her effort a, where a is fully reversible. The unit cost of producing a good of quality a is denoted $\mu(a)$, [$\mu' \geq 0$, $\mu'' > 0$]. There is a minimum quality level L below which sellers cannot, or will not, produce. This may be because such production of lower quality levels is impossible; because they would be observable; or because sellers face an infinitely high penalty for violating L. There is free entry into the industry.

The demand side of the market is described by a set of heterogeneous consumers, who are each interested in purchasing a single unit of the commodity (or nothing), but who differ both in their general willingness to pay for the product ω, and in their taste for quality ν. The distribution of consumers is assumed to be contained within the region $[0,1]\times[0,1]$ in (ν,ω) space. The net benefits to a consumer of type (ν,ω) of a product of quality a selling at price p(a) are given by

$$u(a|\nu,\omega) = \nu a + \omega - p(a).$$

Each consumer will find the quality that maximizes his net benefits and then actually buy a unit of that quality if his net benefits are positive

(and if that quality is available). Thus, given p(a), we can derive aggregate demand.

Consumers can distinguish between sellers and can observe what quality is 'offered' by each seller. However, the *actual* quality provided by a seller is only observed after a 'detection lag' of w periods. There is perfect information transmission between consumers, so that all consumers have the same information about sellers' quality offers and about whether they have cheated.

An entrant into the industry will only be able to sell her product at an entrance price of p^e, whatever its actual quality, since quality is assumed to be unobservable (and also uncontractable) in advance. Thereafter, she will earn $p(R_t)$, where R_t is her quality reputation, which will equal $p(a_t)$ when qualities are stable.

The reputation adjustment process used in this model is simple:

— If a seller's quality offer is less than or equal to the most recent observation of that seller's quality, and if that seller has never been observed to cheat on a quality offer, then consumers will believe the quality offer;

— If a seller's quality offer is greater than the most recent observation of that seller's quality, and if that seller has never been observed to cheat on a quality offer, then consumers will not believe the quality offer, but *will* believe a repetition of the previously observed quality;

— If a seller has ever been observed to cheat on a quality offer (actual a_t < offered a_t), then consumers will only ever believe a quality offer of L thereafter.

This reaction to cheatings is a grim (trigger) strategy and might seem a little strong. Why shouldn't a seller which has milked its reputation just

once be able to promise never to do it again and produce a again? This strategy is, however, very easy to use in modelling[2], and it has the advantage that consumers' expectations of L *will* be rational in equilibrium (since if consumers expect L, they will only ever pay p(L) and sellers will have no incentive to produce any higher a).

Lastly, δ is the discrete time discount rate.

2.4.2 The unregulated competitive equilibrium

The equilibrium price-quality schedule can now be determined. This schedule must be such that sellers have no incentive to destroy their future reputation R_{t+1} by lowering their current quality a_t.

Because sellers require no premium to maintain L, we know that $p(L) = \mu(L)$. For a seller to want to maintain her reputation (which gives her $p(a) - \mu(a)$ forever), instead of milking her reputation (which gives her $p(a)$) - $\mu(L)$ for w periods and $p(L) - \mu(L)$ (which equals 0) thereafter), it must be the case that

$$\sum_{t=0}^{\infty} \delta^t[p(a) - \mu(a)] \geq \sum_{t=0}^{w-1} \delta^t[p(a) - \mu(L)]$$

Rearranging this gives the following seller incentive constraint (SIC)

$$\text{(SIC)} \qquad p(a) \geq [\mu(a) - (1 - \delta^w)\mu_L]/\delta^w$$

where μ_L is shorthand for $\mu(L)$.

---

[2] An alternative strategy (tit-for-tat) is examined in chapter 4 and compared with the trigger strategy. It is shown that, when $p(L) = \mu(L)$ and $w = 1$, then both strategies give identical reputation constraints.

Free entry, meanwhile, requires that the profits of a potential entrant into any quality segment of the market be nonpositive. That is, if p^* is the entry price obtainable, it must be the case that

$$\sum_{t=0}^{w-1} \delta^t[p^* - \mu(a)] + \sum_{t=w}^{\infty} \delta^t[p(a) - \mu(a)] \leq 0$$

which can be rewritten as

$$(p^* - \mu(a))(1 - \delta^w) + (p(a) - \mu(a))\delta^w \leq 0$$

Note, however, that p^* must be equal to μ_L, since otherwise there is a potentially infinite supply of fly-by-night sellers who would overrun the market with items of minimum quality if they could make profits doing so. Substituting this into the above and rearranging gives the following free entry condition (FEC)

$$\text{(FEC)} \qquad p(a) \leq [\mu(a) - (1 - \delta^w)\mu_L]/\delta^w$$

This condition is exactly the same as (SIC) but with the inequality reversed.

Putting (SIC) and (FEC) together we get that, at equilibrium, the following equality must hold.

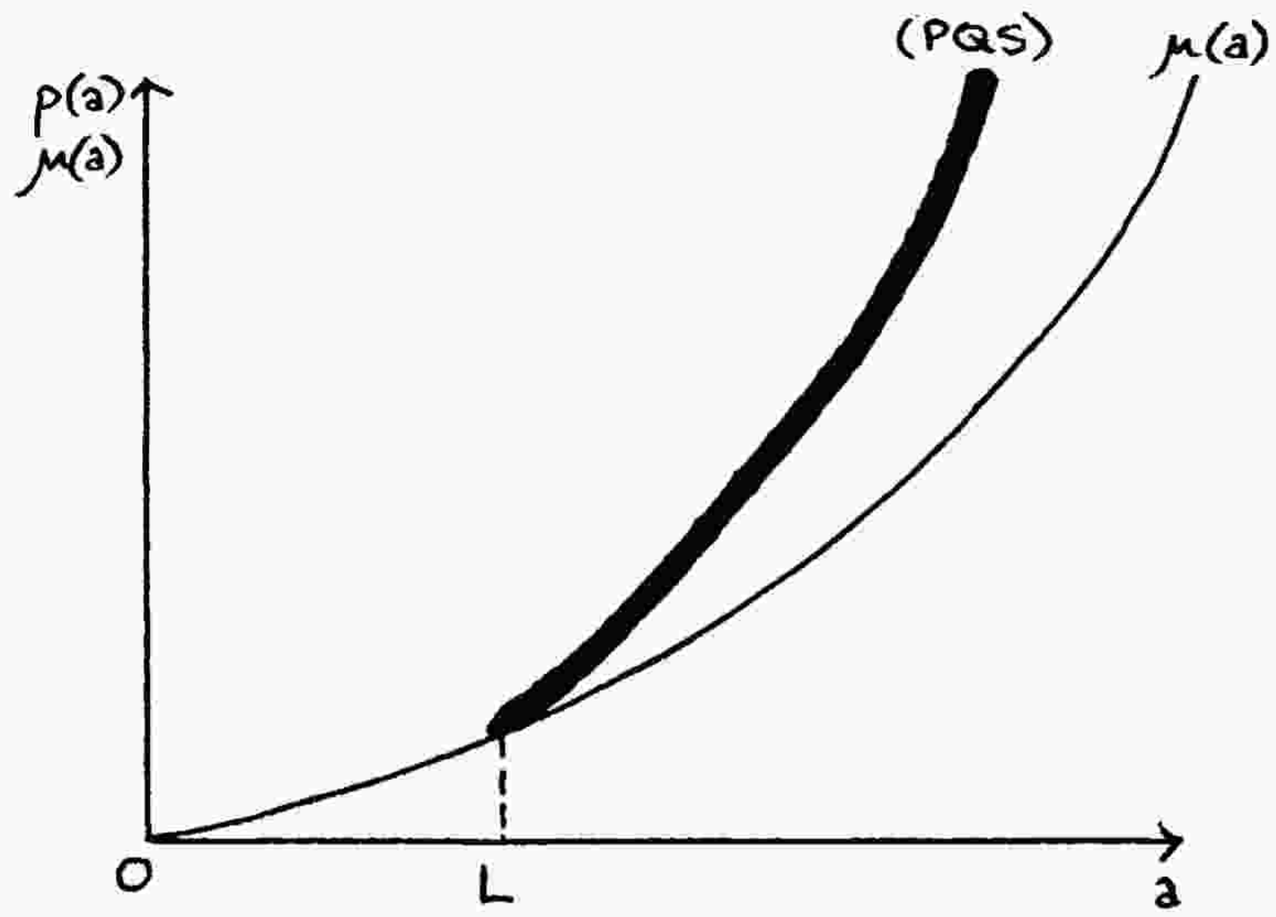
$$p(a) = [\mu(a) - (1 - \delta^w)\mu_L]/\delta^w$$

Let $(\delta^{-w} - 1)$ be denoted by the parameter m. This equality can now be rewritten as the following price-quality schedule (PQS)

$$\text{(PQS)} \qquad p(a) = \mu(a) + m(\mu(a) - \mu_L)$$

Note that m is increasing in w, where w is the detection lag. Thus, a reduction in the detection lag w will reduce m.

Schedule (PQS) is shown below.



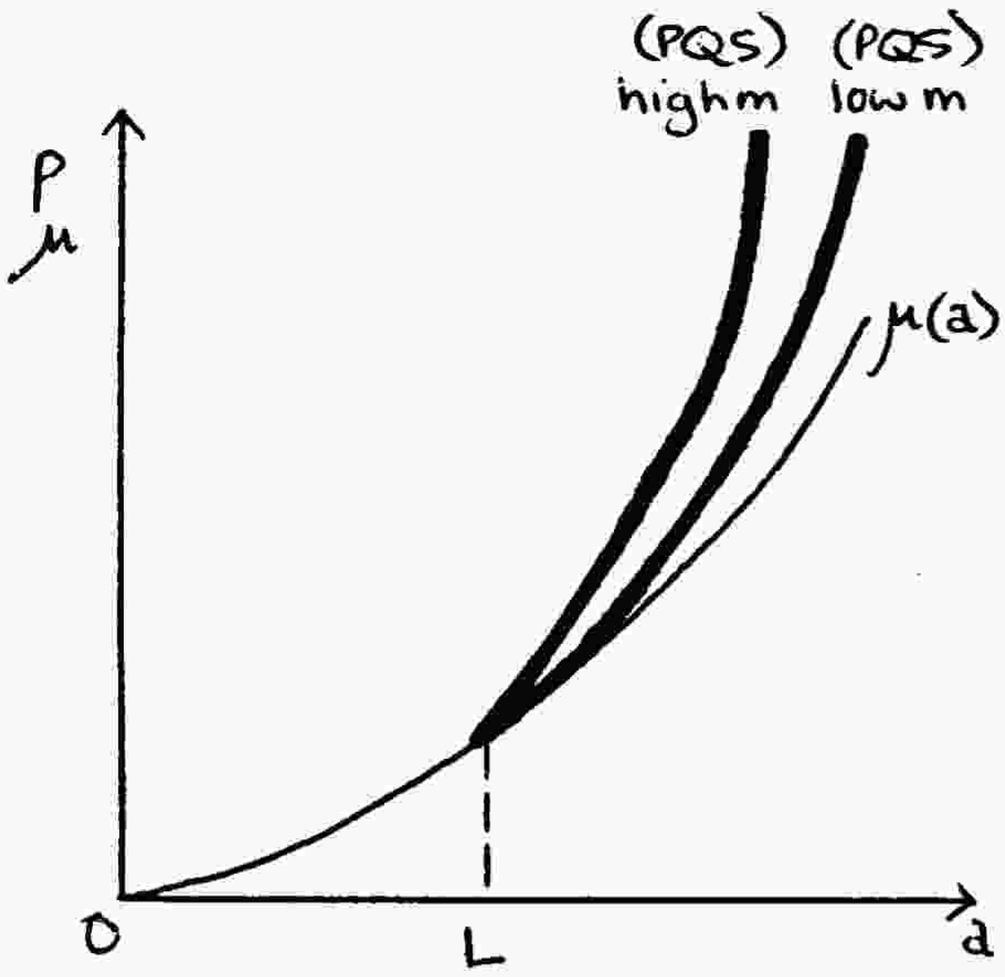
As noted above, each consumer will purchase that quality which maximizes his net benefits, as long as these benefits are positive and *as long as that quality is available*. Because all sellers make identical (zero) profits (which means that they are indifferent as to which quality to produce so long as they can sell it), and because there is only one seller per consumer and vice versa, equilibrium will involve sellers tailoring their quality choices exactly according to consumers' preferences. That is, at equilibrium, sellers will be spread out along (PQS) in such a way that each consumer can purchase his optimal quality.

2.4.3 Welfare analysis of the equilibrium

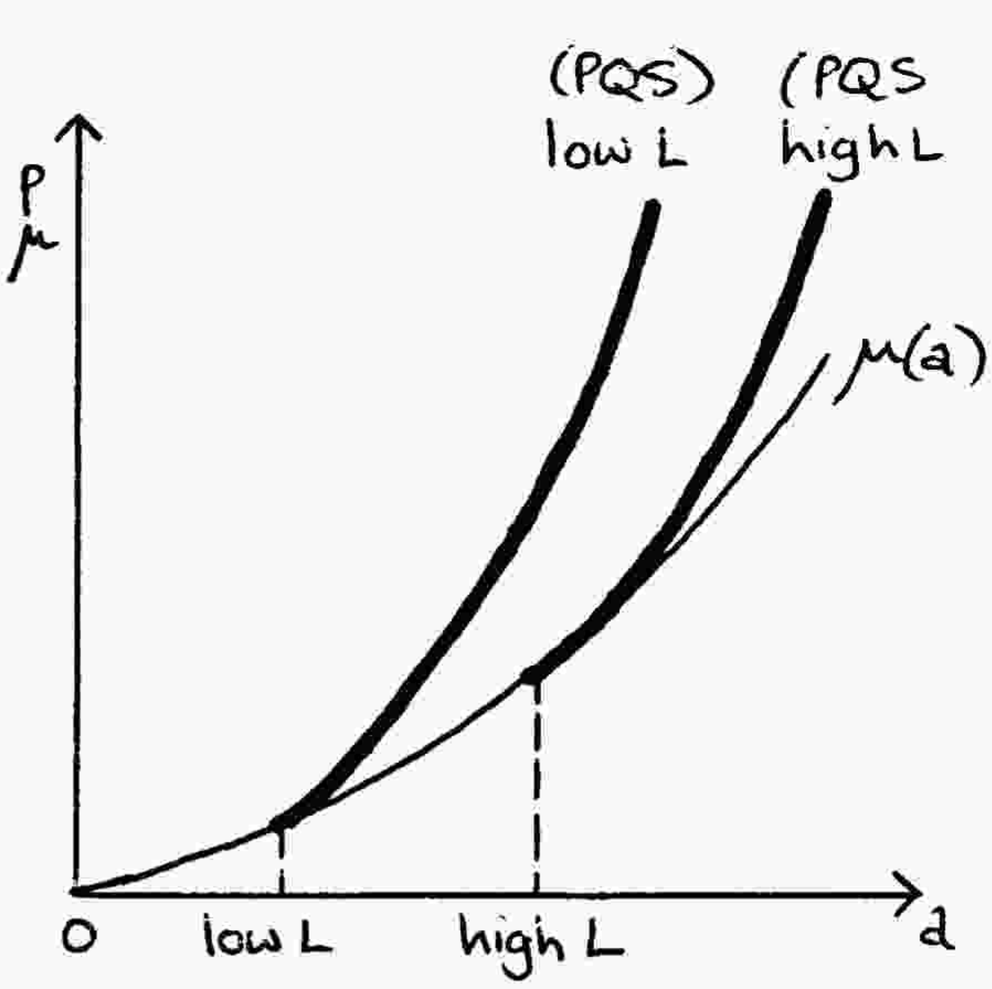
Let us examine (PQS) in more detail. As the lag with which consumers update reputations w gets smaller, m gets smaller and (PQS) rotates clockwise (pivoted on (L, μ_L)), approaching its perfect information position (at which $p(a) = \mu(a)$). This is shown in Figure (1) below.

As the minimum quality level L increases, meanwhile, (PQS) shifts to the right, as is shown in figure (2) above. One interpretation of this is that increases in L raise the price at which entrants can sell their products. This reduces the investment necessary to build up a reputation and hence reduces the premiums for high quality products. Alternatively, raising L

(1)



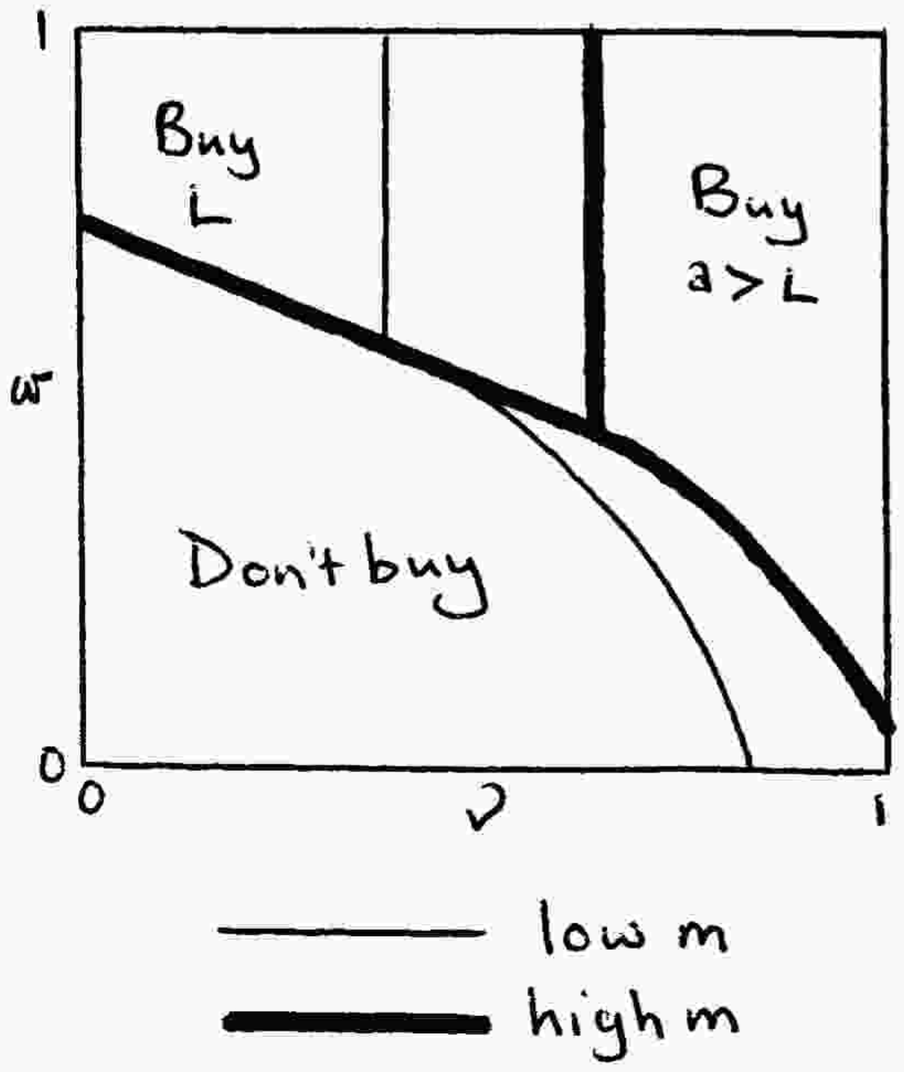
(2)



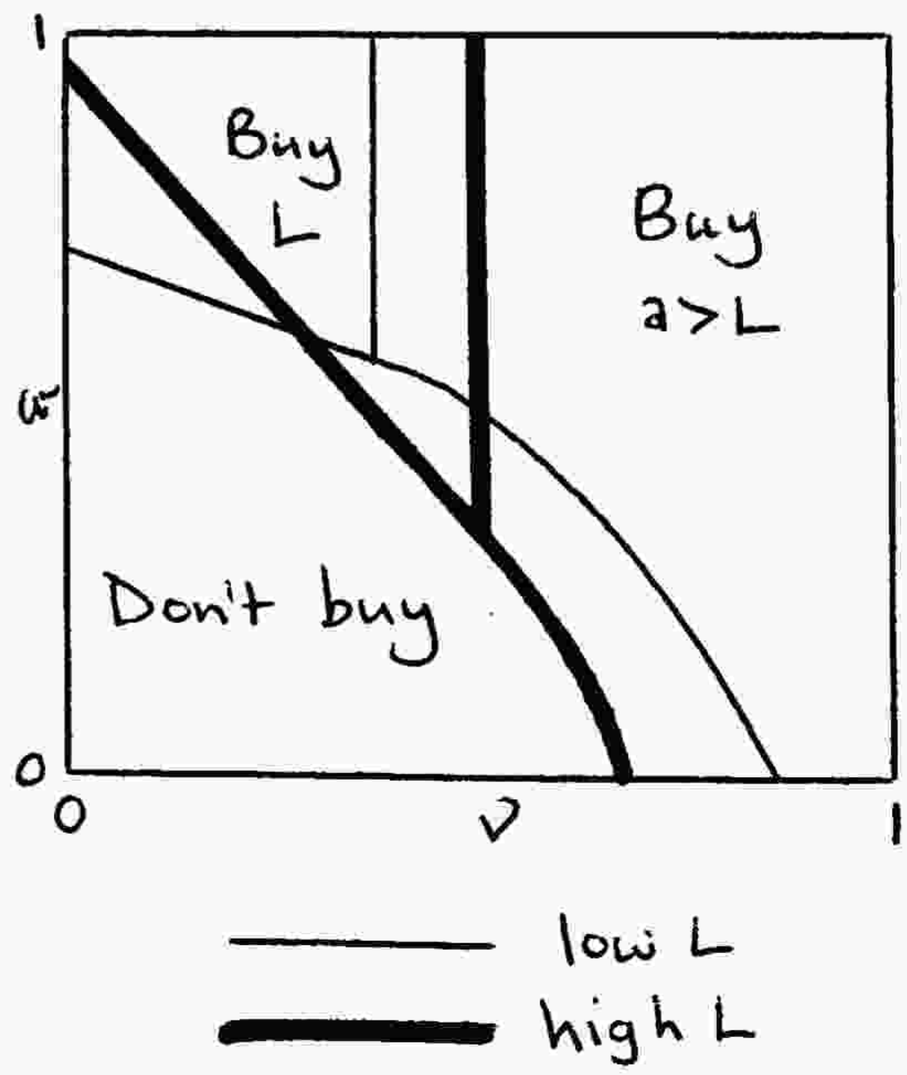
reduces the profits that a seller can earn by milking its reputation and hence reduces the premium necessary to induce quality maintenance.

These movements in (PQS) have important effects on consumer surplus. As might be expected, an increase in m will definitely lead to losses in consumer surplus (i) because consumers who choose qualities above L shift to lower qualities (since the marginal price of quality rises), and (ii) because some consumers leave the market altogether, even though their net benefits from the product exceed its cost of production. This change in consumption patterns is shown in figure (3) below.

(3)



(4)



An increase in L will have more ambiguous consequences. All consumers who choose qualities above L will enjoy a utility gain, since the price they have to pay for their quality choice falls. Those consumers whose optimal quality is lower than L, however, will suffer since the nearest available quality is now even further away from their optimum. This will result in the change in consumption patterns shown in figure (4) [Note that these diagrams are my own, not Shapiro's].

In general, the consumers' optimal L will lie above the optimal quality of the lowest ν consumer, since raising L slightly from this level will have no first-order effects on low ν consumers but will give real benefits to higher ν consumers. For similar reasons, the consumers' optimal L will also lie below the optimal quality of the highest ν consumer. But apart from this we cannot say much about the position of the consumers' optimal L.

These implications for consumer surplus are summarised in the following result.

Result 2.4.1: Increasing m will reduce consumer surplus. Increasing L will raise (reduce) consumer surplus when L is low (high) enough.

2.4.4 Implications for regulation

The importance of this result is clear when m and L are choice variables (rather than exogenous), as could be the case under regulation. At equilibrium, no seller makes positive profits, and so welfare is identical to consumer surplus. An external regulator, therefore, would reduce m by as much as it possibly can, and will set L at the level which maximizes consumer surplus.

Is self-regulation also a feasible form of regulation? One argument for self-regulation might run thus: Since all sellers make zero profits at equilibrium, an SRO will be indifferent between different levels of m and L. Thus, there is no reason why an SRO should *not* be willing to choose the

welfare-maximizing levels of m and L. Moreover, an SRO, with its better knowledge of the industry, may even be able to reduce m to below the minimum level which an external regulator can achieve.

Against this, however, the obvious argument is that there is also no reason why an SRO *should* choose these levels, and if there are even infinitesimal costs of achieving these levels, an SRO will not want to bother.

In addition to this, we should note that those sellers which already exist in the industry will mostly have already done their initial investment in reputation $(p^m - \mu(a))(1 - \delta^w)/(1 - \delta)$. An SRO which is made up of existing sellers will not, therefore, be indifferent between different levels of m and L. If the SRO can shift the price-quality schedule upwards, then it will raise the per seller lifetime profits of its members (that is, it will allow positive per seller lifetime profits to be made). This gives us the following result.

Result 2.4.2: An SRO made up of existing sellers in the industry can raise per seller lifetime earnings of its members by increasing m or reducing L.

From results 2.4.1 and 2.4.2, there is a direct conflict between consumers' interests and per seller lifetime earnings (except when L is high). Thus, if SROs want to maximize per seller lifetime earnings, self-regulation will not be a feasible form of regulation.

The following subsection extends Shapiro's model and, using a different assumption about what SRO's maximize, examines whether this conclusion remains true.

2.4.5 An extension of Shapiro's model

The most obvious objection to the above description of self-regulation is that, as (PQS) shifts upwards, the number of consumers in the industry will fall, and this will necessarily reduce the number of sellers in the industry (since supply equals demand at equilibrium). Thus, if an SRO

cares about total industry profits, rather than profit *per seller remaining in the industry*, then this SRO will not want (PQS) to shift upwards indefinitely.

There is obviously some truth in this. However, I want to attack a different assumption: the assumption that all N potential sellers are identical. In what follows, this assumption is relaxed, and the implications of this change for Shapiro's model are examined.

Let us replace the supply side of Shapiro's model by the following. There are N potential sellers (where N is still greater than the number of potential consumers), and each seller now has unit costs of producing a good of quality a equal to $\mu(a) + \phi$, where sellers' ϕs lie in the interval $[\underline{\phi}, \bar{\phi}]$ and are distributed according to a distribution function $H(\phi)$, $[H(\underline{\phi}) = 0,\ H(\bar{\phi}) = N]$. Note that ϕ is not a sunk cost, so that any seller who leaves the industry will cease to pay it.

The main effect of introducing seller heterogeneity into this model is that the price-quality schedule will now depend on the ϕ of the marginal seller in the industry, $\hat{\phi}$. If this seller is to be prevented from cheating on her quality choice, it must be the case that

$$p(a) \geq (\mu(a) + \hat{\phi}) + m((\mu(a) + \hat{\phi}) - (\mu_{_} + \hat{\phi}))$$

This seller incentive constraint can be rewritten (analogously to (SIC) above) as

$$\text{(SIC')} \qquad p(a) \geq \mu(a) + m(\mu(a) - \mu_{_}) + \hat{\phi}$$

Note that since all inframarginal sellers have $\phi < \hat{\phi}$, (SIC') will always hold (and will be strict) for all inframarginal sellers when it is binding for the marginal seller.

Free entry implies that this marginal seller just makes zero profits. To prevent hit-and-run entry into the entrants' market, p^* must now be equal

to $\mu_L + \hat{\phi}$. So, analogously to (FEC) above, we get a new free entry condition:

(FEC') $\quad p(a) \leqslant \mu(a) + m(\mu(a) - \mu_L) + \hat{\phi}$

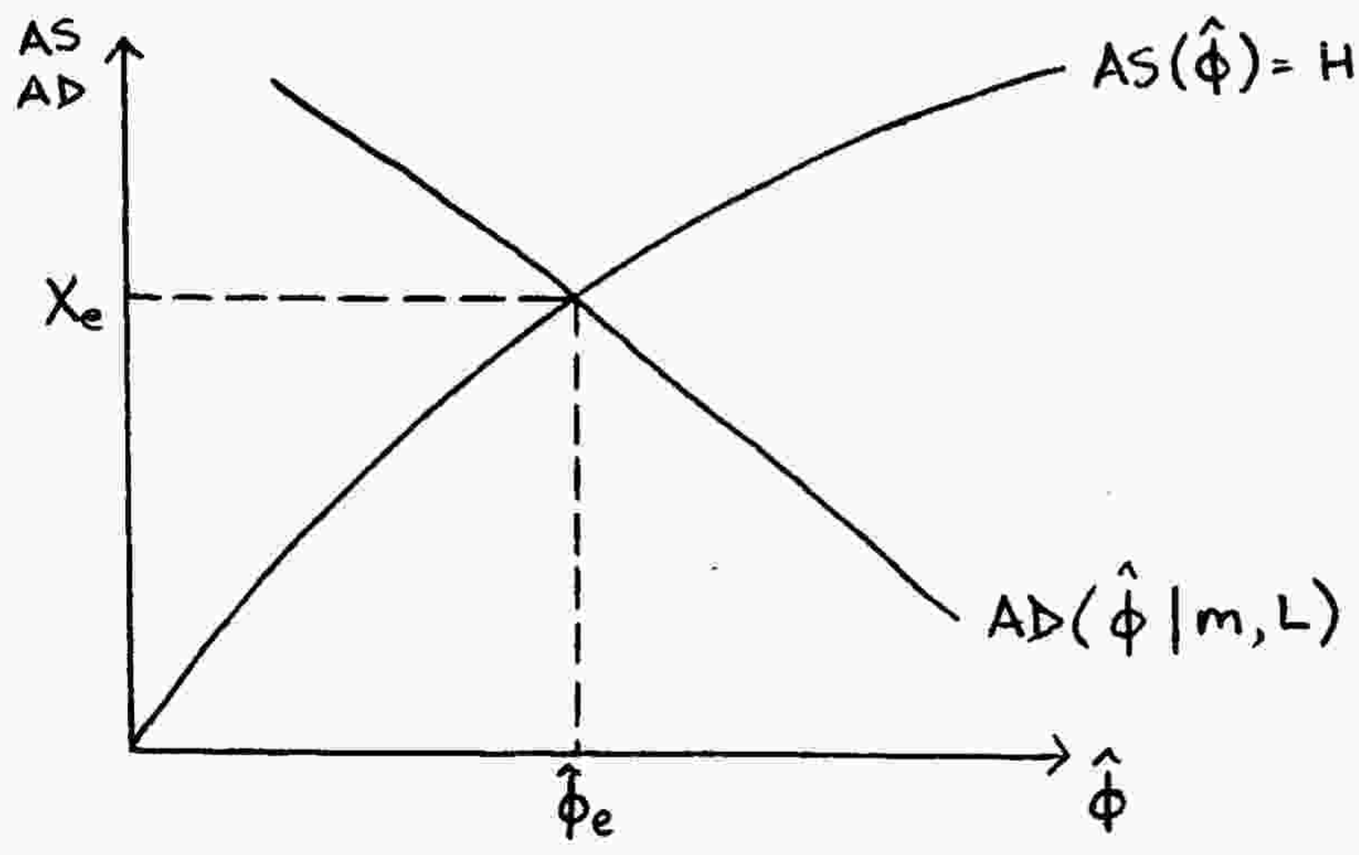
Again, note that all inframarginal sellers will make strictly positive lifetime profits when (FEC') is binding for the marginal seller.

From (SIC') and (FEC'), the price-quality schedule, at equilibrium, will be

(PQS') $\quad p(a) = \mu(a) + m(\mu(a) - \mu_L) + \hat{\phi}$

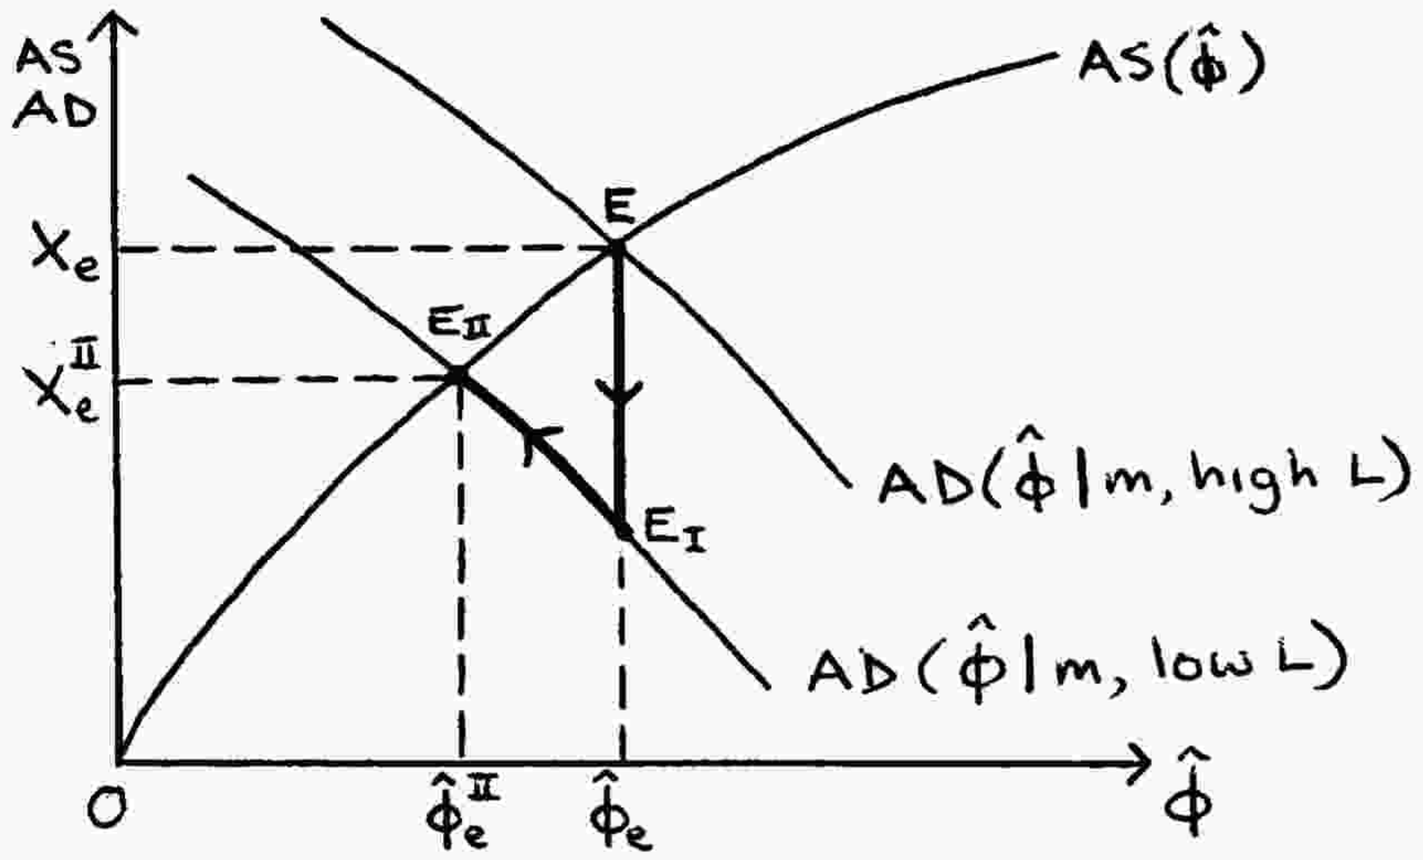
We can now characterise both aggregate demand and aggregate supply in terms of $\hat{\phi}$. The aggregate supply of sellers when the marginal seller has $\phi = \hat{\phi}$, $AS(\hat{\phi})$, is simply $H(\hat{\phi})$, which is increasing in $\hat{\phi}$.

Aggregate consumer demand, meanwhile, is equal to the number of consumers that will stay in the industry given the existing (PQS'), which depends on m, L, $\hat{\phi}$, and the distribution of consumers in (ν, ω) space. Even without making any assumptions about the shape of this distribution, we can say something about this AD function. For given m and L, as $\hat{\phi}$ rises, (PQS') will rise and some consumers will leave the industry. Thus aggregate demand is falling in $\hat{\phi}$. For given $\hat{\phi}$, meanwhile, an increase in m will lower demand, and an increase in L will either raise demand when L is low or lower demand when L is high. Thus $AD(\hat{\phi}, m, L)$ has $AD_{\hat{\phi}} < 0$, $AD_m < 0$ and AD_L is ambiguous.



Equilibrium $\hat{\phi}$, $\hat{\phi}_e$, for given m and L, is determined by the intersection of $AS(\hat{\phi})$ and $AD(\hat{\phi}|m,L)$, as is shown above.

Now, suppose that L is low, so that AD_L is positive. In this case, a fall in L will cause the AD curve to shift downwards, which will reduce both X_e (the equilibrium number of sellers/consumers in the industry) and $\hat{\phi}_e$, as is shown below (i.e. equilibrium shifts from point E to point E_{II}). The same result will hold for an increase in m.



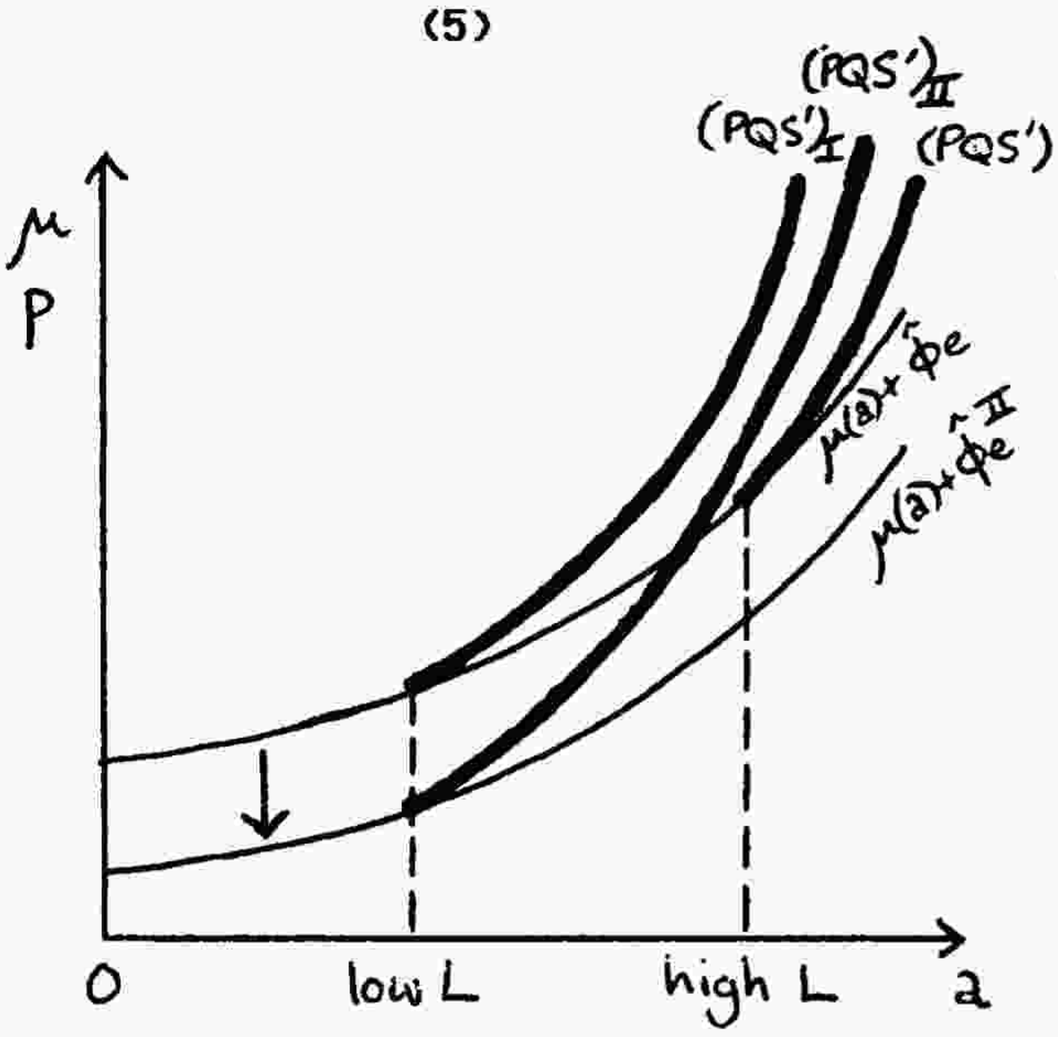
This shift of equilibrium can be separated into two parts:

(i) For given $\hat{\phi}$, the downwards shift of the AD curve will result in a large fall in X_e, and those sellers which are able to remain in the industry at point E_I (i.e. those sellers which are lucky enough to still have customers), will be making higher profits than before. These increased profits are due to the increased price-cost mark-up (remember that, for given m, $\hat{\phi}$ and a, $\partial p/\partial L = -m\mu_e < 0$). This new (PQS'), $(PQS')_I$, is pictured in figure (5) below. The corresponding situation when AD shifts down because of a rise in m, rather than a fall in L, is shown in figure (6).

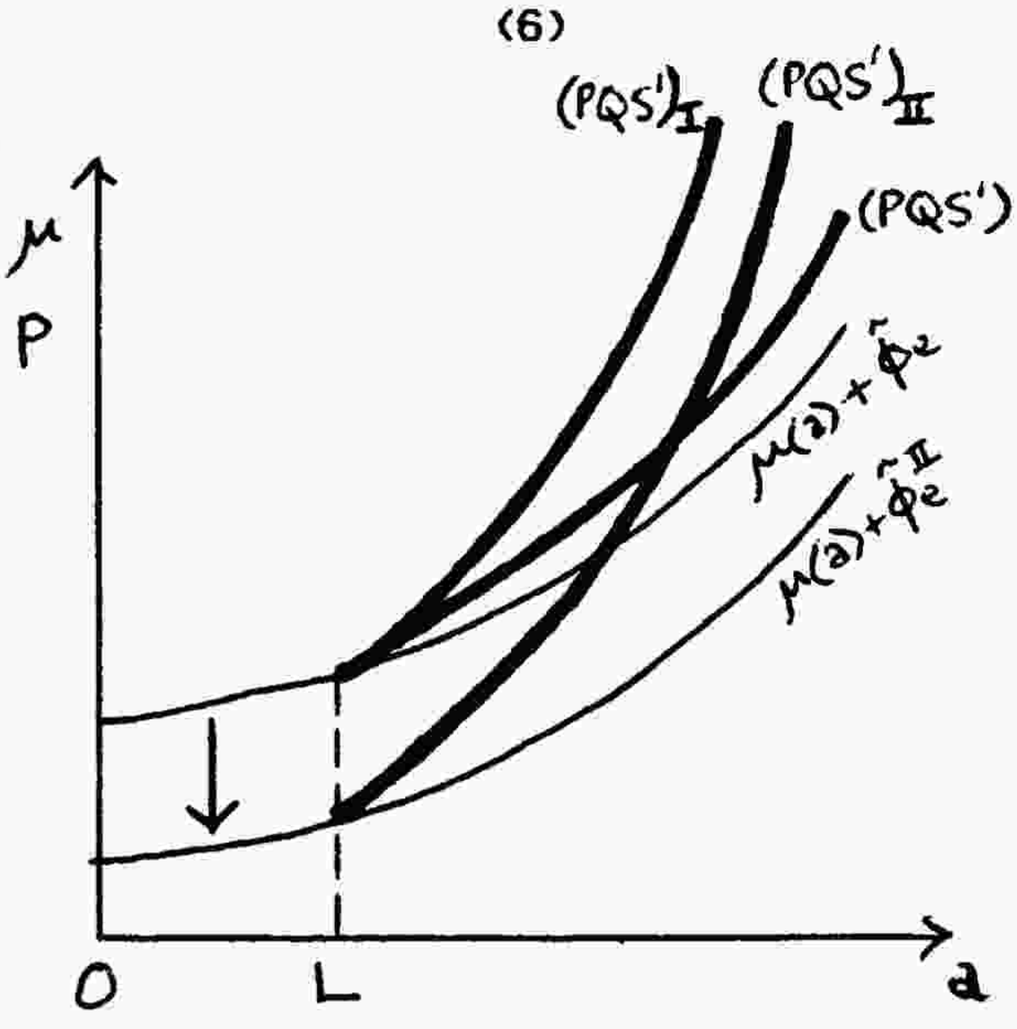
(ii) At E_I, however, there is excess aggregate supply. This will result in prices being bid down, which will in turn lead to some high $\hat{\phi}$ sellers leaving the industry and to some consumers reentering the industry. This process will continue until AS = AD again (at E_{II}), at which

point $\hat{\phi}_-$ is lower than it is at E. Note that (PQS') shifts downwards as $\hat{\phi}_-$ gets lower, so that, at equilibrium, sellers will lie along schedule $(PQS')_{II}$ in figures (5) and (6).

(5)



(6)



2.4.6 Implications for self-regulation

If the first of these two effects was present on its own, then an SRO which wants to maximize per seller profits would definitely like to lower L/raise m. This effect, which can only be bad for consumers, *is* the only effect which is exhibited in the simple Shapiro model.

The second effect, however, acts on per seller profits in the opposite direction to the Shapiro effect. Let us distinguish between pre-entry expected per seller life-time profits $E(\pi(\phi))$ and post-entry actual per seller lifetime profits $\pi(\phi)$, for a type ϕ seller.

In the simple Shapiro model, $E(\pi(\phi))$ is always equal to zero, for all ϕ, but $\pi(\phi)$ may be increased by shifting (PQS) upwards. Thus, a self-regulating industry which cares about per seller profits will generally want to raise m or lower L. In the model described here, this positive effect on $\pi(\phi)$ still holds. However, there is also a negative effect on $E(\pi(\phi))$. To see this, we must use the following result.

Result 2.4.3: $E(\pi(\phi))$ is strictly positive, increasing in $\hat{\phi}_-$, and decreasing in m.

Proof: Examine the pre-entry expected profits of any ϕ seller $E(\pi(\phi))$, when the marginal $\hat{\phi}_-$ seller makes zero profits. Since $p^- = \mu_L + \hat{\phi}_-$, we know that

$$E(\pi(\hat{\phi}_-)) = (\mu_L + \hat{\phi}_- - \mu(a) - \hat{\phi}_-) + (p(a) - \mu(a) - \hat{\phi}_-)/m = 0$$

and

$$E(\pi(\phi)) = (\mu_L + \hat{\phi}_- - \mu(a) - \phi) + (p(a) - \mu(a) - \phi)/m$$

Thus

$$E(\pi(\phi)) = (m + 1)(\hat{\phi}_- - \phi)/m$$

This is clearly increasing in $\hat{\phi}_-$ (and decreasing in m). □

From this result, it is clear that, if m is increased or if L is increased from a low level, so that $\hat{\phi}_-$ falls, then $E(\pi(\phi))$ will fall, for all ϕ. Thus we get,

Result 2.4.4: When a self-regulating industry has to choose m and/or L *in advance* of any investment in reputation (so that only $E(\pi(\phi))$ is important for each ϕ, not $\pi(\phi)$), then m and L will be chosen so as to maximize the number of consumers in the industry, even when the SRO cares only about *per seller* profits. This will in turn maximize consumer surplus. So, in this case, self-regulation will be welfare-optimal.

Proof: For each seller, $E(\pi(\phi))$ is maximized when $\hat{\phi}_-$ is maximized and m is minimized. This requires, however, that L and m be at their welfare-optimal levels. □

Thus, in this case, self-regulation will be welfare-optimal. Moreover, if self-regulation can achieve a lower level of m than an external regulator can, then self-regulation will be welfare-superior to external regulation.

In a more realistic industry, where there has already been some investment in reputation, the Shapiro effect becomes important, and life-time per seller profits will become $\pi(\phi) + E(\pi(\phi))$. How, then, will total per seller profits be affected by a downwards shift in AD? For simplicity, let us consider a rise in m (as described in figure (6) above). Note that examining the implications of a fall in L is complicated by the fact that some sellers will reduce quality to below previously available levels.

At the new equilibrium, the price-cost mark-up will be greater than before for high quality sellers, but it will be *lower* than before for low quality sellers. Similarly, on the consumer side, instead of all consumers being worse off (as in Shapiro's model), we now find that consumers who purchase lower quality goods will actually be *better off* than before.

These results imply that we should consider (i) which part of the industry will be in charge of setting L and m under self-regulation; and (ii) whether we want to give identical weights to the utilities of all consumers. For example, if it will be the lower quality sector of the industry that chooses L and m under self-regulation, the $\hat{\phi}$-effect will outweigh the Shapiro effect for them, so that L will be raised (until $AD_{L} = 0$) and m will be reduced. And this will most benefit consumers of high quality goods. If, however, it will be the higher quality sector of the industry that chooses L and m, the reverse will be true, with the Shapiro effect dominating the $\hat{\phi}$ effect.

2.5 Mayer and Neven (1991)

Mayer and Neven also look at moral hazard and how sellers can be persuaded to behave honestly. However, their model is very different from Shapiro's:

— In Shapiro's model, qualities and per seller profits are endogenous, and can be affected by changes in the minimum quality level L and the detection lag w. The role of the regulator, then, is to set L and w.

— In the Mayer and Neven model, qualities and per seller profits are exogenous (except that a seller of innate quality θ can decide to cheat and provide zero quality). The aim of the regulator is now to give sellers sufficient incentives for honesty, via the use of capital requirements and penalties for cheating, given the exogenous profit levels.

Capital requirements and *ex post* penalties, in this model, are perfect substitutes for each other in the provision of incentives for honesty. However, capital requirements have an additional function of deterring entry by sellers with low capital levels. It is assumed that capital requirements and penalties for cheating cannot be conditioned on θ, for different sellers, even though this would be optimal. If, therefore, capital levels are related to quality in some way, so that capital requirements affect the various θs differently, the optimal use of capital requirements *together with* penalties for cheating will be an improvement on (or at least as good as) the optimal situation using either one on its own.

The following analysis examines this optimal use of capital requirements and penalties for cheating in more detail. It also derives conditions under which self-regulation may be feasible. However, although it provides some interesting implications, some of the assumptions and analysis in Mayer and Neven's model is a little questionable, as is argued in 2.5.4. The model given in chapter 4 can be seen partly as an attempt to set some of Mayer and Neven's ideas within a more economically sound framework.

2.5.1 The model

As in Leland, seller i has innate quality θ_i, where θ is distributed in some way on $(0, 1)$. Each seller's profits π are an increasing concave function of their quality θ $[\pi'(\theta) > 0, \pi''(\theta) < 0]$. However, instead of supplying services honestly, sellers can cheat consumers and earn $\pi_d > \pi(\theta)$, $\forall\theta$, where π_d is earned entirely at the expense of consumers. Sellers can earn a return of μ outside the industry.

Sellers have capital of an amount C, where C is uniformly distributed on $(0, \bar{C}(\theta))$ for each θ, where $\bar{C}'(\theta) \geq 0$. C is assumed to be both appropriable by a regulator and destroyed by expulsion from the industry.

The regulatory framework is composed on a capital requirement CR and an *ex post* penalty J. The penalty could be a fine or imprisonment. it is assumed that all misbehaviour is observed (although this assumption can be relaxed). There is also a probability $(1 - \rho)$ $[0 \leq \rho \leq 1]$ that honest sellers are deemed to have misbehaved, in which case they incur the penalty and lose their capital.

2.5.2 Regulatory contracts

Mayer and Neven first examine external regulation. They assume that an external regulator will regulate optimally, so that only the incentive compatibility and participation constraints *for sellers* need to be examined.

A seller will supply honestly provided the following seller incentive compatibility (SIC) constraint holds.

$$\pi(\theta) - (1 - \rho)(J + C) \geq \pi_d - (J + C)$$

which can be rewritten as

$$\text{(SIC)} \qquad \pi(\theta) \geq \pi_d - \rho(J + C)$$

Assuming that this constraint is satisfied (so that sellers supply honestly), the seller participation constraint (SPC) is

$$(SPC^1) \qquad \pi(\theta) - (1 - \rho)(J + C) \geq \mu$$

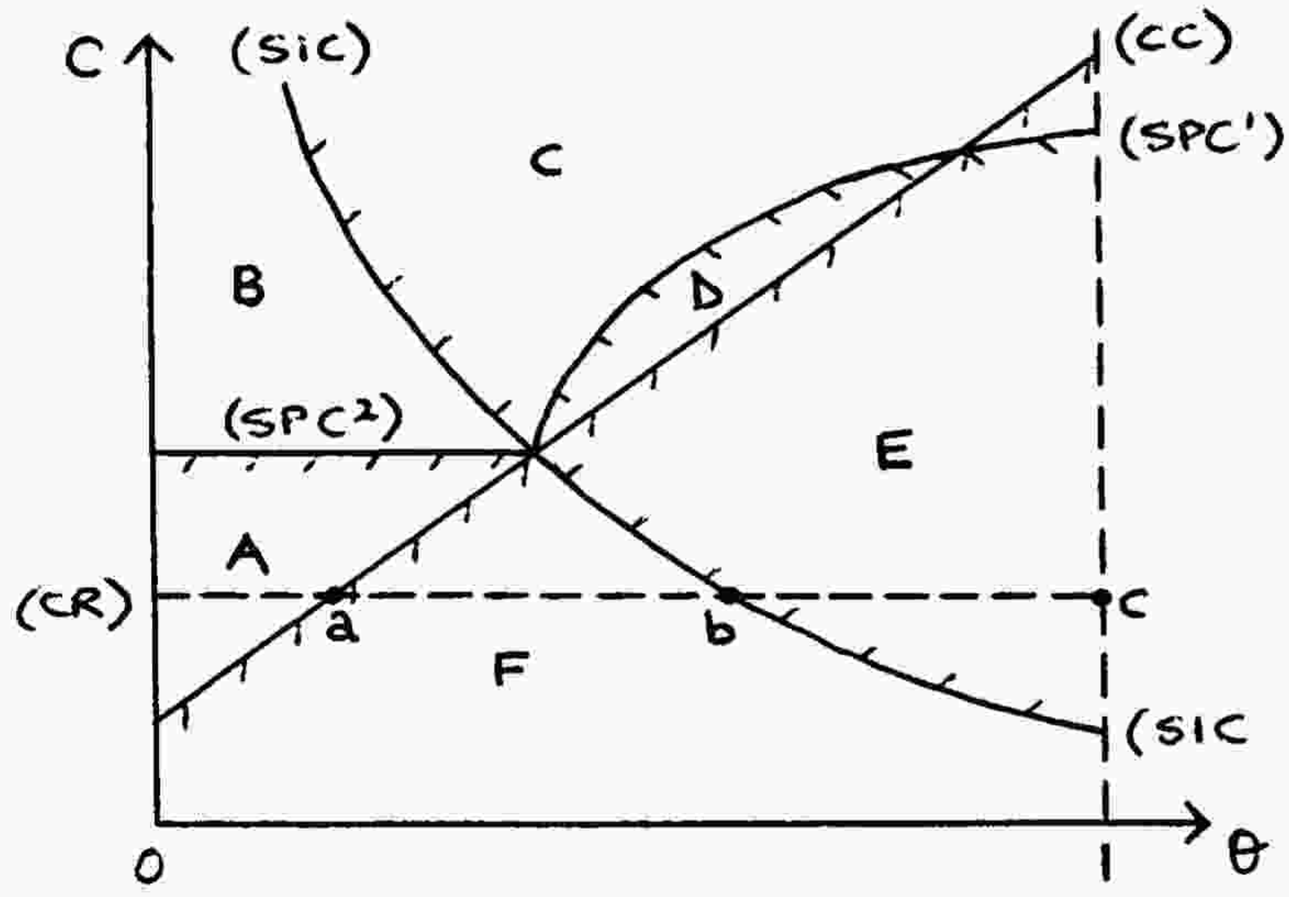
And where (SIC) is not satisfied, the seller participation constraint is

$$(SPC^2) \qquad \pi_{\Box} - (J + C) \geq \mu$$

Lastly, the least upper bound on sellers' capital is also a capital constraint (CC).

$$(CC) \qquad C \leq \bar{C}(\theta)$$

These four conditions can be illustrated (for given J) as below.



Constraints (SIC), (SPC¹) and (SPC²) have a unique common intersection, at which point

$$C_{\bullet} = \pi_{\Box} - J - \mu \quad \text{and} \quad \theta_{\bullet} = \pi^{-1}[(1 - \rho)\pi_{\Box} + \rho\mu]$$

The behaviour of sellers in the above figure is easily summarised. As (CC) has been drawn, there are no sellers in the areas A, B, C and D. If, however, (CC) was not constraining (i.e. if it was shifted sufficiently upwards), sellers in areas B and C would not enter;, sellers in area A

would behave the same as sellers in area F, i.e. they would enter and cheat; and sellers in area D would behave the same as sellers in area E, i.e. they would enter and be honest.

A capital requirement condition CR is now introduced. This gives us a new constraint

(CR) $\quad C \geq CR$

This is shown as a dashed figure in the above diagram.

This condition prevents some sellers from area F, which would cheat, from entering. However, some sellers in area E are excluded. There is therefore a trade-off between allowing dishonest sellers to enter and excluding honest sellers. The first implication of this is that efficent regulation involves setting J such that constraints (SIC), (SPC^1), (SPC^2) and (CC) all intersect. That is, J should be set such that

$$J = \pi_d - \mu - \bar{C}(\theta_-)$$

Mayer and Neven then argue that if equal weight is attributed to cheating and inappropriate exclusion of sellers, CR will be set so that *ab* in the above figure is exactly equal to *bc*. They use this analysis to examine how changes in ρ, μ, π_d, $\bar{C}'(\theta)$ affect the optimal CR, when J is set efficiently.

Consider, for example the effect of raising μ. This leaves constraints (SIC) and (CC) unchanged, but shifts constraints (SPC^1) and (SPC^2) upwards along (SIC) to (SPC^1)' and (SPC^2)'. At this point, however, the capital constraint is not binding. So J will be reduced until all the constraints (SIC), (SPC^1), (SPC^2) and (CC) intersect at the same point again (so that θ_- is raised). This upward shift in (SIC) will also result in CR rising.

This, and similar analysis for for changes in ρ, π_d and $\bar{C}'(\theta)$, gives us

Result 2.5.1: The effects of changes in $\bar{C}'(\theta)$, ρ, μ, and π_d on optimum levels of J, CR and θ_m are set out below.

Increase in	J	CR	θ_m
$\bar{C}'(\theta)$	Same	Rises	Same
ρ	Rises	Falls	Falls
μ	Falls	Rises	Rises
π_d	Rises	Rises	Rises

Two of these results, in particular, seem useful and intuitive:

(1) Capital requirements should be limited when there is a poor relation beween capital and quality (low $\bar{C}'(\theta)$); and

(2) Capital requirements (and so barriers to entry) can be reduced if π_d can be reduced (for example, by reducing the ability of sellers to defraud customers).

2.5.3 Self-regulation

As stated above, Mayer and Neven simply assume that an external regulator will act with the right incentives, so that incentive schemes do not have to determined. When they examine self-regulation, however, they introduce incentive compatibility and participation constraints for the regulator as well. They also simplify their model, by assuming that all sellers are identical (both in terms of quality and in terms of capital).

Suppose that there is a group of sellers, called an SRO, that comprises n identical members. Each member puts up the same capital CR and agrees to the following rules. New members will be admitted only if they too subscribe capital CR. If they cheat, they forego their capital and are subject to an additional penalty J. There is a probability $(1 - \rho)$ that

honest sellers will be penalized (this risk being uncorrelated across sellers).

Consumers can call upon the capital of other members of the SRO to supplement a shortfall in the capital of any deviant member or a failure of the SRO to impose a full penalty J. Other members' capital thus provides coinsurance for the risk that inadequate capital is required of new members or insufficient penalties are imposed on existing members.

When will the SRO enforce its own rules? Consider a two-period model where sellers apply for membership in the first period and in the second period either provide services honestly or cheat. If they cheat, then the SRO may impose penalties.

Mayer and Neven's discussion now becomes unclear. They argue that, by imposing a penalty, the club collectively loses J + CR. This seems to imply that J and CR are both given to the cheated consumer(s). They go on, however, to suggest that, if the SRO fails to penalize, then it loses up to ρnCR (given the assumption that it imposes a capital cost of expected value (1 - ρ)CR on honest sellers). However, if honest sellers *are* penalized in this way, why is the penalty not (1 - ρ)(J + CR)? And if all penalties and foregone capital are passed on to consumers anyway, why does the SRO lose anything by failing to penalize?

Anyhow, if we stick with their formulation, the SRO will respond to the incidence of cheating by imposing the appropriate penalty if and only if

(2.5.1) $J + CR \leq \rho nCR$

Looking at the first period, a deviant seller will want to enter as long as

$$\pi_d - \mu - (J + CR) \geq 0$$

which is identical to (SPC²) above. Thus, entry by deviant sellers is deterred and penalties are imposed if and only if

(2.5.2) $\rho n CR > J + CR > \pi_d - \mu$

Thus, we get the following result.

Result 2.5.2: Self-regulation will be easier (require less high CR) the higher are ρ, n and μ and the lower is π_d.

These results seem reasonable and intuitive.

2.5.4 Problems with this model

Many of Mayer and Neven's results are attractive and intuitive. However, as was described in the last section, some of their analysis is confusing (if not confused). In addition, some of their most important assumptions seem unsatisfactory.

Firstly, the assumption that honest sellers' profits π are increasing in quality θ but that π_d is the same for all sellers is questionable. An honest seller of high quality will have high profits because she has high demand, because she can charge a high price or because she has low costs. It seems unlikely that none of these factors will affect the profits she can make from cheating. The only situation in which π_d might be the same for all sellers is that of fraud. If all sellers have hold the same amount of their consumers cash and then simply run off with it, then the gains they make from cheating will be independent of θ. Even in this case, however, it seems likely that high quality sellers will tend to have access to more of their consumers' cash than low quality sellers.

The implication of π_d being the same for all sellers is that low θ sellers face much worse incentive problems than high θ sellers, since they have less to lose from cheating and the same amount to gain. If, however, π_d is increasing in θ, as I believe to be a more realistic assumption, this result no longer necessarily comes out. In this case, although high θ sellers have a lot to lose from cheating, they may also have so much to gain from cheating that they, too, face large incentive problems.

A second problematic assumption is that there is a probability $(1 - \rho)$ that honest sellers are deemed to have misbehaved, in which case they incur the penalty and lose their capital. The main problem with this assumption is not that it is unrealistic (although it is), but that it drives the whole model. Most importantly, if $\rho = 0$, then the incentive compatibility constraint (SIC) becomes

$$\pi(\theta) \geq \pi_{-1}$$

which is false (by assumption).

Mayer and Neven's justification for this assumption is by example. An announcement by regulators that enquiries are underway into the conduct of a seller often adversely affects public perception of their quality, so that sellers' reputations may suffer even when they are eventually cleared of all misdemeanours. This may be so, but a temporary loss of reputation is very different from expulsion from the industry complete with large fine or imprisonment. Moreover, the implication of this model is that any process which reduces the chance of false penalization of honest sellers to zero will actually be bad, because regulation (of the kind analysed) will become impossible. This seems counter-intuitive.

A third problem with this model is that it is not clear exactly what consumers can observe. If consumers were sophisticated, and if they were able to observe each seller's θ and C, then they would be able tell which sellers are in area E of the above diagram (i.e. which sellers will enter and be honest). So there would be no need for any form of regulation. If, however, consumers cannot see sellers' θs, then it seems strange that sellers' profits depend on their qualities (unless the difference in profits is entirely due to differences in costs, rather than differences in output or price).

Thus, the most promising assumption on observability is that consumers can observe θ, but not C. In this case, though, consumers will never purchase from sellers with $\theta < \theta_{-1}$, since all of these sellers will be dishonest.

Indeed they may not purchase from sellers with θ slightly higher θ_{-}, since the likelihood of being cheated will be so high. However, this change in the model means that the trade-off between allowing dishonest sellers to enter and excluding honest sellers will no longer exist.

Further problematic assumptions involve the distribution of C and how it will affect the optimal choice of CR. It is assumed that C is distributed uniformly on $(0, \bar{C}(\theta))$ for each θ, and that $\bar{C}'(\theta) \geq 0$. This assumption is questionable, in that we might well expect the lower bound on C to be increasing in θ as well. However, suppose that this assumption *is* realistic. No assumption is made on the distribution of sellers over different θs, but suppose that it is also uniform. Combining these two assumptions, we find that there is a greater density of sellers in any given (C', C") interval at low θ than at high θ (since the same number of sellers are distributed uniformly over a smaller interval).

Thus, even if we do attribute equal weight to cheating and inappropriate exclusion of sellers (as is suggested), this doesn't mean setting *ab* and *bc* equal to each other as is argued above. Rather *ab* should be smaller than *bc*, since the density of sellers over *ab* is much higher than it is over *bc*. This will result in CR being higher than in the above results.

Moreover, is it really reasonable to give equal weight to cheating and inappropriate exclusion of sellers? When the costs of fraud are high, it seems that much more weight should be placed on abolishing cheating. More formally, suppose that sellers can appropriate all consumer surplus and that consumers are reimbursed fully for any apparent losses π_{Δ} (even when it turns out that no fraud occurred). In this case, the welfare gain from allowing another honest seller of quality θ to enter will be $\pi(\theta) - \mu$. The loss from allowing a cheating seller to enter, meanwhile, will be $\pi_{\Delta} - J - CR - \mu$.

Consider a marginal increase in CR which causes bc honest sellers (with average profits $\bar{\pi}$) and ab cheating sellers to leave the industry. The total gain to welfare will be

$$ab(\pi_d - J - CR - \mu) - bc(\bar{\pi} - \mu)$$

The first term in this expression will be positive if $\pi_d - J - CR - \mu > 0$, which will be true whenever cheating sellers have an incentive to enter (i.e. for all $CR < C_m$). Welfare will be maximized when ab and bc are such that this expression is zero. Thus, if the costs of fraud (the first term) are large relative to the benefits of honesty (the second term), then bc will be high relative to ab at the welfare optimum (that is, a lot of weight will be placed on abolishing cheating). More generaly, there seems little reason to expect the optimal CR to be that which equates *ab* and *bc*.

Finally, it is not clear to me that an external regulator will regulate properly if given zero incentives. After all, as was mentioned in chapter 1, there is a large literature on how to give external regulators appropriate incentives.

In chapter 4, I consider a model which attempts to overcome some of the above problems. It is still a simplistic model, but it does

— allow π_d to vary with θ;

— let $\rho = 0$;

— allow for proper incentives to be given to external regulators as well as SROs;

— make clear assumptions on the observability of θ and C;

— make more realistic assumptions on the distribution of C; and

— model reputation using an explicitly multi-period set-up.

2.6 Two other models

2.6.1 Shaked and Sutton (1982)

Shaked and Sutton examine a situation in which consumers do not enjoy perfect information, but instead must learn of an SRO's quality through experience, by patronising its members. It is shown that, under such conditions, members of a larger SRO may earn higher incomes than those of a smaller SRO.

The intuition behind their result is as follows. Consumers' expectations of an SRO's quality will converge upon the true quality more quickly as the size of the SRO increases, simply because consumers will have more observations with which to update their expectations in each period. Consumers are assumed to differ in their prior expectations of quality (although the mean prior over the whole population will be accurate). Suppose that those consumers with optimistic expectations of quality will certainly purchase from the SRO, so that the marginal consumer must have a pessimistic expectation of quality. If the price which the SRO can charge is dependent on the quality expectation of this marginal consumer, then price will be higher the closer this consumer's expectations are to the SRO's actual quality. Thus, members of a bigger SRO may receive higher earnings than members of a smaller one.

Shaked and Sutton's analysis is complicated, but the simplified version which follows can get us some way towards their result.

Consider an SRO with n members, each of whose members each supply exactly one unit of a particular good per period. Suppose that this good either works or that it does not work. The actual success rate of members, θ, is exogenous and is a measure of the quality of the SRO.

There are v consumers, where $v > n > v/2$. Consumers are risk neutral. Each consumer forms an expectation $\bar{\theta}$ about θ, and is willing to pay any price less than or equal to $p(\bar{\theta})$ for the good, where the $p(\bar{\theta})$ function is

the same for all consumers and is increasing in $\bar{\theta}$. There can be no price discrimination, however, so that all consumers must in fact pay the same price, even if they have differ according to $\bar{\theta}$.

Consumers have finite lifetimes, and an equal number of consumers enter the industry in each period as depart. Each consumer on entering the market has no information as to the profession's quality, but he does have a prior expectation. Since the observations of "success" versus "failure" are distributed according to a Bernouilli distribution, a conjugate prior distribution is assumed, being a Beta distribution with parameters σ, τ, where

$$\text{mean}(\bar{\theta}) = \frac{\sigma}{\sigma + \tau} = \text{prior expectation of } \theta$$

Consumers have different priors, with σ and τ being distributed according to a probability function $g(\sigma, \tau)$. This probability function is such that an equal number of consumers have pessimistic expectations of θ as have optimistic expectations.

Now consider one consumer, with given σ and τ, who has been alive for t periods. In each period, each consumer has some probability α of hearing of the outcome of each of the n units of service provided by the profession in that period. His prior p.d.f is accordingly updated in the light of this information, the updating rule for expected quality $\bar{\theta}$ being

$$\bar{\theta} = \frac{\sigma + \theta r}{\sigma + \tau + r}$$

where θ is the true quality of the SRO; r is the number of units whose outcome the consumer hears ($= \alpha tn$); and so θr is the corresponding number of these that are successes.

Differentiating this term with respect to r gives

$$\frac{d\bar{\theta}}{dr} = \frac{\theta(\sigma + \tau + r) - (\sigma + \theta r)}{(\sigma + \tau + r)^2}$$

$$= \frac{1}{(\sigma + \tau + r)^2}\left(\theta\tau + \sigma(\theta - 1)\right)$$

This is negative if and only if $\theta \leq \sigma/(\sigma + \tau)$. Thus, if initial expectations of θ are too high, $\bar{\theta}$ will be decreasing in r. Conversely, if initial expectations of θ are too low, $\bar{\theta}$ will be increasing in r. Thus, expectations of quality $\bar{\theta}$ will converge more quickly upon the true quality θ as r (and thus, as n) increases.

Without needing to carry out any explicit integration, we see that

Result 2.6.1: Over the whole set of consumers, who differ according to both (σ, τ) and t, the total number of consumers with expectations of quality greater than or equal to some $\bar{\theta}'$ will be increasing in n, as long as $\bar{\theta}'$ is a pessimistic expectation (i.e. $\bar{\theta}' < \theta$).

Now, because consumers are willing to pay a price less than or equal to $p(\bar{\theta})$, and since p is increasing in $\bar{\theta}$, the marginal consumer will have a lower $\bar{\theta}$ than all other consumers. Moreover, since $v > n > v/2$, and total supply is fixed at n, the marginal consumer must be pessimistic.

Let the quality expectation of the marginal consumer be $\bar{\theta}''$. From result 2.6.1, it is clear that an increase in n will have two effects.

(1) It will cause an equal increase in the equilibrium number of consumers in the industry, which will tend to lower $\bar{\theta}''$, for given r; and

(2) It will increase the total number of consumers with expectations of quality greater than or equal to any given pessimistic $\bar{\theta}'$, which will tend to raise $\bar{\theta}''$.

This is as far as we can get with this simple model. Using explicit integration, and placing some restrictions on $g(\sigma, \tau)$, however, Shaked and Sutton prove the following result.

Result 2.6.2: There exists a size of SRO, $n' > 0$, such that its equilibrium price is strictly greater than that attained by any smaller profession.

Thus, members of a larger SRO may earn higher incomes than those of a smaller SRO. Indeed, Shaked and Sutton show that this result can even hold when actual quality θ is *decreasing* in n.

The above result is interesting, in that it may put an upper bound on the number of SROs that may exist in an industry (this topic is further discussed in chapter 5). Indeed, it seems intuitive that the communi-cations problem faced by small SROs could be an important barrier to entry.

Note that in chapter 4, a second argument for having large SROs is given. An SRO will only have incentives to regulate properly if it cares sufficiently about continuing to receivea subscription fee from each of its members. As n increases, the minimum subscription fee that gives proper incentives will fall, which will be better both for the industry and for consumers.

2.6.2 Shapiro (1986)

Shapiro makes the point that much regulation does not involve minimum quality levels L. Rather, regulators either set minimum qualification/ training levels K, or they run certification schemes which inform consumers about the particular training level of each firm. The assumption behind both of these forms of regulation is that human capital and quality are complements. Or, more specifically, that higher levels of training (human capital investment) reduce the marginal cost of providing quality, and so promote the provision of quality service.

Suppose that there is asymmetric information between sellers and consumers of the following form: consumers cannot observe the quality of a seller's product until w periods after the first purchase is made, but after this they know the seller's quality perfectly (because quality is assumed to be irreversible once chosen). Thus, there is no moral hazard problem ((A3) above), but there is an initial hidden irreversible action problem (A2).

Assume that each seller, after a training period, can elect to produce either a low quality (basic) service q_L, at a cost of μ_L, or a high quality (expert) service q_H, at a cost of μ_H. Initially, the seller's quality is unknown. This is the introductory phase. After w periods of the seller's life, however, her reputation is established and she enters the mature phase.

Consumers are heterogeneous in terms of their marginal valuations of quality. They are unable to observe quality during the introductory phase, and nor can they observe training levels, except insofar as they are informed about them by the regulator. When there is no certification, so that all sellers in this phase appear identical to consumers, they will earn a common introductory price p^*. Once their reputations are established, however, high quality sellers will receive p_H and low quality sellers will receive p_L. Clearly $p_H > p^* > p_L$. A consumer will buy from the high quality market, the low quality market, or the introductory market, depending on his valuation of quality.

It is assumed that an increase in training reduces the marginal cost of quality, so that training and quality are complementary. More formally, for given quality q (= L or H), let the cost of production be $\mu_q(k)$, where $\mu_q' > 0$. Because of the complementarity of quality and training, the optimal choice of training for each quality $\hat{k}_q$ will be such that $\hat{k}_L < \hat{k}_H$.

Now, licensing is the requirement that all sellers obtain some minimum level of training K. If $K < \hat{k}_L$, this will have no effect on equilibrium. Consider, therefore, the introduction of a minimum training level K, such that $\hat{k}_L < K < \hat{k}_H$. This will not affect the cost of producing high quality

μ_H. However, it will increase the cost of producing low quality μ_L and will thus reduce the marginal cost of quality ($\mu_H - \mu_L$).

There are two ways in which this affects equilibrium. First, the direct cost effect (increased μ_L) tends to raise p_L and p^*. Second, by reducing the marginal cost of quality, it lowers p_H. Thus, we get the following result.

Result 2.6.3: Licensing raises p_L and p^*, but lowers p_H. So, licensing benefits those consumers who value quality highly, at the expense of those consumers who value quality less. Licensing also causes the low quality market to shrink and the high quality market to expand (as long as each submarket is initially active), so that average quality increases.

Whether or not licensing will increase overall social welfare is ambiguous. Generally, licensing seems to be most attractive when:

(a) w is neither too low nor too high;

(b) the difference between μ_H and μ_L is small, so that efficiency calls for more sellers of q_H; and

(c) consumers' valuations of quality are high.

The intuitions behind (b) and (c) should be immediately obvious. The intuition behind (a) is as follows: When w is low, there is little reputation problem and so there is little need for licensing. When w is high, the reputation problem will be so huge that licensing will not do much to ease it, since licensing itself relies on reputations to be effective. Thus, given the costs it involves (i.e. the excessive training required of low quality sellers), licensing is unlikely to be socially beneficial in either of these cases.

Now consider a certification scheme. Under certification a seller can provide services without meeting any minimum training standard, but

whatever training a seller does choose will be communicated to the public by the regulator. Training levels may thus be used by sellers to signal their quality.

Shapiro shows that, if an equilibrium exists, then it will be separating. That is, high quality sellers will be able to signal their quality to consumers perfectly.

In some cases, this signalling equilibrium (SE) duplicates the perfect information allocation. This will occur on the condition that low quality sellers find it too costly to mimic high quality sellers by overinvesting, which is most likely to hold if

(i) there is good information about quality (w is small), or if

(ii) quality and training are strong complements, so that the extra training substantially raises the costs of providing the low quality service.

When these conditions are not satisfied, then the SE (if one exists at all) requires high quality sellers to invest in training up to the point, $\tilde{K}$, at which low quality sellers would lose money by attempting to mimic them.

In both of these SE cases, there is no introductory phase, so that low quality sellers get p_L throughout their lives, instead of getting p^* (which is greater than p_L) at first. This means that certification eliminates the cross-subsidy to low quality services, which causes p_L to rise and hurts consumers with low quality valuations.

When the SE duplicates the perfect information allocation, certification benefits consumers with high quality valuations and, since levels of training are efficient, aggregate welfare is raised, even though low quality consumers lose out relative to the equilibrium without certification. When the SE results in overinvestment by high quality

sellers, however, certification may be pareto inferior to the *laissez faire*.

In conclusion, then, licensing leads to overinvestment in training by low quality sellers. Certification can lead to overinvestment in training by high quality sellers. Shapiro shows that neither policy necessarily dominates either the other, or the *laissez faire* outcome. However, he shows that licensing will be most effective when consumers value quality highly, and that certification is most effective when training and quality are strong complements in the production of services.

Shapiro does not examine self-regulation, as opposed to external regulation. Moreover, minimum training levels and certification are not examined anywhere else in this thesis. However, Shapiro's analysis gives us an idea of the kind of implications that such policies would have.

NOTATION - chapter 3

β	$p_{\bar{k}} - \mu_{\bar{k}\times}$ i.e. the marginal effect of k on the equilibrium profit margin $p - \mu_{\times}$
d	distance of buyer from seller
$E(\cdot)$	expectations operator
$F(\bar{u})$	distribution function for reserve utility $\bar{u}$
ϕ	per seller fixed entry costs
$G(\tau)$	distribution function for τ
g_i	relative quality, $k_i - \bar{k}$
γ	position of marginal buyer
$H(\phi)$	distribution function for ϕ
k	training/quality
$\bar{k}$	average quality
λ	proportion of consumer with no observability problems
m	number of sellers regulated
$\mu(k, x)$	cost function
n	number of sellers in the industry
$p(\cdot)$	inverse demand function
$\bar{p}$	average price
Π	total SRO profits
π	per seller profits
$\psi(k)$	fixed costs of quality
$\rho(k)\sigma(x)$	unit costs of quality
t	transport costs
τ	per period search costs
$U(\cdot)$	utility function, $u(p, k) - \bar{u}$
W	welfare
x	seller output
X	total industry output
z_i	relative output

3. SELF-REGULATION UNDER HIDDEN IRREVERSIBLE ACTION

3.1 Introduction

Leland's model of minimum quality levels is good at capturing the potential trade-off between raising quality and restricting supply, the conditions under which this trade-off will exist, and the fact that consumers and industry will weight the two components in this trade-off very differently. However, his model is restrictive in that it assumes:

— Sellers come in a variety (or a continuum) of intrinsic qualities, which they cannot change;

— Each seller sells only one unit of her good (if any), so that total supply is fixed for each quality (or quality interval);

— Consumers can only observe average industry quality (rather than each seller's own quality).

In this chapter, I develop a more general (and hopefully more realistic) model, which relaxes these assumptions, while still capturing the more important characteristics of Leland's model, plus some others.

In this model:

— There is endogenous quality and output choice by each seller, and qualities cannot easily be lowered once chosen;

— Sellers are *ex ante* identical, so that the maximum possible supply at each quality level is equal to the maximum possible total industry supply, which is high;

— Consumers can observe average industry quality and can also observe, to at least some extent, each seller's own quality;

— Equilibrium is symmetric: all sellers end up choosing identical quality and output levels.

— Self-regulation is modelled as collusion on quality, but not on output. Two types of self-regulation are considered:

(1) Sellers play a one-stage game, choosing quality and output simultaneously (see sections 3.3 - 3.5); and

(2) Sellers play a two-stage game in which quality is chosen collusively at stage one, with the knowledge that there will be Cournot competition in output at the stage two. This two-stage set-up allows sellers to manipulate quality so as to restrict output and raise profits (see section 3.6).

— Equilibrium in these two set-ups is then compared with:

(a) the *laissez faire* equilibrium (in which sellers choose quality and output simultaneously without any collusion);

(b) the fully collusive (monopoly) equilibrium (where sellers collude on both quality and output); and

(c) the welfare-optimum.

3.2 The model

Let there be a fixed number n sellers in the industry[1], each of which chooses an output, x_i, and a quality, k_i. Total industry output $X = \sum_I x_i$. Average industry quality $\bar{k} = 1/n \sum_i k_i$.

---

[1] It is assumed that n is small enough to enable all n sellers to make non-negative profits at their profit maximizing choices of quality and quantity. Later in the chapter it is shown that allowing perfect free entry leaves the following results qualitatively unchanged.

Let the costs of each seller i be given by $\mu_i = \mu(k_i, x_i)$. Assume $\mu_x \geq 0$, $\mu_{xk} \geq 0$, $\mu_x > 0$, $\mu_{xx} \geq 0$, and $\mu(0,0) = 0$. Note that the form of $\mu(k_i, x_i)$ is very general in that it makes no assumptions about the importance of the fixed costs of quality relative to the unit costs of quality.

The most important part of this model is the inverse price function which is used. Let g_i be a measure of seller i's *relative* quality, which is defined as

$$g_i \equiv k_i - \bar{k}$$

and let z_i be a measure of seller i's relative output, which is defined as

$$z_i \equiv x_i - \frac{\sum_{j \neq i} x_j}{n-1} \equiv \frac{nx_i - X}{n-1}$$

Each seller i faces an inverse demand function

$$p(g_i, \bar{k}, z_i, X)$$

$[p_g > 0,\ p_{\bar{k}} > 0,\ p_x < p_z < 0,\ p_{gg} < 0,\ p_{\bar{k}\bar{k}} < 0,\ p_{g\bar{k}} < 0,$ and $p_{\bar{k}z} \gtrless p_{\bar{k}x}]$.

The form of this inverse demand function is given some microeconomic justification in the appendix to this chapter. Briefly, the two *absolute* arguments of this function, $\bar{k}$ and X, correspond to income effects. As $\bar{k}$ increases, for given X, the utility of each consumer rises, and so each consumer is willing to pay more. Similarly, in order to expand the total market, X, for given $\bar{k}$, sellers must lower prices. The two *relative* arguments in this function, g_i and z_i, correspond to substitution effects. Consider two otherwise identical sellers: the one which provides the higher quality can charge the higher price, and the one which charges the lower price, for given quality, will get the larger demand.

The reason for using an inverse demand function of this form is that it allows easy analysis of three potential quality externalities which may

result in the *laissez faire* equilibrium being welfare sub-optimal (so that there is a potential role for self-regulation). These three effects are:

(B1) When individual sellers' qualities are only imperfectly observable, a consumer's expectation of quality from seller i will generally be an increasing function of both k_i and $\bar{k}$, lying somewhere between perfect observability (in which there is no role for $\bar{k}$), and zero observability (in which there is no role for k_i, as in Leland). If price is an increasing function of expected quality (as seems reasonable), then an increase in k_i will increase the prices received by all sellers, via its effect on $\bar{k}$. This is a *positive between-seller externality*, and it is captured by $p_{\bar{k}}$. Under *laissez faire*, this positive externality means that sellers have reduced incentives to raise quality, even though an increase in quality would benefit all sellers. This effect was described in chapter 1 as the hidden irreversible action problem (A2), and also as the positive betwen-firm externality (B1).

(B2) When consumers care not only about the absolute quality they receive, but also about how this quality compares to other qualities available (which can be proxied by average quality), then an increase in k_i will, by raising $\bar{k}$, lower the relative quality of, and so the price received by, the other n-1 sellers. This is a *negative between-seller externality* (denoted (B2) in chapter 1), and it is captured by $p_{\bar{k}}$. Under *laissez faire*, this negative externality means that sellers will compete vigorously on quality, which may result in quality being higher than is optimal for the industry. No seller will be willing to reduce quality unilaterally, since it would lose profits to its competitors. But a multilateral quality reduction would improve every seller's profits.

(C1) Quality choice may also be non-optimal because of the Spence effect. When sellers have any market power (that is, they face a downward-sloping demand curve), they only care about the quality valuations of the marginal consumer, whereas social welfare takes into account the

valuations of inframarginal consumers as well. Thus, if quality valuations decrease as the number of consumers increases ($p_{\bar{k}x} < 0$), sellers will tend to provide too little quality for welfare maximization (for given quantity). The converse is true if $p_{\bar{k}x}$ is positive. This is a *seller-consumer externality* (denoted (C1) in chapter 1). Note that this effect is also present in Leland's model.

Lastly, let k be the equilibrium symmetric choice of quality (= $\bar{k}$), and let $\beta(k, x)$ be defined by

$$\beta(k, x) = p_{\bar{k}}(0, k, 0, X) - \mu_{kx}(k, x)$$

So $\beta(k, x)$ is the marginal effect of k on the equilibrium profit margin, $(p - \mu_x)$.

The sign and magnitude of $\beta(k, x)$ will frequently be of importance in what follows. However, all that can be stated categorically about β is that, for given k and x, $\beta(k, x)$ will be higher (i) as price becomes more sensitive to quality and (ii) as the ratio of fixed costs of quality to unit costs of quality rises.[2]

The aim in this chapter is to use this model to make comparisons of equilibria under different regimes, so as to find some conditions under which self-regulation might be expected to be welfare-superior to *laissez*

---

[2] The dependence of $\beta(k, x)$ on the ratio of fixed to unit costs of quality can be proved informally by letting $\mu(k_i, x_i)$ have the form

$$\mu(k_i, x_i) = \lambda\psi(k_i) + (1 - \lambda)\rho(k_i)\sigma(x_i)$$

which gives $\lambda/1-\lambda$ as this ratio. $\lambda/1-\lambda$ is clearly increasing in λ. Thus to find how β varies with this ratio, we just need to find $d\beta/d\lambda$,

$$\frac{d\beta}{d\lambda} = \rho'(k)\sigma'(x)$$

which is positive, given the assumptions made about the cost function.

faire. Various assumptions are made about the relative signs and values of

— $p_{ke} - p_{ss}$

— p_{kex}

— $\beta(k, x)$

The first of these three terms captures the *net* between-seller externality (the difference between the positive and the negative between-seller externalities), and the second captures the Spence effect. The third term is important because it shows how total industry output changes with average industry quality, which captures the output restriction effect (denoted (D1) in chapter 1).

In what follows, it is shown that welfare at equilibrium in the one-stage self-regulation game is most likely to be an improvement on the *laissez faire* equilibrium when

$$p_{ke} - p_{ss} \geq n - p_{kex} \geq n \beta(k, x)$$

And welfare at equilibrium in the two-stage self-regulation game is most likely to be an improvement on the *laissez faire* equilibrium when

$$p_{ke} - p_{ss} \geq n - p_{kex} \geq n - \beta(k, x)$$

The intuitions for these results are given later, at the same time as they are derived.

In section 3.7, it is shown that these results also hold when the number of sellers in the industry is endogenous to the model.

Note that, in those situations where self-regulation is bad for welfare, the government would be justified in banning it. When self-regulation is welfare-superior to no regulation, however, it is not necessarily the case that self-regulation should be encouraged, since it must also be compared with external regulation.

3.3 The one-stage self-regulation game

The profits received by seller i are given by

$$\pi(g_i, \bar{k}, z_i, X) = p(g_i, \bar{k}, z_i, X).x_i - \mu(k_i, x_i)$$

Let us consider three types of market structure.

Fully collusive oligopoly (m): Suppose the n sellers act as a monopoly, colluding on both quality and output choice. If all firms deliberately act identically, then they know that $g_i = 0$ and $X = nx_i$, for all i, even *ex ante*. Profits are maximized by taking this into account, giving the following first order profit maximization conditions:

$$\pi^m_x = np_x x + p - \mu_x = 0$$

and $$\pi^m_k = p_k x - \mu_k = 0$$

If k and x denote symmetric quality and output choices by sellers, these conditions give two loci in (k, x) space: $\pi^m_x = 0$ gives $x^m(k)$ and $\pi^m_k = 0$ gives $k^m(x)$.

***Laissez faire* (c):** Suppose the n sellers compete in both quality and quantity. In this case the first order profit maximization conditions will be

$$\pi^c_x = (p_z + p_x)x + p - \mu_x = 0$$

and $$\pi^c_k = 1/n((n - 1)p_z + p_k)x - \mu_k = 0$$

These conditions give the two loci $x^c(k)$ and $k^c(x)$ in (k, x) space.

Self-regulating oligopoly (r): Suppose the n sellers compete in quantity but are able to collude on quality (although they choose both

simultaneously). In this case the first order profit maximization conditions will be

$$\pi^r_x = (p_x + p_x)x + p - \mu_x = 0$$

and $$\pi^r_k = p_k x - \mu_k = 0$$

These conditions give the two loci $x^r(k)$ and $k^r(x)$ in (k, x) space. Note that, because $\pi^r_x = \pi^c_x$, $x^r(k) = x^c(k)$. And because $\pi^r_k = \pi^m_k$, $k^r(x) = k^m(x)$. Note that, for any given k and x, $(p_x + p_x)$ must exceed np_x (for $n \geq 2$). So $x^m(k) < x^c(k)$, for all k.

Lastly, let us consider the welfare maximizing choices of k and x (given that all sellers will have identical k and x at equilibrium). Welfare *per seller* W is given by

$$W = 1/n \int_0^X p(0, k, 0, X')\, dX' - \mu(k, x)$$

Thus, for welfare maximization,

$$W_x = p - \mu_x = 0$$

and $$W_k = 1/n \int_0^X p_k(0, k, 0, X')\, dX' - \mu_k = 0$$

These conditions give the two loci $x^*(k)$ and $k^*(x)$.

3.4 Quality comparisons, for given output

Consider first the cases of fully collusive oligopoly and self-regulating oligopoly. For given output, k^m will equal k^r and, at the profit maximizing choice of quality,

$$W_k\Big|_{\pi_k=0} = 1/n \int_0^X p_{\bar{k}}(0, k, 0, X')\, dX' - p_{\bar{k}}x$$

This is positive (so that $k^* > k^r$) if and only if

$$1/X \int_0^X p_{\bar{k}}(0, k, 0, X')\, dX' > p_{\bar{k}}$$

That is, $k^* > k^r$ if and only if the average quality valuation over all consumers is greater than the quality valuation of the marginal consumer. Assuming that the sign of $p_{\bar{k}x}$ (whether positive or negative) is the same for all $\bar{k}$ and x, it is clear that this condition will hold if and only if $p_{\bar{k}x}$ is negative. The reverse inequality holds when $p_{\bar{k}x}$ is positive. Thus, we get the following result.

Result 3.4.1: For given output, $k^m = k^r$ and $k^r - k^* \gtrless 0$ as $p_{\bar{k}x}$.

This result is due to the Spence effect described above. When $p_{\bar{k}x} < 0$, inframarginal consumers value quality more than the marginal consumer does. Since profit-maximization takes into account only the quality valuation of the marginal consumer, while welfare maximization takes into account all consumers, it is intuitive that k^m, and k^r, will be less than k^*. Conversely, when $p_{\bar{k}x} < 0$, $k^m = k^r < k^*$.

Now let us compare the qualities k^r (= k^m) and k^c, for given output. By straightforward algebra, we get

$$\pi^r{}_k\Big|_{\pi^c{}_k=0} = x(p_{\bar{k}} - p_{\bar{g}} - 1/n(p_{\bar{k}} - p_{\bar{g}}))$$

Rearranging this gives

Result 3.4.2: For given output, $k^m = k^r$ and $k^r - k^c \gtrless 0$ as $p_{\bar{k}} - p_{\bar{g}}$.

That is, $k^r > k^c$ when the *net* between-seller externalities are positive,

and $k^r < k^c$, when the net between-seller externalities are negative. This is intuitive. In the *laissez faire* set-up, positive net between-seller externalities mean that sellers have insufficient incentives to raise quality up to its industry profit-maximizing level. Self-regulation, which allows the internalization of these externalities, results in an increase in quality, so that $k^r > k^c$. The converse is true when $p_E < p_S$.

Let us define 3 cases.

Case A: $p_E > p_S$, so that $k^r(x) > k^c{}_A(x)$, for any given x.

Case B: $p_E = p_S$, so that $k^r(x) = k^c{}_B(x)$, for any given x.

Case C: $p_E < p_S$, so that $k^r(x) < k^c{}_C(x)$, for any given x.

Modified versions of result 3.4.1 are applicable in each of these cases:

Result 3.4.3: For given output, we get the following inequalities:

In Case A: $p_{Ex} < 0 \Rightarrow k^c{}_A < k^*$ and $k^c{}_A > k^* \Rightarrow p_{Ex} > 0$.

In Case B: $p_{Ex} < 0 \Leftrightarrow k^c{}_B < k^*$ and $k^c{}_B > k^* \Leftrightarrow p_{Ex} > 0$.

In Case C: $p_{Ex} > 0 \Rightarrow k^c{}_C > k^*$ and $k^c{}_C < k^* \Rightarrow p_{Ex} < 0$.

3.5 Quality comparisons when output is allowed to vary

Let $k^m(x^m)$, $k^r(x^r)$, $k^c(x^c)$, and $k^*(x^*)$, be the equilibrium quality levels under full collusion, self-regulation, *laissez faire* and welfare-maximization, respectively. And let $x^m(k^m)$, $x^r(k^r)$, $x^c(k^c)$ and $x^*(k^*)$, be the corresponding equilibrium output levels.

From the first order π- and W-maximization conditions, with respect to output, it is clear that, for given k, $x^m < x^r = x^c < x^*$. Thus, when comparing $k^m(x^m)$, $k^r(x^r)$ or $k^c(x^c)$ with $k^*(x^*)$, we need to know not only how $k^m(x)$, $k^r(x)$, $k^c(x)$ and $k^*(x)$ compare for given x (as described in section 3.4), but also how these qualities change as outputs vary, and how outputs change as qualities vary. That is, we need to find dx(k)/dk and dk(x)/dx, for each of these market structures.

Total differentiation of the first-order π- and W-maximization conditions with respect to *output*, gives:

$$\frac{dx^c(k)}{dk} = \frac{dx^r(k)}{dk} \gtrless (p_{Rx} + p_{Rz})x + \beta \; ; \qquad \frac{dx^m(k)}{dk} \gtrless p_{Rx}X + \beta;$$

and $$\frac{dx^*(k)}{dk} \gtrless \beta$$

Total differentiation of the first-order π- and W-maximization conditions with respect to *quality*, gives:

$$\frac{dk^r(x)}{dx} = \frac{dk^m(x)}{dx} \gtrless p_{Rx}X + \beta \; ; \qquad \frac{dk^*(x)}{dx} \gtrless \beta;$$

and $$\frac{dk^c(x)}{dx} \gtrless p_{Rx}X + \beta + \frac{(n-1)}{n}.[(p_g - p_R) + X(p_{gx} - p_{Rx})]$$

In addition to these slopes, we also know that

$$x^m(k) < x^r(k) = x^c(k) < x^*(k), \text{ for any given } k;$$

$$k^c{}_C(x) > k^c{}_B(x) = k^r(x) = k^m(x) > k^c{}_A(x), \text{ for any given } x;$$

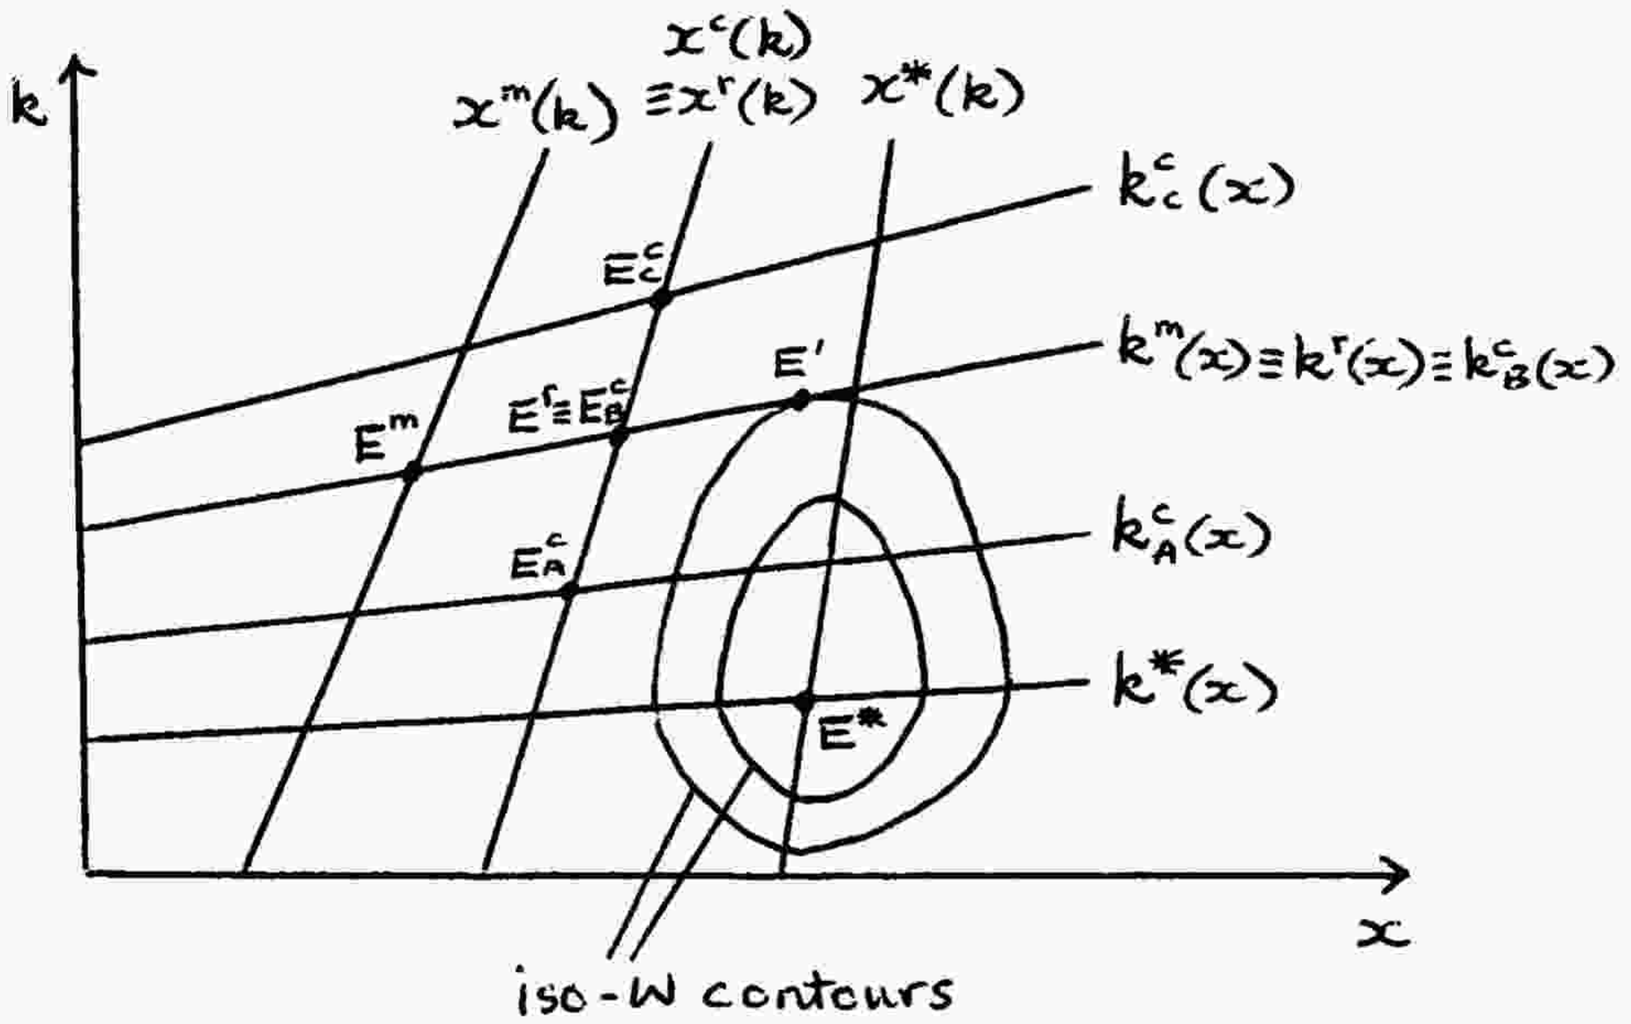
and $$k^m(x) - k^*(x) \gtrless p_{Rx}, \text{ for any given } x;$$

where the subscripts A, B and C relate to the sign of $p_g - p_R$, as described in section 3.4.

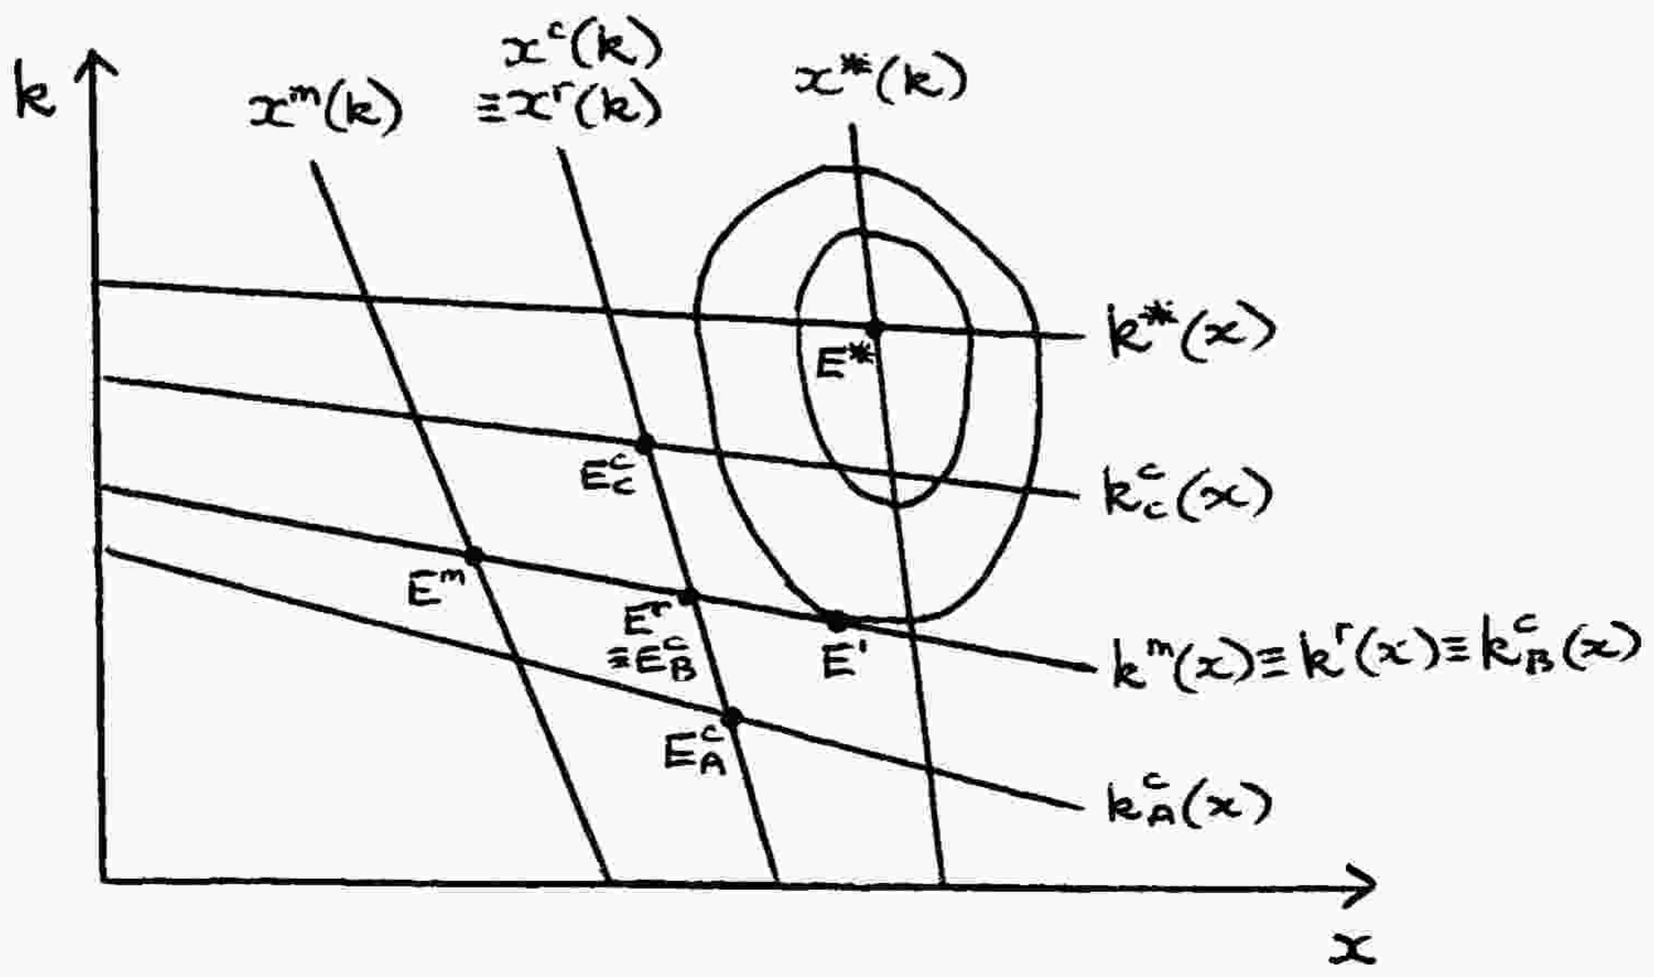
Before making any comparisons between different equilibria, it is useful to draw the various k(x) and x(k) loci in (k, x) space for various permutations of p_{Rx} and β, using the above information about their slopes and relative positions.

In these diagrams, it is assumed, for simplicity, that when $p_{Ex} > 0$, then $k^c_A(x) > k^*(x)$ for all x, and that when $p_{Ex} < 0$, then $k^c_C(x) < k^*(x)$ for all x. It is also assumed that the slopes (and, especially, the *signs* of the slopes) of these various k(x) and x(k) curves do not vary with x or k (although the way in which these slopes *do* vary with x and k is clearly very important). E^m, E^r, E^c_A, E^c_B, E^c_C and E^* are the various equilibria, and E' shows the point at which welfare is maximized *given that* there is collusion on quality.

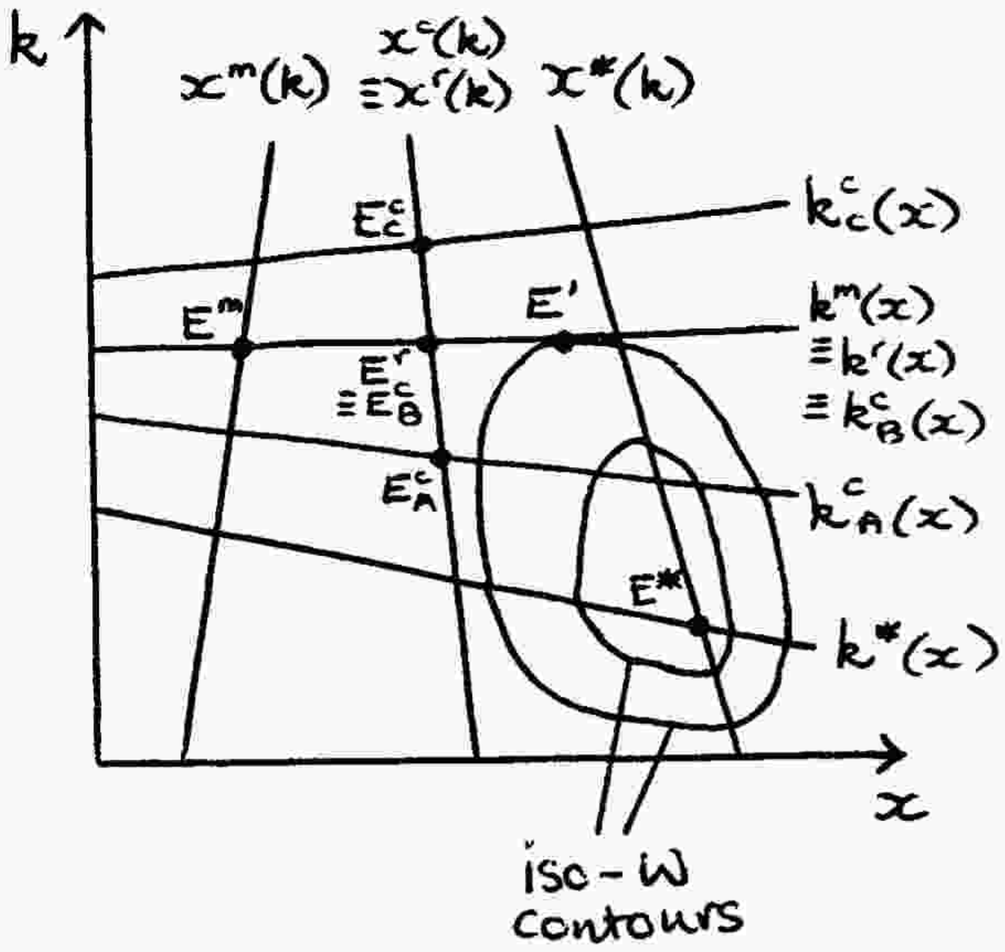
$p_{Ex} > 0$ and $\beta \geq 0$



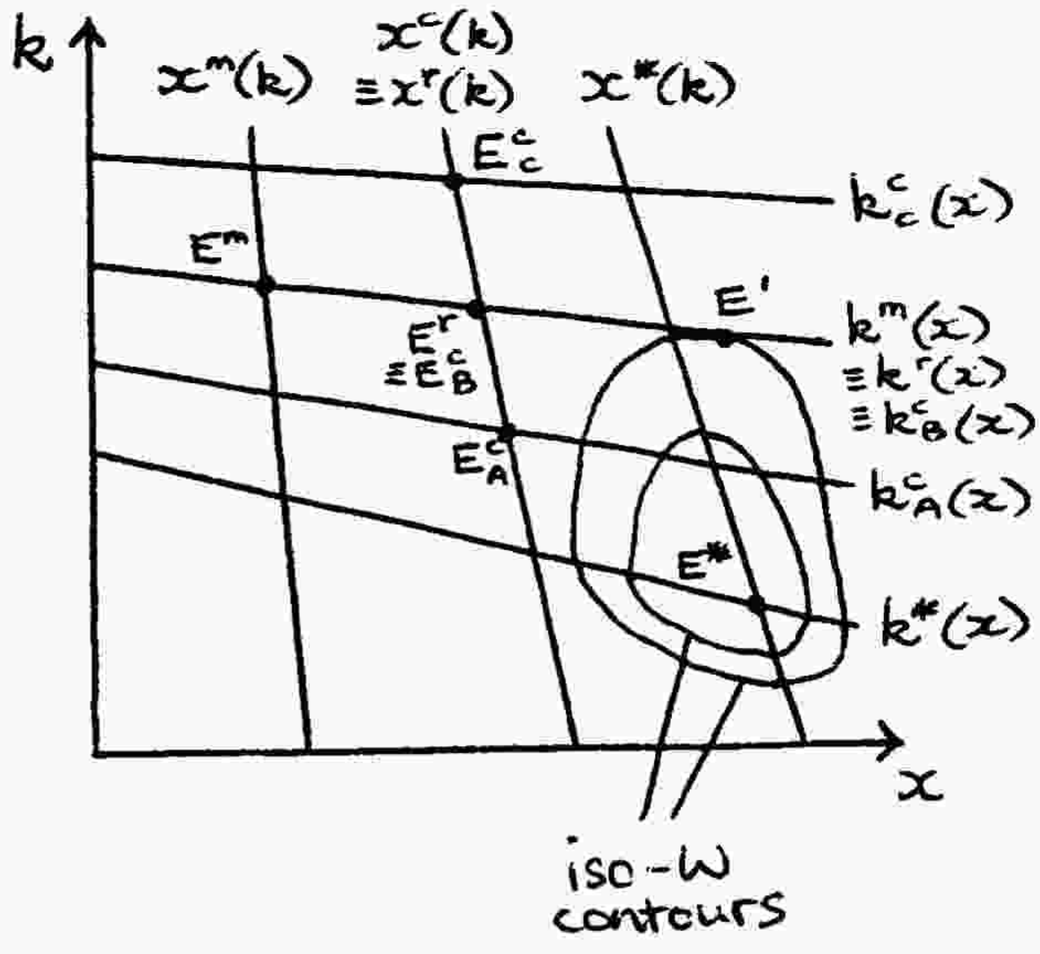
$p_{Ex} < 0$ and $\beta \leq 0$



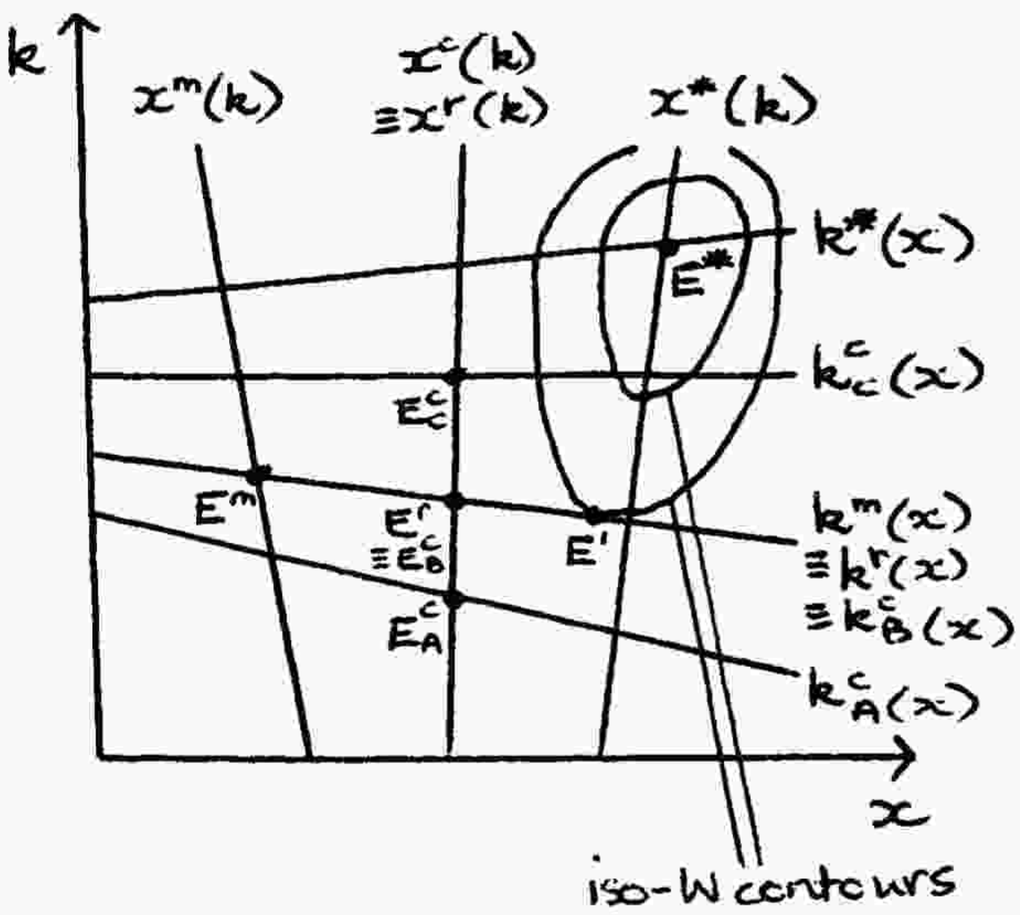
$p_{\bar{k}x} \gtrless p_{\bar{k}x}X + \beta \gtrless -\beta \quad > 0$



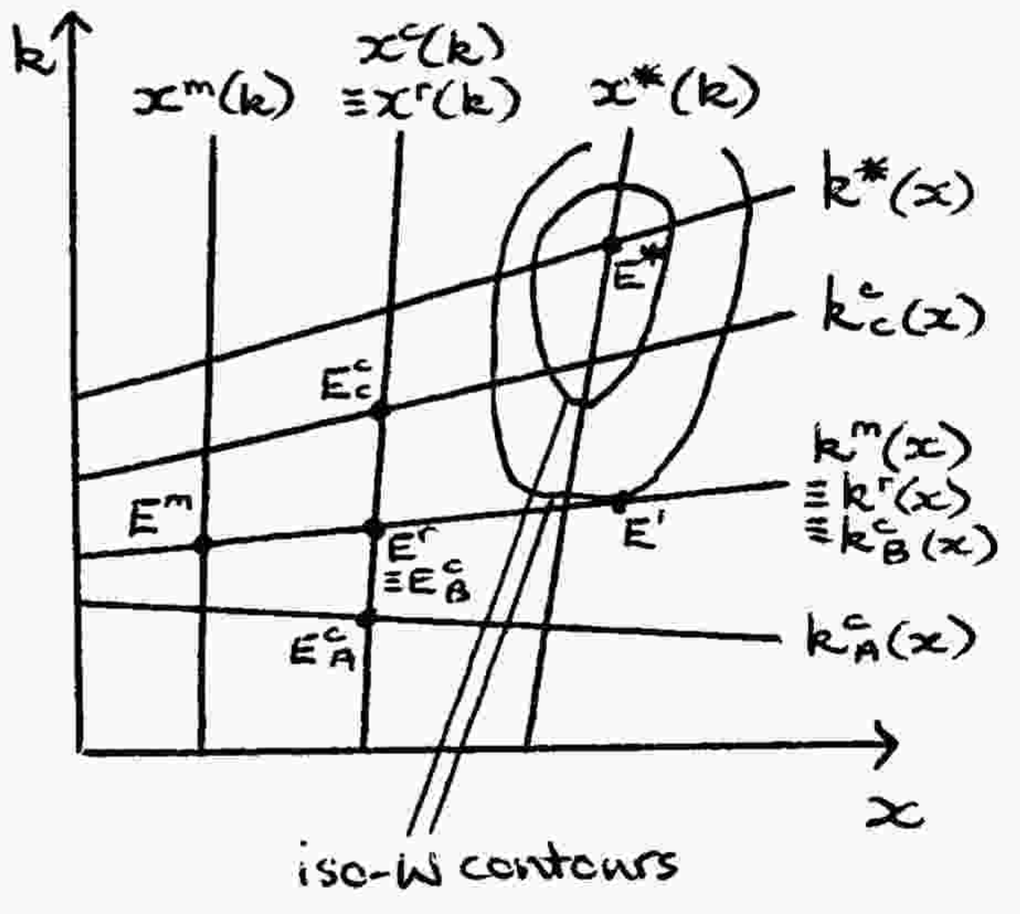
$p_{\bar{k}x} \gtrless -(p_{\bar{k}x}X + \beta) \gtrless -\beta \quad > 0$



$p_{\bar{k}x} \gtrless p_{\bar{k}x}X + \beta \gtrless -\beta \quad < 0$



$p_{\bar{k}x} \gtrless -(p_{\bar{k}x}X + \beta) \gtrless -\beta \quad < 0$



3.5.1 Comparing equilibria under quality collusion

From the fact that $dk^r(x)/dx = dk^m(x)/dx \gtrless p_{\bar{k}x}X + \beta$, we can derive the following result.

Result 3.5.1: $k^r(x)$ and $k^m(x)$ will become closer to $k^*(x^*)$ as x falls from $x^*(k^*)$ if and only if $p_{\bar{k}x}$ and $(p_{\bar{k}x}X + \beta)$ are of the same sign.

Proof: Suppose that $p_{\bar{x}x}$ is positive. From proposition 3.4.1, we know that $k^m(x) = k^r(x) > k^*(x)$, when $x = x^*(k^*)$. If $(p_{\bar{x}x}X + \beta)$ is also positive, then $dk^m(x)/dx$ and $dk^r(x)/dx$ must be positive. Thus, a reduction in x from $x^*(k^*)$ will result in a reduction of $k^m(x)$ and $k^r(x)$, which must bring them both closer to $k^*(x^*)$. The converse argument holds when $p_{\bar{x}x}$ and $(p_{\bar{x}x}X + \beta)$ are both negative. □

This result states that, as x contracts from its socially optimal level $x^*(k^*)$, the quality chosen under monopoly and self-regulation will become closer to the globally welfare-optimal quality $k^*(x^*)$ if and only if $p_{\bar{x}x}$ and $(p_{\bar{x}x}X + \beta)$ are of the same sign. Thus, under these conditions, it is possible that the contraction in x that occurs under monopoly and self-regulation (to $x^m(k^m)$ and $x^r(k^r)$) will bring quality levels nearer to their globally welfare-optimal level. Note, however, that the more x contracts, the more likely it becomes that the sign of $(k^r(x) - k^*(x^*))$ will change, and from this point any further contraction of x will move $k^r(x)$ away from $k^*(x^*)$. This result is clear in the above diagrams.

Now, profit maximization under both self-regulation and full collusion involves setting $p_{\bar{x}} = \mu_k/x$. Thus,

$$\beta(k, x) = \mu_k/x - \mu_{kx}$$

So, whenever there is quality collusion, $\beta(k, x)$ depends only on the ratio of fixed to unit costs of quality, and not on $p_{\bar{x}}$. If all costs of quality are fixed (e.g. $\mu(k_i, x_i) = \psi(k_i)$), then $\beta(k, x)$ is definitely positive, since $\mu_{kx} = 0$. If all costs of quality are unit costs (e.g. $\mu(k_i, x_i) = \rho(k_i)\sigma(x_i)$), then $\beta(k, x)$ will generally be nonpositive, given that $\mu_{xx} > 0$.

This point can be used to explain the following (which is really just a corollary to result 3.5.1).

Result 3.5.2: (i) If $p_{\bar{x}x}$ $\gtrless$ β, then a reduction in x from $x^*(k^*)$ will (at least at first) bring $k^r(x)$ and $k^m(x)$ closer to $k^*(x^*)$

(ii) If $p_{\bar{k}x} \lessgtr_{sgn} -\beta$, then as $|\beta| \to 0$ (for given k and x), it becomes more and more likely that a reduction in x from $x^*(k^*)$ will (at least at first) bring $k^r(x)$ and $k^m(x)$ nearer to $k^*(x^*)$.

The intuition behind this result is as follows. Suppose $p_{\bar{k}x} > 0$, so that $k^r(x) = k^m(x) > k^*(x)$, for given x. As β increases (becomes less negative ($|\beta| \to 0$), or more positive, for given k and x), the fixed costs of quality become more important relative to the unit costs of quality. As x falls, the fixed costs of quality are spread over fewer units. Thus, as the proportion of fixed costs increases, the rise in the overall (per unit) costs of quality, resulting from a reduction in x, will be increased, which means that the incentive to improve quality is reduced, and there is an increased tendency for k(x) to fall towards $k^*(x^*)$. The converse of this intuition holds when $p_{\bar{k}x} < 0$. This result is clear in the above diagrams.

We can go further this. The fact that $k^m(x)$ and $k^r(x)$ get nearer to $k^*(x^*)$, as x falls from $x^*(k^*)$, doesn't necessarily imply that welfare increases, since the beneficial effects of the changes in k^m and k^r must be weighed up against the detrimental effects of the reduction in x. Remember that $k^m(x) = k^r(x)$, for all x, and let $x^*(k^r(x^*))$ denotes the value of x at the intersection of $x^*(k)$ and $k^r(x)$[3]. We can then derive the following.

Result 3.5.3: A reduction in x from $x^*(k^r(x^*))$, under self-regulation or full collusion, will (at least at first) increase welfare if and only if $p_{\bar{k}x}$ and $(p_{\bar{k}x}X + \beta)$ are of the same sign.

Proof: To prove this, we need to find the derivative of $W(k^r(x), x)$ with respect to x (which shows how W changes as we move along the $k^r(x)$ locus). This is

$$\frac{dW(k^r(x), x)}{dx} = \frac{\partial W}{\partial x} + \frac{\partial W}{\partial k}\cdot\frac{dk^r(x)}{dx}$$

---

[3] Note that $x^*(k^r(x^*)) > x^*(k^*)$, when $p_{\bar{k}x}$ and β are of the same sign (or β is zero), because $(k^r(x) - k^*(x)) \lessgtr_{sgn} p_{\bar{k}x} \lessgtr_{sgn} \beta \lessgtr_{sgn} dx^*(k)/dk$.

This must then be evaluated at $x^*(k^r(x^*))$: At $x^*(k^r(x^*))$, $\partial W/\partial x = 0$ and, from result 3.4.1, $\partial W/\partial k \gtrless p_{\bar{x}x}$, so

$$\left.\frac{dW(k^r(x),x)}{dx}\right|_{x^*(k^r(x^*))} \gtrless p_{\bar{x}x}\cdot(p_{\bar{x}x}X + \beta)$$

When $p_{\bar{x}x}$ and $(p_{\bar{x}x}X + \beta)$ are of the same sign, $dW(k^r(x),x)/dx$ must be negative. Thus, as x falls from $x^*(k^r(x^*))$, W increases (at least at first). □

The intuition for this result is as follows. As x falls from $x^*(k^r(x^*))$, the beneficial effects of the change in k^r on welfare are first order, while the detrimental effects of the reduction in x on welfare are only second order.

We also know that $x^*(k^r(x^*)) < x^*(k^*)$ if and only if $p_{\bar{x}x} \gtrless -\beta$, because $(k^r(x) - k^*(x)) \gtrless p_{\bar{x}x}$ and $dx^*(k)/dk \gtrless \beta$, Thus we get.

Result 3.5.4: When $p_{\bar{x}x} \gtrless (p_{\bar{x}x}X + \beta) \gtrless -\beta$, a reduction in x from $x^*(k^*)$, under self-regulation or full collusion, will (at least at first) increase welfare.

Again, the truth of results 3.5.3 and 3.5.4 is clear in the above diagrams.

Lastly, note that the first order profit maximizing conditions (with respect to x) imply that $x^m(k) < x^r(k) < x^*(k)$, for given k. Thus, comparing quality at $x^*(k^*)$ with that at equilibrium, the overall change in k^m, $|k^m(x^m) - k^m(x^*(k^*))|$, will always exceed the overall change in k^r, $|k^r(x^r) - k^r(x^*(k^*))|$. More generally, when there is total collusion on quality choice, this overall change in quality, $|k(x) - k(x^*(k^*))|$, will be increasing in the degree of output collusion. This result is helpful in what follows.

3.5.2 Comparing self-regulation with *laissez faire*

The main aim of this chapter is to compare self-regulation with the *laissez faire*, and to find conditions under which we might expect self-regulation (collusion on quality) to be welfare-optimal to the *laissez faire* (so that self-regulation is a feasible form of regulation).

To do this, we need to compare $k^c(x^c)$, $k^r(x^r)$ and $k^*(x^*)$. The first result we get is:

Result 3.5.5: As long as

$$k^r(x^r) - k^*(x^*) \gtreqless k^r(x^*) - k^*(x^*) \quad (\gtreqless p_{ex}),$$

quality under self-regulation $k^r(x^r)$ will be closer to $k^*(x^*)$ than will be quality at the *laissez faire* equlibrium $k^c(x^c)$, if and only if

$$p_e - p_s \gtreqless - p_{ex}.$$

Proof: Suppose that $p_{ex} > 0$, so that $k^r(x^*) > k^*(x^*)$. When there are *net* negative externalities ($p_e < p_s$), self-regulation (which involves internalizing these externalities) will lower quality. Thus, as long as $k^r(x^r) > k^*(x^*)$ (i.e. has same sign as $k^r(x^*) - k^*(x^*)$), we know that

$$k^c(x^c) > k^r(x^r) > k^*(x^*).$$

That is, self-regulation will bring quality closer to $k^*(x^*)$. The argument for the case where $p_{ex} < 0$ is the converse of this. □

Again, we can go further than this. Examining the effects of self-regulation on total welfare, we get:

Result 3.5.6: When $p_K - p_\sigma \gtrless p_{Kx}$, then as long as

$$k^r(x^r) - k^*(x^*) \gtrless k^r(x^*) - k^*(x^*) \quad (\gtrless p_{Kx}),$$

self-regulation will become more likely to be welfare superior to the *laissez faire* equilibrium as

(i) β falls for $p_{Kx} > 0$; and

(ii) β increases for $p_{Kx} < 0$.

Moreover, under these conditions, self-regulation will definitely be welfare-superior to *laissez faire* when $p_{Kx}x + \beta \gtrless p_{Kx}$.

Proof: First, we need to find the derivative of $W(k, x^r(k))$ with respect to k (which shows how W changes as we move along the $x^r(k)$ $(= x^c(k))$ locus). This is

$$\frac{dW(k, x^r(k))}{dk} = \frac{\partial W}{\partial k} + \frac{\partial W}{\partial x} \cdot \frac{dx^r(k)}{dk}$$

Evaluating this at $k^r(x^r)$ can be ambiguous since, while $\partial W/\partial x$ is definitely positive; $\partial W/\partial k \gtrless p_{Kx}$; and $dx^r(k)/dk \gtrless ((p_{Kz} + p_{Kx})x + \beta)$.

Now consider the case where $p_{Kx} > 0$, $p_K < p_\sigma$, and $k^r(x^r) > k^*(x^*)$, so that

$$k^c(x^c) > k^r(x^r) > k^*(x^*).$$

Welfare will be higher under self-regulation than under *laissez faire*, if $dW(k, x^r(k))/dk$ is negative at $k^r(x^r)$. This must be the case when

$$((p_{Kz} + p_{Kx})x + \beta) \gtrless p_{Kx} \quad (< 0).$$

More generally, $dW(k, x^r(k))/dk$ will be more likely to be negative, in this set-up, as β falls (becomes less positive/more negative). The converse argument proves the proposition when $p_{Kx} < 0$. □

To see the intuition behind this result, consider again the case where $p_{Rx} > 0$, $p_R < p_S$ and $k^r(x^r) > k^*(x^*)$. When the slope of $dx^r(k)/dk$ is positive then, although self-regulation brings quality closer to $k^*(x^*)$ (relative to $k^c(x^c)$), it also moves output away from $x^*(k^*)$ (relative to $x^c(k^c)$). Whether or not there is an overall welfare gain from the introduction of self-regulation into the *laissez faire* set-up depends on the trade off between these two effects.

Now, as β falls, the slope of $dx^r(k)/dk$ becomes less positive/more negative. So the price paid for improved quality (k closer to $k^*(x^*)$), in terms of reduced x, is lowered. Indeed, when $(p_{Rx} + p_{Rx})x + \beta < 0$, dx^r/dk^r is negative and there is no trade off at all. Not only is $k^c(x^c) > k^r(x^r) > k^*(x^*)$, but also $x^c(k^c) > x^r(k^r) > x^*(k^*)$. Thus, self-regulation is definitely welfare-superior to *laissez faire* under these conditions. The intuition for the case where $p_{Rx} < 0$ is the converse of this.

More generally, we can simplify the above analysis to give the following result.

Result 3.5.7: Equilibrium in the one-stage self-regulation game is most likely to be a welfare-improvement on the *laissez faire* equilibrium when

$$p_R - p_S \text{ is } \bar{>} \text{ or } \bar{=} \text{ or } \bar{<} \; p_{Rx} \text{ is } \bar{>} \text{ or } \bar{=} \text{ or } \bar{<} \; \beta$$

3.6 The two-stage self-regulation game

It is perhaps more realistic to consider two effects of introducing self-regulation into the *laissez faire* set-up. The internalization of externalities effect was analysed above. In addition, however, there may be an output manipulation effect, whereby sellers use quality as an instrument for achieving a degree of implicit output collusion, even though explicit output collusion is impossible.

Analysis of this second effect involves looking at a two-stage game. Knowing that there will be Cournot competition in outputs in the second

stage of the game, sellers will attempt to manipulate the outcome of this second stage game via the collusive quality choice that is made in the first stage.

Given that there will be Cournot competition in the second stage of the game (so that equilibrium must lie along the $x^r(k)/x^c(k)$ locus), the first stage profit maximization condition is

$$\frac{d\pi(k, x^r(k))}{dk} = \frac{\partial\pi}{\partial k} + \frac{\partial\pi}{\partial x}\cdot\frac{dx^r(k)}{\partial k}$$

To find out how the two-stage game equilibrium differs from that of the one-stage game, we need to evaluate this expression at E^r_1, the self-regulation equilibrium in the one-stage game. At this point, $\partial\pi/\partial k$ must be zero; $\partial\pi/\partial x$ must be negative; and (from the last section) $dx^r(k)/dk$ [illegible] $(p_{kx} + p_{kx})x + \beta$.

Therefore, at the quality which maximizes profits in the one stage game, k^r_1,

$$\left.\frac{d\pi}{dk}\right|_{k^r_1} \text{[illegible]} - ((p_{kx} + p_{kx})x + \beta) \text{[illegible]} - \frac{dx^r(k)}{dk}$$

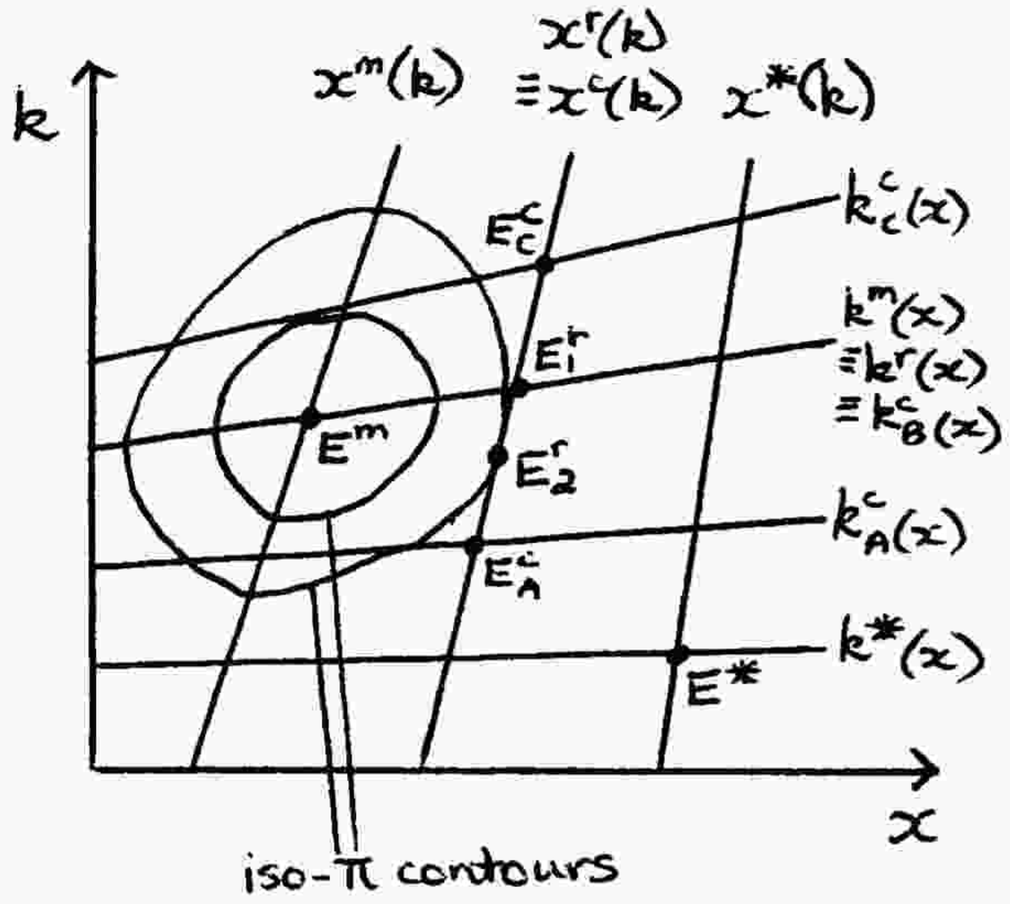
This gives us

Result 3.6.1: For a self-regulating oligopoly, quality in the one-stage game k^r_1 is greater than quality in the two-stage game k^r_2 if and only if $dx^r(k)/dk > 0$ (i.e. if and only if $(p_{kx} + p_{kx})x + \beta > 0$).

This result is easy to see in the figures below. The industry's isoprofit curves, which are centred on the point E^m, are horizontal along the $x^m(k)$ locus and vertical along the $k^m(x)$ locus. Given that there is to be Cournot competition in stage two of the game, the two-stage equilibrium must lie on the $x^r(k)$ locus. Thus profits will be maximized when this

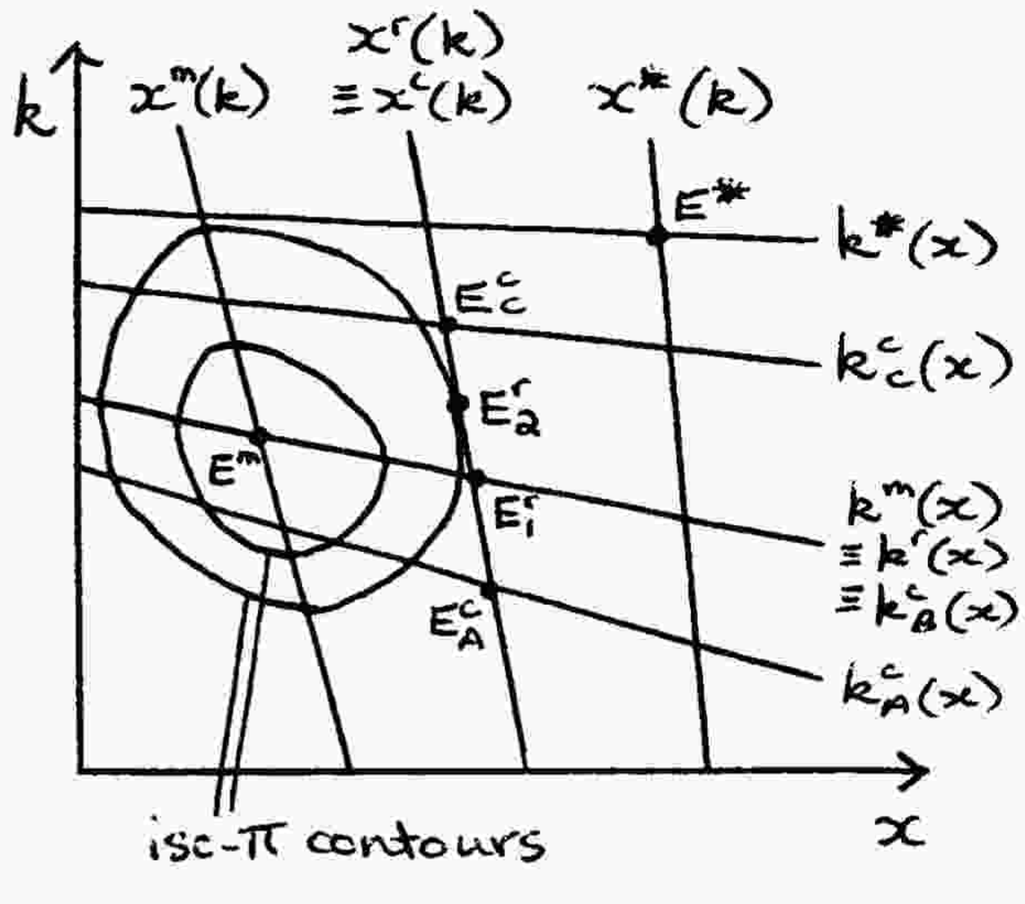
locus is just tangential to an isoprofit curve, which implies that $k^r_1 > k^r_2$ if and only if $dx^r(k)/dk > 0$.

$p_{\bar{k}x} > 0$ and $\beta \geq 0$

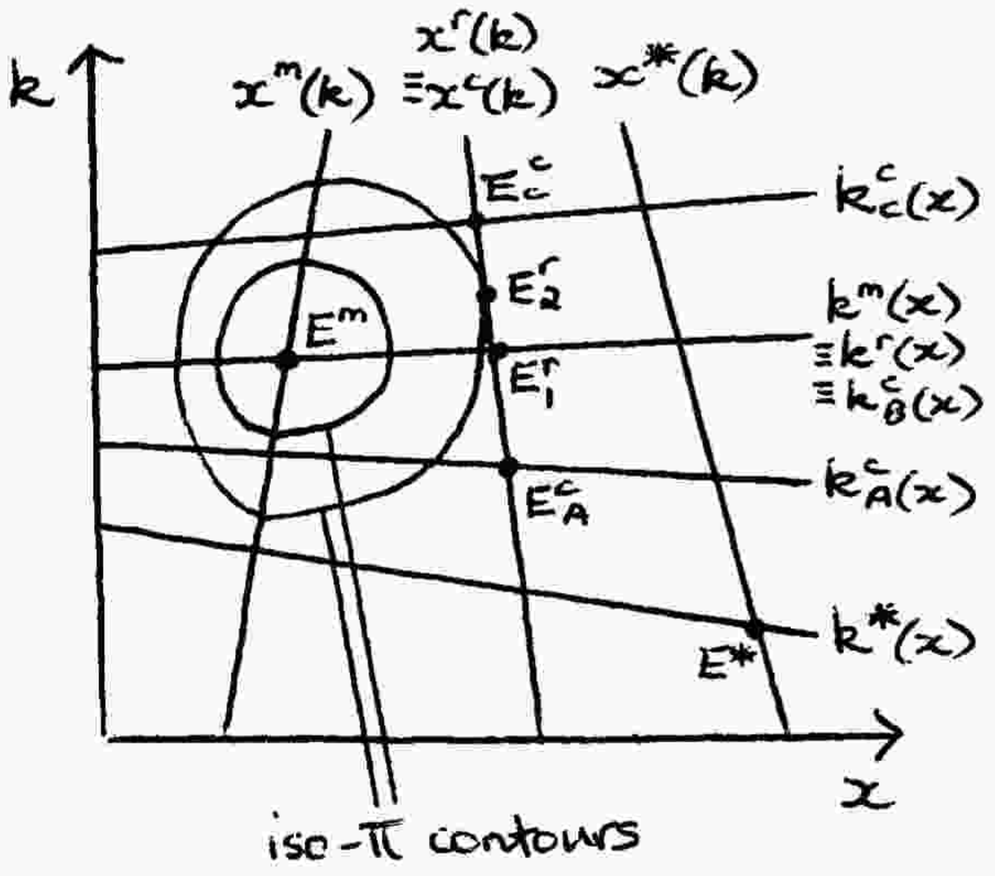


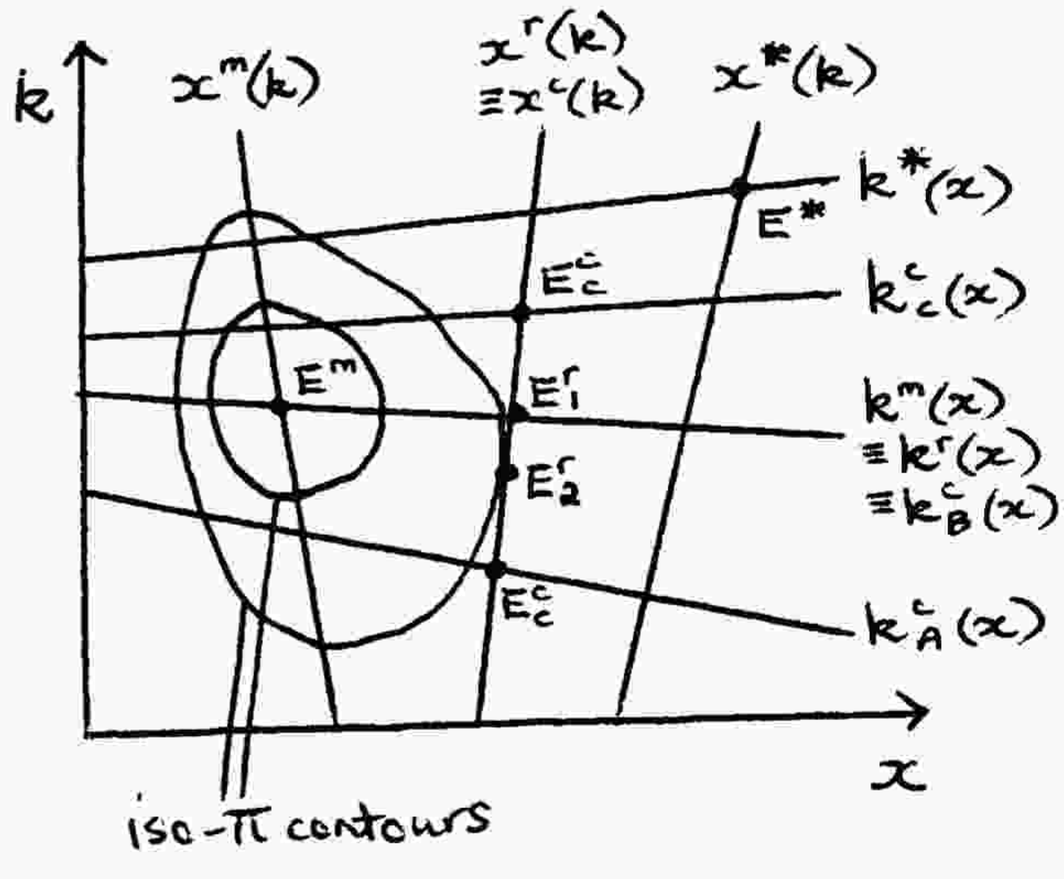
$p_{\bar{k}x} \bar{sgn} - (p_{\bar{k}x}x + \beta) \bar{sgn} - \beta > 0$

$p_{\bar{k}x} < 0$ and $\beta \leq 0$



$p_{\bar{k}x} \bar{sgn} - (p_{\bar{k}x}x + \beta) \bar{sgn} - \beta < 0$





Result 3.6.1 is very intuitive. The sellers want to use quality as an instrument to lower quantity. If $dx^r/dk^r > 0$, then lowering quantity will involve lowering quality, and so $k^r_2 < k^r_1$. The converse of this intuition holds for $dx^r(k)/dk < 0$.

It is clear from the above diagrams that the welfare consequences of allowing self-regulation to be a two-stage game are not unambiguous. However, we can derive the following result.

Result 3.6.2: As long as

$$k^r_2(x^r_2) - k^*(x^*) \gtreqless k^r(x^*) - k^*(x^*),$$

$k^r_2(x^r_2)$ will be closer than $k^r_1(x^r_1)$ to $k^*(x^*)$ if and only if

$$(p_{Rz} + p_{Rx})x + \beta \gtreqless p_{Rx}$$

And so, two-stage self-regulation will become more likely to improve quality (bring quality nearer to $k^*(x^*)$) relative to the *laissez faire*,

(i) as β increases for $p_{Rx} > 0$; and
(ii) as β falls for $p_{Rx} < 0$.

Proof: This proof combines result 3.6.1 with result 3.4.1. Suppose p_{Rx} is positive. Result 3.4.1 tells us that $k^r(x^*) > k^*(x^*)$. Result 3.6.1 tells us that $k^r_1(x^r_1) > k^r_2(x^r_2)$ if and only if $(p_{Rz} + p_{Rx})x + \beta > 0$. Thus, as long as

$$k^r_2(x^r_2) - k^*(x^*) \gtreqless k^r(x^*) - k^*(x^*)$$

we get

$$k^r_1(x^r_1) > k^r_2(x^r_2) > k^*(x^*)$$

if and only if $(p_{Rz} + p_{Rx})x + \beta > 0$. The converse holds for $p_{Rx} < 0$. □

Again, we can simplify this result to give

Result 3.6.3: Equilibrium in the two-stage self-regulation game is most likely to be a welfare-improvement on the *laissez faire* equilibrium when

$$p_R - p_S \cdot \bar{g}_n - p_{Rx} \cdot \bar{g}_n - \beta.$$

Compare this with result 3.5.7, which states that equilibrium in the one-stage self-regulation game is most likely to a welfare-improvement on the *laissez faire* equilibrium when

$$p_R - p_S \cdot \bar{g}_n - p_{Rx} \cdot \bar{g}_n \beta.$$

The difference between these two results is due to the fact that β, via its effect on the slope of $dx^r(k)/dk$, has two contrasting effects.

(i) As quality moves towards $k^*(x^*)$ (which depends on $p_R - p_S \cdot \bar{g}_n - p_{Rx}$), output will be driven away from $x^*(k^*)$ unless $dx^r(k)/dk \cdot \bar{g}_n - p_{Rx}$, which requires $\beta \cdot \bar{g}_n - p_{Rx}$.

(ii) When sellers use quality as an instrument to reduce output, quality will be driven away from $k^*(x^*)$ unless $dx^r(k)/dk \cdot \bar{g}_n p_{Rx}$, which is more likely to be the case when $\beta \cdot \bar{g}_n p_{Rx}$.

Under one-stage self-regulation, only the first of these effects is important, while under two-stage regulation, the second effect will outweigh the first.

3.7 Endogenizing entry

The above model assumes that:

— the number of sellers in the industry n is exogenous;

— n is small enough to enable all n sellers to make non-negative profits under each of the market structures considered;

— when collusion occurs, all n sellers take part.

These assumptions are clearly unrealistic. However, relaxing them is not straightforward. Various questions need to be asked:

—— Are sellers absolutely homogeneous?

—— Is membership of a self-regulatory organization (SRO) compulsory for practising in the industry?

—— Are SROs able to control membership numbers directly, or must they accept all sellers which fulfill the relevant quality requirements (in which case SROs can only control membership numbers via their quality choices, if at all)?

Let m denote the number of sellers belonging to an SRO, and let there be a maximum of one SRO in the industry. In this section, we look at three possible self-regulation set-ups:

(1) The SRO cannot control m separately from n, so that m is simply identical to n (perfect free entry).

(2) The SRO can control m, and membership is compulsory, so that n is constrained to be equal to m (SRO membership as an entry barrier).

(3) The SRO can control m, but membership is not compulsory, so that m could be less than n (threat-constrained entry).

In addition to $\mu(k_i, x_i)$, it is now assumed that each seller i faces a fixed cost of functioning in the industry ϕ_i, where ϕ is distributed according to some distribution function $H(\phi)$.[4] If all potential and actual sellers are placed in ascending order of ϕ, $\phi(n)$ is defined as the fixed cost of functioning of the nth practitioner. If all sellers are homogeneous *ex ante*, then $\phi(n)$ will be the same for all n. Otherwise, $\phi'(n) \geq 0$.

---

[4] This is similar to the assumption made in my extension of Shapiro's (1983) model above (see section 2.4.5).

3.7.1 Perfect free entry

Suppose that the SRO cannot control m separately from n. If membership of the SRO is compulsory, then this lack of control means that the SRO must let in every seller who wants to practise in the industry. If membership of the SRO is not compulsory, it means that the SRO must let in every seller who wants to join it. The framework in this chapter is unsuitable for analysing an industry which has some sellers inside an SRO and some sellers outside it. So, let us assume that an SRO will form only if all sellers agree to join it, and that penalties for leaving it can be designed such that no member will ever want to cheat and provide $k_i \neq k$. In this case, all sellers will join, and remain in, the SRO, since profits are higher under self-regulation than under *laissez faire*. Thus, whether or not membership is compulsory, m will simply be identical to n.

The aim of this subsection is to show that most of the results derived above for fixed n will be *qualitatively* unchanged under perfect free entry. *Quantitative* comparisons between similar equilibria in the fixed n and the free entry set-ups are difficult, however, because there is a whole continuum of fixed n equilibria, corresponding to the continuum of possible n choices.

Perfect free entry implies that n will increase until the last seller to enter the industry (the marginal seller) just makes non-negative profits.

When all sellers are homogeneous *ex ante*, and when equilibrium is symmetric, the profits of the marginal seller will be equal to the profits of all inframarginal sellers. So free entry drives total industry profits down to zero. In this case, there is little incentive for the sellers in an industry to form an SRO, since any resulting profits will be dissipated by new entry.

Suppose, then, that sellers are heterogeneous *ex ante*. In this case, the profits received by seller i are given by

$$\pi(g_i, \bar{k}, z_i, X, \phi_i) = p(g_i, \bar{k}, z_i, X).x_i - \mu(k_i, x_i) - \phi_i$$

We know that, at equilibrium, the marginal (nth) seller in the industry just makes non-negative profits. Thus, n is given by

$$p(0, k, 0, nx).x - \mu(k, x) - \phi(n) = 0$$

This equation gives us the implicit function n(k, x). Using total differentiation, we can derive

$$n_x = \frac{np_x x + p - \mu_x}{\phi' - x^2 p_x} \equiv \frac{np_x x + p - \mu_x}{-\phi' + x^2 p_x}$$

and

$$n_k = \frac{p_k x - \mu_k}{\phi' - x^2 p_x} \equiv \frac{p_k x - \mu_k}{-\phi' + x^2 p_x}$$

Total industry equilibrium profits are given by

$$\Pi(k, x, n) = p(0, k, 0, nx).nx - n.\mu(k, x) - \int_0^n \phi(n')\, dn'$$

And total welfare is given by

$$W(k, x, n) = \int_0^{nx} p(0, k, 0, X')\, dX' - n.\mu(k, x) - \int_0^n \phi(n')\, dn'$$

Suppose that entry decisions are made *after* the k and x decisions. Collusion on a variable (k and/or x) now involves considering the *indirect* effect of that variable on profits, due to its effect on n, as well as the direct effects considered above. For collusion on k and x, in a simultaneous choice game, the first order profit maximization conditions become:

$$\Pi^m_x = n(np_x x + p - \mu_x) + (px + nx^2 p_x - \mu - \phi)n_x = 0$$

and $$\Pi^m_k = \Pi^r_k = n(p_k x - \mu_k) + (px + nx^2 p_x - \mu - \phi)n_k = 0$$

Because $(px - \mu - \phi)$ equals zero at equilbrium by definition, and substituting in for n_x and n_k, these conditions become

$$\Pi^m_x = \frac{\phi' n}{\phi' - x^2 p_x} \cdot (np_x x + p - \mu_x) = 0$$

$$\Pi^m_k = \frac{\phi' n}{\phi' - x^2 p_x} \cdot (p_k x - \mu_k) = 0$$

which are satisfied if

$$(np_x x + p - \mu_x) = 0$$

and $$(p_k x - \mu_k) = 0$$

respectively. These last conditions are identical to the first order profit maximization conditions derived in section 3.3 (which don't take $n(k,x)$ into account). Thus, allowing perfect free entry leaves the one-stage collusive choices of k and x qualitatively the same (although, since $n(k,x)$ need not be identical to the fixed n assumed in section 3.3, these choices will not generally be *quantitatively* the same).

Now, consider total welfare. When $n = n(k,x)$, the first order welfare maximization conditions become:

$$W_x = n(p - \mu_x) + (px - \mu - \phi)n_x = 0$$

and $$W_k = \int_0^{nx} p_k(0, k, 0, X')\, dX' - n\mu_k + (px - \mu - \phi)n_k = 0$$

Because $(px - \mu - \phi)$ equals zero at free entry equilibrium by definition, these conditions become

$$W_x = n(p - \mu_x) = 0$$

and
$$W_k = \int_0^{nx} p_{\bar{x}}(0, k, 0, X')\, dX' - n\mu_k = 0$$

which are again qualitatively identical to the first order welfare maximization conditions in section 3.3. So allowing perfect free entry leaves the welfare maximizing choices of k and x qualitatively the same, too.

When there is competition on either k or x, it is less clear whether perfect free entry changes things. One possible assumption is that, since n depends on k and x, rather than k_i and x_i, no seller will think that she can affect n, and the analysis will be exactly as it is with fixed n. That is

$$\pi^{r}_{x} = \pi^{e}_{x} = (p_{\bar{x}} + p_x)x_i + p - \mu_x = 0$$

and
$$\pi^{e}_{k} = 1/n((n - 1)p_{\bar{x}} + p_{\bar{k}})x_i - \mu_k = 0$$

Let us call this *myopic* competition.

An alternative assumption is that we can rewrite $n(k, x)$ as $\bar{n}(\sum k_i, \sum x_i)$, where the partial derivatives of $\bar{n}$ with respect to k_i and x_i are equal to 1/n times the partial derivatives of $n(k, x)$ with respect to k and x, respectively. This is a *forward-thinking* assumption. The profits of an individual forward-thinking seller i, who expects all other sellers to produce x, are given by

$$\pi(g_i, \bar{k}, x_i, x, \bar{n}, \phi_i) = p(g_i, \bar{k}, (x_i - x), (x_i + (\bar{n} - 1)x)).x_i - \mu(k_i, x_i) - \phi_i$$

If there is competition in x, this gives a first order profit maximization condition of

$$\pi^{r}_{x} = \pi^{c}_{x} = (p_{z} + p_{x})x_{i} + p - \mu_{x} + xx_{i}p_{x}\bar{n}_{x} = 0$$

Comparing this with the first order profit maximization condition under myopic competition, there is one extra term $xx_{i}p_{x}\bar{n}_{x}$. Since p_{x} is always negative, and $\bar{n}_{x}$ is negative for all $x(k) > x^{m}(k)$, this extra term will be positive at the myopic competition equilibrium output. So, when sellers are forward thinking, equilibrium x will be higher, for given k, than it is under myopia. This is because raising output lowers n, which raises the price each seller receives for given k_{i} and x_{i}. Note that total industry profits are lower in this equilibrium than they are when sellers are myopic.

If there is competition in k, the analysis is the same, except that the extra term is now $xx_{i}p_{x}\bar{n}_{k}$. We know that p_{x} is always negative and that $\bar{n}_{x}$ is negative if and only if $k(x) > k^{m}(x)$. For case A *laissez faire*, where $p_{\bar{k}} > p_{\bar{x}}$, we know that $k^{c}(x) < k^{m}(x)$ in the myopic competition case (see result 3.4.2). When sellers are forward-thinking, therefore, $\bar{n}_{x}$ is positive, which means that the extra term is negative. So equilibrium $k^{c}_{\theta}(x)$ is even lower under forward-thinking than under myopia. The converse holds for case C *laissez faire*, where $p_{\bar{k}} < p_{\bar{x}}$. Again, total industry profits are lower in this equilibrium than they are when sellers are myopic.

This discussion is summarised in the following result.

Result 3.7.1: Allowing perfect free entry instead of fixed n

(a) has no qualitative effect on the equilibrium k and x choices made under one-stage collusion;

(b) has no qualitative effect on the welfare maximizing levels of k and x;

(c) has no qualitative effect on the equilibrium k and x choices made under competition, as long as sellers *do not* consider the effect of their k and x choices on n (myopic competition);

(d) drives the equilibrium k and x choices made under competition even further away from their collusive levels, when sellers *do* consider the effect of their k and x choices on n (forward-thinking competition).

Remember that self-regulation is taken to involve collusion on k and competition on x. Thus, under myopic competition, one-stage self-regulation is most likely to be welfare-superior to *laissez faire* under exactly the same conditions as were given in section 3.5. That is, when

$$p_k - p_q = \bar{g}_n - p_{kx} = \bar{g}_n \beta$$

Moreover, when one-stage self-regulation *is* welfare-superior to *laissez faire*, then the welfare improvement associated with the introduction of self-regulation will be even greater when competing sellers are forward-thinking than when they are myopic.

Turning to the two-stage collusion game. Given that there will be Cournot competition in the second stage of the game, the first stage profit maximization condition with free entry is

$$\frac{d\pi(k, x^r(k), n(k, x^r(k)))}{dk} = \frac{\partial\pi}{\partial k} + \frac{\partial\pi}{\partial n} n_k + \left[\frac{\partial\pi}{\partial x} + \frac{\partial\pi}{\partial n} n_x\right] \frac{dx^r(k)}{dk}$$

At the one-stage equilibrium, the first two terms in this expression must be zero, and since $x^r(k) > x^m(k)$, for all k, the bracketted term must be negative at any point along $x^r(k)$. Thus, just as under fixed n, $d\pi/dk = \bar{g}_n - dx^r(k)/dk$, and so $k^r_1 > k^r_2$ if and only if $dx^r(k)/dk > 0$. Moreover, by total differentiation of the first order profit maximization condition with respect to x, and using the fact that n_k equals zero at the one-stage equilibrium, it can also be shown that $dx^r(k)/dk = \bar{g}_n p_{kx} x + \beta$, as in the fixed n case.

Thus, the results comparing the one-stage with the two-stage self-regulation equlibrium are qualitatively the same under perfect free entry

as under fixed n. That is, two-stage self-regulation is most likely to be welfare-superior to *laissez faire* when

$$p_{\bar{k}} - p_{g\,\bar{g}n} - p_{\bar{k}x\,\bar{g}n} - \beta$$

3.7.2 Compulsory SRO membership as an entry barrier

When the SRO can control m, and when membership of the SRO is compulsory, then the SRO can effectively choose n, separately from its choice of k and x (as long as the n choice made is lower than the free entry n, n(k, x)). In this section, the self-regulation and full collusion equilibria where n is a choice variable are compared with the corresponding equilibria under perfect free entry. Under free entry, n is positively related to profits (because higher profits induce further entry). Since profits are higher under self-regulation and full collusion than at the welfare-optimum, it must be the case that perfect free entry involves too much entry for social welfare maximization[5] (a familiar result for Cournot equilibrium - see Mankiw and Whinston (1986)).

This section examines an implication of this statement: that a restriction on entry into the industry will, at least at first, improve welfare. Only the one-stage game game is considered. The first order welfare maximization conditions for this will be exactly as in section 3.3, with the addition of

$$W_n = p(0, k, 0, nx)x - \mu(k, x) - \phi(n) = 0$$

This condition will always hold under perfect free entry. Thus $n^*(k, x)$ is identical to $n(k, x)$ (the number of sellers at given k and x under free entry), and the welfare optimal choices of k and x will be the same as under free entry.

---

[5] Although notice that *for given k and x* free entry results in exactly the optimal number of sellers, since $W_n = (px - \mu - \phi)$ which must equal zero under free entry.

For profit maximization in the fully collusive oligopoly and the SRO, the first order conditions with respect to k and x are as in section 3.3. That is, collusion on k gives

$$\pi^m_k = \pi^r_k = p_{\bar{x}}(0, k, 0, nx)x - \mu_k(k, x) = 0$$

Collusion on x gives

$$\pi^m_x = p_x(0, k, 0, nx)nx + p(0, k, 0, nx) - \mu_x(k, x) = 0$$

And competition on x gives

$$\pi^r_x = [p_{\bar{x}}(0, k, 0, nx) + p_x(0, k, 0, nx)].x + p(0, k, 0, nx) - \mu_x(k, x) = 0$$

In addition, when n is also a choice variable, Total industry profits will be maximized when

$$\Pi^m_n = \Pi^r_n = p_x(0, k, 0, nx)nx^2 + p(0, k, 0, nx)x - \mu(k, x) - \phi(n) = 0$$

Since p_x is negative, $(px - \mu - \phi)$ must be positive under profit maximization. Thus $n^r(k, x) = n^m(k, x) < n(k, x)$, for any given k and x. That is, *profit maximization requires that the industry restricts entry.* This is intuitive.

What effects will this restriction of entry have on k, x and welfare?

For given x, total differentiation of $\pi^m_k = 0$ (and so $\pi^r_k = 0$) gives

$$\frac{dk^m(x, n)}{dn} = \frac{dk^r(x, n)}{dn} = \frac{p_{\bar{x}x}x^2}{\mu_{kk} - p_{\bar{x}\bar{x}}} \equiv \bar{\bar{g}}_n \; p_{\bar{x}x}$$

When $p_{\bar{x}x}$ is positive (negative), a reduction in n causes the $k^m(x)/k^r(x)$ schedule to shift downwards (upwards). From result 3.4.1, a positive (negative) $p_{\bar{x}x}$ means that, under fixed n (and therefore also under perfect free entry), the $k^m(x)/k^r(x)$ schedule lies above (below) the $k^*(x)$

schedule. Thus, the effect of restricting entry will bring the $k^m(x)/k^r(x)$ schedule closer to the $k^*(x)$ schedule (although it may overshoot).

For given k, total differentiation of $\pi^m_x = 0$ and $\pi^r_x = 0$ gives

$$\frac{dx^m(k,n)}{dn} = \frac{p_{xx}nx^2 + 2p_xx}{\mu_{xxx} - 2np_x - p_{xx}n^2x}$$

$$\frac{dx^r(k,n)}{dn} = \frac{p_{xx}x^2 + p_xx}{\mu_{xxx} - (n+1)p_x - p_{xx}nx}$$

both of which are negative. Thus, a reduction in n causes the $x^m(k)$ and $x^r(k)$ schedules to shift to the right, which means that they move towards $x^*(k)$ (although, again, they might overshoot).

Under perfect free entry, n is increasing in $px - \mu$, the per seller profits before fixed costs ϕ are taken into account. Because $px - \mu$ is maximized with respect to both x and k at $(k^m(x^m), x^m(k^m))$, and maximized with respect to k at $(k^r(x^r), x^r(k^r))$, it is clear that

$$n(k^*(x^*), x^*(k^*f)) < n(k^r(x^r), x^r(k^r)) < n(k^m(x^m), x^m(k^m))$$

Using this inequality, and remembering that a reduction in n from its free entry level causes the $k^m(x)/k^r(x)$ schedule to move towards the $k^*(x)$ schedule *and* the $x^m(k)$ and $x^r(k)$ schedules to shift towards the $x^*(k)$ schedule, we get the following result:

Result 3.7.2: When n is a choice variable (as under compulsory SRO membership), then as n falls from its free entry level, welfare will generally (although not necessarily) be improved (at least at first), since

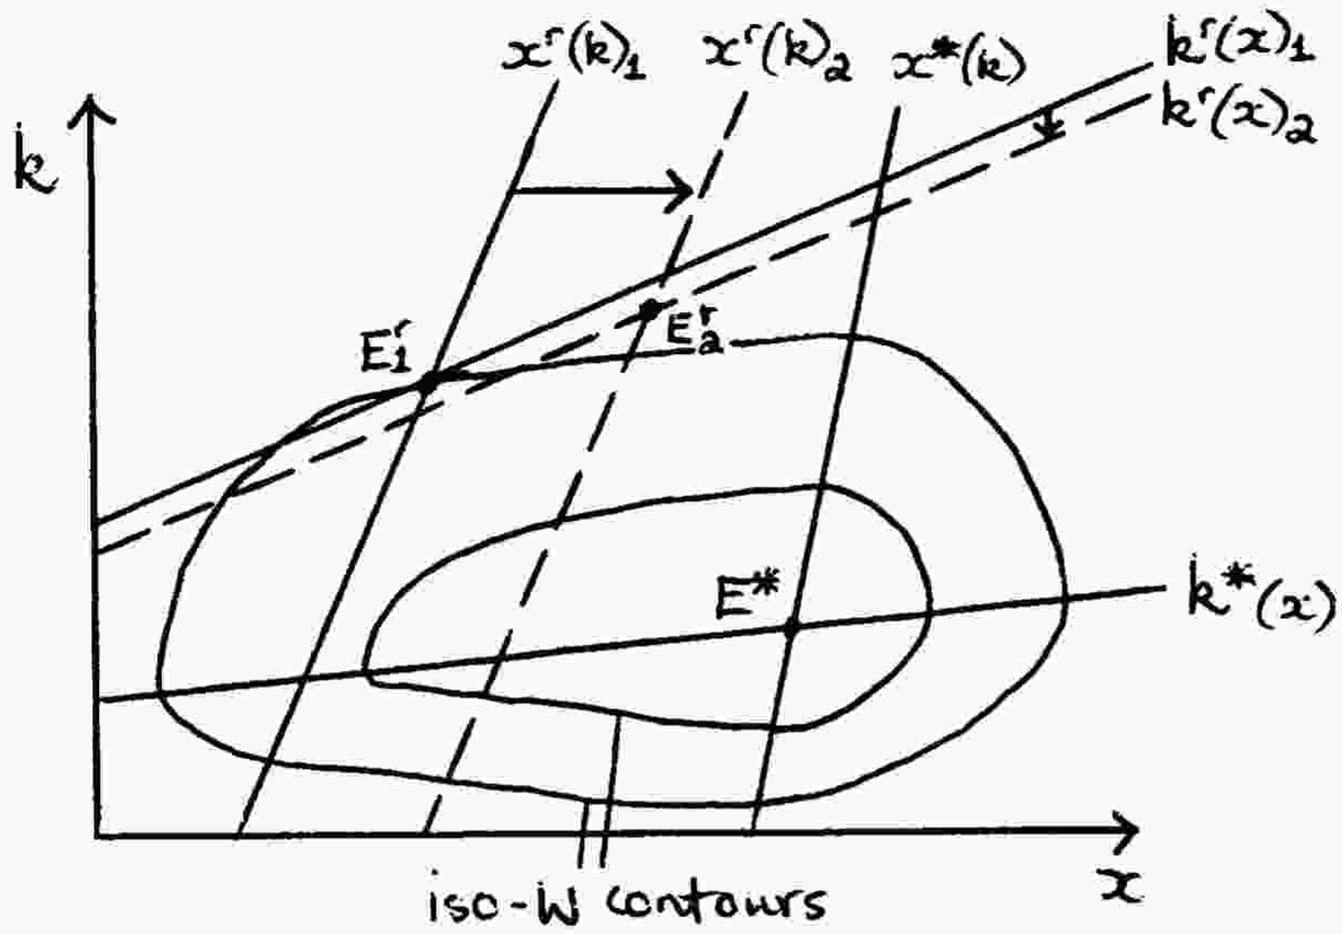
— $k^m(x^m)$ and $k^r(x^r)$ move towards $k^*(x^*)$;

— $x^m(k^m)$ and $x^r(k^r)$ move towards $x^*(k^*)$; and

— $n^m(k^m, x^m)$ and $n^r(k^r, x^r)$ move towards $n^*(k^*, x^*)$.

However, this result does not imply that compulsory membership is welfare-superior to free entry. There are two possible reasons for this.

The first reason is simply that even when loci move towards the welfare optimal loci, their intersection need not become welfare superior. A straightforward example where this is true is shown below.



The second reason why the above argument does not imply that compulsory membership is necessarily welfare-superior to free entry is that, at equilibrium, the profit-maximization loci may overshoot and jump over the welfare-maximization loci. Let us compare self-regulation and the welfare optimum. For given k and x, we get

$$W_n\big|_{n^r} = - p_x n^r x^2 > 0$$

Thus, $n^*(k, x) > n^r(k, x)$, which means that, for any given k and x, n is restricted by *more* than is socially optimal.

Next, for given k, we get

$$W_x\big|_{\Pi^r{}_x=0} = n^*(k, x).[\,p(0, k, n^*(k, x)x) - p(0, k, n^r(k, x)x)]$$
$$- n^*(k, x)[\,p_x(0, k, n^r(k, x)x) + p_x(0, k, n^r(k, x)x)]$$

The second line of this expression is positive. However, $n^*(k, x) > n^r(k, x)$, which means that the first line of this expression is negative.

Thus, it is no longer obvious that W_{x} is positive when $\Pi^{r}_{x} = 0$, and so the $x^{*}(k)$ schedule need no longer lie to the right of the $x^{r}(k)$ schedule, as it did under free entry.

Similarly, for given x, we get

$$W_{k}\,|_{\Pi^{r}_{k}=0} = \int_{0}^{n^{*}(k,x)x} p_{\bar{k}}(0,k,X')\,dX' - p_{\bar{k}}(0,k,n^{*}(k,x)x) + [p_{\bar{k}}(0,k,n^{*}(k,x)x) - p_{\bar{k}}(0,k,n^{r}(k,x)x)]$$

The first line of this expression has the same sign as $-p_{\bar{k}x}$ (as under free entry). But, because $n^{*}(k,x) > n^{r}(k,x)$, the second line has the same sign as $p_{\bar{k}x}$. It no longer clear, then, that the $k^{r}(x)$ schedule will lie on the same side of $k^{*}(x)$ as it did under free entry.

Thus, even if a small restriction on membership does improve welfare, relative to the free entry equilibria, the actual restriction chosen by an SRO or a fully collusive oligopoly may be so great as to be welfare inferior to free entry.

3.7.3 Threat-constrained entry

The point to be made in this subsection is straightforward. It was shown above that collusive choice of n may not be welfare superior to free entry for two reasons. The second (and arguably the more important) of these reasons derives from the fact that n will be restricted beyond its welfare optimal level, and that this might cause the profit maximization output and quality loci to overshoot and jump over the welfare maximization loci. If, then, we can find a form of self-regulation which allows *some* restriction of n, from its free entry level, but which prevents any overshooting, then this form of self-regulation will generally be welfare-superior to both the set-ups described above.

This subsection looks at such a set-up. Industry profits are always higher under self-regulation than under *laissez faire*. Thus, when there is

perfect free entry, there will be more sellers under self-regulation than under *laissez faire*. Now, suppose

(i) that an SRO can choose m, but that membership is not compulsory, so that m could be less than n.

(ii) that an SRO can threaten to abolish itself (and so cause a reversion to the *laissez faire* equilibrium) in the face of any entry into the industry by non-members.

In this case, a non-member will enter the industry if and only if it would still be able to function under *laissez faire*. An SRO will do best by preventing any non-members from entering the industry (since it must be able to do at least as well when it can control what new entrants do, as when it cannot). Thus, an SRO will do best by setting m equal to the number of sellers which would exist under *laissez faire* and perfect free entry, and allowing the m lowest ϕ sellers to become members. Its threat to revert to *laissez faire* in the face of further entry will then succeed in preventing any further entry[6].

Because m (and therefore n) cannot be restricted beyond its free entry *laissez faire* level, $n(k^c(x^c), x^c(k^c))$, it is no longer the case that m under self-regulation (or even under full collusion) can end up lower than $n^*(k, x)$, for any given k and x. And this means that the profit maximization output and quality loci cannot overshoot and jump over the welfare maximization loci.

In this threat-constrained entry set-up, then, the second of our two explanations for why collusive choice of n may not be welfare superior to free entry no longer holds. Thus, self-regulation (and full collusion) in this set-up is likely to be welfare superior to that under free entry.

---

[6] Note that it is costly for the SRO to carry out its threat of returning to *laissez-faire*, *ex post*. Thus, for this threat to be credible *ex ante*, the self-regulation game must, implicitly, be a multi-period one, providing the SRO with sufficient reputation incentives to carry out the threat.

The conditions under which self-regulation is likely to be welfare-superior to *laissez faire* remain the same in this set-up as under free entry (since this set-up is effectively identical in welfare terms to the fixed n case, which also gives these same conditions (see subsection (a) above)). However, because the actual welfare improvement achieved will be generally be higher under this set-up, the benefits of self-regulation are more likely to outweigh any costs if threat-constrained entry is allowed.

3.7.4 Endogenizing entry - conclusions

To briefly conclude this section, we have seen that:

— Endogenising entry does not drastically change the analysis of sections 3.2 - 3.6, either under perfect free entry or under threat-constrained entry;

— Self-regulation under free entry results in too much entry for welfare-maximization;

— Self-regulation when n is a choice variable may result in too little entry for welfare-maximization;

— Thus, self-regulation under threat-constrained entry (which results in too much entry for welfare-maximization but less entry than under free entry), may well be welfare-superior to either of the other set-ups.

Appendix 3.A. Microfoundations for the inverse demand function $p(g_i, \bar{k}, z_i, X)$

The rationale for using an inverse demand function which has g_i and $\bar{k}$ as separate arguments is given in section 3.2 above. Briefly, the benefit of this formulation (over more usual formulations - like $p(\bar{k}, X)$) is that it captures two potential between-seller quality externalities:

— p_g captures a negative between-seller externality. When sellers compete in relative quality, to some extent, then a unilateral

increase in k_i will, by raising $\bar{k}$, lower the price received by all sellers other than i.

— $p_{\bar{k}}$ captures a positive between-seller externality. When quality observability is imperfect, a consumer's expectation of seller i's quality will tend to depend, at least partly, on $\bar{k}$. So a unilateral increase in k_i will, by raising $\bar{k}$, raise the price received by all sellers other than i.

In order to provide microfoundations for this inverse demand function, it is necessary to find microfoundations for each of these externalities, and also to find a way of putting them together.

In this appendix I suggest two possible models that might be used to achieve this. The first is a search model, and I do not solve it explicitly. I do give and solve a simple example of the second model, which is essentially a variant of the Hotelling model. My aim, however, is not so much to give concrete examples, as to give some hints as to how examples could, in theory, be derived.

3.A.1 The negative externality

Let us look at $p_{\bar{k}}$ first. Let each consumer get utility

$$U(p, k, \bar{u}) = u(p, k) - \bar{u}$$

from a good of expected quality k and expected price p, where $u_p < 0$, $u_k > 0$, $u_{kk} < 0$, and where the reserve utility $\bar{u}$ is different for each consumer, and is distributed over the interval $[0, \infty]$, with a cumulative distribution function $F(\bar{u})$.

When $\bar{u}$ is the only factor which differentiates consumers, but sellers are differentiated by both p and k (which are observable by consumers), then every consumer will buy from the seller which offers them the highest $u(p, k)$. That is, the demand faced by seller i is discontinuous in

$u(p_i, k_i)$. If there are n sellers, and if all the other (n - 1) sellers charge price p and provide quality k, then seller i's demand function will be

$$x_i = 0 \qquad \text{if } u(p_i, k_i) < u(p, k)$$

$$x_i = F(u(p, k))/n \qquad \text{if } u(p_i, k_i) = u(p, k)$$

$$x_i = F(u(p_i, k_i)) \qquad \text{if } u(p_i, k_i) > u(p, k)$$

In this set-up, there will effectively be Bertrand-style competition in $u(p, k)$, with $u(p, k)$ rising until all sellers make zero profits.

If we want to allow Cournot-style competition in k, then, demand x_i needs to be continuous in $u(p_i, k_i)$. This requires the introduction of a further dimension of differentiation between consumers. Two commonly used dimensions are:

—— differentiation by marginal search costs

—— differentiation by total transport costs

These are described in turn below:

(a) A search model

Consider the following set-up. There are X consumers in the industry and n sellers. Initially, these X consumers are spread equally over the n sellers, so that each seller's demand is X/n.

While consumers know $\bar{k}$ and $\bar{p}$, and the k_i and p_i of their particular seller, they need to carry out search in order to find out about the qualities and prices of other sellers. One period of search costs τ, where τ is different for every consumer and is distributed by $G(\tau)$ over the interval $(0, \infty]$.

Search is assumed to function in the following way. A consumer who decides to search in any period has an equal probability (1/n) of arriving at each of the n sellers, whereupon he learns the price and quality of that seller[7]. He may then stay with this new seller, or decide to search again. There is no limit to the number of periods for which search may continue, but a cost of τ is incurred every time.

A consumer will decide to search if the expected utility improvement from doing so exceeds the search cost. The expected utility from one period of search, EU, will be increasing in $\bar{k}$ and decreasing in $\bar{p}$[8]. Thus, a consumer who currently pays p_i for quality k_i will search if and only if

$$EU(\bar{p}, \bar{k}) - u(p_i, k_i) > \tau$$

As $u(p_i, k_i)$ increases from zero, the number of consumers who leave it to search will be falling, and the number of consumers who stay at it, after search from a lower utility seller, will be rising. Thus, we get

$$x_i = x(u(p_i, k_i), EU(\bar{p}, \bar{k}))$$

where x_i is increasing in its first argument and decreasing in its second argument.

From this, it should be possible, although arduous, to obtain an inverse demand function of the form

$$p_i = \tilde{p}(g_i, \bar{k}, z_i, \bar{p}))$$

for each seller (in which $\tilde{p}_g$ will be greater than $\tilde{p}_k$).

---

[7] This set-up means that a searcher has a 1/n chance of finding himself back where he started. This is obviously unrealistic, but makes the model a lot more tractable mathematically.

[8] It is assumed that consumers know nothing else about the distribution of prices and qualities (other than $\bar{p}$ and $\bar{k}$). Otherwise, this information would also affect the expected utility of search.

(b) A Hotelling model with vertical differentiation

Consider the following simple model. Two sellers (A and B) are located at opposite ends (0 and 1) of a unit horizontal differentiation line, and consumers are located uniformly along this line. Each consumer faces transport costs td, where d is the distance of that consumer from the seller whose good he purchases. Thus, utility is now

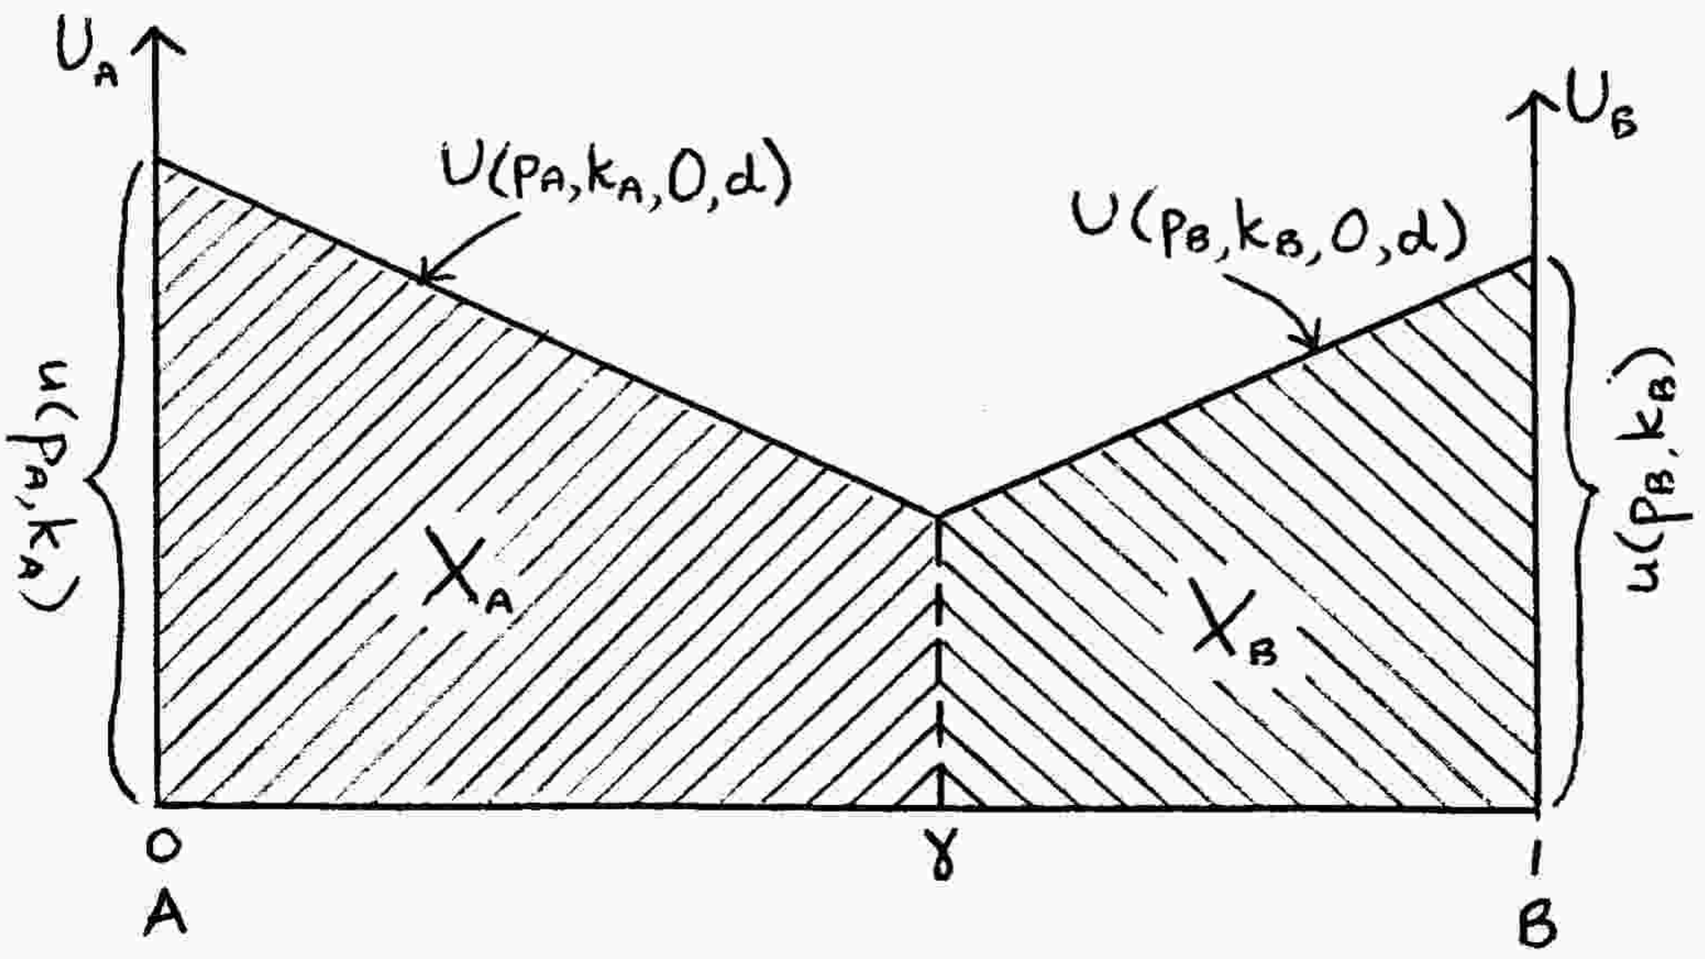
$$U(p, k, \bar{u}, d) = u(p, k) - \bar{u} - td$$

The consumer who is just indifferent about which seller he buys from will be positioned at γ, where γ is given by

$$\gamma = \frac{u(p_A, k_A) - u(p_B, k_B) + t}{2t}$$

If $\bar{u}$ is distributed uniformly on $[0, \infty]$, with a constant density of one, then the total demand facing seller A (and B respectively) will be the area under its $U(p, k, 0, d)$ line over the interval $[0, \gamma]$ (and $[\gamma, 1]$ respectively), as shown below.

Thus, if seller B chooses p_B and k_B, seller A's demand X_A is given by



$$\tilde{x}_A = \frac{(u(p_A, k_A) - u(p_B, k_B) + t).(3u(p_A, k_A) + u(p_B, k_B) - t)}{8t}$$

In this simple example, $\tilde{x}_i$ is continuous and increasing in $u(p_i, k_i)$, as required. It should now be possible, although again arduous, to obtain an inverse demand function $p(g_i, \bar{k}, \tilde{z}_i, \tilde{X})$ for each seller from this framework (in which p_B will be greater than p_A)[9].

There are various ways of extending this two seller model to the multi-seller case. These tend, however, to get extremely complex and intractable. Thus, when discussing the positive externality, and in the simple detailed example given below, we stick to the two-seller case.

3.A.2 The positive externality

There are various ways of modelling this externality, which can involve different assumptions about what exactly consumers can observe. In this subsection, we consider two possible assumptions, each of which fits quite nicely into one of the above negative externality models

(a) The search model

In the search model above, the number of consumers in the industry X was taken as given. Suppose now that search (as described above) occurs in the second stage of a two stage game, with entry decisions being made at the first stage.

Assume that potential consumers into an industry can only observe $\bar{p}$ and $\bar{k}$. They enter the industry if and only if their expected utility from doing so, given $\bar{p}$ and $\bar{k}$, is non-negative. Thus $X = X(\bar{p}, \bar{k})$. It will generally be possible to rearrange this into to an industry inverse demand function $\bar{p}(\bar{k}, X)$, where $\bar{p}_{\bar{k}} > 0$. Substituting this into the inverse demand function from the second part of the game $\tilde{p}(g_i, \bar{k}, z_i, \bar{p})$, and rearranging, gives

---

[9] Note that $\tilde{X}$ and $\tilde{z}_i$ are to $\tilde{x}_i$ exactly as X and z_i are to x_i.

$$p_i = p(g_i, \bar{k}, z_i, X)$$

And while β_g is definitely greater than $\beta_{\bar{k}}$, the sign of $(p_g - p_{\bar{k}})$ is ambiguous, since $p_{\bar{k}} = \beta_k + \beta_B \bar{p}_{\bar{k}}$ and $\beta_B \bar{p}_{\bar{k}}$ is definitely positive.

In this case, then, $(p_g - p_{\bar{k}})$ will be more likely to be positive, the lower are β_B and $\bar{p}_{\bar{k}}$.

(b) The Hotelling model

Let us change the two seller Hotelling model described above, in the following simple way: There is a proportion λ of all potential and actual consumers who face no observability problems. They can observe p_A, p_B, k_A and k_B, and they will thus behave exactly as above. Call their demand x^{λ}_A, where $x^{\lambda}_A \equiv \lambda x_A$. The rest of the actual and potential consumers are assumed to be less well informed. They can observe each seller's price, but they can only observe average quality $\bar{k}$. These consumers decide between sellers purely on the basis of their price. Thus, demand for A's product from these consumers will be

$$x^{1-\lambda}_A = (1 - \lambda).\frac{(u(p_A, \bar{k}) - u(p_B, \bar{k}) + t).(3u(p_A, \bar{k}) + u(p_B, \bar{k}) - t)}{8t}$$

Total differentiation of this shows that, for given p_B and $x^{1-\lambda}_A$, $dp_A/d\bar{k}$ is positive. Thus, an increase in k_B will, via its effect on $\bar{k}$, increase p_A. This is the positive externality we aimed to model.

Aggregating over all consumers, total demand for A's product becomes

$$x_A = x^{\lambda}_A(p_A, p_B, k_A, k_B) + x^{1-\lambda}_A(p_A, p_B, \bar{k})$$

Again, it should be possible, if arduous, to obtain an inverse demand function $p(g_i, \bar{k}, z_i, X)$ for each seller from this framework. However, it is no longer obvious that p_g will be greater than $p_{\bar{k}}$. All we can say is:

— When $\lambda = 1$, then $p_g > p_E$.

— When $\lambda = 0$, then $p_g < p_E$ (indeed, $p_g = 0$).

— When $0 < \lambda < 1$, the inequality may go either way.

Generally, then, $(p_g - p_E)$ will be more likely to be positive, the higher is λ (i.e. the higher is the proportion of well informed consumers in the population).

This result should extend to a multi-seller version of this model.

3.A.3 A simple 2-seller example

In order to show in more detail how one might go about finding an inverse demand function for such a Hotelling model, let us make the following simplifying assumption about the utility function:

$$u(p, k) = \theta k - p$$

Let i denote either seller A or B and let j denote the other seller. In this case x_i is given by

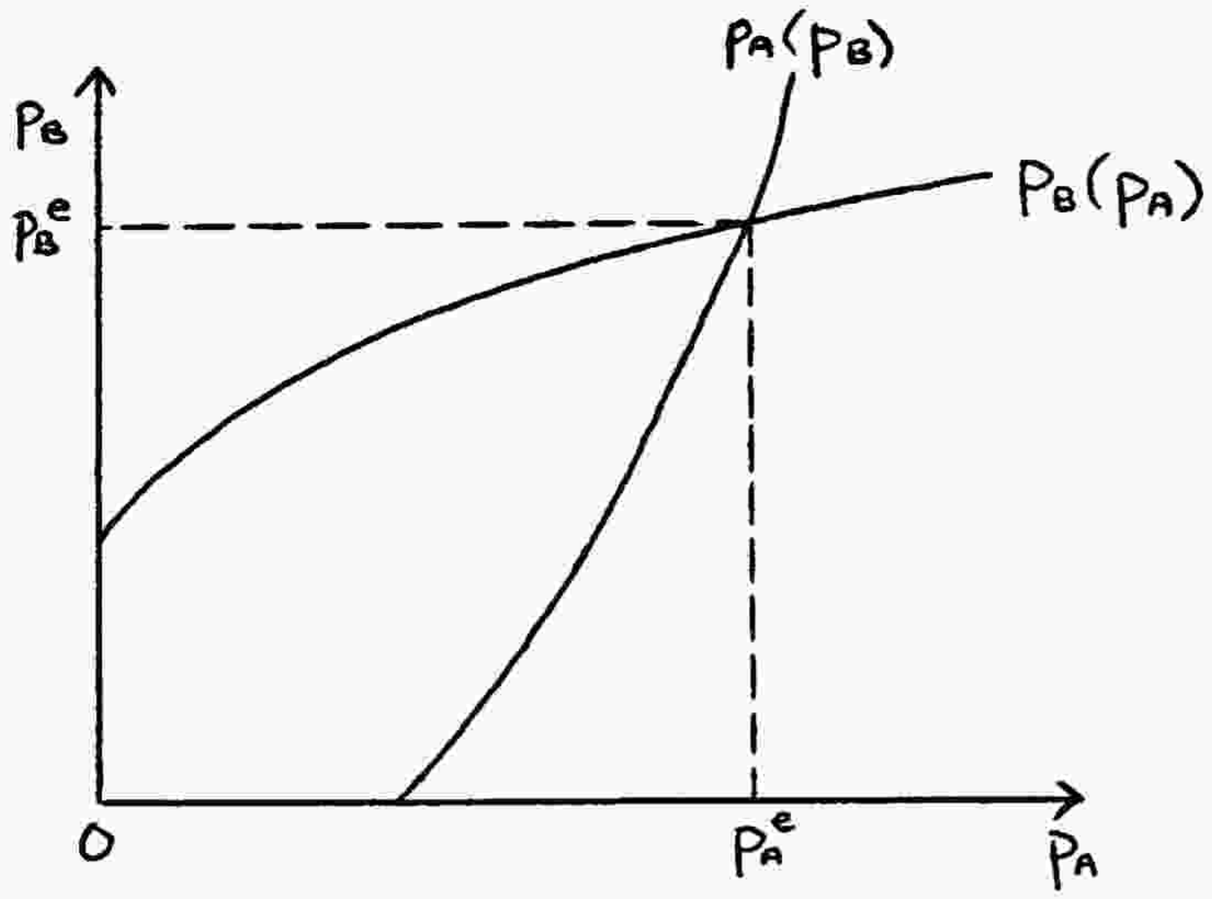
$$(3.A.1) \quad 8x_i t = 4\lambda\theta g_i(2\theta\bar{k} + \theta g_i - 2p_i) + (p_j - p_i + t)(4\theta\bar{k} - 3p_i - p_j - t)$$

Solving this explicitly for p_i involves a complicated square root sign. However, by total differentiation, we can derive

$$(3.A.2) \quad \frac{dp_i}{dp_j} = \frac{2\theta\bar{k} - p_i - p_j - t}{4\lambda\theta g_i + 2\theta\bar{k} - 3p_i + p_j + t}$$

which must be positive at equilbrium for the model to be meaningful (since if either the numerator or the denominator is negative, there will be some points along the horizontal quality line at which no consumers buy from either seller).

Thus we can draw the two loci $p_A(p_B)$ and $p_B(p_A)$ in (p_A, p_B) space, both of which will be positively sloped at least at the point where they cross (the equlibrium).



Now, consider how an increase in $\bar{k}$ will affect equilibrium p_A and p_B. Total differentiation of equation (3.A.1) gives

$$(3.A.3) \qquad \frac{\partial p_i}{\partial \bar{k}} = \frac{2\theta(2\lambda\theta g_i - p_i + p_j + t)}{4\lambda\theta g_i + 2\theta\bar{k} - 3p_i + p_j + t}$$

for given p_j. This is certainly positive. Thus, if g_A and g_B remain the same, an increase in $\bar{k}$ will shift both $p_A(p_B)$ and $p_B(p_A)$ outwards, and both prices will increase in equilibrium. More formally,

$$\frac{dp_i}{d\bar{k}} = \frac{\partial p_i}{\partial \bar{k}} + \frac{dp_i}{dp_j} \cdot \frac{\partial p_j}{\partial \bar{k}}$$

which, using the fact that $g_i \equiv -g_j$, comes out as

$$\frac{dp_i}{d\bar{k}} = \frac{4\theta[(2\lambda\theta g_i - p_i + p_j + t)(2\lambda\theta g_j + p_i - p_j + t) + t(2\bar{k}\theta - p_i - p_j - t)]}{(4\lambda\theta g_i + 2\theta\bar{k} - 3p_i + p_j + t)(4\lambda\theta g_j + 2\theta\bar{k} + p_i - 3p_j + t)}$$

which is definitely positive and, indeed, greater than $\partial p_i/\partial \bar{k}$.

Now consider an increase g_i. Total differentiation of equation (3.A.1) gives

$$(3.A.4) \qquad \frac{\partial p_i}{\partial g_i} = \frac{4\lambda\theta(\theta g_i + \theta\bar{k} - p_i)}{4\lambda\theta g_i + 2\theta\bar{k} - 3p_i + p_j + t}$$

for given p_j, which is positive as long as $\lambda > 0$ (i.e. as long as there are at least some informed consumers). Note that this expression is also identical to dp_i/dg_i, since

$$\frac{dp_i}{dg_i} = \frac{\partial p_i}{\partial g_i} + \frac{dp_i}{dp_j} \cdot \frac{\partial p_j}{\partial g_i}$$

and g_i does not appear as an argument in p_j.

Simple subtraction can now give us $(p_{\bar{u}} - p_{\bar{k}})$. When $\lambda = 0$, $(p_{\bar{u}} - p_{\bar{k}})$ is exactly equal to $- p_{\bar{k}}$. Thus, when no consumers are informed about quality $(p_{\bar{u}} - p_{\bar{k}})$ is negative, as expected. When $\lambda = 1$, we get

$$(3.A.5) \qquad p_{\bar{u}} - p_{\bar{k}} = \frac{4\theta(\theta g_j + \theta\bar{k} - p_j)(2\theta\bar{k} - p_i - p_j - t)}{(4\theta g_i + 2\theta\bar{k} - 3p_i + p_j + t)(4\theta g_j + 2\theta\bar{k} + p_i - 3p_j + t)}$$

which is definitely positive. Thus, when all consumers are informed about quality, $(p_{\bar{u}} - p_{\bar{k}})$ is positive, as is also expected.

This model is very simplistic. Most importantly, the additive separability of $u(p, k)$ means that $p_{\bar{k}x}$ and $p_{\bar{u}x}$ must be zero. However, it does show clearly how the two externalities described above might interact.

NOTATION - chapter 4

A	lowest quality that can enter for J = 0
B	highest quality that can enter for J = 0
C	capital ϵ $[0, \bar{C}]$
c_H	cost of high quality
c_L	cost of low quality
δ	discount rate
$F(\theta)$	fraud function
$[G, H]$	quality interval supplied by SRO
I	quality interval supplied by external regulator
J	punishment for cheating
$\hat{J}(X)$	constrained version of $J(\theta, X)$
λ	weighting parameter (for theorem 4.3.3)
$\mu(\theta)$	cost function
N	maximum number of sellers
n	number of sellers regulated
$p(\theta, X)$	price function
p_H	price for high quality
p_L	price for low quality
Π	Total industry profits
$\pi(\theta, X)$	profit function for honest sellers
$\pi_d(\theta, X)$	profit function for dishonest sellers
Q	compensation paid to cheated consumers
$\hat{Q}(X)$	constrained version of $Q(\theta, X)$
r	rate of interest
S	subscription fee (per period)
σ	shift parameter for constraint (4.3.2')
T	length of punishment
t	time
θ	quality ϵ $[\underline{\theta}, \bar{\theta}]$
X	actual number of sellers (= output)
$x(\theta, \sigma)$	number of sellers of quality θ
w	quality observation lag
Z	consumer surplus from transacting

4. SELF-REGULATION UNDER MORAL HAZARD

4.1 Introduction

Chapter 3 examined the quality problems which can arise when (i) the price that each seller receives depends partly on the qualities of other sellers (and not just on that seller's own quality) and when (ii) the marginal effect of quality on price varies with total output. In that chapter, quality was modelled as irreversible (or 'not easily reversible'), but the problems examined may also be present when quality is reversible. There is, however, a potentially more serious problem in the reversible quality case: moral hazard. In this chapter, I abstract from the quality problems of chapter 3 so as to concentrate on the pure moral hazard problem. In reality, though, it should be remembered that these various problems may co-exist.

The problem of moral hazard was described in section 1.3.1, and was analysed in more detail in sections 2.4 and 2.5. Essentially, it is assumed that

— Consumers cannot see sellers' qualities prior to purchase (although they may be able to see them at some point after purchase);

— Sellers' quality choices are fully reversible, so that past quality performance is (in itself) no proof of future quality;

— At the time of purchase, consumers know exactly what quality each seller is offering, but these offers need not correspond to the actual qualities that would be received;

— Consumers prefer higher qualities to lower qualities, but higher qualities are more costly to produce than are lower qualities. Overall, the benefits of high quality outweigh the costs.

In a simple one-shot game, if a consumer believes that a seller's good is of a high quality and decides to buy it, then that seller will have an incentive to cheat the consumer by producing a low quality (since lower quality is cheaper to produce). A rational consumer, though, will deduce that this will happen, so he will only believe a low quality offer in the first place. Thus, the one-shot Nash equilibrium involves trade of low quality goods only, even when both buyer and seller would be better off (under full information) trading in higher quality goods.

This one-shot result can be avoided in a repeated game. For example, Shapiro (1983) shows that if consumers can refuse to buy from a seller that has ever cheated on its quality offer, then there will exist some price at which the future stream of profits received by a non-cheating seller will be just high enough to give that seller sufficient incentives not to cheat (although this price may equal infinity, for particularly bad moral hazard). This is a form of reputation model.

As was shown in section 2.4, one problem with such a reputation model is that, for any high quality to be credible, its price must exceed its marginal cost. This has two effects:

— An 'income' effect: as prices increase for given qualities, consumers will tend to buy *fewer (or even zero) goods* of each quality; and

— A 'substitution' effect: as the slope of the price-quality budget line increases (pivots upwards from the locus along which price = marginal cost), consumers will tend to buy *lower qualities* than is efficient.

These two effects mean that equilibrium in the reputation model will be inefficient compared with the full information free entry equilibrium.

This chapter approaches the idea of reputation slightly differently. Instead of allowing price to be endogenous (determined by the reputation constraint and free entry/Bertrand competition), price is taken to be exogenous in this model (except for the brief discussion in appendix 4.A2).

When price is exogenous, the income and substitution effects of Shapiro's model do not occur. Instead, inefficiency exists if there are sellers, who *would* enter the industry and make positive profits under full information, but who are prevented from entering the industry at the reputation equilibrium.

In a *laissez faire* reputation model (with or without price exogeneity), the degree of inefficiency which results will depend, amongst other things, on:

(a) How much sellers have to gain from cheating;

(b) How quickly cheatings are detected;

(c) How well information about cheatings spreads amongst potential consumers;

(d) How much sellers can be punished once cheating has been detected;

(e) Sellers' discount rates.

When the inefficiency resulting from these effects under *laissez faire* is high, then there may be a role for market intervention/regulation. Regulation can work by changing any of (a) to (d) above (factor (e) is more difficult to affect). More specifically,

(a) Gains from cheating can be reduced (e.g. via higher minimum quality or qualification levels);

(b) The speed of detection can be increased (e.g. by monitoring, and by educating consumers about what to expect of sellers);

(c) The transmission of information about instances of cheating can be improved (either explicitly via better access to the media, or implicitly via a formal complaints procedure combined with a licensing scheme (this is the information collation role for regulation));

(d) Sanctions against cheating sellers can be increased (consumers can only threaten future nonpurchase, but a regulator may also be able to impose fines and losses of tied-up capital (this is the penalization role for regulation));

The discussion of Shapiro's (1983) model in section 2.4 looked at (a) raising minimum quality levels and (b) increasing the speed of detection of cheating. Factors (c) and (d) were ignored. It was simply assumed that (c) transmission of information between consumers about instances of cheating is perfect[1] and (d) cheating sellers face future non-purchase but no further punishment.

The model in this chaper concentrates on factors (c) and (d). That is, it focusses on two roles for regulation: the *information collation* role and the *penalization* role. It then examines whether (and when) self-regulation will be a feasible/optimal form of regulation. The model in this chapter can also be seen as a response to Mayer and Neven (1991). Their model deals with (c) the use of punishments for cheating sellers, but was shown (in section 2.5) to be unnattractive in some ways.

In the following model, two separate parts of effective regulation are distinguished:

(i) Determination of the regulatory rules; and

(ii) Proper application of the regulatory rules.

The determination of regulatory rules takes place in period zero of the model. These rules are formulated by some welfare-maximizing body in the case of external regulation and by the industry in the case of self-regulation. The chosen rules are then applied by a separate regulatory

---

[1] This implies that the multi-player quality choice game is formally identical to the corresponding two-player game (see section 4.2 for a more detailed discussion of this point).

body for all of the following periods. That is, the actual regulation is assumed to be at 'arms-length'. Rules must thus be formulated so as to give regulatory bodies appropriate incentives to regulate properly, taking into account any restrictions on the ability of these regulatory bodies to distinguish between different sellers.

4.1.1 Comparing different forms of regulation

Once the distinction between determination and application of regulatory rules has been made, it is easier to focus on the potential differences between external regulation and self-regulation. In this model, three of these are examined:

(1) When the regulatory body is either an SRO or a non-benevolent external regulator, it will regulate properly only if it is given sufficient incentives to do so. This issue must be taken into account at the rule formulation stage of the model (and it may well affect the optimal regulatory rules). No such incentive scheme is needed, however, for a benevolent external regulator.

(2) When quality is observable (with a lag) but non-verifiable, an external regulator, whose regulatory rules tend to be set out very formally and legalistically, may not be able to condition these rules on sellers' qualities. However, an SRO, whose rules tend to be less formal, may be able to condition its regulation on sellers' qualities. Thus, the set of rules which an external regulator can apply may be a restricted subset of those which an SRO can apply. Note, however, that an external regulator must be able to observe and verify the *existence* of a cheating (even if it cannot verify the extent of the cheating). Otherwise, external regulation is not possible at all.

(3) At the rule formation stage, external regulation will involve welfare being maximized (given constraints (1) and (2) above). For self-regulation, however, the industry will set rules so as to maximize its

own profits (in particular, rules may be chosen which restrict the number of sellers in the industry to below the welfare-optimal level).

Summary: External regulation by a benevolent regulator involves welfare-maximization given constraint (2) above; external regulation by a non-benevolent regulator involves welfare-maximization given constraints (1) and (2); and self-regulation involves industry profit-maximization given constraint (1). These differences allow us to compare the three forms of regulation explicitly.[2]

4.1.2 Layout and main conclusions of chapter

Section 4.2 examines a very simple two-player quality choice game, and shows how an increase in the punishment faced by cheating sellers can improve on the efficiency of the reputation equilibrium. This simple exposition demonstrates the penalization role of regulation.

The information collation role is then examined. Consider a multi-player version of the two-player quality choice game, in which any particular buyer and seller will meet each other only rarely . If there is perfect transmission of information between consumers about cheatings, then the multi-player quality choice game will be formally identical to the corresponding two-player game and have the same reputation-based solution.

When the transmission of information is very imperfect, however, it is unlikely that there will be a reputation solution to this multi-player game. Under these circumstances, it is shown that a formal consumer complaints procedure combined with a seller licensing system can make the industry work *as if* there is perfect information transmission. This is the *information collation* role of regulation.

---

[2] In Mayer and Neven's model, difference (1) only is examined. This leads them to conclude that self-regulation is possible in only a subset of the situations in which external regulation (by a benevolent regulator) is possible. Once differences (2) and (3) are introduced, though, it is no longer obvious which of the two forms of regulation will be superior.

Note that, in this set-up, there are essentially two games going on: a consumer/regulator game and a regulator/seller game. The usual form of reputation model consists of a consumer/seller game, which may be affected in some way by a regulator. When reputation formation is replaced by a licensing system, consumers no longer directly affect sellers' behaviour. Rather, sellers are given appropriate incentives by the regulator (which many involve increased penalties: the *penalization* role of regulation), and the regulator is given appropriate incentives by the consumers.

Sections 4.3 and 4.4 model this 'two game' set-up more explicitly and examine the regulatory rules which will be chosen under the three forms of regulation described above. In both sections it is assumed that the consumer/regulator game is played with trigger strategies. However, the two sections make different assumptions about the strategies played in the regulator/seller game, with trigger strategies being used in section 4.3 and tit-for-tat strategies being used in section 4.4. It is shown that the necessary conditions for seller honesty under tit-for-tat strategies are more stringent than the corresponding conditions under trigger strategies.

In this model, self-regulation is certainly a *feasible* form of regulation. It is also shown that self-regulation becomes more likely to be welfare-superior to external regulation as fraud becomes relatively more important, as sellers become more *homogeneous* in terms of capital, as the observation lag becomes greater, and as the negative effect of output on price becomes smaller.

4.2 The quality choice game with moral hazard

4.2.1 The simple two-player quality choice game

Appendix 4.A1 describes a very general two-player moral hazard game. Let δ denote the discrete time discount rate. A trigger strategy involves both players playing Honest until one player cheats, after which both play Cheat for evermore. A tit-for-tat strategy, on the other hand, involves player i

playing Cheat in period t if and only if player i played Honest and player j played Cheat in period t-1.

The appendix shows how repetition of a one-shot moral hazard game can achieve the (Honest, Honest) outcome, under either of these strategies, if δ is sufficiently high. It also shows that the condition on δ for subgame perfect equilibrium (SPE) with tit-for-tat strategies will be at least as stringent as the condition on δ for SPE with trigger strategies. That is, if there is an SPE in tit-for-tat strategies, then there must also be an SPE in trigger strategies, but the converse need not be true. This simple result is reflected in result 4.4.2 for my, more complex, reputation model.

A further result shown in 4.A1 is that the condition on δ for SPE will be exactly the same, for either type of strategy, if and only if the sum of the payoffs when either player is fooled (i.e. when one plays Honest while the other plays Cheat) is equal to the sum of the payoffs when neither player is fooled (i.e. both play Honest or both play Cheat).

Let us now examine the simple quality choice game below, where the seller chooses between high quality and low quality and the buyer decides whether to believe high quality (and pay p_H) or low quality (which costs p_L):

		Consumer	
		p_H	p_L
Seller	H	$p_H - c_H$	$p_L - c_H$
	L	$p_H - c_L$	$p_L - c_L$

where $p_H - c_H > p_L - c_L$

Only the seller's payoffs are shown, because only the seller faces moral hazard. In this game, it is clear that the condition on δ for SPE *will* be the same for either type of strategy (because $(p_H - c_L) + (p_L - c_H) = (p_H - c_H) + (p_L - c_L)$). Using result 4.A1.1 (see 4.A1), this condition is

$$\delta \geq \frac{c_H - c_L}{p_H - p_L}$$

Suppose that the price-cost mark-up for the lowest quality level L is zero, so that $p_L = c_L$. In this case, the above condition becomes

(4.2.1) $\quad p_H \geq (c_H - (1 - \delta)c_L)/\delta$

More generally, when there is a continuous range of available qualities q, rather than just a binary choice, and when L is the minimum quality level, this inequality turns into the price-quality schedule

(PQS) $\quad p(q) \geq (c(q) - (1 - \delta)c_L)/\delta$

as in Shapiro (1983).[3] The fact that the rhs of this inequality is greater than c(q), for all $q > L$, means that there must be a premium to producing quality in order for sellers to want to maintain high quality.

If prices are endogenous, this premium may have a detrimental effect on the efficiency of the industry: via an income effect and/or a substution effect. If, on the other hand, prices are exogenous, there is no reason to expect them to satisfy (PQS), so that sellers, who could enter the industry under full information, may be prevented from doing so. This is also inefficient. In either case, then, there may be a role for regulation to lower the required high quality price premium to more acceptable (and perhaps even efficient) levels.

Consider, for example, the effect of imposing a fine J on cheating firms, so that the normal form (binary choice) game becomes

---

[3] Note that, when there is a continuous quality range, the strategies of consumers need to be formulated more carefully than above. One possibility involves a cheated consumer reacting exactly as though he had been given quality L, even if his seller actually provided a quality only slightly lower than the quality offer. In equilibrium, a cheating seller then does best by providing quality L, and consumers expectations are rational.

		Consumer	
		p_H	p_L
Seller	H	$p_H - c_H$	$c_L - c_H$
	L	$p_H - c_L$	$-J$

For an SPE in trigger strategies, we require that

$$\delta \geq \frac{c_H - c_L}{p_H - c_L + J} \quad \text{or} \quad p_H \geq (c_H - (1 - \delta)c_L)/\delta - J$$

Thus, by setting $J = (1 - \delta)(c_H - c_L)/\delta$, a regulator can induce a desire to maintain reputation without requiring any price premium for high quality. Similar (if more complicated) results hold for tit-for-tat strategies.

In principle, then, the use of additional penalties for cheating may achieve a more efficient reputation equilibrium than the free market can, for a two-player simple moral hazard game. This is the penalization role of regulation. In the next section, a multi-player version of the above game is analysed, to see whether (and when) these results will still hold.

4.2.2 The multi-player quality choice game

When there are many consumers and many sellers in an industry, three new factors become important in determining how well the reputation mechanism will work. These are

(a) the matching rule (how individual sellers are matched up with individual buyers in each period);

(b) the degree of communication among consumers; and

(c) the incentives that non-cheated buyers have for carrying out retribution against cheating firms and the incentives that cheated buyers have for carrying out retribution against non-cheating firms.

Suppose that each seller has one unit of the good to sell in each period, and that each buyer only wants to buy one unit per period. The multi-player game is identical to the two-player game in two special cases:

(1) If there is a constant matching rule (so that the same seller always meets the same buyer);[4] and

(2) If there is perfect costless communication between buyers.[5]

4.2.3 An information collation role for regulation

Apart from these two special cases, however, the multi-player game varies from the simple two-player game. Consider a case where there is no communication amongst buyers (so that only a cheated consumer can ever take retribution against the seller who cheated him). The effect of changing

---

[4] Note that the game may change when one seller sells to several buyers, even if there is a constant matching rule. This is because these buyers may be able to get together and threaten that, if the seller ever cheats *any one* of them, then they will all, multilaterally, stop buying from him.

[5] Note that this requires a slight change in the descriptions of players' strategies. These are now

Trigger: Every player i plays Honest in every period t, unless both of the following conditions hold: (a) player i (which could be a buyer or a seller) has always played as specified by his equilibrium strategy and (ii) player i's match has, at least once, *not* played as was specified by his equilibrium strategy. When these two conditions hold in period t, player i plays Cheat (pays p_L or provides low quality).

Tit-for-Tat: Every player i plays Honest in every period t, unless both of the following conditions hold: (i) player i (which could be a buyer or a seller) made a play in period t-1 that was specified by his equilibrium strategy and (ii) player i's match did *not* make the play in period t-1 that was specified by his equilibrium strategy. When these two conditions hold in period t, player i plays Cheat (pays p_L or provides low quality).

These strategies ensure that any seller who cheats will be punished by the next buyer (or buyers, for trigger strategies), even if that buyer is not (those buyers are not) the one who was cheated.

the matching rule, in this case, is formally the same as changing the discount rate δ. A more frequent matching of a given seller and buyer can be denoted by a higher discount rate. Now, recall that for the seller to have adequate incentives to maintain quality in a free market, we need

$$(4.2.1) \quad p_H \geq (c_H - (1 - \delta)c_L)/\delta$$

If this inequality is made into an equality which can be used to find the minimum necessary p_H, it is clear that

$$\frac{dp_H}{d\delta} = \frac{c_L - c_H}{\delta^2} < 0$$

Thus, the less frequently a given buyer and seller are matched (the lower is δ), the higher is the minimum necessary quality premium. This is definitely bad for the buyer, and if p_H is greater than the buyer's reservation price (so that no purchase is made), the seller is also worse off than she would be with more frequent matching.

Now, consider an extreme case of the above situation, in which there is no communication amongst buyers, and no buyer ever meets the same seller more than once. This matching rule is called the 'Townsend Turnpike', after Townsend (1981) who suggests that one way to think of this is as two infinitely long sets - one of buyers and one of sellers - moving in opposite directions. In this case, reputation formation will be impossible in the free market, because a cheated buyer will never have the chance to punish a cheating seller by refusing to purchase from her again.

It is evident that this situation could be improved by the introduction of an institution that improves the information transmission between buyers about the behaviour of each seller. However, efficient trade does not require that every buyer and seller know the full history of every other buyer and seller. For example, for the multi-player tit-for-tat strategy game, each buyer and seller only needs to know his/her own history of

behaviour, and whether his/her current (period t) trading partner played according to his/her equilibrium strategy in period t-1.

Milgrom, North and Weingast (1990) look at a situation which has many similarities with the buyer-seller situation described above: the system of long-distance trading which built up in Europe over the 12th and 13th centuries. Prior to the revival of trade in the early middle ages, there was little or no policing of commercial activity. When trading communities grew larger, however, it became increasingly difficult for traders within each community to monitor one another's behaviour. So, as trade between different regions increased (due to higher levels of geographical specialization), there grew a need to prevent reneging by traders who might cheat in one location, never to be seen again.

Milgrom et al argue that Law Merchants, who were effectively in charge of collating information about cheating traders and passing this information on to other traders dealing with them, evolved to solve exactly these problems. Moreover, this Law Merchant solution essentially involved giving traders just enough information to enable the reputation mechanism to function. Their reasoning is just as above: In a large community, it would be too costly to keep everyone informed about what transpires in all trading relationships, as a simple reputation system might require. So the system of Law Merchants was designed to promote private resolution of disputes and otherwise to transmit *just enough* information to the right people in the right circumstances, so as to give traders sufficient incentives to maintain their reputations.

In the Milgrom et al model, traders question the Law Merchant (LM) about their new trading partner at the start of every period, and the LM tells them whether that partner has his backing. The LM will give any trader his backing unless he has received valid complaints about that trader's behaviour (and maybe even then, if the trader has already been penalized in some way). In this case, as long as the LM is considered trustworthy, the only information traders need to know is whether or not their partners have LM backing. Meanwhile, the only information an LM needs is whether he has

received a complaint about a trader in the past (and even this may not be needed if that trader was penalized).

It is straightforward to see how this information collation idea can be applied to a buyer-seller game. Suppose that there exists a regulator, which gives out licences to sellers. Suppose also that, in order to be given a licence, sellers must agree to the regulator's penalization scheme which will give sellers sufficient incentives for honesty. The regulator will remove a seller's licence if and only if (i) it has received valid complaints about that seller's behaviour and (ii) the penalization scheme dictates that the licence should be removed.

Assume now that the penalty for falsely pretending to have a licence is very high, so that no seller will claim that she has one when she doesn't, even for very low δ. This means that any seller which has a licence will have sufficient incentives to behave honestly, as long as the regulator is doing its job properly. Thus, the information requirement for consumers is greatly reduced: a consumer only needs to know (i) whether a seller is licensed and (ii) whether the regulator is trustworthy.

Moreover, because many sellers are regulated by a single regulator, consumers will have repeated dealings with that regulator (even though they never meet the same seller more than once). This means that the regulator may be able to develop a reputation for behaving properly (in the usual way), even though individual sellers cannot.

This set-up effectively places the regulator *between* consumers and sellers, so that instead of a buyer/seller quality choice game, we now have a regulator/seller quality choice game and a buyer/regulator honesty game. The following model elaborates on these ideas more formally.

4.3 A model of regulation under moral hazard

In this section, I examine the buyer/regulator and the regulator/seller games described above:

— In the buyer/regulator game, buyers play a trigger strategy. If the regulator is ever observed to act improperly, buyers will no longer trust it, and will thus no longer trust the quality assurances of its member sellers. These sellers will leave the regulator, since they no longer gain anything from membership, and the regulator will lose all of its income. If this income is high enough relative to the gains to be made from acting badly, then the regulator will behave properly.

— In the regulator/seller game, the regulator can play a trigger strategy (this section) or a tit-for-tat strategy (section 4.4). If a seller cheats on its quality offer, the cheated buyer will report this to regulator and the regulator will punish the seller. If this punishment is high enough relative to the gains to be made from cheating, then each seller will behave honestly.

Note that, for this regulation scheme to work, it is important to check that consumers have the right incentives (i) to complain if and only if cheated and (ii) to refuse to buy either from non-regulated sellers or from sellers whose regulator has been seen to act improperly.

Given the constraints on regulators and sellers derived from these games, I then examine how regulation may be set up at period zero (before these games commence). Within this model, our three potential difference between different forms of regulation become:

(1) When the regulatory body is either an SRO or a non-benevolent external regulator, it will regulate properly only if the constraints derived from the buyer/regulator game all hold. A benevolent external regulator, however, faces no such constraints.

(2) The constraints resulting from the regulator/seller game are likely to be different for sellers of different qualities, with higher punishments for cheating being imposed on those sellers who have bigger incentives to cheat. When quality is observable (with a lag) but non-verifiable, the punishments imposed upon cheating sellers by an external regulator may have to be the same for all sellers, while an SRO may be able to fine-tune punishments for individual sellers.

(3) At the rule formation stage, welfare will be maximized (given constraints (1) and (2)) under external regulation, but profits will be maximized (given constraint (1) only) under self-regulation.

It is shown that self-regulation is certainly feasible in this model (that is, there exists a set of regulation schemes which satisfy both the buyer/regulator and the regulator/seller game). These regulation schemes are characterized, and the choice of scheme under self-regulation is compared with the choice of scheme under external regulation. Conditions are then derived under which self-regulation is likely to be welfare-superior to external regulation.

4.3.1 Framework

This model bears similarities to the Mayer and Neven (1991) model (see section 2.5). Seller i has an innate quality θ_i, where θ is distributed in some way on the interval $[\underline{\theta}, \bar{\theta}]$, with $\bar{\theta} > \underline{\theta} \geq 0$. The total number of sellers (potential and actual) equals N. Each seller provides one unit of the good, or nothing, so that maximum industry supply is N.

A seller of innate quality θ has a binary quality choice to make: she can provide quality θ and receive profits $\pi(\theta, X)$, where X is total industry output, or she can cheat her consumers (by providing zero quality and perhaps also by direct fraud) and receive profits $\pi_{\sqsupset}(\theta, X)$, where $\pi_{\sqsupset}(\theta, X) > \pi(\theta, X)$, for all θ and X.

This model's assumption of innate (honest) qualities and exogenous profits is a useful simplification. One benefit is that it is allows us to abstract from all of the quality choice problems described in chaper 3 and in Shapiro (1983), so that we can concentrate on the simple moral hazard problem. A second benefit of the quality assumption is that, because each seller just has a binary choice to make, it is easy to formulate strategies for players. That is, we don't need to specify what the regulator will do if a seller cheats only slightly on her quality offer, etc.

The profit function for honest sellers $\pi(\theta, X)$ has the following characteristics:

$$\pi_x \leqslant 0, \quad \pi_{\theta x} = 0, \quad \pi_{\theta\theta} \leqslant 0, \quad \pi_{xx} = 0, \quad \text{and } \pi(\theta, N) > 0, \quad \forall\theta.$$

Note that, unlike Mayer and Neven, no assumption is made on the sign of π_θ (i.e. it may well be negative for at least some θ). This reduced form profit function is derived from a simple price function $p(\theta, X)$ [$p_x \leqslant 0$, $p_\theta \geqslant 0$, $p_{\theta x} = p_{xx} = 0$, $p_{\theta\theta} \leqslant 0$] and an (opportunity) cost function $\mu(\theta)$ [$\mu' > 0$, $\mu'' \geqslant 0$]. That is,

$$\pi(\theta, X) = p(\theta, X) - \mu(\theta)$$

Sellers have zero fixed and marginal costs of production.

Note that this exogenous price function is a special case of the price function used (and rationalized) in chapter 3. Price *can* be modelled as endogenous in this set-up (assuming either free entry, as in Shapiro (1983), or Bertrand competition), but the analysis quickly becomes very complex (see the second appendix to this chapter, 4.A2, for some results).

The profits of cheating sellers $\pi_\sqsupset(\theta, X)$ are earned entirely at the expense of consumers. This reduced form profit function is made up of two components, which derive from the two different ways in which sellers can cheat consumers:

(i) Sellers can lower quality to zero. Suppose that the costs of providing zero quality are zero and that no consumer would pay anything for such quality. The profit to be made from lowering quality from θ to zero is thus $p(\theta, X)$ (the price consumers pay for the product), whereas the profit to be made from maintaining quality θ is $p(\theta, X) - \mu(\theta)$.

(ii) Sellers can commit direct fraud (stealing consumers' cash/belongings). The amount that a seller can defraud the consumer in terms of direct stealing is denoted by $F(\theta)$ $[F' \geq 0]$.

Total profits from cheating are thus given by

$$\pi_{\Box}(\theta, X) = p(\theta, X) + F(\theta)$$

Note that $F(\theta)$ is likely to be important in industries where sellers have control of consumers' cash (such as in some areas of the financial services industry). In other industries, cheating on quality offers is likely to be the larger part of $\pi_{\Box}$.

Note that since profits for all θ and X ($\leq N$) are non-negative, welfare will improve as the number of sellers in the industry increases (for any given quality spread). Thus, any form of regulation that can increase the number of sellers in the industry, without reducing the quality spread available, will be socially beneficial.

Each seller i has capital C_i, where C differs between sellers. Various assumptions can be made about the distribution of C. Two extremes are examined below. In the first (Case I), C is distributed independently from θ and uniformly on $[0, \bar{C}]$. In the second (Case II), C is perfectly correlated with with θ.

Three components of regulation are considered, the subscription (or licence fee) per period S, the punishment for cheating J, and the amount of compensation paid to cheated consumers Q. It is argued that, under trigger

strategies, J must be paid upfront to the regulator, on joining the regulation scheme, since no seller will have an incentive to pay any punishment *after* it has cheated. Thus, J acts as a capital requirement. It is also argued that S will be chosen to be negligible (as long as the number of sellers regulated by each regulatory body is fairly large).

Lastly, unregulated sellers are assumed to face prohibitively large penalties for falsely claiming to be regulated; there is a lag of w periods before qualities become observable; and δ is the discrete time discount rate.

4.3.2 The *laissez faire* equilibrium under perfect information transmission

Under full information, all N sellers will function in the industry and will behave honestly, since $\pi(\theta, N) \geq 0$ and $\pi(\theta, N) \geq \pi(0, N)$ for all θ. When consumers cannot see sellers' qualities in advance of purchase, however, all sellers will cheat in a one-shot game, since $\pi_d(\theta, X) > \pi(\theta, X)$.

Now suppose that consumers can observe a seller's quality after a lag of w periods, and that there is then perfect information transmission between consumers. If consumers play a trigger strategy whereby no consumer will ever purchase again from a seller who has cheated, then the moral hazard problem may be partly overcome even under *laissez faire*. A seller who cheats will do best by cheating in all w periods, before she is found out. Thus, a seller will cheat if and only if she would rather earn $\pi_d(\theta, X)$ for w periods than earn $\pi(\theta, X)$ forever. That is, a seller will be honest if and only if

$$\sum_{t=0}^{w-1} \delta^t \pi_d(\theta, X) \leq \sum_{t=0}^{\infty} \delta^t \pi(\theta, X)$$

which can be calculated and rearranged to give

$$\pi(\theta, X) \geq (1 - \delta^w)\pi_d(\theta, X) \qquad (4.3.1)$$

If $\delta = 0$ or if $w = \infty$, then $(1 - \delta^{w}) = 1$ and so condition (4.3.1) will not hold for any θ or X, since $\pi_{d}(\theta, X) > \pi(\theta, X)$. That is, when quality is never observed, or if sellers do not care about the future, then equilibrium is as in the one-shot game. On the other hand, if $w = 0$ or if $\delta = 1$, then $(1 - \delta^{w}) = 0$ so that condition (4.3.1) will hold for all θ and X (and the market will be efficient). That is, if quality is observed immediately, or if sellers care about the future just as much as the present, then equilibrium is as under full information.

Lying between these two extremes, however, $1 > (1 - \delta^{w}) > 0$, and so there may exist θ and X such that (4.3.1) does not hold. In these cases, there may be role for a regulator to impose further punishments on cheating sellers (the penalization role of regulation).

Moreover, even if a *laissez faire* reputation equilibrium does exist, this equilibrium might not be the *unique* Nash Equilibrium of the seller/consumer game. When no compensation is given to cheated consumers, a consumer who buys from a cheating seller loses π_{d}, while a consumer who buys from an honest seller is assumed to receive some consumer surplus from the transaction (denoted Z, where $Z > 0$). Thus, the above extensive game can be simplified to the following normal form game:

		Consumer	
		Buy	Don't buy
Seller	Honest	$\frac{\pi}{1 - \delta}$, Z	0 , 0
	Cheat	$\frac{1 - \delta^{w}}{1 - \delta}\pi_{d}$, $-\pi_{d}$	0 , 0

When condition (4.3.1) holds, {Honest, Buy} is one Nash Equilibrium of this game, but since $-\pi_{d} < 0$, {Cheat, Don't buy} is also a Nash Equilibrium.

One way of ensuring that {Honest, Buy} is the unique Nash Equilibrium is to commit to giving cheated consumers an amount of compensation Q, where $Q > \pi_d$. However, it is not possible for individual sellers, acting on their own, to credibly commit to such a compensation scheme, since *ex post* a cheating seller has no incentive to pay the compensation. The need for a credible compensation scheme thus provides a further role for a regulator.

Finally, when information transmission is *not* perfect, the seller may continue to be able to find uninformed consumers who will buy from her, even after she has cheated other consumers. In this case, she will have reduced incentives not to cheat, and there is a role for a regulator to act as a collator of information, to prevent any consumers from buying from a seller which has cheated.

Thus, under moral hazard, regulation may be necessary (superior to the *laissez faire*) for various reasons. The following section examines one plausible form which such regulation might take.

4.3.3 The regulator/seller game

Consider first the punishment for cheating J (which may or may not differ for different sellers). Under trigger strategies, a seller who has cheated will be ejected from the regulatory system (which effectively means ejection from the industry). At this point, it will be difficult to force a cheating seller to pay J (at least without the aid of the legal system). Thus, in this model, it is assumed that every regulated seller must give the regulatory body an initial payment or 'capital requirement', equal to J, at period zero. If a seller is ever found to be cheating, then not only will she be ejected from the industry, but she will also forfeit her initial payment J, which is appropriated by the regulator.

If this threat to appropriate J is to have any force in changing a seller's behaviour, it is further assumed that if any seller chooses to terminate the contract voluntarily, then J will be returned to her (although she will

have to wait w periods before getting it, so that the regulator can be sure that no cheating has occurred).

Each seller also pays the regulator a subscription fee S in each period that she remains regulated (the need for this will be made clear later). Under this set-up, a seller of quality θ will be honest if and only if she would rather earn $\pi(\theta, X) - S$ forever than earn $\pi_{cd}(\theta, X) - S$ for w periods and then lose J.[6] That is, a seller will be honest if and only if

$$\sum_{t=0}^{w-1} \delta^t [\pi_{\sqsupset}(\theta, X) - S] - J \leq \sum_{t=0}^{\infty} \delta^t [\pi(\theta, X) - S]$$

This simplifies out to the seller incentive constraint (SIC)

$$\text{(SIC)} \qquad \delta^w S \leq \pi(\theta, X) - (1 - \delta^w)\pi_{\sqsupset}(\theta, X) + (1 - \delta)J$$

which will be identical to (4.3.1) whenever S and J are zero.

The seller participation constraint (SPC) for each period is

$$\text{(SPC)} \qquad S \leq \pi(\theta, X) \qquad \text{and} \qquad C \geq J$$

Let us turn now to the regulator side of this game. One reason for having a subscription fee is that, without it, if the regulator is successful in giving sellers incentives not to cheat, then it will never appropriate any seller's J. Thus, the regulator will not have any incentive to participate in the game in the first place. Mayer and Neven get around this problem by making the unsatisfactory assumption that there is a small possibility of penalizing honest sellers. Suppose, though, that we simply allow the

---

[6] Note that, since all the terms in this model are analysed using their discounted present values, the seller will actually lose her fund multiplied by δ^{w-1}. However, it is assumed that the regulator lets sellers' funds collect interest (and that $\delta \equiv 1/(1 + r)$, so that by period w-1, the fund will actually be J/δ^{w-1}, the pdv of which is clearly J.

regulator to charge a subscription fee S in each period, as payment for its regulatory function. As long as S is slightly bigger than its per seller costs of regulation (which, for simplicity, are assumed to be zero in this model), the regulator's participation constraint (RPC) will be satisfied.

A second reason for having a subscription fee is that it helps to provide the regulator with incentives to behave properly. Some of these incentives are provided by buyers in the buyer/regulator game. However, sellers also play a role. Suppose that the regulator regulates n sellers, and that sellers play a trigger strategy, so that if the regulator is observed to treat any seller unfairly, all the sellers will leave the regulator, cease their subscriptions S, and take their capital (their Js) with them.

Now consider the possibility of a regulator accusing one seller of cheating (so as to appropriate her J), when she has not. It seems reasonable to assume that all of the n sellers can tell if a regulator appropriates some seller's J without paying out the related compensation, even if they cannot observe whether an actual cheating ever happened. Because of the sellers' trigger strategy, the regulator will behave in this way if and only if it would rather receive J now than receive (n - 1)S forevermore. That is, the regulator will behave properly in this regard if and only if the following regulator incentive constraint (RIC) holds.

$$(RIC_1) \qquad \frac{(n-1)S}{1-\delta} \geq J$$

This does not, however, prevent the possibility of a regulator accusing a non-cheating seller of cheating (so as to appropriate her J) *and* then paying out the relevent compensation (Q each) to w lucky consumers (note that the present discounted value of total compensation paid is $(1 - \delta^w)Q/(1 - \delta)$). Suppose that this action is only observed by the particular seller affected. A regulator which does this will no longer receive that seller's S, and so it will behave properly in this regard if and only if

$$(RIC_2) \qquad \frac{S}{1-\delta} \geq J - \frac{(1-\delta^{\omega})Q}{1-\delta}$$

If the total required compensation is greater than J, then this condition will hold for all non-negative S.

4.3.4 The buyer/regulator game

Let us now examine the buyer/regulator game. Given that there is an honesty equilibrium in the regulator/seller game, any consumer who wants to know whether to expect honesty from a particular seller only needs to know (i) whether the seller he faces is regulated and (ii) whether the regulator is regulating properly.

The buyer/regulator game described below is concerned with part (ii) of this condition. Suppose that there *is* perfect information transmission about any cheating by a regulator. This is a reasonable assumption, since many consumers will deal with the same regulator and also because there is likely to be media/government interest in dishonest regulators.

Consumers are then assumed to play a trigger strategy: no consumer will ever purchase again from any seller whose regulator has been seen to behave badly in any way (e.g. by failing to punish cheating sellers or by failing to compensate cheated customers). These (n - 1) sellers will then terminate their dealings with the regulator, cease their payments S and take back their funds J.

Under this set-up, the regulator will admit to a cheating and pay out the appropriate compensation if and only if it would rather get (n - 1)S forevermore than save itself $(1 - \delta^{\omega})Q/(1 - \delta) - J$. That is, a regulator will behave properly in this regard if and only if

$$(RIC_3) \qquad \frac{(n-1)S}{1-\delta} + J - \frac{(1-\delta^{\omega})Q}{1-\delta} \geq 0$$

Finally, we need to check that consumers face appropriate incentives. Firstly, we need to check that they will complain to the regulator if and only if they are cheated. This will be true as long as there is some small cost involved in complaining to a regulator (which prevents false complaints), but this cost is smaller than the compensation received if the complaint is valid (which ensures true complaints).

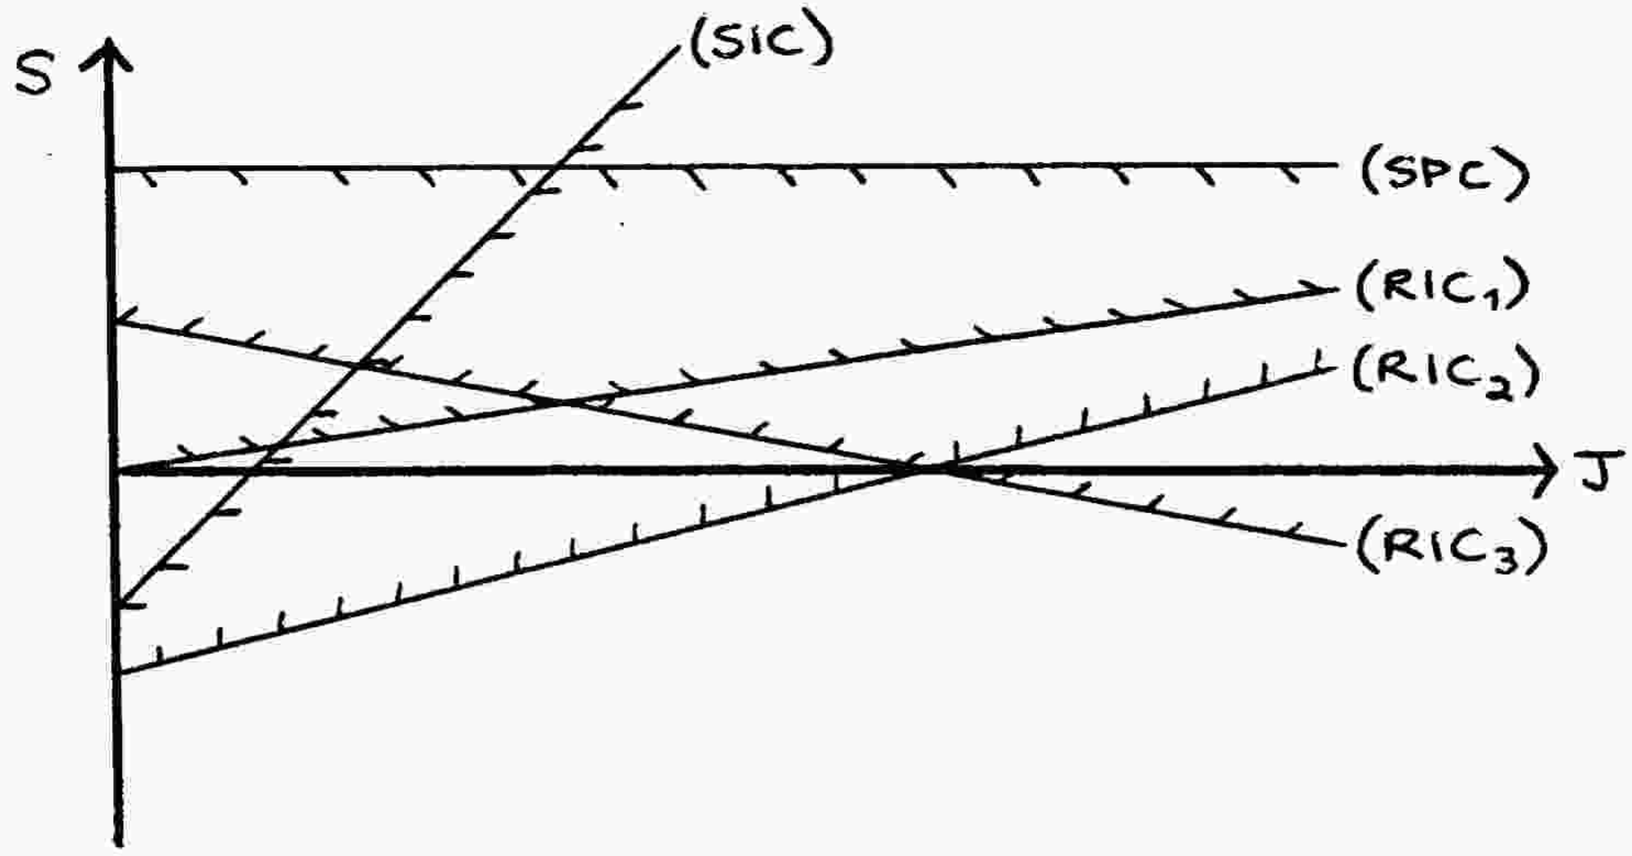
Secondly, we need to check that a consumer will purchase from a seller if and only if she is regulated by a trustworthy regulator. It is clear that, at equilibrium, no consumer will purchase from a seller whose regulator has been shown to act improperly, since that seller will definitely cheat, given the strategies being played. It is less clear that consumers *will* purchase from an honest seller at equilbrium, since there is again a non-uniqueness of equilibrium problem.

This problem was described above (in section 4.3.2). It was shown that, in order for {Honest, Buy} to be the unique subgame perfect equilibrium of this game, consumers must receive (overall) a positive utility from purchasing, even when cheated. Thus, it must be the case that the following consumer purchase condition (CPC) holds.

(CPC) $\quad Q > \pi_{\square}(\theta, X)$

4.3.5 A characterization of the regulatory system

In this section, the above constraints are used to describe some aspects of the expected regulation scheme. The first point to make is that, when the regulatory rules are being formulated (either by government or by the industry), S will be set as low as possible (for given J). This is because neither the public good nor industry profits will be improved by allowing the regulator to make profits. Assume that (CPC holds). The diagram below shows the other constraints faced by sellers and regulator for low n. The binding constraint on S, for given J, is clearly ($RIC_{\ni}$). The exact choices of S and J are not obvious, though, since there is a trade-off between the two along ($RIC_{\ni}$).



Now, as n increases, constraints (RIC$_3$) and (RIC$_1$) both converge upon the S = 0 axis. That is, as n increases, the amount S that each seller has to pay the regulator, for given J, decreases. This is because, as n increases, the regulator has more subscriptions to lose from behaving improperly. Thus, the more sellers that the regulator deals with, the better it will be for all the other sellers engaged in contracts with the regulator. One can view this as an argument for saying that there should be one regulatory body for a large number of sellers, rather than lots of small regulatory bodies (Shaked and Sutton (1982) provide another argument, see section 2.6).

In the limit as n gets large, the periodic payment S that each seller has to pay will be negligibly greater than zero. In this case, constraints (RIC$_1$) and (RIC$_3$) will be easiily satisfied and (SIC), (SPC) and (RIC$_2$) become

$$\text{(SIC')} \qquad J(\theta, X) \geq \frac{(1 - \delta^w)\pi_3(\theta, X) - \pi(\theta, X)}{(1 - \delta)}$$

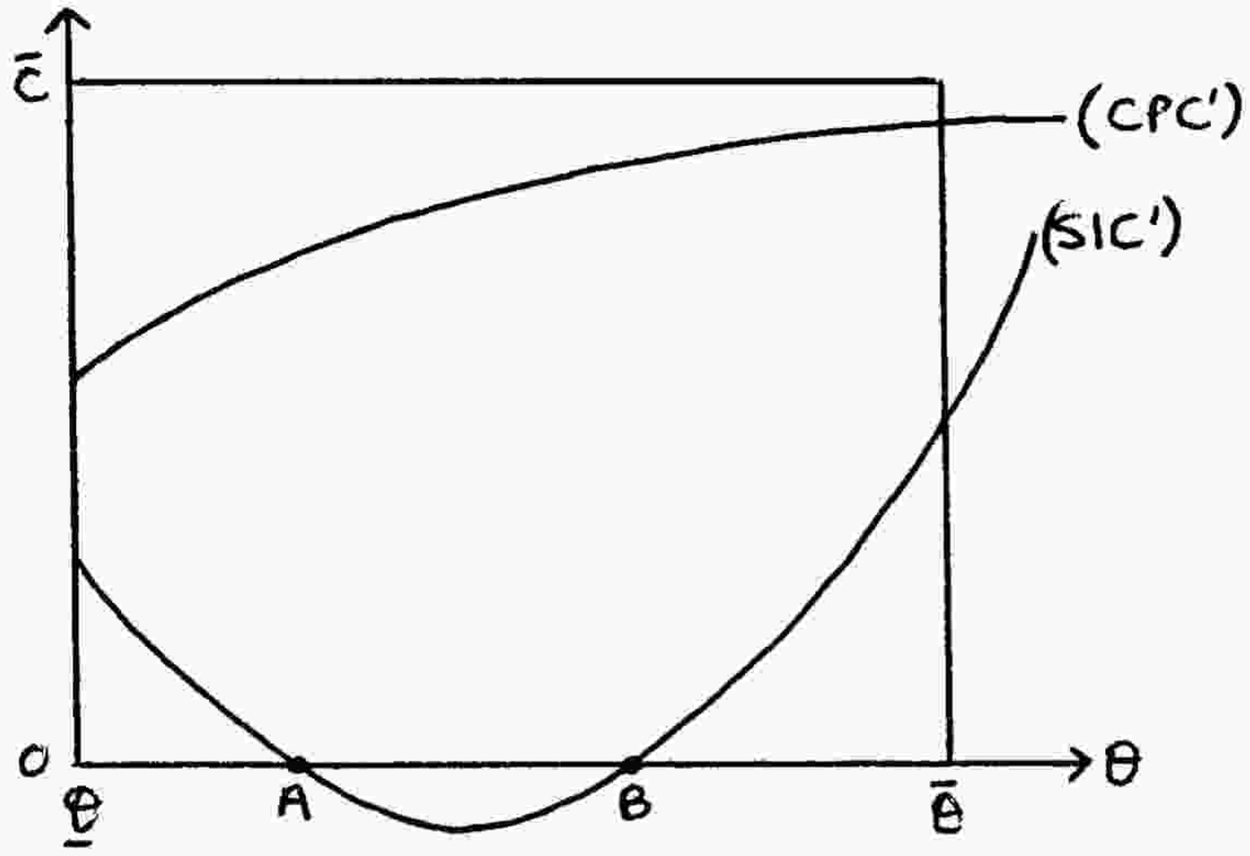
$$\text{(SPC')} \qquad C \geq J(\theta, X)$$

$$\text{(RIC}_2\text{')} \qquad J(\theta, X) \leq \frac{(1 - \delta^w)Q(\theta, X)}{(1 - \delta)}$$

Also, (CPC) can be rewritten

$$\text{(CPC')} \qquad \frac{(1 - \delta^{\omega})Q(\theta, X)}{(1 - \delta)} \geq \frac{(1 - \delta^{\omega})\pi_{\sqsupset}(\theta, X)}{(1 - \delta)}$$

Constraints (SIC') and (CPC') are drawn in (C, θ) space below.



Now, designing the regulatory rules involves setting $J(\theta, X)$ and $Q(\theta, X)$. In this diagram, the sellers which enter will be those sellers with θ and C such that, for given X,

(a) $J(\theta, X)$ lies above constraint (SIC') and below $(1 - \delta^{\omega})Q(\theta, X)/(1 - \delta)$ (constraint (RIC_2'));

(b) $(1 - \delta^{\omega})Q(\theta, X)/(1 - \delta)$ lies above constraint (CPC'); and

(c) $C \geq J(\theta, X)$.

The importance of constraint (SIC') in the above diagram is clear: A seller with quality θ and capital C will only be able to enter the industry if it lies above $J(\theta, X)$, and $J(\theta, X)$ must lie on or above constraint (SIC'). Thus, no seller which lies below constraint (SIC') will ever be able to enter the industry, even under optimal regulation.

The intuition for this is straightforward. The height of constraint (SIC') shows the punishment J that is necessary to ensure honesty by sellers of each quality θ. That is, it corresponds the amount that a θ type seller has to gain by cheating. Those sellers that have the most gain (either because $\pi(\theta, X)$ is relatively low, or because $\pi_d(\theta, X)$ is relatively high) will face the highest punishments, and vice versa. Any seller who has less capital than this required $J(\theta, X)$ will be unable to enter the industry, because she will be unable to commit to an appropriate incentive scheme.

Note that, for sellers with θs such that constraint (SIC') falls below the C = 0 line (i.e. θs between A and B on the above diagram), punishments are not needed to ensure honesty. For these sellers, therefore, the important role played by the regulator is the collation of information (and also the running of the compensation scheme), rather than the penalization role.

The importance of constraint (CPC') is less obvious. Remember that (CPC') essentially shows, for each θ and X, the total losses sustained by cheated buyers. The total compensation paid must be at least slightly greater than these losses to avoid a nonuniqueness of equilibrium. However, consider the implications of total compensation being significantly greater than total losses for some θs. The possibility of consumers gaining from cheatings may change consumers purchasing patterns (even if cheatings are unlikely to happen at equilibrium). A consumer may be willing to buy slightly different qualities or slightly greater quantities than otherwise, if he perceives some chance of making a gain (in terms of compensation minus losses) from the move.

Thus, it will be optimal from a demand side point of view if total compensation is negligibly bigger than total losses (so that $(1 - \delta^w)Q/(1 - \delta)$ lies an insignificant amount above constraint (CPC')). However, from constraint (RIC_{π}'), $J(\theta, X)$ must lie below $(1 - \delta^w)Q/(1 - \delta)$. Thus, *either* $J(\theta, X)$ must lie below (CPC') *or* total compensation must be significantly larger than total losses for at least some θs.

We can derive two results from these constraints. First:

Result 4.3.1: For each θ and X, there exist $J(\theta, X)$ and $Q(\theta, X)$ such that constraints (SIC'), (RIC_2'), and (CPC') all hold. That is there exist $J(\theta, X)$ and $Q(\theta, X)$ such that $Q(\theta, X) > \pi_{\square}(\theta, X)$ and

$$\frac{(1 - \delta^{\omega})Q(\theta, X)}{(1 - \delta)} \geq J(\theta, X) \geq \frac{(1 - \delta^{\omega})\pi_{\square}(\theta, X) - \pi(\theta, X)}{(1 - \delta)}$$

Moreover, the gap between constraints (SIC') and (CPC') is independent of $F(\theta)$, the level of potential fraud.

Proof: Let $Q(\theta, X)$ be infinitessimally larger than $\pi_{\square}(\theta, X)$. In this case the above inequality becomes

$$\frac{(1 - \delta^{\omega})\pi_{\square}(\theta, X)}{(1 - \delta)} \geq J(\theta, X) \geq \frac{(1 - \delta^{\omega})\pi_{\square}(\theta, X) - \pi(\theta, X)}{(1 - \delta)}$$

Clearly the final term in the above inequality is just equal to the first term minus $\pi(\theta, X)/(1 - \delta)$. Thus, since $\pi(\theta, X)$ is always positive, the first term is strictly larger than the final term, which means that there must be potential values for $J(\theta, X)$ which lie between the two terms. Moreover, since $F(\theta)$ only appears in the $\pi_{\square}(\theta, X)$ term, the gap between the two constraints will be independent of $F(\theta)$ (that is, any change in $F(\theta)$ will affect both constraints identically). □

The main implication of this result is that, if regulatory rules can be conditioned on sellers' qualities, then regulatory rules exist such that all sellers can enter the industry - except those lying below constraint (SIC') - and such that total compensation can be negligibly larger than total losses. That is, *there does exist a set of feasible regulation schemes.*

The second result examines how fraud and cheating on quality each affect this set of feasible schemes. First of all, notice that constraint (SIC') can be expanded to give

$$J(\theta, X) \geq \frac{(1 - \delta^{w})F(\theta) - \delta^{w}p(\theta, X) + \mu(\theta)}{(1 - \delta)}$$

Now, consider a special case, in which $F' = 0$ (so that the level of potential fraud is the same for all θ). Assume also that $\mu(\underline{\theta})$ is zero. Now, compare two polar cases: When $F = 0$ and $\mu(\bar{\theta}) > \delta^{w}p(\bar{\theta}, X)$, it is clear that constraint (SIC') will only become positive for high θ sellers. That is, when fraud is not a problem but cheating on quality is, it is the high quality sellers that will need the highest levels of regulation. On the other hand, assume now that $\mu(\theta) = 0$, for all θ, and that $(1 - \delta^{w})F \geq \delta^{w}p(\theta, X)$. In this case, constraint (SIC') will always be downward sloping and will be highest for low θ. That is, when cheating on quality is not a problem, but fraud is, it is the low quality sellers that will need the highest levels of regulation.

More generally, we get:

Result 4.3.2: Suppose that potential fraud $F(\theta) = \lambda\hat{F}(\theta)$ and that costs of quality $\mu(\theta) = (1 - \lambda)\hat{\mu}(\theta)$. That is, for given $\hat{F}(\theta)$ and $\hat{\mu}(\theta)$, an increase in λ implies that the threat of fraud becomes more important relative to the threat of cheating on quality. Suppose also that $\hat{\mu}(\bar{\theta}) > (1 - \delta^{w})\hat{F}(\bar{\theta})$ and that $(1 - \delta^{w})\hat{F}(\underline{\theta}) > \hat{\mu}(\underline{\theta})$. In this case, an increase in λ will increase the required regulation (raise J) for low θ sellers, and will reduce the required regulation (lower J) for high θ sellers.

Proof: In this case, constraint (SIC') can be rewritten

$$\text{(SIC')} \qquad J(\theta, X) \geq \frac{\lambda(1 - \delta^{w})\hat{F} - \delta^{w}p(\theta, X) + (1 - \lambda)\hat{\mu}(\theta)}{(1 - \delta)}$$

The marginal effect of an increase in λ will cause the constraint to shift by

$$\frac{(1 - \delta^{w})\hat{F} - \hat{\mu}(\theta)}{(1 - \delta)}$$

Since $\mu(\bar{\theta}) > (1 - \delta^{w})\hat{F}(\bar{\theta}) \geq (1 - \delta^{w})\hat{F}(\underline{\theta}) > \mu(\underline{\theta})$, this will be positive for low θ sellers and negative for high θ sellers. Thus an increase in the relative importance of fraud will increase the regulation required by low θ sellers and reduce the regulation required by high θ sellers. □

Note that $\mu(\bar{\theta}) > (1 - \delta^{w})\hat{F}(\bar{\theta}) \geq (1 - \delta^{w})\hat{F}(\underline{\theta}) > \mu(\underline{\theta})$ may not hold. In this case, an increase in λ may cause regulation to increase or to decrease for all θ sellers. However it remains the case that it will increase more/decrease less for low θ sellers than for high θ sellers.

Stated less formally, this result shows that an increase in the relative importance of fraud will raise the *relative* regulation of low θ sellers vis-à-vis that of high θ sellers. This accords with empirical evidence that the relative regulation of low θ sellers tends to be high in industries where fraud is a possibility (such as some areas of financial services), but lower in more quality based industries (such as electrical services/plumbing).

4.3.6 External regulation versus self-regulation

Recall that there are three potential differences between external regulation and self-regulation in this set-up:

(1) Both a non-benevolent external regulator and an SRO will need to be given sufficient incentives to regulate properly, whereas a benevolent external regulator requires no such incentives.

(2) When quality is observable, but non-verifiable, it is possible that external regulations cannot be conditioned on sellers' qualities, while self-regulatory rules can.

(3) At the rule formation stage, external regulation will involve welfare being maximized (subject to relevant constraints), whereas a self-regulating industry will aim to maximize industry profits.

The first of these differences principally relates to constraint (RIC_2'), since constraints (RIC_1) and (RIC_3) certainly hold. Within this model, then this difference translates as:

(1) Both a non-benevolent external regulator and an SRO must satisfy constraint (RIC_2'), but a benevolent external regulator need not.

In what follows, this effect is ignored (although its implications are straightforward to trace). Instead, it is assumed that all forms of regulation set $J(\theta, X)$ without regard for its effects on $Q(\theta, X)$, and then set $Q(\theta, X)$ to be as low as it can be and still satisfy constraints (RIC_2') and (CPC'). This assumption effectively supposes that consumers' demand patterns are not greatly affected when total compensation to cheated consumers exceeds their total losses.

In what follows, therefore, we are only concerned with constraints (SIC') and (SPC'). Note that (SPC') is identical to $J(\theta, X)$. In terms of these, the second difference above can be stated thus:

(2) Under self-regulation, constraint (SPC') can be of any shape (and, in particular, it can replicate constraint (SIC') along some or all of its length); whereas under external regulation, constraint (SPC') is constrained to be linear and horizontal.

On the other hand, the third difference above is now

(3) The external regulator will set (SPC') so as to maximize welfare, subject to this horizontality constraint; whereas under self-regulation, the industry, which has free choice over the shape of constraint (SPC'), will set it so as to maximize industry profits.

The question to be addressed in this subsection is whether the self-regulation situation will ever be welfare-superior to the external regulation situation.

Case I

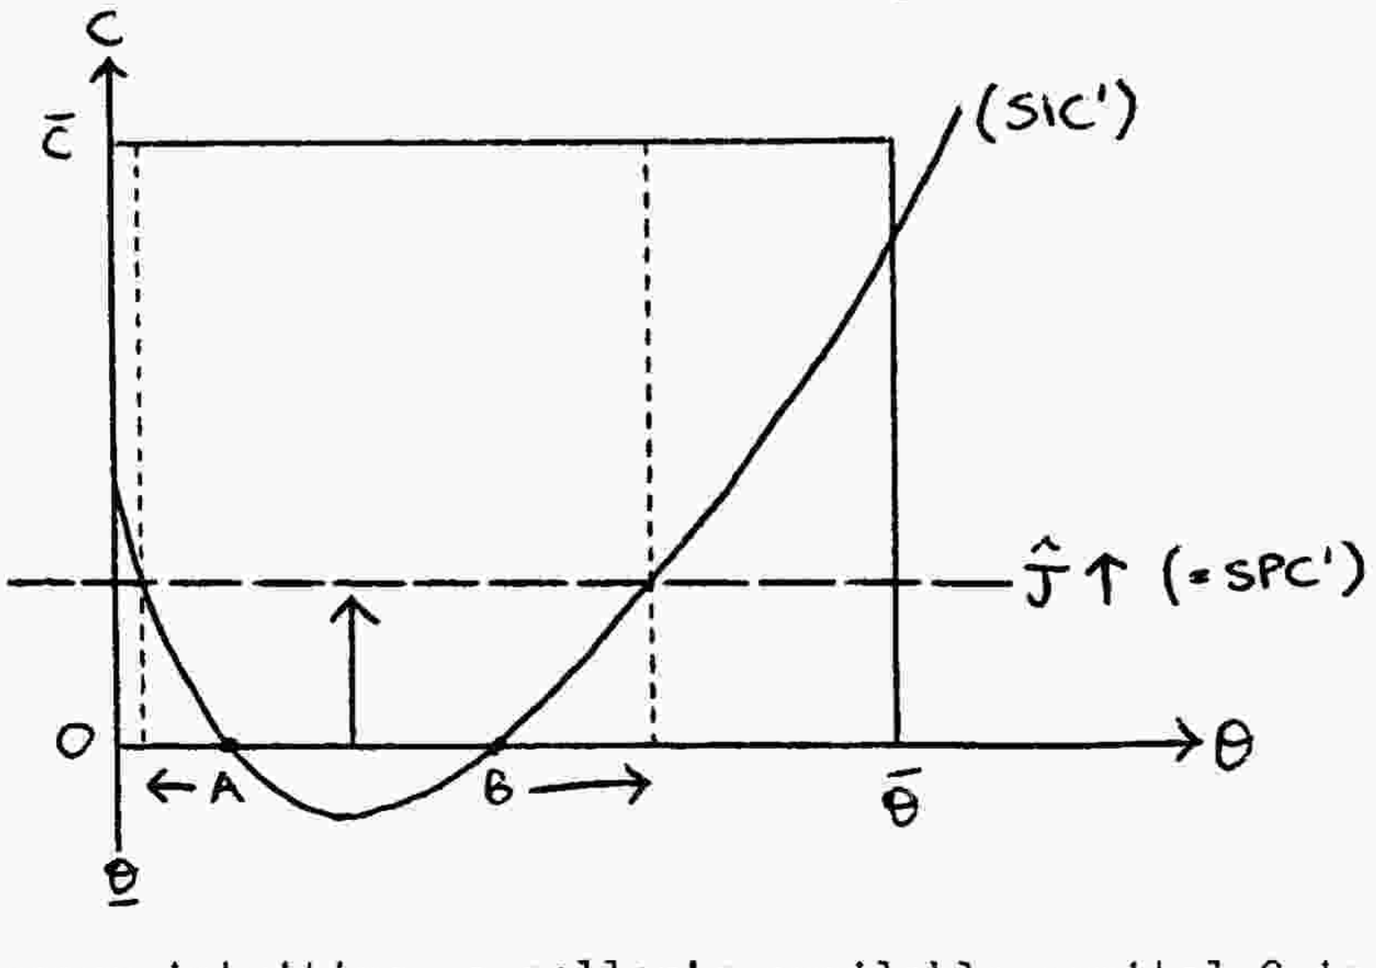
Suppose that θ is spread uniformly over $[\underline{\theta}, \bar{\theta}]$ and, initially, that C is spread uniformly on $[0, \bar{C}]$ for each θ (Case I). Let us consider the case in which constraint (SIC') is U-shaped and cuts the C = 0 axis twice between $\underline{\theta}$ and $\bar{\theta}$ (at points A and B).[7] Under external regulation, we know that $J(\theta, X)$ becomes $\hat{J}(X)$ and $Q(\theta, X)$ becomes $\hat{Q}(X)$. When $\hat{J}(X)$ is zero for all X (so that the regulator just acts as a collator of information and provider of compensation), the quality spread available will be $[A, B]$, and the total number of sellers in the industry will be $(B - A)N/(\bar{\theta} - \underline{\theta})$.

As $\hat{J}(X)$ shifts upwards from zero, there are two effects:

(1) The spread of qualities that are able to enter increases at both ends until the full range of qualities are available; but

(2) The number of sellers of each θ that can enter decreases (because they have $C \leq J$).

These two effects are obvious in the diagram below.



If, as seems intuitive, a seller's available capital C is directly related to her size, then this trade-off translates into the following result:

---

[7] Extending the following analysis to different characterisations of constraint (SIC') is straightforward.

Result 4.3.3: Consider the case in which constraint (SIC') is U-shaped and cuts the $C = 0$ axis twice between $\underline{\theta}$ and $\bar{\theta}$. As an external regulator increases $\hat{J}(X)$ from zero, a greater variety of qualities is able to enter the industry, but at the cost of smaller sellers (of all qualities) being prohibited from entry.

Thus, there is a trade-off between two factors. The socially optimal level of J will depend on the relative importance of these two factors for social welfare. Note that this trade-off seems a lot more realistic than Mayer and Neven's trade-off between allowing dishonest sellers to enter and excluding honest sellers. In this model, there may be a small chance that any seller will be dishonest, but almost all sellers entering the industry will be honest.

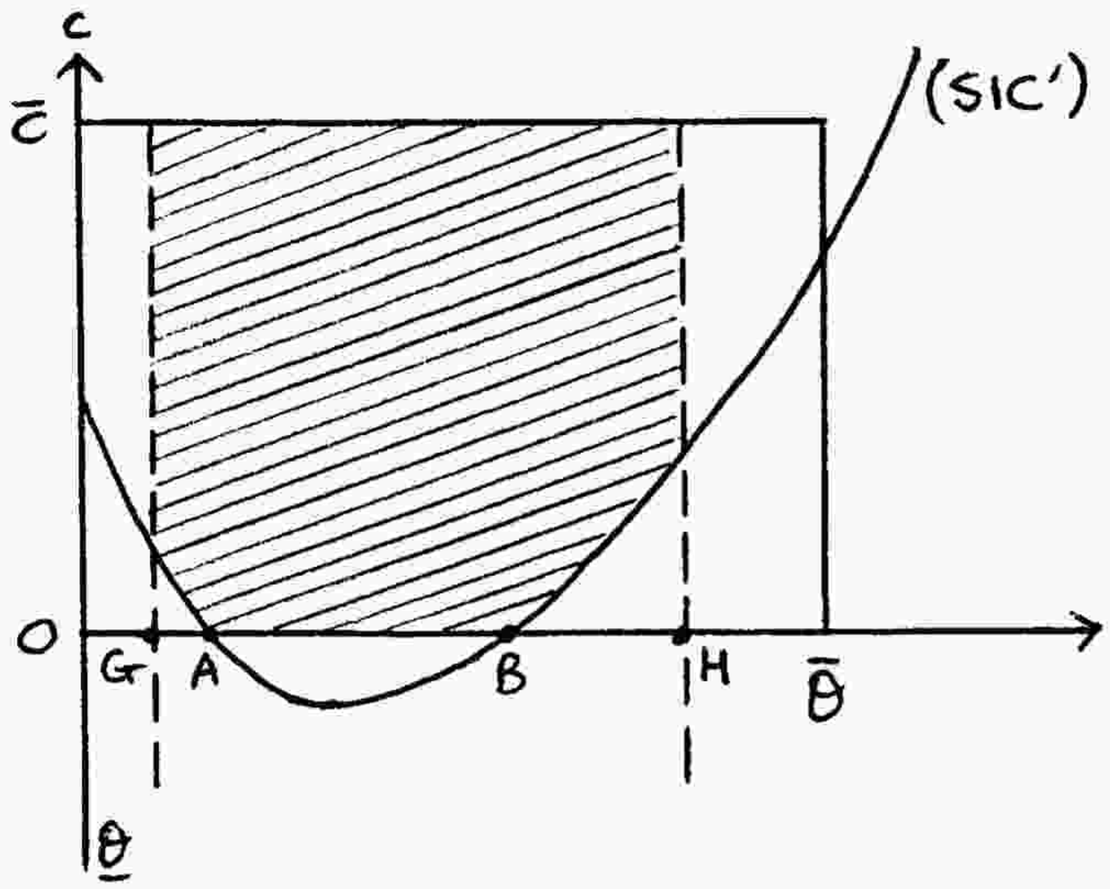
Without making more specific assumptions about the demand side of the economy, it is hard to say at which point along this trade-off welfare will be optimal. If consumers care little about quality, it seems likely that welfare will be maximized when the total number of sellers in the industry is maximized. However, an increased desire for variety (or for high and low qualities) will tend to result in variety being higher and total supply being lower, at the welfare-optimum, than this.

Turning to self-regulation, the advantage of self-regulation over external regulation, in this model, is that an SRO *can* make $J(\theta, X)$ depend on θ, since verifiability is not required for self-regulatory contracts. When C is distributed uniformly over $[0, \bar{C}]$ for every θ (Case I), the maximum amount of entry can be achieved by setting $J(\theta, X)$ such that it exactly follows constraint (SIC'). That is, X is maximized when J is set such that

$$J(\theta, X) = \frac{(1 - \delta^{\omega})\pi_{\sigma}(\theta, X) - \pi(\theta, X)}{(1 - \delta)}$$

This will achieve the maximum welfare possible (since the number of available qualities is also maximized), given that there is not full information (i.e. this gives us the second-best).

An SRO, however, may have incentives to reduce output to lower than this maximum X, since the industry price (for given quality) is decreasing in output. Since $\pi_{ex} = 0$, the θs which give the highest markups will be the same for all X. Thus, an SRO will maximize industry profits by letting in the maximum number of sellers with high mark-ups (which involves setting $J(\theta, X)$ as above for those sellers), and letting in zero sellers with low mark-ups (which involves setting $J = \bar{C}$ for those sellers). This is shown in the diagram below, which is otherwise identical to the diagram above.



As argued above, it is likely to be the high θ and low θ consumers which have the lowest mark-ups (and which the SRO will want to keep out of the industry). Thus, in this diagram, the quality spread available under self-regulation will be something like [G, H] and the total industry output will be $N/\bar{C}(\bar{\theta} - \underline{\theta})$ times the shaded area.

To compare external regulation with self-regulation (for case I), we need to compare the two situations described above. One result follows immediately.

Result 4.3.4: When profits are not decreasing in X (i.e. $\pi_x = 0$), self-regulation will certainly be welfare-superior to external regulation.

Proof: When $\pi_x = 0$ and when $J(\theta, X)$ can vary with θ, industry profits are maximized by allowing the maximum number of sellers to enter the industry (which involves setting $J(\theta, X)$ equal to constraint (SIC'), for all θ). In

this situation, since output and variety are both maximized, they must be higher than under optimal external regulation.[8] Thus, in this case, self-regulation will be welfare-superior to external regulation. □

The same result may hold when there is free entry of regulators. Consider an industry containing one SRO, which chooses $J(\theta, X)$ so as to prohibit sellers of some θs from entering the market. As long as some of these prohibited sellers have Cs which lie above constraint (SIC'), there is a role for a second SRO to enter the market and to set its $J(\theta, X)$ such that these prohibited sellers can now function in the industry (although note that there may be a minimum possible size (n) for such an SRO, which may prevent its entry). If free entry of SROs thus makes it impossible (or at least a lot harder) for one SRO to control X, SROs will act *as if* $\pi_x = 0$, which means that welfare will be maximized.

A second result is also straightforward to derive.

Result 4.3.5: When the quality interval available under optimal external regulation is a sub-interval of the quality interval [G, H] supplied under self-regulation, self-regulation will be welfare-superior to external regulation.

Proof: A self-regulating industry will supply at least as much (and usually more) of each available quality than will an externally regulated industry. Thus, even if the quality interval supplied is the same for both forms of regulation, self-regulation will be welfare-superior (because total supply is higher). If external regulation provides only a subsection of the SRO's quality interval, this welfare-superiority will be increased. □

The following result is a little more complex to prove.

---

[8] There is one exception: when constraint (SIC') is horizontal, the two forms of regulation will be identical.

Result 4.3.6: If constraint (SIC') shifts upwards, without the $\pi(\theta, X)$ function being affected, then an SRO will allow a larger quality interval to enter the industry than before (until the maximum quality interval is achieved), but an external regulator will allow a smaller quality interval than before. Thus, from result 4.3.5, it becomes more likely that self-regulation will be welfare-superior to external regulation.

Proof: Self-regulation involves $J(\theta, X)$ lying along constraint (SIC') between G and H and equalling $\bar{C}$ outside of this interval. Let $x(\theta, \sigma)$ denote the number of sellers of each quality θ that lie above constraint (SIC'), so that total output is given by

$$X(G, H, \sigma) = \int_{G}^{H} x(\theta, \sigma)\, d\theta$$

σ denotes the height of constraint (SIC'), so that $x_\sigma < 0$ and $X_\sigma < 0$. In this case, total profits under self-regulation are given by

$$\Pi(G, H, \sigma) = \int_{G}^{H} x(\theta, \sigma).\, \pi(\theta, X(G, H, \sigma))\, d\theta$$

Under self-regulation, $\Pi_G = \Pi_H = 0$, $\Pi_{GG} < 0$, and $\Pi_{HH} < 0$. Thus,

$$\frac{dG}{d\sigma} \gtrless 0 \text{ as } \Pi_{G\sigma} \quad \text{and} \quad \frac{dH}{d\sigma} \gtrless 0 \text{ as } \Pi_{H\sigma}$$

Simple differentiation gives

$$\Pi_{G\sigma} = -\, x_\sigma(G, \sigma).\, [\pi(G, X) + \pi_X X] - 2x(G, \sigma)\pi_X.\, \int_{G}^{H} x_\sigma(\theta, \sigma)\, d\theta$$

$$\Pi_{H\sigma} = x_\sigma(H, \sigma).\, [\pi(H, X) + \pi_X X] + 2x(H, \sigma)\pi_X.\, \int_{G}^{H} x_\sigma(\theta, \sigma)\, d\theta$$

At G, however, $\Pi_{G} = 0$, which implies that $\pi(G, X) + \pi_{x}X = 0$. Similarly at H, we get $\pi(H, X) + \pi_{x}X = 0$. Thus, we get

$$\frac{dG}{d\sigma} \gtrless 0 \quad - \quad 2x(G, \sigma)\pi_{x} . \int_{G}^{H} x_{\sigma}(\theta, \sigma)d\theta \quad < 0$$

$$\frac{dH}{d\sigma} \gtrless 0 \quad 2x(H, \sigma)\pi_{x} . \int_{G}^{H} x_{\sigma}(\theta, \sigma)d\theta \quad > 0$$

So an increase in σ (a shift upwards of constraint (SIC')) will lower G and raise H, enlarging the available quality interval at each end.

Turning to external regulation, let I denote the quality interval available, given $\hat{J}$, $\bar{C}$ and constraint (SIC'). Under welfare-maximization, $\hat{J}$ will be increased until the marginal loss of mid-quality sellers is just offset (in welfare-terms) by the marginal gain of sellers at either end. More formally, let I_{J} be the marginal increase in I due to an increase in $\hat{J}$ (note that I_{J} is equal to the sum of the moduluses of the slopes of constraint (SIC') at each end of I, given $\hat{J}$). Welfare-maximization involves weighing up $I_{J}(\bar{C} - \hat{J})$ against I.

Now, starting from a welfare-optimal situation, consider an upwards shift in constraint (SIC') and suppose, first of all, that $\hat{J}$ simply moves sufficiently upwards to maintain the initial I. At this point, I remains unchanged, but $(\bar{C} - \hat{J})$ has fallen. Thus, unless I_{J} has increased substantially, welfare-maximization requires a reduction in $\hat{J}$ from this point, which implies that I will be lower than previously. Thus, an upwards shift in (SIC') will reduce the quality interval available under external regulation. □

The economic implication of this result is straightforward. Constraint (SIC') may shift (relatively) upwards, without the $\pi(\theta, X)$ function being affected, for various reasons:

—— $F(\theta)$ may shift upwards;

—— $\bar{C}$ may fall, while N remains the same (so that the density of sellers at each C increases;

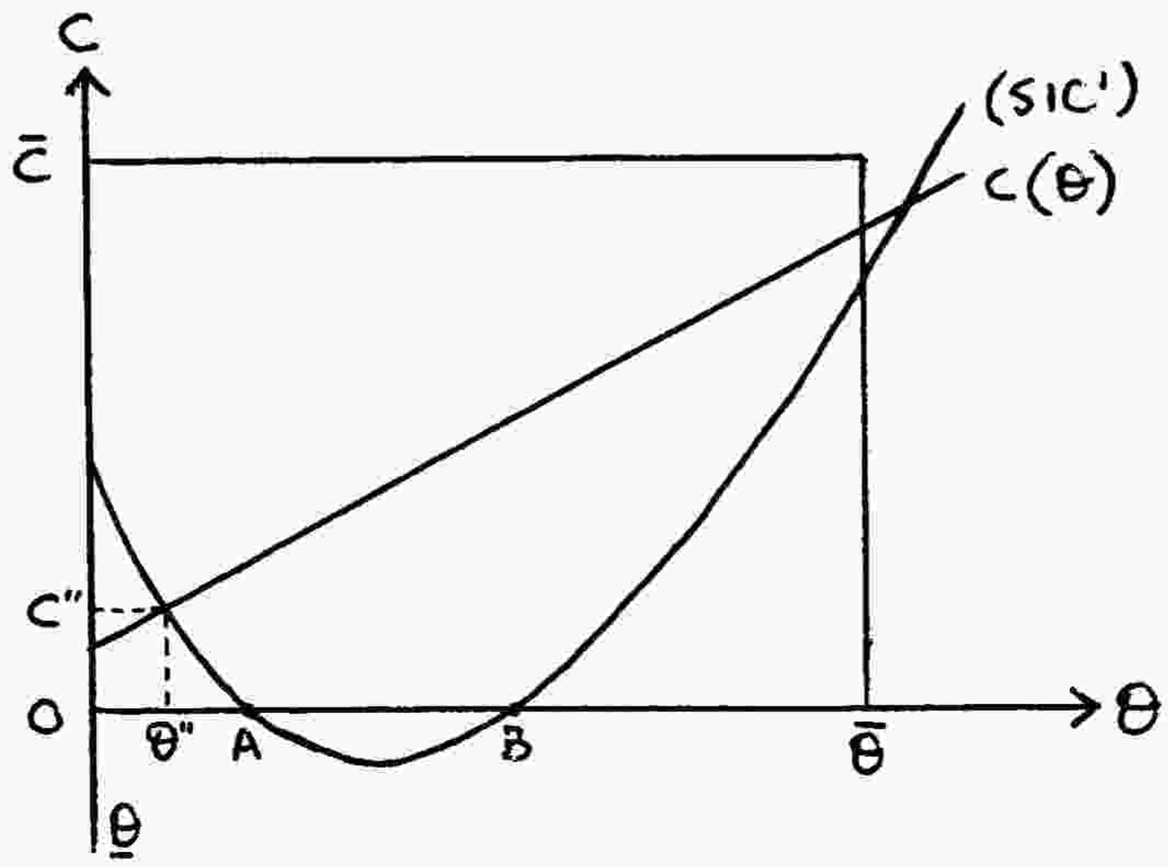
—— w may increase;

Thus, self-regulation will become more likely to be welfare-superior to external regulation when fraud is more important, when sellers are more homogeneous in terms of capital, and when the lags involved in the moral hazard game are higher.

This result is interesting in that the intuition runs against the usual intuition about self-regulation. We generally think that, the harder it is to maintain proper incentives, the less likely self-regulation is to achieve them. In this model, though, both forms of regulation are equally able to ensure proper behaviour, but an SRO will have fewer incentives to try to restrict industry output when output is already low (i.e. when high levels of regulation (high J) are required).

Case II

If C is perfectly correlated with θ (Case II), then the analysis is similar to the above, except that the total number of sellers in the industry is directly proportional to the spread of qualities available. This case is drawn below.



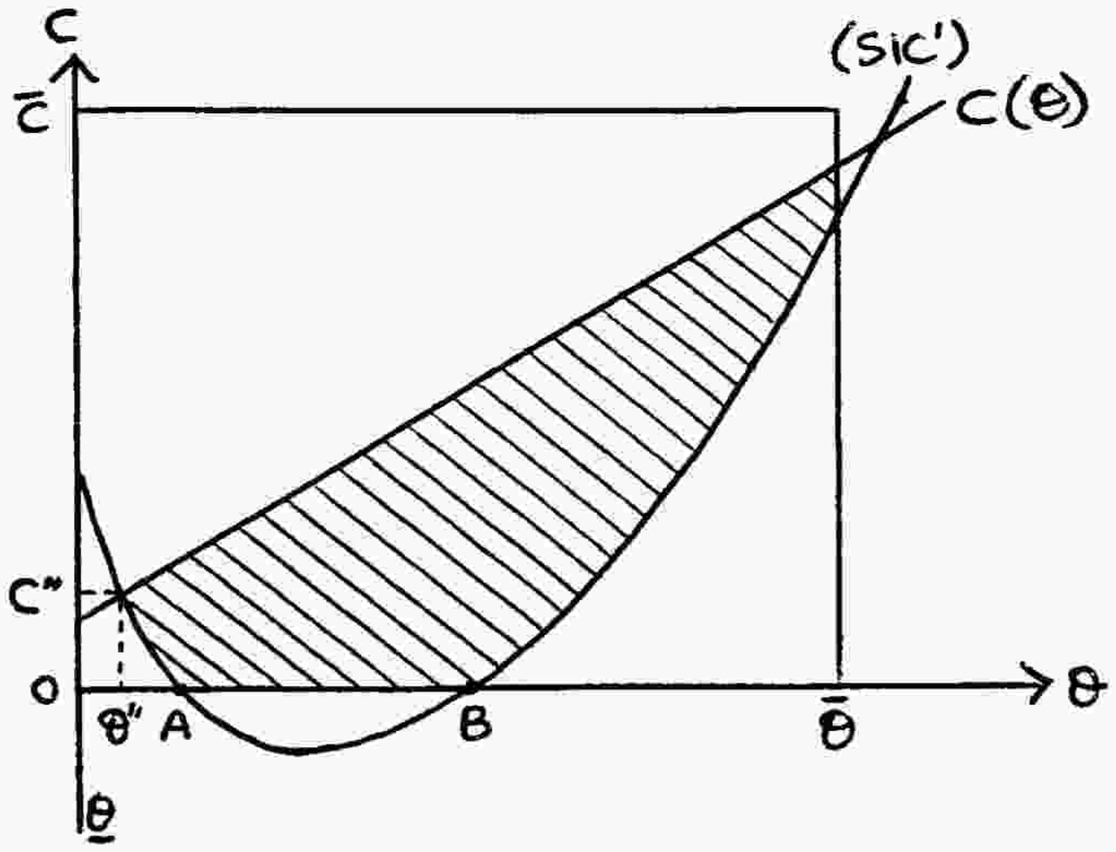
Consider external regulation in this set-up. When there there are no capital requirements for entry ($\hat{J}(X) = 0$, $\forall X$), then the number of sellers in the industry is again $(B - A)N/(\bar{\theta} - \underline{\theta})$ Let C" be the point at which $C(\theta)$ and constraint (SIC') intersect. As $\hat{J}(X)$ increases from zero to C", the spread of qualities increases, as does (by implication) the number of sellers in the industry. This must be welfare improving. Thus the welfare-optimal $\hat{J}$ must be at least as big as C".

As $\hat{J}$ goes up beyond C", constraint (SPC') becomes binding and the lowest quality available will be increasing in $\hat{J}$ (along $C(\theta)$). That is, both the upper *and lower* ends of the available quality interval are increasing with $\hat{J}$, for $\hat{J} > C"$. Instead of facing a trade-off between losing middle quality sellers and gaining sellers at either end of the quality interval, therefore, we now face a trade-off between losing low quality sellers and gaining high quality sellers. Note that it is not clear whether the total length of the quality interval (the total number of sellers in the industry) will be increasing or decreasing in $\hat{J}$.

The welfare-optimal choice of $\hat{J}$ will depend on consumers' valuations of quality, and on how different consumers' valuations are weighted. If consumers care little about quality, or if they value variety over absolute quality, then $\hat{J}$ should be chosen to maximize the available quality interval. Otherwise, as J increases, those consumers who desire higher qualities do better, while those consumers who desire lower qualities do worse.

Turning to self-regulation, denote the quality at which $C(\theta)$ and constraint (SIC') intersect as $\theta"$. When J can be conditioned on θ, it is possible to design regulation such that all sellers with $\theta \geq \theta"$ can enter the industry. Indeed, any $J(\theta, X)$ lying between $C(\theta)$ and constraint (SIC') (that is, in the shaded area below) will allow this.

Note, though, that a self-regulator may eliminate some of these potential sellers from the industry (by setting $J(\theta, X) > \bar{C}$) so as to maximize profits, as before.



The comparison of external with self-regulation is much the same as before. An SRO can choose any quality interval (whose lower end is at least θ"), but will choose it so as to maximize profits. An external regulator is severely constrained in its choice of quality interval, but will choose it so as to maximize welfare. Results 4.3.4 and 4.3.6 remain true. That is, when $\pi_x = 0$, self-regulation will certainly be welfare-superior to external regulation. And, for $\pi_x \neq 0$ and for unchanged $\pi(\theta, X)$, it becomes more likely that self-regulation will be welfare-superior to external regulation as constraint (SIC') shifts upwards, since the quality interval achievable will shrink far less (if at all) under self-regulation than under external regulation.

Result 4.3.5, meanwhile, can be sharpened. Since output X is directly proportional to the length of quality interval available, we get

Result 4.3.7: Suppose $C = C(\theta)$. If the quality interval available under one form of regulation is a subinterval of the quality interval available under the other form of regulation, then the form of regulation with the larger quality interval will be the welfare-superior.

The next section considers how these various results will change if tit-for-tat (rather than trigger) strategies are used in the regulator/seller game.

4.4 The seller/regulator game with tit-for-tat strategies

A potential problem with the above model is that trigger strategies may be unrealistically grim: a seller who cheats, but then undergoes some punishment, seems likely to be given a second chance. Moreover, the concept of putting up an initial amount of capital $J(\theta, X)$ is not only a barrier to entry for many sellers, but is also rarely seen in the real world, except in industries where fraud and/or systemic risk are major problems.

The model in this section is identical to the above model except that the strategies played in the regulator/seller game are tit-for-tat and $J(\theta, X)$ denotes a punishment which is paid *after* a cheating has been detected (rather than a capital sum which is paid on entering the industry). Note that, J can only be paid *after* cheating, with tit-for-tat strategies, because sellers have an incentive to pay J so that they can re-enter the industry. This incentive is absent with trigger strategies, since a cheating seller can never re-enter the industry.

The punishment $J(\theta, X)$ also acts as (partial) compensation, since it is paid directly to the cheated consumer. This tit-for-tat model is similar to the model of medieval merchant traders by Milgrom, North and Weingast (1990).

4.4.1 The model

For simplicity, let us assume that $w = 1$. That is, by the time consumers and buyers are matched in a period, all quality observations and punishments from the previous period will have been executed.

Consider a regulator/seller game with the following sequence of play within each period.

(a) Sellers and buyers are matched in some way. Each seller tells her current buyer whether or not she is regulated (by a *credible* regulator

- see the buyer/regulator game above). Any seller who is regulated pays a subscription fee of S to her regulator.

(b) Each seller chooses whether to be honest and provide quality θ or to cheat, which involves providing zero quality and stealing $F(\theta)$ directly from the customer ($F(\theta)$ may be zero). The buyer she is matched with forms a belief about this quality and pays a price (for one unit of the good) which is commensurate with this belief.

(c) The buyer may file a complaint with the regulator at minimal personal cost, but only if the seller was regulated in the first place.

(d) If the buyer files a complaint, then the regulator can award a judgment $J(\theta, X)$ to the buyer (to be paid by the seller) and can give the buyer further compensation $Q(\theta, X)$.

(e) If a judgment $J(\theta, X)$ is awarded, the seller may pay it, or she may refuse to pay it.

(f) If a seller refuses to pay a judgment, she may be expelled from the regulation system.

Consider the following Regulation System Strategies (RSS):

At stage (a), a seller will pay a fee S to her regulator if and only if she has not been expelled.

At stage (b), if a seller is regulated, then she will provide quality θ and the buyer she is matched with will pay $p(\theta, X)$ for one unit of the good. If the seller is unregulated, no sale will take place.

At stage (c), if any seller is regulated but provides low quality, then the cheated buyer files a complaint with her regulator. Otherwise, no complaint is filed.

At stage (d), if and only if a valid complaint was filed, the regulator awards damages of $J(\theta, X)$ to the buyer and gives the buyer a further compensation of $Q(\theta, X)$.

At stage (e), the seller pays any judgment $J(\theta, X)$ made against her.

At stage (f), if the seller doesn't pay a judgment, then she is expelled from the regulation system until she *is* willing to pay it.

In this game, there is no question of the regulator appropriating funds, since a cheating seller pays its punishment $J(\theta, X)$ directly to the cheated buyer. The regulator must, however, have sufficient incentives to award judgements fairly and to pay out any required further compensation $Q(\theta, X)$. This can be easily achieved in a buyer/regulator game with trigger strategies as in the previous section. As before, it can be shown that when n is large, S can be reduced towards zero without adversely affecting the regulator's incentives. In what follows, let us suppose that S is negligibly small.

Note that $Q(\theta, X)$ is required (as discussed in 4.3.2) to eliminate the alternative sequential equilibrium in which sellers cheat and consumers do not buy. Elimination of this equilibrium requires that consumers do better by purchasing, even if they end up being cheated. More formally, this can be written

$$J(\theta, X) + Q(\theta, X) > p(\theta, X) + F(\theta)$$

which corresponds to condition (CPC') above. It is assumed, in what follows, that $Q(\theta, X)$ is always chosen by the regulator so as to fulfil this condition.

We can now establish the following result.

Result 4.4.1: The RSS are symmetric sub-game perfect equilibrium strategies of the above regulation game if and only if the following inequality holds

$$(4.4.1)\qquad \frac{(p(\theta,X) - \mu(\theta))\delta}{1-\delta} \geq J(\theta,X) \geq F(\theta) + \mu(\theta)$$

Proof: The Optimality Principle of dynamic programming is used. This says that, if we show that there is no point at which a single change in any agent's current action (followed by later adherence to the RSS) can raise that agent's expected payoff, then there is no point at which some more complicated deviation can be profitable, either. Using this, we must examine:

(1) When it "pays to pay judgments": In terms of lifetime average payoff, paying the judgment gives $(-(1-\delta)J(\theta,X) + \delta(p(\theta,X) - \mu(\theta)))$, and refusing to pay the judgment (given the RSS) gives zero. So it "pays to pay judgments" when $J(\theta,X) \leq (p(\theta,X) - \mu(\theta))\delta/(1-\delta)$.

(2) When it pays the victim of cheating to file a complaint at stage (c): Given the RSS, a cheated customer expects the judgment to be paid. So he will file a complaint if and only if $J(\theta,X) + Q(\theta,X)$ is greater than his minimal cost of appealling (which it will be).

(3) When it pays a seller of quality θ to be honest, given that there are no unpaid judgments and that she is currently regulated: If she plays honest, the seller gets $p(\theta,X) - \mu(\theta)$ this period. If she cheats and later adheres to the strategy, her payoff will be $p(\theta,X) + F(\theta) - J(\theta,X)$. So she will play honest if and only if $J(\theta,X) \geq \mu(\theta) + F(\theta)$.

(4) When it pays the sellers to be regulated: A regulated seller of quality θ gets $(p(\theta,X) - \mu(\theta))$ this period. Otherwise (given the RSS) she gets zero. So she will participate in the system if $p(\theta,X) - \mu(\theta) \geq 0$ (but this condition is redundant in view of the others). □

Corollary: There is a judgment $J(\theta,X)$ for each (θ,X) which makes the RSS a symmetric sequential equilibrium strategy if and only if

$$(4.4.2)\qquad p(\theta,X) \geq (\mu(\theta) + (1-\delta)F(\theta))/\delta$$

This corollary leads us to a second result.

Result 4.4.2: The condition for sellers to be honest is stricter (holds in fewer circumstances) in the tit-for-tat version of this game than in the trigger version with $w = 1$. More precisely, a seller can only be given sufficient incentives for honesty in the tit-for-tat game if it is possible to regulate that seller properly in the trigger game without using capital requirement (i.e. with $J(\theta, X) = 0$).

Proof: Condition (4.4.1) is the tit-for-tat equivalent to condition (SIC') in the above trigger model. We can rewrite condition (SIC') for $w = 1$ as

$$\text{(SIC')} \qquad p(\theta, X) \geqslant (\mu(\theta) + (1 - \delta)F(\theta))/\delta - (1 - \delta)J(\theta, X)/\delta$$

Thus, since $J(\theta, X) \geqslant 0$, condition (SIC') will certainly hold whenever condition (4.4.2) holds. That is, condition (4.4.2) is more strict than condition (SIC'). Conversely, condition (4.4.2) will hold only if condition (SIC') holds for $J(\theta, X) = 0$. □

This result has two interesting implications. Firstly, trigger strategies may be able to achieve sufficient incentives for sellers to be honest, even when this is impossible under tit-for-tat strategies. Thus, when incentives to cheat are very high (i.e. $\mu(\theta)$ and $F(\theta)$ are high relative to $p(\theta, X)$), initial capital requirements may be needed, whereas when incentives to cheat are lower, *ex post* punishments for cheating may be sufficient.

Secondly, one obvious potential benefit of using tit-for-tat strategies and *ex post* punishments, instead of initial capital requirements, is that sellers are not prevented from entering the industry on grounds of lack of capital. However, the above result shows that it is only those sellers which could enter the industry under trigger strategies *without* depositing a capital fund which are able to be regulated in this way. Thus, this potential benefit turns out to be fallacious.

4.4.2 External regulation versus self-regulation

This section returns to the two essential differences between external and self-regulation analysed above:

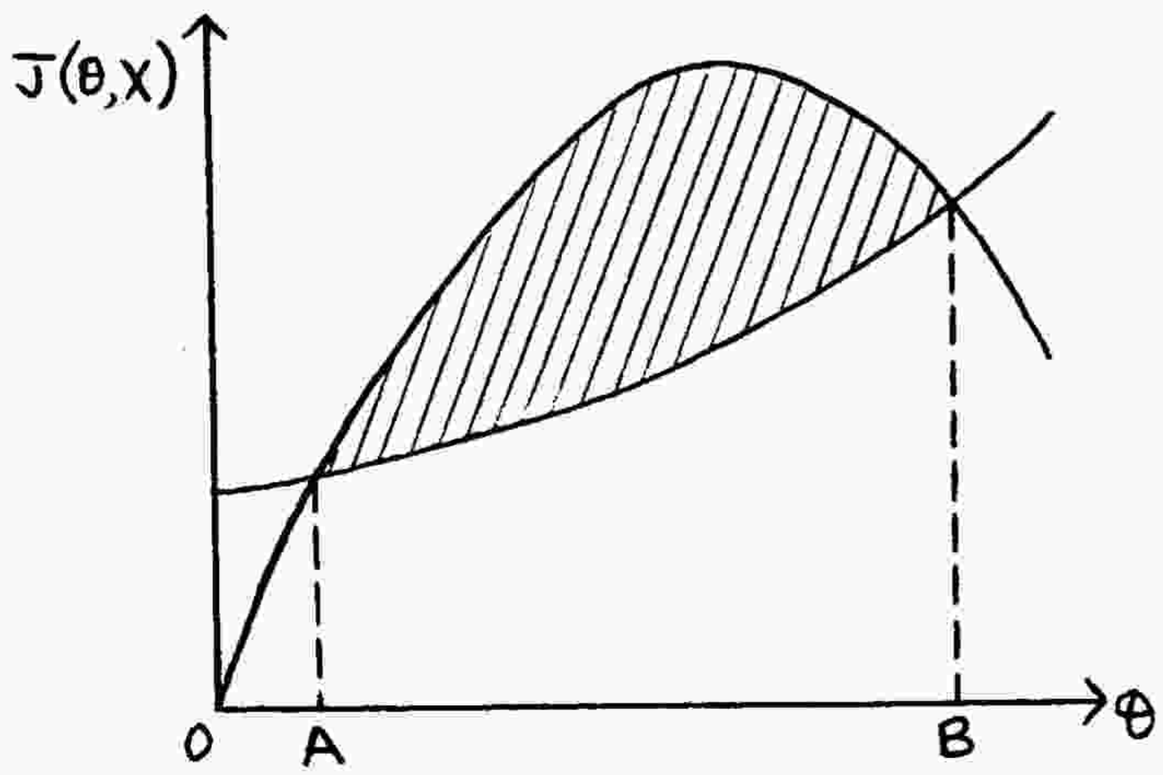
—— An external regulator will aim to maximize welfare, whereas an SRO will aim to maximize industry profits; and

—— An SRO can set $J(\theta, X)$ however it wants, but an external regulator is constrained by being unable to condition $J(\theta, X)$ on θ, due to nonverifiability.

The implications of this model are similar to those in the trigger strategies model. The equilibrium condition for honesty in this model is

$$(4.4.1) \qquad \frac{(p(\theta, X) - \mu(\theta))\delta}{1 - \delta} \geq J(\theta, X) \geq \mu(\theta) + F(\theta)$$

Let us define A and B such that condition (4.4.2) holds for all $\theta \in [A, B]$, where $\underline{\theta} \leq A < B \leq \bar{\theta}$.[9] There are effectively two constraints on $J(\theta, X)$ in the above condition. Call the lhs constraint (4.4.1a) and the rhs constraint (4.4.1b). Clearly these two constraints will cross at $\theta = A$ and $\theta = B$, as in the diagram below.



[9] Note that condition (4.4.2) only holds when (SIC') would hold for C = 0. Thus, So A and B are defined here in exactly the same way as in section 4.3.

$J(\theta, X)$ must lie somewhere in the shaded area between the two constraints. Note that an SRO has no incentives to set $J(\theta, X)$ higher than its minimum level for any θ sellers which it wants to allow into the industry. Thus, for those sellers which it allows into the industry, an SRO will set $J(\theta, X)$ equal to $\mu(\theta) + F(\theta)$. For any other sellers, the SRO will either simply refuse membership to them, or it will set $J(\theta, X)$ to lie outside the shaded area (which means that rational buyers will not purchase from such sellers). An external regulator, meanwhile, is constrained to set $J(\theta, X) = \hat{J}(\theta)$.

Clearly, the above situation is formally similar (although not identical) to the situation in section 4.3.4 where $C = C(\theta)$, with constraint (4.4.1b) corresponding to (SIC') and with constraint (4.4.1.a) corresponding to (SPC'). As in that set-up, the number of sellers in the industry in this model will be $N/(\bar{\theta} - \underline{\theta})$ times the available quality interval. That is, if any seller of quality θ can enter the industry, then all sellers of the same quality can enter, regardless of C.

Unsurprisingly, then, the results comparing self and external regulation for $C = C(\theta)$ (i.e. results 4.3.4, 4.3.6, and 4.3.7) will also hold for this model. The main difference, is that $[A, B]$ is the *maximum* quality interval achievable under tit-for-tat strategies, whereas it was the *minimum* quality interval achievable (even by an external regulator) under trigger strategies.

Appendix 4.A1 Tit-for-tat and trigger strategies

Consider the following simple two-player normal-form game.

		Player 2: Honest	Player 2: Cheat
Player 1	Honest	A_1 A_2	C_1 B_2
	Cheat	B_1 C_2	D_1 D_2

where $B_1 > A_1 > D_1 > C_1$

and $A_2 > D_2 > C_2$.

Outcome $\{A_1, A_2\}$ is pareto-superior to outcome $\{D_1, D_2\}$. Thus, both players would prefer that they both play Honest than that they both play Cheat. Because $B_1 > A_1$ and $D_1 > C_1$, however, player 1 will certainly want to play Cheat, whatever player 2 is doing; and since $D_2 > C_2$, player 2 will thus choose Cheat. So {Cheat, Cheat} is the Nash equilibrium of the one-shot version of this game. If $A_2 \geq B_2$, then the game is one of one-sided moral hazard, since player 2 *would* play Honest, given that player 1 was going to. Otherwise (when $B_2 > A_2$), both players face moral hazard incentives.

When this game is repeated, it may become possible for players to achieve {Honest, Honest}, due to a desire not to damage their good reputations. Two possible strategies that are often examined are

(i) Trigger - in which both players play Honest until one player cheats, after which both play Cheat forevermore.

(ii) Tit-for-tat - in which a player i plays Cheat in period t if and only if player i played Honest and player j played Cheat in period t-1.

Let a_{it} represent the action taken by player i in period t; and let $\pi_i(a_{1t}, a_{2t})$ represent the resulting payoff earned by player i in period t; and let δ be the discount factor applied to compute the present value of a stream of payoffs. A player's time weighted average payoff over the whole sequence of trades is given by:

$$\bar{\pi}_i = (1 - \delta) \sum_{t=0}^{\infty} \delta^t \pi_i(a_{1t}, a_{2t})$$

We can now derive some results about the ability of the free market to overcome the moral hazard problem.

Result 4.A1.1: For δ near enough to one - specifically if

$$(4.A1.1)\quad \delta \geq \frac{B_i - A_i}{B_i - D_i}$$

- then Trigger strategies form a subgame perfect equilibrium (SPE) of the above game.

Proof: By the Optimality Principle of dynamic programming it suffices to show that there is no point at which player i can make a one-time play different from the equilibrium play that raises his total payoff. If neither player has so far cheated, this requires that

$$A_i \geq (1 - \delta)B_i + \delta D_i$$

which holds if and only if the condition in the result 4.A1.1 holds. And if either player has already cheated, this requires that

$$D_i \geq (1 - \delta)C_i + \delta D_i$$

which holds if and only if $D_i \geq C_i$, which we have already assumed. □

Result 4.A1.2: For δ near enough to one - specifically if

$$(4.A1.2)\quad \delta \geq \text{Max}\left[\frac{B_i - A_i}{A_i - C_i}, \frac{D_i - C_i}{A_i - C_i}\right]$$

- then Tit-for-tat strategies form an SPE of the above game.

Proof: Again the Optimality Principle can be used, but now there are four possible decision situations

(i) If both players played honest at t-1 or if both players cheated at t-1 - a current period deviation by player i is unprofitable if and only if

$$A_i \geq (1 - \delta)(B_i + \delta C_i) + \delta^2 A_i$$

which holds if and only if $\delta \geq (B_i - A_i)/(A_i - C_i)$.

(ii) If player j was honest at t-1, but player i cheated - a deviation by i is unprofitable if and only if

$$(1 - \delta)C_i + \delta A_i \geq (1 - \delta)(D_i + \delta C_i) + \delta^2 A_i$$

which holds if and only if $\delta \geq (D_i - C_i)/(A_i - C_i)$.

(iii) If player i was honest at t-1, but player j cheated - a deviation by i is unprofitable if and only if

$$(1 - \delta)B_i + \delta A_i \geq A_i$$

which holds if and only if $B_i \geq A_i$, which has already been assumed to be true. □

Result 4.A1.3: When $B_i = A_i + D_i - C_i$, then the condition for SPE is same for both trigger and tit-for-tat strategies. That is, whichever strategies are used, {Honest, Honest} will be an SPE if and only if

$$(4.A1.3) \quad \delta \geq \frac{B_i - A_i}{B_i - D_i}$$

Otherwise, the condition for SPE with tit-for-tat strategies is always more stringent that that for SPE with trigger strategies.

Proof: The first part of this result is proved by simple substitution. To prove the second part, two cases need to be considered.

(i) $B_i > A_i + D_i - C_i$ implies that the tit-for-tat condition becomes $\delta \geq (B_i - A_i)/(A_i - C_i)$, and this is greater than the trigger condition $(B_i - A_i)/(B_i - D_i)$ if and only if $(B_i - D_i) > (A_i - C_i)$, which is true by assumption.

(ii) $B_i < A_i + D_i - C_i$ implies that the tit-for-tat condition becomes $\delta \geq (D_i - C_i)/(A_i - C_i)$, and this is greater than the trigger condition $(B_i - A_i)/(B_i - D_i)$ if and only if

$$
\begin{aligned}
& (D_i - C_i)(B_i - D_i) > (B_i - A_i)(A_i - C_i) \\
\Leftrightarrow \quad & A_i^2 - D_i^2 > A_iB_i + A_iC_i - D_iB_i - D_iC_i \\
\Leftrightarrow \quad & (A_i - D_i)(A_i + D_i) > (A_i - D_i)(B_i + C_i) \\
\Leftrightarrow \quad & (A_i + D_i) > (B_i + C_i)
\end{aligned}
$$

which is true by assumption. □

More generally, we might like to look at strategies that lie midway between tit-for-tat and trigger, in which i's response to j's cheating in period t-1 will be to cheat for T periods before returning to acting honestly. The condition on δ with this type of strategy will be more stringent than the trigger condition but less stringent than the tit-for-tat condition. Moreover, I postulate that, for any δ such that

$$\frac{B_i - A_i}{B_i - D_i} < \delta < \text{Max}\left[\frac{B_i - A_i}{A_i - C_i}, \frac{D_i - C_i}{A_i - C_i}\right]$$

there will be a length of punishment T* such that any shorter punishments will not achieve the honesty SPE and any longer punishments are wasteful.

Appendix 4.A2 The model with endogenous price

The main section of this chapter assumes an exogenous inverse demand function $p(\theta, X)$ and derives, from it, some implications about punishment and compensation schedules, $J(\theta, X)$ and $Q(\theta, X)$.

An alternative assumption is that there is Bertrand competition between sellers in the industry, so that prices are competed down to their minimum levels, without regard to their effects on demand. This idea is similar to that in Shapiro (1983), although he assumes free entry and perfect competition, rather than restricted entry and Bertrand competition.

When price is endogenous, we must initially derive the constraints on prices that must be fulfilled at an {honest, buy} regulation equilibrium. By its choice of $J(\theta)$, a regulator can now affect $p(\theta)$, and thus total demand (since demand will be dependent on prices and qualities, in some way).

It is difficult to derive any explicit implications about regulation in the endogenous price versions of the above models. This appendix just contains a few preliminary observations.

Consider first the trigger version of the above game (described in section 4.3). The principal constraints which must hold at equilibrium in this game are

$$\text{(SIC')} \qquad J(\theta) \geq \frac{(1 - \delta^w)F(\theta) - \delta^w p(\theta) + \mu(\theta)}{(1 - \delta)}$$

$$\text{(SPC')} \qquad C \geq J(\theta)$$

$$\text{(CPC')} \qquad \frac{(1 - \delta^w)Q(\theta)}{(1 - \delta)} \geq \frac{(1 - \delta^w)(p(\theta) + F(\theta))}{(1 - \delta)}$$

From these formulations, it is clear that a shift upwards of $p(\theta)$ will shift constraint (SIC') downwards and will shift constraint (CPC') upwards. This concurs with intuition. Raising the price that sellers receive gives sellers extra incentives for honesty, which means that fewer incentives need to be given via $J(\theta)$ (i.e. (SIC') shifts downwards). At the same time, raising $p(\theta)$ raises the lower bound on compensation to be given to cheated consumers (i.e. (CPC') shifts upwards).

Bertrand competition in prices means that prices will fall until they are only just high enough to give sellers incentives for honesty (that is, until constraint (SIC') just holds). In this case, the following result is straightforward to derive.

Result 4.A2.1: When $p(\theta)$ is endogenous to the trigger model, there is a trade-off (along constraint (SIC')) between lowering prices (towards costs) and raising minimum capital requirements.

Lowering prices towards costs is beneficial, partly because it lowers minimum compensation levels, but more importantly because it increases demand in the industry (and brings consumers' quality and quantity choices nearer to their full information levels). However, raising minimum capital requirements is detrimental to the supply side of the economy (since sellers must have $C \geq J(\theta)$ in order to function in the industry).[10]

These two effects essentially describe a demand function and a supply function, for each quality θ. Raising prices reduces demand in the usual way (via a downward-sloping demand curve). On the supply side, raising prices reduces the minimum capital requirements necessary to fulfill constraint (SIC'), which allows new sellers to enter the industry. Thus, raising prices raises supply, giving us an upward-sloping supply function.

As with normal supply and demand schedules, there is likely to be a price and output which equates supply and demand. Unlike a normal supply and demand model, however, there are no outside forces ensuring that this equilibrium point is reached. Instead, prices are determined by the regulator's choice of $J(\theta)$, and output is then given by either demand or supply - whichever is smallest - for that price.

Turning to the tit-for-tat model, condition (4.4.1) gives us the minimum $p(\theta)$ that is consistent with the {honest, buy} sequential equilibrium:

$$p(\theta) = \mu(\theta) + (1 - \delta)J(\theta)/\delta$$

where $J(\theta) \geq \mu(\theta) + F(\theta)$.

---

[10] Note that when $p(\theta)$ is actually lowered to $\mu(\theta)$ (its full information level), constraints (SIC') and (CPC') coincide.

Thus, we get the following result.

Result 4.A2.2: When $p(\theta)$ is endogenous to the tit-for-tat model, $p(\theta)$ will be increasing in $J(\theta)$.

In this model, an increase in $J(\theta)$ has no effect on supply of given θ, since C is not constraining in this model. Since demand for each quality θ will be decreasing in $p(\theta)$, it is clear that welfare is maximized when there is complete entry ($X = N$), and $p(\theta)$ is as low as possible. This is done by setting $J(\theta)$ equal to $F(\theta) + \mu(\theta)$, for each θ.

Again, neither self-regulation or external regulation will achieve this welfare optimum. An external regulator must set $J(\theta) = \hat{J}$ for all θ. Thus, those θ with $\hat{J} < \mu(\theta) + F(\theta)$ will be unable to enter the industry; and those θ with $\hat{J} > \mu(\theta) + F(\theta)$ will have greater than minimal price. Raising $\hat{J}$ has two effects under external regulation: it raises the price which must be paid for any given quality by $(1 - \delta)\hat{J}/\delta$; and it raises the maximum available quality, $\bar{\theta}(\hat{J})$, which is given by

$$\bar{\theta}(\hat{J}) = [\theta \mid \mu(\theta) + F(\theta) = \hat{J}]$$

That is,

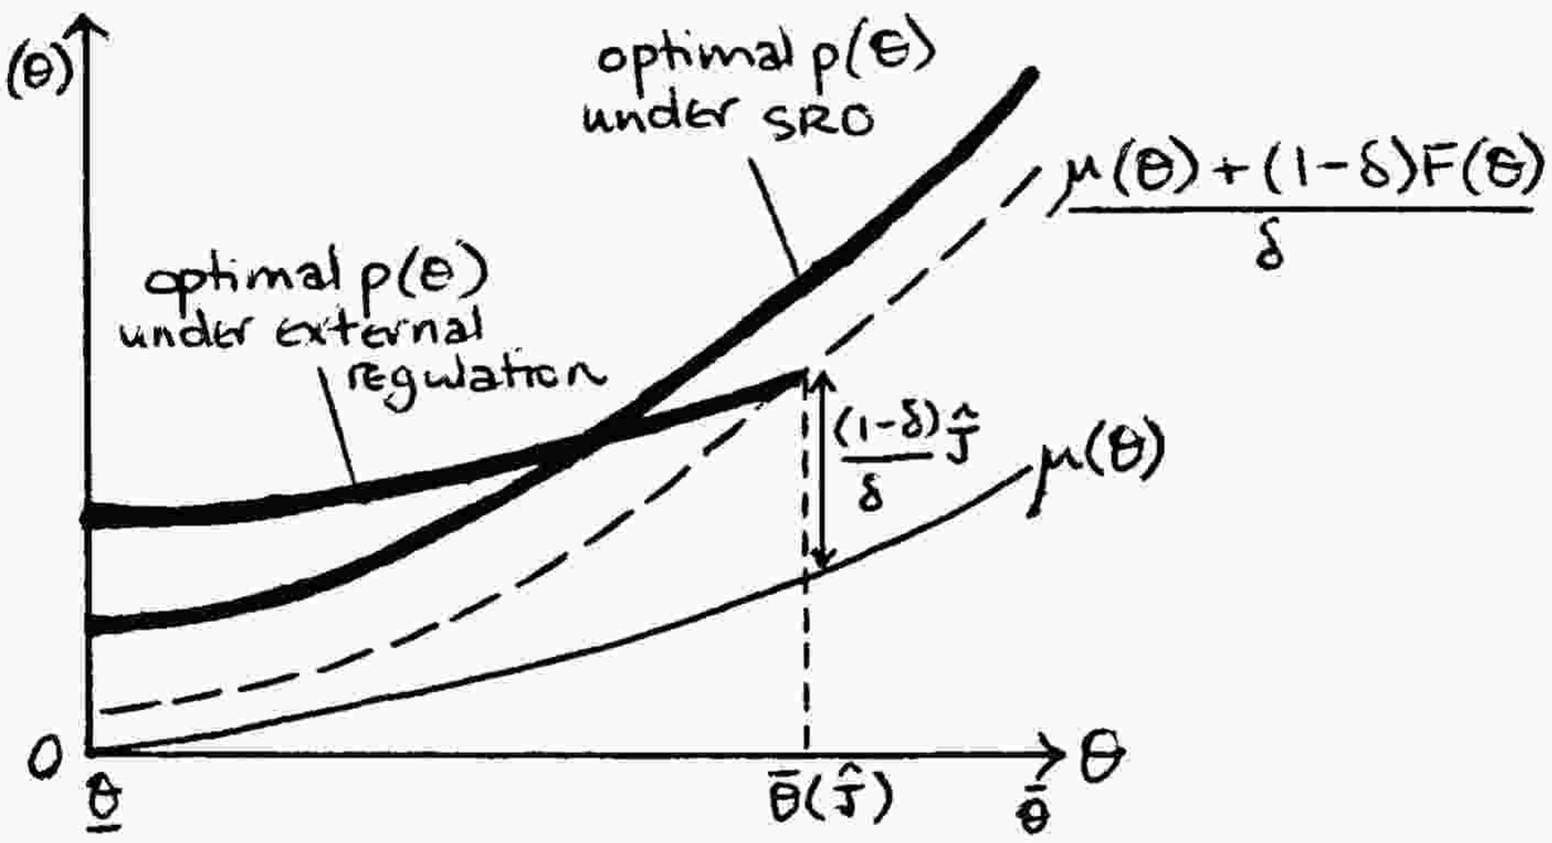
Result 4.A2.3: When $p(\theta)$ is endogenous to the tit-for-tit model, an external regulator faces a trade-off between increasing the highest quality available in the industry and raising the prices of all other qualities.

An increase in $\hat{J}$ will thus tend to be good for those consumers who value high quality, but bad for those who prefer low qualities at low prices.

An SRO, on the other hand, *can* condition $J(\theta)$ on θ, but will do this so as to maximize industry profits, which will tend to involve $J(\theta)$ (and so $p(\theta)$) being higher than is welfare-optimal.

It is unclear which of the two forms of regulation will be welfare-superior in this set-up. Suppose that optimal external regulation sets $\hat{J}$ such that some high θ sellers (with $\theta > \bar{\theta}(\hat{J})$) are kept out of the industry and such that $p(\theta)$ for some low θs is greater than any consumers wish to pay. In this case, those consumers which value high and low qualities will prefer self-regulation, since an SRO will tend to set $J(\theta)$ such that the higher qualities can enter while, at the same time, the prices charged for low θs are reduced. However, consumers who prefer mid-range qualities are likely to prefer external regulation, since price will close to its minimal level, whereas it will tend to be raised above this minimal level by an SRO.

A situation like this is shown in the diagram below.



To derive further results about these endogenous $p(\theta)$ models would require additional specification of the demand side of the models.

NOTATION - chapter 5

c	unit costs
Φ	fixed set-up cost for each SRO
h	labourers' services consumed per buyer
k	quality/training
L	minimum quality level
$\lambda\tau k^2$	cost of training
$\lambda\tau\theta$	cost of innate quality
$M(h, \bar{q})$	quantity of labourers' services a consumer is willing to forego for one unit of quality $\bar{q}$ legal services
μ	simplification parameter $\equiv \tau/\underline{u}$
N	total number of sellers
n	number of sellers in SRO
p	price
ω	most preferred quality
$\bar{q}(\alpha, \beta)$	average quality over quality interval $[\alpha, \beta]$
Π	Total industry profits
π	per seller profits
ρ	weighting on Π in welfare function
S	consumer surplus
T	$\equiv n_1 n_2 + 2n_1 + 2n_2 + 2$
τ	transport costs
θ	(innate) quality ϵ $[0, 1]$
U	utility
$\underline{u}$	difference between a consumer's maximum utility and his reserve utility u_o
u_o	reserve utility
v	number of consumers
W	welfare
X	total SRO output
x	per seller output
Z	the number of SROs

5. THE NUMBER OF SROS

5.1 Introduction

In previous chapters, self-regulation was generally assumed to be carried out by a single SRO, an assumption which facilitated comparisons between self-regulation and other possible regimes. This chapter considers whether (and when) this assumption is really justified. That is,

— When will the optimal number of SROs be equal to one?

— Will it ever be socially beneficial for an SRO to face the threat of entry by another SRO even if entry never actually occurs?

— Should membership of an SRO be compulsory for all the sellers in an industry?

This chapter also examines whether an industry, *acting on its own*, will have the welfare-optimal number of SROs, or whether the number of SROs should be influenced by public policy, even though the regulation itself is internal.

5.1.1 Changing the number of SROs

There are essentially five potential effects of increasing the number of SROs, Z, in an industry. The relative signs and magnitudes of these effects will be important for answering the above questions. Briefly, these effects are:

— The competition effect: As Z increases, competition between SROs will increase. This causes prices to fall and total output to rise, which will be beneficial to consumers and detrimental to profits.

— The output manipulation effect: As Z increases, SROs will become less able to manipulate quality so as to restrict total output (as in the

sequential self-regulation model in chapter 3). This will tend to result in reduced manipulation of quality (although it could result in more), and less manipulation will cause average quality to move back towards its socially optimal level. Overshooting is possible, however.

— The quality externality effect: As Z increases, any quality externalities between SROs (such as the positive and negative externalities modelled in chapter 3) will become less internalized, and so each SRO's quality will diverge further from its collusive profit-maximizing value. This will certainly be bad for industry profits. It is not clear, however, whether it will be good or bad for consumers.

— The variety effect: As Z increases, the number of qualities on offer may increase. This will

  — reduce average 'transport' costs per consumer; and may

  — reduce variances (around expected quality) which will improve the utilities of risk averse consumers.

both of these effects will improve welfare.

— The duplication of costs effect: There are likely to be significant costs involved in setting up and running an SRO. If the average costs of each SRO are decreasing with the SRO's size, then an increase in Z will result in an increase in the industry's overall costs of self-regulation, which will be bad for welfare.

5.1.2 Layout and main conclusions of chapter

In this chapter, I present a simple model which exhibits all of these effects (with the exception of the variance reduction effect which I have found difficult to incorporate). This model is a variant of the Hotelling

(1929) model, in which consumers differ in their preferences for vertical quality, rather than in their preferences for horizontal quality. More precisely, consumers are assumed to differ according to their 'most preferred' vertical qualities. Transport costs are incurred if the quality a consumer receives is less than his 'most preferred' quality (but not if it is greater - for the rationale behind this see section 5.3).

It is assumed that there are n_j sellers in each SRO j. Each seller makes output choices and (more or less) irreversible quality/training choices. Members of an SRO all collude on quality levels, and either compete or collude on output levels. The SROs compete with each other on both output and quality. The industry contains either one or two SROs.

This model is simple, but it gives relatively strong results, many of which seem likely to hold for more general set-ups. Of these results, five are especially interesting, the intuitions for which are given in sections 5.6 and 5.7.

(1) When entry is member-stealing (e.g. SRO 2 steals half of SRO 1's members on entry), then entry will result in SRO 1 raising its quality level. However, when entry is member-augmenting (e.g. SRO 2 enters at the same scale as SRO 1 but with completely new members), then entry will result in SRO 1 lowering its quality level.

(2) Consumer surplus will always be higher when there are two SROs than when there is only one, except (for certain parameter values) when there is output competition within SROs and entry is member-augmenting.

(3) When entry is member-augmenting, total industry profits will always be lower when there are two SROs than when there is only one. When entry is member-stealing, the results are more complex: Total industry profits will be higher when there are two SROs than when there is only one, except under full collusion within SROs for certain parameter values. Another interesting result is that, when entry is member-

stealing and when there is output competition within SROs, SRO 1's own profits will be higher after entry for certain parameter values.

(4) When there is full collusion within SROs, welfare is definitely higher when there are two SROs, even when total profits are lower. That is, even with the maximum weighting on profits, the beneficial effect of entry on consumer surplus always dominates. This result also holds when there is output competition within SROs for the member-stealing case, but not for the member-augmenting case.

(5) SRO 1 will need to lower its quality level in order to deter entry by SRO 2. At this point, however, SRO 1's profits may be lower than they are if entry is accomodated. Similarly, if SRO 2 wants to push SRO 1 out of the industry, it must raise its quality level from its accomodation level (note, though, that there are global stability problems in this case).

This model can be extended by introducing Leland's (1979) model into the demand side, which makes total output and average quality supplied interdependent. This extension is examined in section 5.8. When applied to simple one-SRO and two-SRO frameworks, this model gives qualitatively similar results to the simple Hotelling model. It is useful, however, in that it can be used to examine how a single SRO will react to the introduction of a non-SRO sector. The most interesting results are that

(1) It is ambiguous whether the SRO will raise or lower L in response to entry by a non-SRO sector, but it will become more likely to raise it as transport costs and costs of quality become higher.

(2) The SRO will certainly make lower profits when there is entry.

(3) Consumer surplus (and, in fact, total welfare) will certainly be higher when there is entry, so that the SRO's preferences with respect to entry will be diametrically opposed to those of consumers.

This extension of my model can also be used to examine the optimal number of SROs (and whether this number will emerge automatically in an industry). The analysis is relatively simple because total output must equal one and SROs turn out not to change their L-choices in the face of further entry. So there is a simple trade-off between extra fixed set-up costs Φ and reduced transport costs. One conclusion in this section is that the optimal number of SROs will be increasing in the magnitude of transport costs and decreasing in that of fixed costs. A less obvious conclusion is that, for given transport costs, there will tend to be too many SROs in the industry for welfare-maximization when Φ is low, but too few SROs when Φ is high.

Before considering these models, however, I examine the two-SRO extension of Shaked and Sutton's (1981) model. The main conclusion in Shaked and Sutton's one-SRO model is that an SRO which is allowed to choose its own minimum quality level L will choose too high a level of L for welfare maximization (result 2.3.2 in section 2.3).

In this case, the idea behind the introduction of a second SRO into an industry is simple. If the introduction of a second SRO results in a reduction of the industry minimum quality level from its one-SRO level, then welfare will (at least at first) be improved. Shaked and Sutton examine two polar cases. Firstly, they show that if SRO 1 doesn't change its choice of L_1 in response to entry, then the entry of a second SRO will certainly not reduce welfare, and will probably raise it. Secondly, they show that if SRO 1 reduces L_1 so as to deter entry by the second SRO, L_1 may overshoot its optimal level, and welfare may actually be reduced.

This model is presented in the next section. It is a simple and appealing model, but its assumptions are not unproblematic. Some reservations about these were stated in chapter 2, and section 5.2.4 considers two more. These reservations provide a partial motivation for my model.

5.2 Shaked and Sutton (1981) - the two-SRO extension

In this section, I examine the two-SRO extension of Shaked and Sutton's (1981) model. This model was presented in chapter 2, section 2.3. The main conclusion, in the one-SRO model, was that an SRO which is allowed to choose its own minimum quality level L will choose too high a level of L for welfare maximization (result 2.3.2).

In this case, the idea behind the introduction of a second SRO into an industry is simple. If the introduction of a second SRO results in a reduction of the industry minimum quality level from its one-SRO level, then welfare will (at least at first) be improved. Shaked and Sutton consider three possible ways in which the incumbent SRO (SRO 1) might react to entry by a second SRO:

(1) If SRO 1 doesn't change its choice of L_1 in response to entry, then SRO 2 will have to choose $L_2 < L_1$ to get any members. Thus the industry minimum quality level is lower than in the one-SRO set-up.

(2) If SRO 1 decides to deter entry, it will generally have to lower L_1 to eliminate any potential profits for SRO 2.

(3) More generally, if the SRO 1 chooses to allow entry but to change its choice of L_1 after entry, it will tend to choose a lower L_1 than it did before entry. This is because SRO 1 will no longer be able to restrict output directly via its use of L_1 (although it may have some ability to control output indirectly, if it knows SRO 2's reaction curve). Thus, SRO 1 will have less incentive to set $L_1 > L^*$.

In all of these cases, the industry minimum quality level will be lowered, so that welfare will (at least at first) be improved. The qualification 'at least at first' is needed in cases (2) and (3) because there may be overshooting, so that the industry minimum quality level ends up *lower* than is welfare-optimal. This is especially true of case (2), where L_1 might have to be lowered a great deal in order to prevent entry. Note also that

the costs involved in setting up a second SRO need to be set against any benefits.

5.2.1 The framework

There are N workers who can be employed either as "lawyers" or as "labourers". Workers are identical when employed as labourers, but differ from each other in the quality of service they offer when employed as lawyers, and they are ranked along the interval [0, N] on this basis, in ascending order of potential quality.

There are two SROs, which choose minimum quality levels L_1 and L_2, where $L_1 \geq L_2$, and where SRO 2 does not exist when $L_1 = L_2$. The number of lawyers in SRO 1 will be $X_1 = N - L_1$; the number of lawyers in SRO 2 will be $X_2 = L_1 - L_2$; and there will be L_2 labourers. The expected qualities of the two SROs will be $\bar{q}^1$ and $\bar{q}^2$ respectively, where $\bar{q}^1 = \bar{q}(L_1, N)$ and $\bar{q}^2 = \bar{q}(L_2, L_1)$, where $\bar{q}(\alpha, \beta)$ represents a preference ordering which is complete, reflexive, transitive and continuous.

There are v consumers, each of which can purchase one unit of legal services from a member of SRO 1, or one unit of legal services from a member of SRO 2, or none. The services of labourers are a residual category of expenditure, denoted h. Define $M(h, \bar{q})$ as the quantity of labourers' services the consumer is willing to forego in return for 1 unit of lawyers services of quality $\bar{q}$, where $M_{\bar{q}}\bar{q}_{\alpha}$ is positive. Let M^1 denote $M(h, \bar{q}^1)$ and let M^2 denote $M(h, \bar{q}^2)$.

5.2.2 The two-SRO equilibrium

At equilibrium, the v consumers are divided into three groups:

— there are X_1 consumers who buy one unit of legal services from a member of SRO 1 and enjoy $h - M^1$ units of labourers' services;

— there are X_2 consumers who buy one unit of legal services from a member of SRO 2 and enjoy $h - M^2$ units of labourers' services; and

— there are $v - X_1 - X_2$ consumers who enjoy h units of labourers' services only.

The total supply of labourers, meanwhile, is L_2. Thus, equilibrium in the market for labourers occurs when

$$(v - X_2 - X_1)h + X_1(h - M^1) + X_2(h - M^2) = L_2$$

which can be rewritten as

$$(5.2.1) \qquad vh = X_1M^1 + X_2M^2 + L_2$$

This corresponds to equation (2.3.1) in the one-SRO model.

Keeping X_1 constant, this equation can be differentiated with respect to L_2 to obtain

$$(5.2.2) \qquad vh_{L_2} = X_1M^1_{h}h_{L_2} + X_2M^2_{h}h_{L_2} + X_2M^2_{q}\bar{q}^2_{x} - M^2 + 1$$

which, on rearranging, gives

$$(5.2.3) \qquad h_{L_2} = \frac{X_2M^2_{q}\bar{q}^2_{x} - (M^2 - 1)}{v - X_1M^1_{h} - X_2M^2_{h}}$$

Note that the denominator of this term is positive since $v > X_1 + X_2$ and $1 > M^2_{h}, M^1_{h}$.

5.2.3 Results

At equilibrium, M^1 is equal to the ratio of per seller earnings of SRO 1 members to those of labourers. Similarly, M^2 denotes the ratio of per seller earnings of SRO 2 members to those of labourers. An SRO will be

viable only if this ratio is greater than or equal to 1. That is, SRO 2 is viable if and only if, at equilibrium, $M^2 \geq 1$. When this viability condition is satisfied, let us say that SRO 2 can "threaten" SRO 1. Using this we can derive the following result.

Result 5.2.1: If an SRO of quality $\bar{q}^2$ can threaten an SRO of quality $\bar{q}^1$, then any SRO of quality $\bar{q} > \bar{q}^2$ can threaten $\bar{q}^1$.

Proof: Viability requires that $M^2 \geq 1$. From (5.2.3), we know that when $M^2 = 1$,

$$h_{L_2} = \frac{X_2 M^2_{\bar{q}} \bar{q}^2_\alpha}{v - X_1 M^1_h - X_2 M^2_h} > 0$$

When $h_{L_2} > 0$, however, it must be the case that $M^2_{L_2} > 0$ (by total differentiation of $M(h, \bar{q}^2)$). Thus, once M^2 exceeds unity for some $\bar{q}^2$, it will never fall below unity as $\bar{q}^2$ increases. □

Corollary: If an SRO of quality $\bar{q}^1$ can be threatened by a second SRO, it can certainly be threatened by an SRO of quality $\bar{q}^{2'} = \bar{q}(L_1, L_1)$, that is, by an SRO consisting of a single individual just excluded from SRO 1.

Now, let $\bar{q}^{2''}$ be the largest viable size of SRO 2 (which corresponds to $M^2 = 1$), and suppose that $\bar{q}^1$ is such that $\bar{q}^{2'} > \bar{q}^{2''}$ (to ensure the viability of SRO 2). It is clear that the equilibrium condition (5.2.1) coincides with the equilibrium condition for a single SRO (condition (2.3.1) in chapter 2), when $\bar{q}^2$ equals $\bar{q}^{2'}$ or $\bar{q}^{2''}$. We can also derive the following.

Result 5.2.2: Welfare as a function of $\bar{q}^2$, $W(\bar{q}^2)$, is a single peaked function, taking an internal maximum at a point $\bar{q}^{2*}$, where $\bar{q}^{2'} > \bar{q}^{2*} > \bar{q}^{2''}$.

Proof: As is explained in section 2.3.2, welfare can be identified with h, the number of labourers at equilibrium. Result 5.2.2 will thus hold as long as $h(L_2)$ is a single peaked function, taking an internal maximum

between $\bar{q}^{2"}$ and $\bar{q}^{2'}$. It is clear from equation (5.2.3) that for $\bar{q}^2 = \bar{q}^{2"}$ (and so $M^2 = 1$), $h_{L2} > 0$, whereas for $\bar{q}^2 = \bar{q}^{2'}$, $h_{L2} < 0$. All that is required to complete the proof is to show that when $h_{L2} = 0$, $h_{L2L2} < 0$, which is clear below

$$h_{L2L2}\Big|_{h_{L2}=0} = \frac{X_2(M^2_\xi \bar{q}^2_\alpha)'_\alpha - 2M^2_\xi \bar{q}^2_\alpha}{v - X_1 M^1_n - X_2 M^2_n}$$

which is negative by the assumptions made on $M(h, \bar{q})$. □

From results 5.2.1 and 5.2.2, Shaked and Sutton's central result comes out easily.

Result 5.2.3: If SRO 1 does not change its choice of L_1 in response to the entry of SRO 2, welfare will certainly not decrease and will definitely increase for any choice of $\bar{q}^2$ between $\bar{q}^{1'}$ and $\bar{q}^{1"}$.

This result is important, as it suggests that the introduction of a second SRO will always tend to improve welfare. However, the assumption that SRO 1 will not change its choice of L_1 in response to this entry seems unrealistic. Indeed, in the Hotelling model in this chapter, it will be shown that SRO 1 will *raise* its average quality in response to entry when entry is member-stealing, and it will *lower* it when entry is member-augmenting. The intuition for this will be discussed later. The point here is simply that, in the present model, if SRO 1 were to react to entry by changing its choice of L_1, then the above result might fail to hold.

Moreover, if the profession faces the 'threat' of entry rather than its actuality, Shaked and Sutton themselves show that the introduction of such a threat will have ambiguous consequences for welfare. Their result follows.

Result 5.2.4: If the threat of entry is a feasible one (that is if entry is not blockaded due to L_1 being less than $\bar{q}'$ in the unthreatened single SRO

equilibrium), then SRO 1 will need to expand (lower L_1) if it is to deter entry by a second SRO.

Proof: If SRO 1 can be threatened, then it can be threatened by its last member, which has quality L_1. Thus a second SRO can pose a threat if and only if $M^2 = M(h, L_1) \geq 1$, where h is the equilibrium value for the single SRO. Now, recall from result 2.3.1 that there exists some quality $\bar{q}'$ (given by $M(h, \bar{q}') = 1$) such that all single SROs with quality $\bar{q}^1 > \bar{q}'$ are viable, while no single SRO with quality $\bar{q}^1 < \bar{q}'$ is viable. Since $M(h, \bar{q}') = 1$, a second SRO can pose a threat if and only if $L_1 > \bar{q}'$. Thus, if SRO 1 wants to prevent entry, it must lower L_1 to $\bar{q}'$. □

This reduction in L_1 will improve welfare, at first. Thus, the threat of entry may improve welfare, even if no entry actually occurs. However, there may be overshooting, with SRO 1 reducing L_1 to below its socially optimal level, in which case the overall effect of the threat on welfare will be indeterminate.

5.2.4 Discussion

This model is simple and appealing, but its assumptions are not completely unproblematic. In addition to the reservations noted in chapter 2, there are new problems involved the extension of this model to a two-SRO regime.

Firstly, Shaked and Sutton discuss two polar cases: the case where SRO 1 does not change its choice of L_1 in response to entry by a second SRO; and the case where SRO 1 reduces L_1 by enough to deter entry. More generally, they suggest (but do not model) that we might expect SRO 1 to allow entry but to lower L_1 after entry.

However, this implication is not obvious. In a more simple framework, it is a standard result that an incumbent monopolist will need to (irreversibly) expand output in order to deter entry, but that in the face of actual entry she will reduce output to below its monopoly level (ref. Spence (1977), Dixit (1980)). If this is true for such a simple case,

might we not want to consider the possibility that SRO 1 would in fact *raise* L_1 in the face of actual entry?

It seems necessary, then, to model explicitly SRO 1's choice of L_1 in the face of entry by a second SRO. This is a partial motivation for my model in this chapter.

A second reservation with this model is that consumers are assumed to be identical, while multiple SROs seem to be especially important and likely to exist when consumers have differing preferences for quality (at least when each SRO is only able to offer a single q (or $\bar{q}$)).

For example, suppose that there is just one profession of health workers, all of whose members have to be qualified to surgeon level. Prices would be very high and most simple treatments would go undone (since, for example, a person with a wart would not want to pay anything approaching the same price per hour as would a person who requires open heart surgery). The introduction of a para-profession, which offers simple treatments at lower prices, seems almost certain to increase social welfare in this scenario.

The idea here is that multiple SROs will result in there being a variety of qualities on offer, which will reduce consumers' transport costs and so increase welfare. Again, the need to model this effect explicitly provides a partial motivation for my model in this chapter. This model now follows.

5.3 The number of SROs - a 'Hotelling' model

In this section, I introduce a model which cannot be accused of either of the two objections made above to Shaked and Sutton's two-SRO model: the reaction of an incumbent SRO to entry by a second SRO is modelled explicitly; and explicit assumptions are made about the quality preferences of different consumers. My model is simple, but it gives relatively strong results, many of which seem likely to hold for more general set-ups.

5.3.1 The model

On the demand side, the model is a variant of the Hotelling (1929) model. The Hotelling model used to justify the inverse demand function in chapter 3 was relatively conventional. All consumers had the same valuation of vertical quality, but they differed in their preferences for horizontal quality (and also in their reserve utilities). The model in this chapter considers a different case. Here, consumers differ in their preferences for vertical quality. More precisely, consumers are assumed to differ according to their 'most preferred' vertical qualities.

Suppose that consumer i has a 'most preferred' quality of product ω_i, where ω is distributed uniformly along a line from quality k = 0 to some quality $k = k_{max}$, with a density equal to one. This 'most preferred' quality is assumed to fulfil all the functions the consumer needs of the product. For qualities in excess of a consumer's 'most preferred' quality, that consumer will only care about the price of this quality, and not about how much better it is than his optimal quality. That is, he will face zero transport costs. For qualities which fall below a consumer's 'most preferred' quality, however, that consumer will care both about the price of that quality *and* about how far this quality is from his optimal quality (i.e. he faces positive transport costs).

These assumptions on consumers' transport costs are made simply because they seem intuitively to be correct, at least for situations in which consumers are concerned with receiving certain characteristics from a

product. To take an earlier example, a person with a wart (low ω_i) will be indifferent between being treated by a low-level nurse (low k) and being treated by a top surgeon (high k), if each charges an identical fee, since both would be able to perform this simple task perfectly well. A person who needs open heart surgery (high ω_i), however, will be willing to pay a sustantial premium in order to be operated on by a top surgeon (high k) rather than by a low-level nurse (low k).

More formally, consumer i, who has a 'most preferred' quality ω_i and who buys a good of expected quality $\bar{k}$ at price p, will have expected utility

$$U_i = \underline{u} - p - \tau(\omega_i - \bar{k}) \quad \text{if } \omega_i > \bar{k}$$

and $$U_i = \underline{u} - p \quad \text{if } \omega_i \leq \bar{k}$$

where τ are (linear) transport costs, and $\underline{u}$ is the difference between the utility consumers get from their ideal quality $U(\omega_i)$ and their reserve utility u_o. Note that $\underline{u}$ is assumed to be identical for all consumers. If given a choice of qualities, consumer i will choose the quality which gives him the highest U_i (or nothing if the highest U_i is negative).

The supply side of this model is very simple. It is assumed that there are n_j sellers in each SRO j. Each seller makes (more or less) irreversible quality (training) choices, and faces a quadratic cost of quality function μk^2. Without any loss in generality, we can let $\mu \equiv \lambda\tau$. In what follows, various conditions are placed upon λ. These are required to ensure that a stable equilibrium exists. Sellers also choose outputs, given that they face a unit cost c.

Qualities and outputs are chosen at stages I and II, respectively, of a two stage game. Two alternative assumptions are made about the stage II output choice game

(r) There is no collusion on output choice, so that every seller chooses output by maximizing her own individual profits, given what all the other sellers are doing (Cournot-Nash competition). Note that I only

consider semi-symmetric equilibria, in which every member of each SRO j chooses an identical output x_j.

(m) There is full collusion on output *within* each SRO, but Cournot-Nash competition on output *between* SROs. That is, outputs are modelled as being chosen by the SROs (rather than their individual members), so as to maximize SRO profits, given what any other SROs are doing.

In both cases, the stage I quality choice game involves all the members of each SRO j colluding on quality k_j, but the SROs competing with each other. That is, each SRO j is modelled as choosing its own k_j level to maximize its own profits, given the k choices of any other SROs.

Note that the two regimes (r) and (m) correspond to the identically labelled regimes in chapter 3.

Consumer surplus S is the sum of U_i over all i, and welfare W is a weighted sum of consumer surplus and total industry profits profits Π

$$W = S + \rho\Pi$$

where $0 \leq \rho \leq 1$. There are zero fixed set-up costs for SROs.

In sections 5.4 and 5.5, I analyse the equilibrium (for both (m) and (r)) first when there is only one SRO, and then when there are two SROs. These two sections are relatively dry and mathematical (and could be skipread), their aim being to prepare the groundwork for results sections 5.6 and 5.7. In these sections, these various equilibria are compared, so that we can determine the effects of entry by a second SRO on quality, profits, consumer surplus, and welfare. Intuitions for these effects are also given. Two polar forms of entry are considered:

(S) Member-stealing entry: the second SRO takes half of SRO 1's members on entry (so that the total number of sellers does not increase).

(A) Member-augmenting entry: the second SRO does not steal any members, and enters on the same scale as SRO 1 (so that the total number of sellers doubles).

5.4 Case I - The industry with one SRO

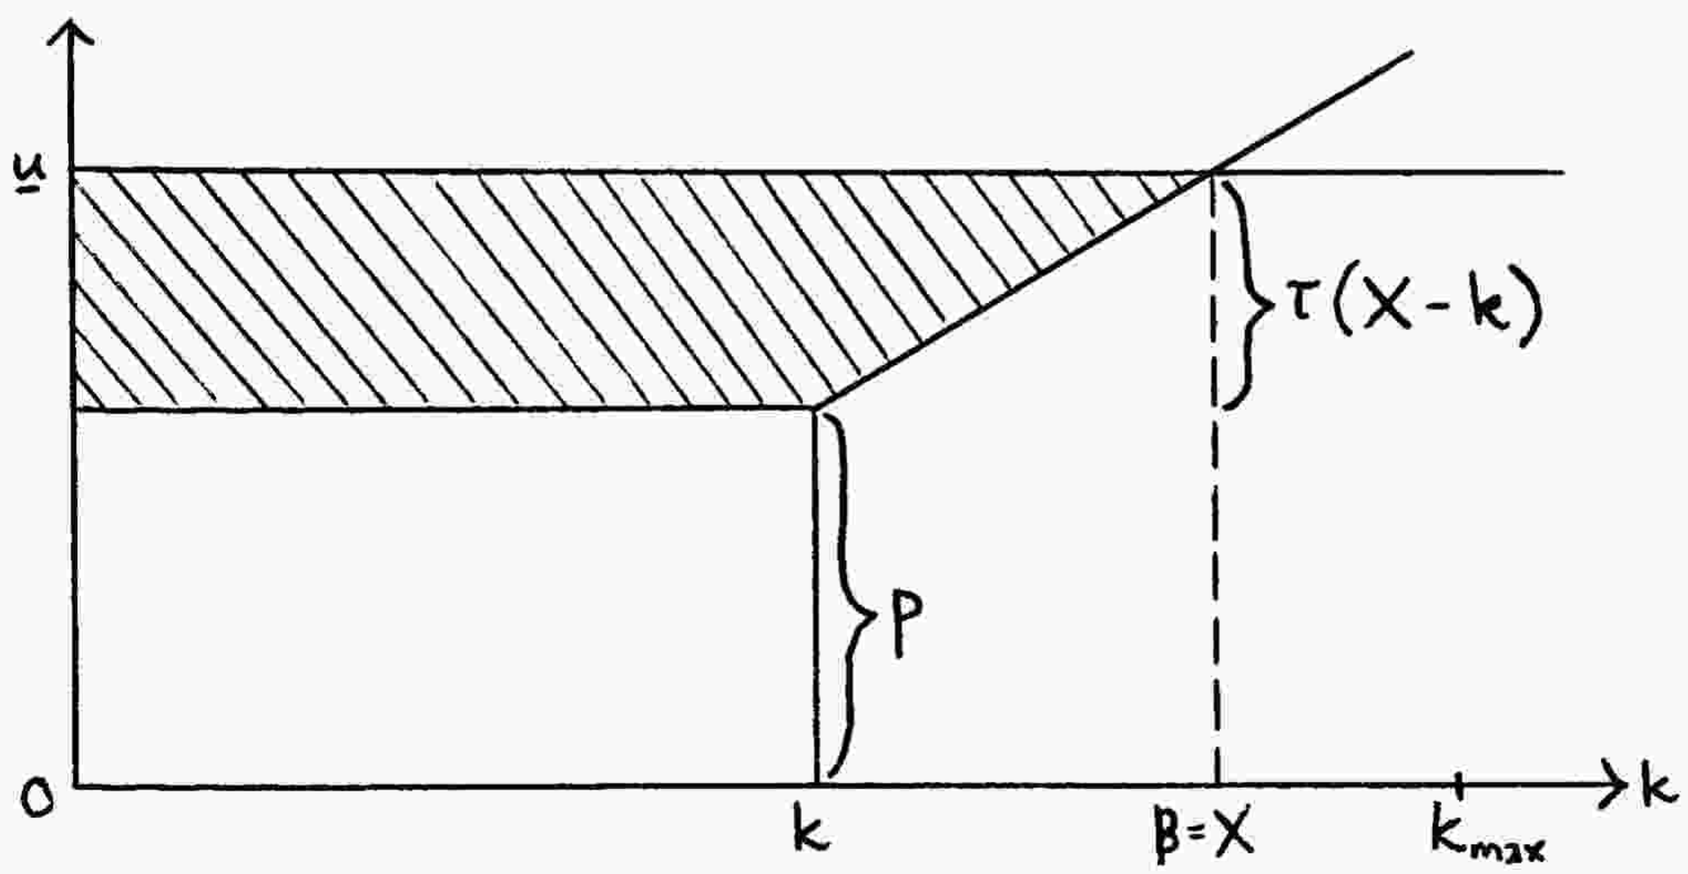
Consider an industry in which there is only one SRO to which all of the n firms in the industry belong. The consumer who is just indifferent between buying a good of quality k and not consuming anything has an optimal value of k, β given by

$$p + \tau(\beta - k) = \underline{u}$$

Since consumers with optimal quality levels below k only care about p, then, as long as $p < \underline{u}$ and $\beta < k_{max}$, the total industry demand X will equal β. Thus we can derive a simple inverse demand function.

$$p = \underline{u} + \tau(k - X)$$

where τ are the linear transportation costs and p is the price of the good. This situation is shown in the diagram below.



Consumer surplus is given by the shaded area in this diagram. More formally, total consumer surplus in case I is given by

$$S_I = k(\underline{u} - p) + \tfrac{1}{2}\tau(X - k)^2$$

$$\equiv \qquad S_I = \tau(X^2 - k^2)/2$$

Using these formulations, we can now find the equilibrium values of p, k, X, Π and S_I, under full collusion (m) and under quality collusion only (r).

5.4.1 Full collusion within each SRO (regime m)

In regime (m), members of an SRO collude on both quality k and output x, choosing k in the first stage and x in the second stage of a two-stage game. At the second stage, the SRO chooses X (= nx) to maximize industry profits which are given by

$$\Pi^m = (\underline{u} - c + \tau(k - X)).X - n\lambda\tau k^2$$

Substituting the first-order Π-maximization condition back into this expression gives us the first stage maximand:

$$\Pi^m = \frac{(\underline{u} + \tau k - c)^2}{4\tau} - n\lambda\tau k^2$$

From the first-order Π^m-maximization condition (with respect to quality), we get

$$k^m = \frac{\underline{u} - c}{\tau(4\lambda n - 1)} \qquad X^m = \frac{2\lambda n(\underline{u} - c)}{\tau(4\lambda n - 1)}$$

$$p^m - c = \frac{2\lambda n(\underline{u} - c)}{4\lambda n - 1} \qquad \Pi^m = \frac{\lambda n(\underline{u} - c)^2}{\tau(4\lambda n - 1)}$$

$$S_r^m = \frac{(\underline{u} - c)^2(4\lambda^2 n^2 - 1)}{2\tau(4\lambda n - 1)^2}$$

and $$W_I^m = \frac{(\underline{u} - c)^2.[(4\lambda^2n^2 - 1) + 2\rho\lambda n(4\lambda n - 1)]}{2\tau(4\lambda n - 1)^2}$$

Note that, for profits to be maximized (rather than minimized) at this point, $\partial^2\Pi^m/\partial k^2$ must be negative, which requires that $\lambda n > 1/4$.

5.4.2 Collusion on quality only (regime r)

In regime (r), members of an SRO still collude on quality k, but there is Cournot competition in output. This can be modelled as a two-stage game, with Cournot competition at the second stage of the game. At this stage, each firm i chooses x^i to maximize its profits π^i, which are given by

$$\pi^i = (\underline{u} + \tau(k - \sum_i x^i) - c).x^i - \lambda\tau k^2$$

Substituting the first order π-maximization condition into the term for total industry profits Π gives us the first stage maximand:

$$\Pi^r = \frac{n(\underline{u} + \tau k - c)^2}{\tau(n + 1)^2} - n\lambda\tau k^2$$

From the first-order Π^r-maximization condition (with respect to quality), we get

$$k^r = \frac{\underline{u} - c}{\tau.[\lambda(n + 1)^2 - 1]} \qquad X^r = \frac{\lambda n(n + 1)(\underline{u} - c)}{\tau.[\lambda(n + 1)^2 - 1]}$$

$$p^r - c = \frac{\lambda(n + 1)(\underline{u} - c)}{[\lambda(n + 1)^2 - 1]} \qquad \Pi^r = \frac{\lambda n(\underline{u} - c)^2}{\tau.[\lambda(n + 1)^2 - 1]}$$

$$S_I^r = \frac{(\underline{u} - c)^2.[\lambda^2n^2(n + 1)^2 - 1]}{2\tau.[\lambda(n + 1)^2 - 1]^2}$$

and $$W_I^r = \frac{(\underline{u} - c)^2.[(\lambda^2n^2(n + 1)^2 - 1) + 2\rho\lambda n.(\lambda(n + 1)^2 - 1)]}{2\tau.[\lambda(n + 1)^2 - 1]^2}$$

Again, for profits to be maximized (rather than minimized) at this point, $\partial^2\Pi^r/\partial k^2$ must be negative, which requires that $\lambda(n + 1)^2 > 1$. This will certainly be true if $\lambda n > 1/4$

5.4.3 A brief comparison of regimes (m) and (r) for the one-SRO case

Because there is output collusion in regime (m), but not in regime (r), we would expect to find that $p^m - c > p^r - c$, that $X^m < X^r$, and that $\Pi^m > \Pi^r$. Simple algebra confirms that these inequalities do hold.

Less obviously, it is also true that $k^m > k^r$. There are two reasons for this. Firstly, the higher price-markup in regime (m) means that there are increased returns to quality, so that the optimal quality will be higher. Secondly, when output cannot be controlled directly, quality may be chosen strategically at stage one so as to reduce the output choices made at stage two. Since equilibrium output at stage two is increasing in quality, strategic output restriction will require a lowering of quality.

In addition to this, since $p^r < p^m$ and $X^r > X^m$, it must be the case that $S^r > S^m$. That is, consumers are better off under output competition, even though quality is lower.

5.5 Case II - The industry with two SROs

Turning now to the situation where there are two SROs in the industry, of true qualities k_1 and k_2, where $k_1 > k_2$. The consumer who is just indifferent between the SRO of quality k_1 and not consuming anything has an optimal value of k, β, given by

$$p_1 + \tau(\beta - k_1) = \underline{u}$$

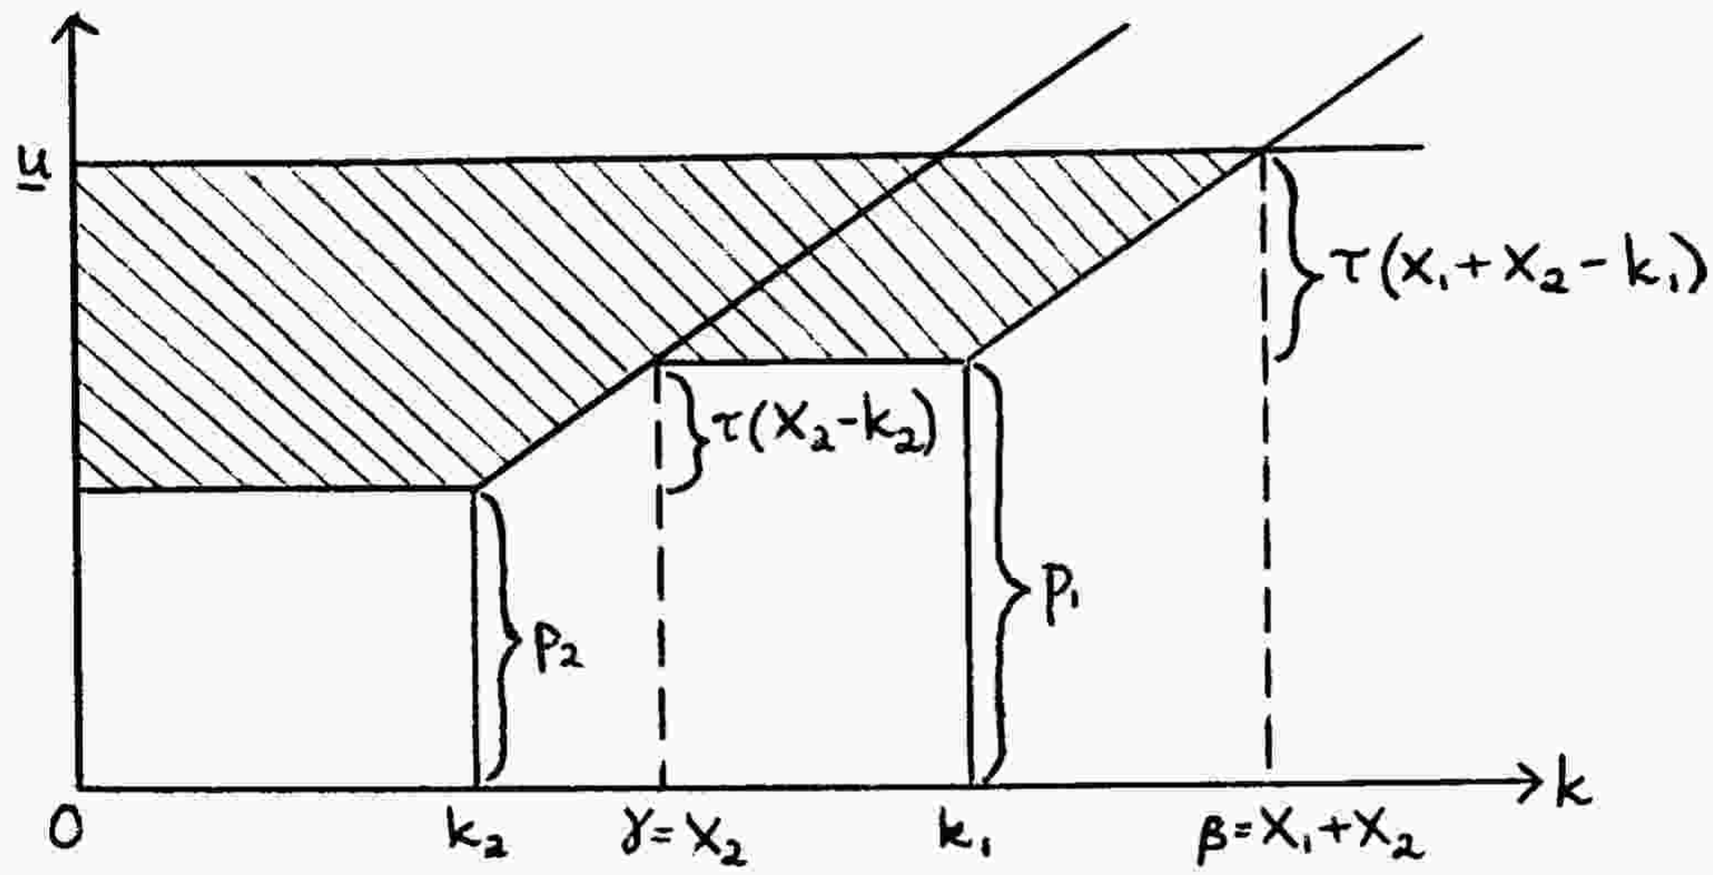
i.e.
$$\beta = \frac{\underline{u} - p_1}{\tau} + k_1$$

The consumer who is just indifferent between the two SROs, of qualities k_1 and k_2, has an optimal value of k, γ, given by

$$p_2 + \tau(\gamma - k_2) = p_1$$

i.e. $$\gamma = \frac{p_1 - p_2}{\tau} + k_2$$

For two SROs, then, the situation is as pictured below.



As long as p_2, $p_1 < \underline{u}$ and $\beta < k_{max}$, the outputs of each SRO are given by

$$X_1 = [1/\tau].[\underline{u} - 2p_1 + p_2 + \tau(k_1 - k_2)]$$
$$X_2 = [1/\tau].[p_1 - p_2 + \tau k_2]$$

These can be rearranged to give inverse demand functions

$$p_1 = \underline{u} + \tau(k_1 - X_2 - X_1)$$
$$p_2 = \underline{u} + \tau(k_2 + k_1 - 2X_2 - X_1)$$

Consumer surplus is given by the sum over all consumers, up to the marginal consumer β, of the reserve utility $\underline{u}$ minus the actual utility each receives. Consumer surplus is thus the shaded area in the diagram above. More formally, consumer surplus in case II is given by

$$S_{II} = k_2(\underline{u} - p_2) + \tfrac{1}{2}\tau(\omega - k_2)^2 + (k_1 - k_2)(\underline{u} - p_1) + \tfrac{1}{2}\tau(\beta - k_1)^2$$

$$= \tfrac{1}{2}\tau[(X_1 + X_2)^2 - k_1^2 + X_2^2 - k_2^2]$$

Using these formulations, we can now find the equilibrium prices, qualities, outputs, profits and consumer surplus under full collusion (m) and under quality collusion only (r).

5.5.1 Full collusion within each SRO (regime m)

In regime (m), members of an SRO collude on both k and X. However, SROs do not collude with each other on either variable. At the second stage of the game, SRO 1 will choose X_1 to maximize its total profits

$$\Pi_1 = [\underline{u} - c + \tau(k_1 - X_2 - X_1)].X_1 - n_1\lambda\tau k_1^2$$

and SRO 2 will choose X_2 to maximize its total profits

$$\Pi_2 = [\underline{u} - c + \tau(k_2 + k_1 - 2X_2 - X_1)].X_2 - n_2\lambda\tau k_2^2$$

By putting together the first order profit-maximization conditions for both SROs, we find that, at the stage two equilibrium,

$$X_1^m = \frac{3(\underline{u} - c + \tau k_1) - \tau k_2}{7\tau}$$

and $$X_2^m = \frac{\underline{u} - c + \tau(k_1 + 2k_2)}{7\tau}$$

Note that $(p_1^m - c) = \tau X_1^m$ and $(p_2^m - c) = 2\tau X_2^m$. Substituting these terms back into SRO 1's profits gives us the stage one SRO maximands:

$$\Pi_1^m = \frac{[3(\underline{u} - c + \tau k_1) - \tau k_2]^2}{49\tau} - n_1\lambda\tau k_1^2$$

and $$\Pi_2^m = \frac{2[\underline{u} - c + \tau(k_1 + 2k_2)]^2}{49\tau} - n_2\lambda\tau k_2^2$$

Putting together the first order profit-maximization conditions (with respect to quality) for both SROs now gives us, at equilibrium,

$$k_1^m = \frac{21(\underline{u} - c)(21\lambda n_2 - 4)}{\tau.[(49\lambda n_2 - 8)(49\lambda n_1 - 9) + 12]}$$

$$k_2^m = \frac{49 \times 4\lambda n_1(\underline{u} - c)}{\tau.[(49\lambda n_2 - 8)(49\lambda n_1 - 9) + 12]}$$

$$X_1^m = \frac{7\lambda n_1}{3}.k_1^m \qquad X_2^m = \frac{7\lambda n_2}{4}.k_2^m$$

$$\Pi_1^m = \frac{\lambda n_1\tau(49\lambda n_1 - 9)}{9}.k_1^{m2} \qquad \Pi_2^m = \frac{\lambda n_2\tau(49\lambda n_2 - 8)}{8}.k_2^{m2}$$

and $$S_{II}^m = \frac{(\underline{u} - c)^2[7^3\lambda^3 n_1^2 n_2(119\lambda n_2 - 32) - 9(21\lambda n_2 - 4)^2]}{2\tau.[7^3\lambda^2 n_1 n_2 - 56\lambda n_1 - 63\lambda n_2 + 12]^2}$$

Note that for any Nash equilbrium to be stable, it must be the case that

$$\frac{\partial^2\Pi_1^m}{\partial k_1^2}.\frac{\partial^2\Pi_2^m}{\partial k_2^2} > \left|\frac{\partial^2\Pi_1^m}{\partial k_1\partial k_2}.\frac{\partial^2\Pi_2^m}{\partial k_2\partial k_1}\right|$$

In this case, the following condition must hold

$$(49\lambda n_1 - 9)(49\lambda n_2 - 8) > 12$$

$$\equiv \quad 7^4\lambda^2 n_1 n_2 - 7^2\lambda((9n_2 + 8n_1) + 60 > 0$$

This condition is more stringent than either second order condition for profit maximization. In the special case where $n_1 = n_2 = \bar{n}$, this condition holds if $\lambda\bar{n} > 12/49$, which is certainly true if $\lambda\bar{n} > 1/4$.

Two further conditions on λ come from the fact that, for the model to make sense, it must be the case that $p_2^m < p_1^m < p_2^m + \tau(k_1^m - k_2^m)$, which comes out to be

$$- 49\lambda^2 n_1 n_2 + 63\lambda n_2 - 12 > 0 \text{ and } \lambda n_2 > 4/7$$

When $n_1 = n_2 = \bar{n}$, these conditions both hold if $0.571 < \lambda\bar{n} < 1.053$.

5.5.2 Collusion on quality only (regime r)

In regime (r), sellers choose outputs under Cournot assumptions in the second stage of the game. In the first stage of the game, the two SROs (rather than individual sellers) choose quality levels, with each taking the other SRO's choices as given when it solves its profit-maximization problem. That is, the solution concept used is Nash Equilibrium.

First consider the second stage Cournot game. Seller i, who is a member of SRO 1 and has quality k_1 will have profits

$$\pi^i_1 = [\underline{u} + \tau(k_1 - X_2 - \sum_i x^i_1) - c].x^i_1 - \lambda\tau(k_1)^2$$

and seller j, who is a member of SRO 2 and has quality k_2 will have profits

$$\pi^j_2 = [\underline{u} + \tau(k_2 + k_1 - 2\sum_j x^j_2 - X_1) - c].x^j_2 - \lambda\tau(k_2)^2$$

By putting together the first order profit-maximization conditions for both SROs, we find that, at the stage two equilibrium,

$$X_1^r = \frac{n_1.[(n_2 + 2)(\underline{u} - c + \tau k_1) - n_2\tau k_2)]}{\tau.[n_1 n_2 + 2n_1 + 2n_2 + 2]}$$

and

$$X_2^r = \frac{n_2.[\underline{u} - c + \tau k_1 + \tau(n_1 + 1)k_2)]}{\tau.[n_1 n_2 + 2n_1 + 2n_2 + 2]}$$

Note that $(p_1^r - c) = \tau X_1^r/n_1$ and $(p_2^r - c) = 2\tau X_2^r/n_2$. Substituting these terms back into SRO 1's profits gives the stage one SRO maximands:

$$\Pi_1^r = \frac{n_1.[(n_2 + 2)(\underline{u} - c + \tau k_1) - n_2\tau k_2]^2}{\tau.[n_1n_2 + 2n_1 + 2n_2 + 2]^2} - n_1\lambda\tau k_1^2$$

and $$\Pi_2^r = \frac{2n_2.[\underline{u} - c + \tau k_1 + \tau(n_1 + 1)k_2]^2}{\tau.[n_1n_2 + 2n_1 + 2n_2 + 2]^2} - n_2\lambda\tau k_2^2$$

Define T as

$$T \equiv n_1n_2 + 2n_1 + 2n_2 + 2$$

Putting together the first order profit-maximization conditions (with respect to k) for both SROs now gives us, at equilibrium,

$$k_1^r = \frac{T(n_2 + 2)(\underline{u} - c).[(n_2 + 2)\lambda T - 2(n_1 + 1)]}{\tau.[\lambda T^2 - (n_2 + 2)^2].[\lambda T^2 - 2(n_1 + 1)^2] + 2\tau n_2(n_1 + 1)(n_2 + 2)}$$

$$k_2^r = \frac{2\lambda T^2(n_1 + 1)(\underline{u} - c)}{\tau.[\lambda T^2 - (n_2 + 2)^2].[\lambda T^2 - 2(n_1 + 1)^2] + 2\tau n_2(n_1 + 1)(n_2 + 2)}$$

$$X_1^r = \frac{n_1\lambda T}{(n_2 + 2)} k_1^r \qquad X_2^r = \frac{n_2\lambda T}{2(n_1 + 1)} k_2^r$$

$$\Pi_1^r = \frac{n_1\lambda\tau.[\lambda T^2 - (n_2 + 2)^2]}{(n_2 + 2)^2}.k_1^{r2}$$

$$\Pi_2^r = \frac{n_2\lambda\tau.[\lambda T^2 - 2(n_1 + 1)^2]}{2(n_1 + 1)^2}.k_2^{r2}$$

and $$S_{II}^r = \frac{\tau}{2}.\left[\left[\frac{\lambda T n_1 k_1^r}{(n_2 + 2)} + \frac{\lambda T n_2 k_2^r}{2(n_1 + 1)}\right]^2 + \left[\frac{\lambda T n_2 k_2^r}{2(n_1 + 1)}\right]^2 - (k_1^r)^2 - (k_2^r)^2\right]$$

The stability condition for this Nash equilibrium is

$$(\lambda T^2 - 2(n_1 + 1)^2)(\lambda T^2 - (n_2 + 2)^2) > 2n_2(n_1 + 1)(n_2 + 2)$$

Again this condition is more stringent than either first order condition. In the special case where $n_1 = n_2 = \bar{n}$, this condition can be written as

$$\lambda > \frac{2[\bar{n}^2 + 3\bar{n} + 2]}{T^2}$$

which is certainly true if $\lambda\bar{n} > 1/4$.

Two further conditions on λ come from the fact that, for the model to make sense, is must be the case that $p_2^r < p_1^r < p_2^r + \tau(k_1^r - k_2^r)$, which comes out to be

$$n_2\lambda^2T^2 - (n_2 + 2)^2\lambda T + 2(n_1 + 1)(n_2 + 2) < 0$$

and

$$n_2\lambda T > 2(n_1 + 1)$$

When $n_1 = n_2 = \bar{n} > 1$, this condition holds if $0.428 < \lambda\bar{n} < 0.799$.

5.6 Results for regime (m): full collusion within SROs

We are now in a position to make comparisons between cases I and II. That is, we can examine how each of the two SROs will react to the actions of the other. There are three things to look at here:

— How will a single incumbent SRO react to the entry, into the industry, of a second SRO, and how will the entrant itself act?

— What will the reaction curve of each SRO look like, in (k_1, k_2) space, when there are two SROs in the industry?

— What would an incumbent SRO need to do to deter entry by a second SRO, and will this ever be a profitable strategy to adopt?

5.6.1 The effect of actual entry (Comparing cases I and II)

To answer the first of these questions, two polar forms of entry are considered:

(S) The member-stealing case, in which $n_1 = n_2 = \frac{1}{2}n = \bar{n}$. That is, the second SRO takes half of SRO 1's members on entry.

(A) The member-augmenting case, in which $n_1 = n_2 = n = \bar{n}$. That is, the second SRO doesn't steal any members, and enters on the same scale as SRO 1.

The absence of subscripts on variables means they come from case I (the one-SRO case). Note that, in order for this model to make sense when $\bar{n}$ is divided by 2, it must be the case that $\bar{n} \geq 2$.

Remember that the two-SRO full collusion model only makes sense for values of $\lambda\bar{n}$ betwen 0.571 and 1.053. These allowable values of $\lambda\bar{n}$ are used in proving the following results.

First let us compare k_1^m, k_2^m and k^m. Simple algebraic manipulation gives us the following result.

Result 5.6.1: Suppose there is full collusion within SROs. When entry is member-augmenting (A), the incumbent SRO, SRO 1, will lower k_1 in response to actual entry, as is suggested by Shaked and Sutton. When entry is member-stealing (S), however, the incumbent SRO will *raise* k_1 in response to actual entry.

Proof: In both case (S) and case (A),

$$k_1^m = \frac{21(\underline{u} - c)(21\lambda\bar{n} - 4)}{\tau.[(49\lambda\bar{n} - 8)(49\lambda\bar{n} - 9) + 12]}$$

and $$k_{\geq}^{m} = \frac{49 \times 4\lambda\tilde{n}(\underline{u} - c)}{\tau.[(49\lambda\tilde{n} - 8)(49\lambda\tilde{n} - 9) + 12]}$$

while, in cases (S) and (A) respectively,

$$k^{mS} = \frac{\underline{u} - c}{\tau(8\lambda\tilde{n} - 1)} \quad \text{and} \quad k^{mA} = \frac{\underline{u} - c}{\tau(4\lambda\tilde{n} - 1)}$$

It is clear that $k^{mA} > k^{mS}$. Using simple algebraic manipulation, we get that $k^{mA} > k_1^{m}$ if $\lambda\tilde{n} > 8/91$, which is true for all allowable $\lambda\tilde{n}$. Thus, when entry is member-augmenting, the incumbent SRO will lower its quality level in the face of actual entry.

Similarly, $k_1^{m} > k^{mS}$ if $\lambda\tilde{n} > 40/161$, which is also true for all allowable $\lambda\tilde{n}$. Thus, when entry is member-stealing, the incumbent SRO will raise its quality level in the face of actual entry. □

The intuition behind this result is as follows. There are three effects going on:

— A differentiation effect: When a new, lower quality SRO enters the industry, an incumbent SRO will have an incentive to raise quality so as to differentiate itself from the new SRO (and so lower the degree of output competition between them).

— A returns to quality effect: As is shown below, prices will be lower when there are two SROs in the industry. This means that the returns to quality are reduced, which will tend to lower k_1^{m}.

— A duplication of costs effect: When the number of sellers in the industry increases, the slope of the cost of quality curve (= $2n\psi k$) increases. This means that there is a tendency for quality to fall as the number of sellers rises.

The duplication of costs effect is not present in the member-stealing case, since the total number of sellers in the industry is the same before and after entry. When only the first two effects only are present, the first will outweigh the second, so that $k_1^m > k^{ms}$. In the member-augmenting case, however, all three effects are present, and the addition of the duplication of costs effect tips the balance, so that $k^{mA} > k_1^m$.

We can also say something about the quality chosen by the entrant, relative to that chosen by a single SRO. Since $k_2^m < k_1^m$ and $k_1^m < k^{mA}$, it must be the case that $k_2^m < k^{mA}$. Algebraic manipulation gives $k_2^m < k^{ms}$ if and only if

$$119\lambda^2\hat{n}^2 - 91\lambda\hat{n} + 12 > 0$$

which holds as long as $\lambda\hat{n} > 0.595$. Thus we get

Result 5.6.2: Suppose there is full collusion within SROs. In the member-augmenting case (A), SRO 2 will certainly set quality lower than will a single SRO. This will also be true for the member-stealing case (S), as long as $\lambda\hat{n} > 0.595$. For $\lambda\hat{n} \in [0.571, 0.595]$, however, the quality set by SRO 2 in the member-stealing case will actually exceed that chosen by a single SRO.

The intuition for this result is as follows. There are two effects going on in the member-stealing case:

— An output competition effect: In this model, the choices made by SRO 1 makes effectively determine the total size of the market, while the choices made by SRO 2 simply determine SRO 2's *share* of that total market. Thus, SRO 2 can raise its output in stage two of the game only by raising its market share. And this is made easier if it raises k_2^m.

— A returns to quality effect: As described above, the fact that prices will be lower when there are two SROs in the industry means that the returns to raising quality are reduced, which will tend to lower k_2^m.

When λ is low, so that costs of quality are low, the first of these two effects will be the more important, so that $k_2^m > k^{ms}$. As λ rises, however, the second effect becomes more important until $k^{ms} > k_2^m$. For the member-augmenting form of entry, there is a third effect: the duplication of costs effect described above. This will tip the balance so that k^{mA} always exceeds k_2^m.

These results are summarised below.

Result 5.6.1/2: Under full collusion within SROs (regime m), relative qualities are

$$k^{mA} > k_1^m > k_2^m > k^{ms} \quad \text{if } \lambda\tilde{n} \in [0.571, 0.595], \text{ and}$$

$$k^{mA} > k_1^m > k^{ms} > k_2^m \quad \text{if } \lambda\tilde{n} \in [0.595, 1.053].$$

for all allowable values of $\lambda\tilde{n}$ (i.e. $\lambda\tilde{n} \in [0.571, 1.053]$).

The results which follow are all derived in the same way as above. All the relevant terms are found in terms of $\tilde{n}$, and are then compared using algegraic manipulation (and, in some case, Taylor's approximation).

For prices, this method gives the following result.

Result 5.6.3: Under full collusion within SROs (regime m), prices will be higher with one SRO than with two SROs, for both the member-stealing and the member-augmenting cases. More formally, relative prices are

$$p^{mA} > p^{ms} > p_1^m > p_2^m$$

for all allowable values of $\lambda\tilde{n}$.

The explanation for $p^{mA} > p^{mS}$ is simply that $k^{mA} > k^{mS}$. The intuition for why both p^{mS} and p^{mA} are greater than either p_1^m or p_2^m is simply that there is inter-SRO competition when there are two SROs which raises outputs and lowers prices.[1]

For profits, the above method gives the following result.

Result 5.6.4: Under full collusion within SROs (regime m), relative profits are

$$\Pi^{mA} > (\Pi_1^m + \Pi_2^m) > \Pi^{mS} > \Pi_1^m > \Pi_2^m \quad \text{if } \lambda\text{ñ} \in [0.571, 0.821] \text{ and}$$
$$\Pi^{mA} > \Pi^{mS} > (\Pi_1^m + \Pi_2^m) > \Pi_1^m > \Pi_2^m \quad \text{if } \lambda\text{ñ} \in [0.821, 1.053]$$

which covers all allowable values of λñ.

There are various things to notice in this result. First, the incumbent's profits will definitely be lowered in the face of actual entry. This is essentially due to the introduction of competition, and concords with Shaked and Sutton's intuition that an incumbent SRO will generally prefer no entry to entry, if it doesn't have to do anything to prevent entry.

---

[1] There is indeed a corollary result to result 5.6.3.

Result 5.6.3': Under full collusion within SROs (regime m), relative outputs are

$$X^{mA} > (X_1^m + X_2^m) > X^{mS} > X_1^m > X_2^m \quad \text{if } \lambda\text{ñ} \in [0.571, 0.652]$$
$$(X_1^m + X_2^m) > X^{mA} > X^{mS} > X_1^m > X_2^m \quad \text{if } \lambda\text{ñ} \in [0.652, 1.053]$$

for all allowable values of λñ.

The intuition for why X^{mA} might be greater than $(X_1^m + X_2^m)$ is that there are two effects working in opposite directions. The first, a competition effect, means that output tends to rise when output competition between SROs is introduced. But the second, a quality effect, says that total output tends to increase with the top quality on offer. So, because k^{mA} is bigger than both k_1^m and k_2^m, there is a tendency for output to fall when a second SRO enters the industry.

Second, for certain parameter values, total industry profits will be *higher* under what is effectively a duopoly than under what is effectively a monopoly. This seems slightly counter-intuitive. Indeed, in the member-augmenting form of entry (A), the duplication of costs of quality means that Π^{mA} is definitely greater than $(\Pi_1^m + \Pi_2^m)$. But there are two other effects going on, which make the total effect on profits ambiguous for the member-stealing form of entry (S). These are:

— A competition effect: This tends to lower prices and profits when entry occurs. This usual effect of competition *is* present in this model, in so far as both p_1^m and p_2^m are lower than either p^{mS} or p^{mA}.

— A price discrimination effect: When there is only one quality on offer, the total transport costs in the economy will be larger than when there are two qualities on offer. Thus, because a monopoly SRO can only provide one quality, it is at a profit disadvantage relative to a duopoly set-up, in which there is some discrimination between customers.

The competition effect and the price discrimination effect will work in opposite directions on industry profits, and which effect outweighs the other will, in general, be ambiguous for the member-stealing form of entry. As λ falls, however, so that transport costs become larger relative to quality costs, this price discrimination effect will rise. Thus, Π^{mS} becomes more likely to exceed $\Pi_1^m + \Pi_2^m$ as λ rises, for given $\bar{n}$.

Similarly, as $\bar{n}$ rises, for given λ, there will be more duplication of quality costs, and so transport costs will become smaller relative to quality costs. The price discrimination effect will thus be reduced, and Π^m will become more likely to exceed $(\Pi_1^m + \Pi_2^m)$.

For consumer surplus, the above method gives the following result.

Result 5.6.5: Under full collusion within SROs (regime m), relative consumer surpluses are

$$S_{II}^{m} > S_{I}^{mS} > S_{I}^{mA} \qquad \text{if } \lambda\tilde{n} \in [0.571, 0.777], \text{ and}$$

$$S_{II}^{m} > S_{I}^{mA} > S_{I}^{mS} \qquad \text{if } \lambda\tilde{n} \in [0.777, 1.053].$$

which covers all allowable values of $\lambda\tilde{n}$.

This result states that, for all allowable $\lambda\tilde{n}$, consumer surplus is higher when there are two SROs than when there is just one. This is partly because transport costs are lowered and partly because output is increased (and prices decreased).

Turning to welfare, it is clear that for $\lambda\tilde{n} \in [0.571, 0.821]$ in the member-stealing case (S), total welfare will certainly be higher when there are two SROs than when there is only one, since total profits are higher. To see whether this result carries over for other $\lambda\tilde{n}$, or for the member-augmenting case, let us consider the situation in which our welfare function accords equal weight to profit and consumer surplus. That is, profit is given its maximum weighting, $\rho = 1$. In this case, the above method gives us the following result.

Result 5.6.6: Under full collusion within SROs (regime m), relative welfares when $\rho = 1$ are

$$W_{II}^{m} > W_{I}^{mA} > W_{I}^{mS}$$

for all allowable values of $\lambda\tilde{n}$.

Since this result holds true for $\rho = 1$, and since consumer surplus is always higher with two SROs than with one, this result must also hold for all $\rho \in [0, 1]$. Thus, when there is full collusion within SROs (regime m), it is certainly true that having two SROs in an industry will be welfare superior to one. That is, any loss in profits is more than made up for by the increase in consumer surplus.

Note that, in order for two SROs to exist in an industry, the choice of the number of SROs may have to be made externally, and not left up to the

industry, since the industry as a whole does better when there is only one SRO, except (for certain parameters) when entry is member-stealing.

5.6.2 Reaction curves

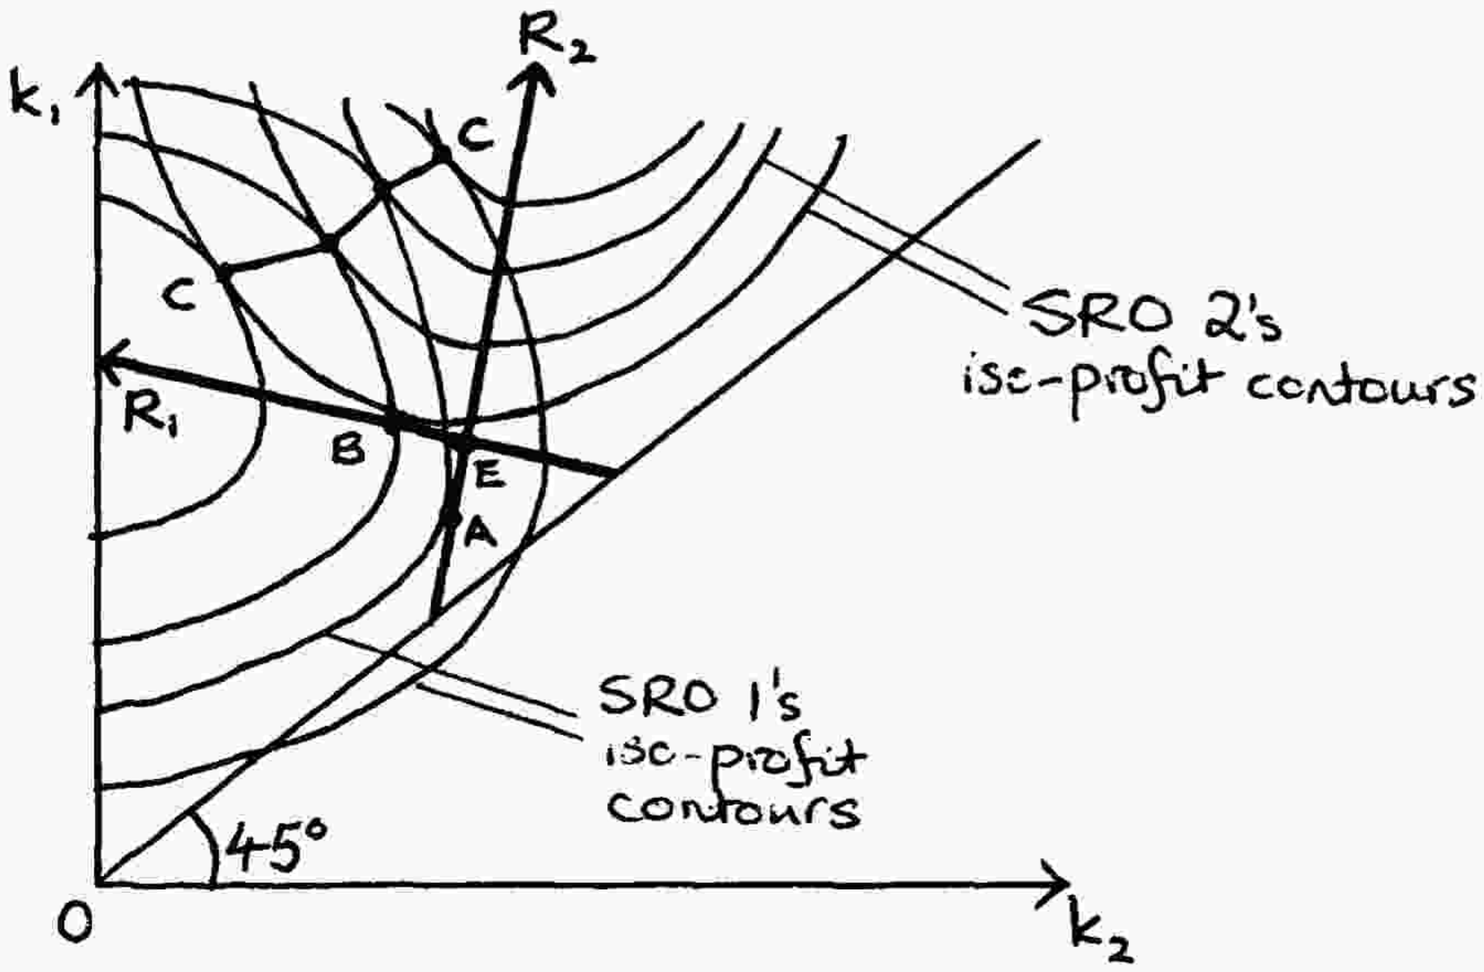
Let us now turn to the question of reaction curves. By total differentiation of the first order profit-maximization (with respect to k_1) condition, we get

$$\frac{dk_1^m(k_2^m)}{dk_2^m} = -\frac{3}{49\lambda n_1 - 9}$$

$$\frac{dk_2^m(k_1^m)}{dk_1^m} = \frac{4}{49\lambda n_2 - 8}$$

Thus the slope of SRO 1's reaction curve in (k_1^m, k_2^m) space at equilibrium will be negative and, for the case where $n_1 = n_2 = \bar{n}$, it will lie somewhere in the range [- 0.070, - 0.158]. And the slope of SRO 2's reaction curve in (k_2^m, k_1^m) space at equilibrium will be positive and will lie somewhere in the range [0.092, 0.200].

These reaction curves are shown in the following diagram.



This diagram is of interest because it tells us a lot about how the two SROs will act if it turns out that the straight Cournot equilibrium E,

which we have so far been analysing, is not the appropriate equilibrium concept. For example:

(a) Points A and B show how the equilibrium choices of k_1^m and k_2^m will change if either SRO has first mover advantage. If SRO 1 is the Stackelberg leader, equilibrium will be at point A, where SRO 1 does better and SRO 2 does worse than at E. The intuition for this result is obvious: By lowering k_1, SRO 1 leaves less of the market open to sales from SRO 2 and so makes higher profits. Similarly, if SRO 2 is the Stackelberg leader, equilibrium will be at point B. This equilibrium is particularly interesting since both SROs are better off at this point. Thus, it seems possible that SRO 1 might do best by deliberately giving SRO 2 first mover advantage.

(b) The contract curve CC shows how equilibrium will change if the SROs are able to collude on their k choices since, along CC, both SROs do better than they do at E. Thus, if collusion is possible we would expect to see k_1^m higher and k_2^m lower than otherwise.

(c) The arrows on the reaction curves show in which direction profits are increasing. This means that, to force the other SRO to leave the industry, SRO 1 would have to lower k_1^m (as in Shaked and Sutton) and SRO 2 would have to raise k_2^m.

5.6.3 Entry deterrence

This last point brings us to the question of entry deterrence. In this model, entry will only ever be deterred if, in addition to the costs mentioned above, there is a fixed cost of functioning in the industry (either a fixed cost of production by sellers or a fixed cost of setting up an SRO). Let this fixed cost be Φ (note that it might be better to let it be $n\Phi$ when it is a fixed cost of production, but the analysis is pretty much the same).

To show the first of the two points in (c) more formally, note that when SRO 2 is on its reaction curve (i.e. when it is profit-maximizing), its profits can be given as a function of k_1

$$\Pi_2^m(k_1) = \frac{2\lambda n_2(\underline{u} - c + \tau k_1)^2}{\tau(49\lambda n_2 - 8)} - \Phi$$

This is clearly increasing in k_1, so that k_1 must be reduced if entry is to be deterred. Entry will, in fact, be deterred only when SRO 1's quality reaches

$$\hat{k}_1^m = \left[\frac{\Phi(49\lambda n_2 - 8)}{2\tau\lambda n_2}\right]^{1/2} - \frac{\underline{u} - c}{\tau}$$

At which point, SRO 1's profits $\hat{\Pi}_1$ are

$$\hat{\Pi}_1^m(\hat{k}_1^m) = \frac{(\underline{u} - c + \tau\hat{k}_1^m)^2}{4\tau} - n\lambda\tau(\hat{k}_1^m)^2 - \Phi$$

One interesting question is whether or not $\hat{k}_1^m$ is lower than k^m, the profit-maximizing quality choice when there is only one SRO. It is straightforward to show that $\hat{k}_1^m < k^m$ if and only if

$$\Phi < \bar{\Phi} = \frac{32\lambda^3 n^2 n_2(\underline{u} - c)^2}{\tau(49\lambda n_2 - 8)(4\lambda n - 1)^2}$$

If $\Phi > \bar{\Phi}$, then the one-SRO profit-maximizing choice of k^m will be sufficient to deter entry. This is blockaded entry. If $\Phi < \bar{\Phi}$, then the incumbent SRO will have to make an effort if it is to deter entry. Clearly, whether or not entry is worth deterring depends on whether or not $\hat{\Pi}_1^m$ is bigger than Π_1^m. When Φ is only a little smaller than $\bar{\Phi}$, then the incumbent SRO's quality will not have to be lowered by very much from in order to deter entry. For Φ close to $\bar{\Phi}$, therefore, it is likely that $\hat{\Pi}_1^m$ is bigger than Π_1^m, and so that entry will be deterred. But for lower Φ, this becomes less likely.

Note also that, while the entry deterrence equilibrium is unlikely to be welfare superior to the two-SRO equilibrium, it may well be welfare superior to the one-SRO equilibrium when Φ is only a little lower than $\bar{\Phi}$. To see this, remember that $d\Pi/dk$ is zero at k^m, so that $dW_I/dk \equiv dS_I/dk$, which can be shown to be negative for all $\lambda\tilde{n} < 1$. Thus, for $\lambda\tilde{n} \in [0.571, 1]$, welfare will be improved (at least at first) when k is lowered from its profit-maximizing level k^m.

5.7 Results for regime (r): Collusion within SROs on quality only

The results for regime (r), where there competition within SROs on output, are stated briefly in this section. Most of these results have been proved by rigorous numerical simulation (using Mathematica, Version 2 (1991)), rather than by analytical methods. Intuitions are only given when the results differ qualitatively from the results for full collusion.

It was shown above that the model makes sense for values of $\lambda\tilde{n}$ between 0.428 and 0.799 only. We also know that $n \geq 2$. The following results make use of these allowable parameter values. Remember that case (S) is that of member-stealing entry, while case (A) is that of member-augmenting entry.

Result 5.7.1/2: Under collusion within SROs on quality only (regime r), relative qualities are

$$k^{rA} > k_1^r > k_2^r$$

and

$$k_1^r > k^{rS}$$

for all allowable λ and $\tilde{n}$. In addition, for low λ and $\tilde{n}$, $k_2^r > k^{rS}$, but as λ and $\tilde{n}$ rise, it becomes more likely that $k^{rS} > k_2^r$.

This result is essentially the same as result 5.6.1/2, and the same intuitions apply.

Result 5.7.3: Under collusion within SROs on quality only (regime r), relative prices are

$$p^{rA} > p_1{}^r > p_2{}^r$$

and $$p_1{}^r > p^{rs}$$

for all allowable λ and $\bar{n}$. In addition, for low λ and $\bar{n}$, $p_2{}^r > p^{rs}$, but as λ and $\bar{n}$ rise, it becomes more likely that $p^{rs} > p_2{}^r$.

More precisely,

— for $\bar{n} = 2$, $p_2{}^r > p^{rs}$;
— for $\bar{n} \geq 5$, $p_2{}^r < p^{rs}$; and
— for $\bar{n} = 3$ or 4, $p_2{}^r > p^{rs}$ for low λ, but $p_2{}^r < p^{rs}$ for high λ.

This result differs from result 5.6.3. Under full collusion within SROs (regime m), we had

$$p^{mA} > p^{ms} > p_1{}^m > p_2{}^m$$

There were two main effects at work behind these results. A quality effect (which tends to relate prices to qualities) and an output competition effect (whereby prices fall as output competition increases). In the full collusion regime, the output competition effect for two SROs outweighs the quality effect, so that $p^{ms} > p_1{}^m$. Under regime (r), however, there is no longer a big increase in output competition for the two SRO case (indeed, output competition might even be reduced). Thus, prices are more likely to relate to qualities, as is seen to be the case by comparing result 5.7.3 with result 5.7.1/2.

Result 5.7.4: Under collusion within SROs on quality only (regime r), relative profits are

$\Pi^{rA} > \Pi_1{}^r + \Pi_2{}^r > \Pi_1{}^r > \Pi^{rs} > \Pi_2{}^r$ if $\bar{n} \geq 3$, and if $\bar{n} = 2$ when λ is high; and

$\Pi^{rA} > \Pi_1{}^r + \Pi_2{}^r > \Pi^{rs} > \Pi_1{}^r > \Pi_2{}^r$ if $\bar{n} = 2$ when λ is low.

The main difference between this result and result 5.6.4 is that Π^{rS} and Π^{rA} are a lot lower, relative to Π_1^r and Π_2^r, than were their full collusion equivalents. The reason for this is straightforward. Under full collusion, the entry of a second SRO introduces output competition where there was none before. So profits are bound to fall by a large amount. When there is already output competition, however, the entry of a second SRO will have a much smaller negative effect on profits, and total profits are actually higher with entry because of the positive price-discrimination effect described above.

Another interesting difference is that $\Pi_1^r > \Pi^{rS}$ for $\bar{n} \geq 3$ (and for high λ when $\bar{n} = 2$), so that SRO 1 might actually prefer member-stealing entry to no entry. This is partly due to the price-discrimination effect and partly due to the fact that output competition between sellers is actually lessened under member-stealing entry.

Result 5.7.5: Under collusion within SROs on quality only (regime r), relative consumer surpluses are

$S_I^{rA} > S_{II}^r > S_I^{rS}$ if $\bar{n} \geq 5$, and if $\bar{n} = 3$ or 4 and λ is low;
$S_{II}^r > S_I^{rA} > S_I^{rS}$ if $\bar{n} = 3$ or 4 and λ is high; and
$S_{II}^r > S_I^{rS} > S_I^{rA}$ if $\bar{n} = 2$

The main difference between this result and result 5.6.5 is that, just as Π^{rS} and Π^{rA} are a lot lower, relative to Π_1^r and Π_2^r, so also S_I^{rS} and S_I^{rA} are a lot higher, relative to S_{II}^r, than are their full collusion equivalents. This means that consumer surplus is no longer always higher when there are two SROs than when there is just one. For example, consider the mose realistic case, where $\bar{n} \geq 5$. Consumer surplus is highest under entry if entry is the member-stealing kind, because prices are not much higher and transport costs are reduced. But consumer surplus is lower under entry if entry is the member-augmenting kind, even though prices are lower, since the duplication of quality costs causes quality to fall.

Turning to welfare, it is clear that $W_I^{rA} > W_{II}^{r} > W_I^{rS}$ when $\tilde{n} \geq 5$, for any ρ, since both consumer surplus and profits are ordered in this way. For lower $\tilde{n}$, the results are not so clear-cut. However, when $\rho = 1$ (so that profits are given their maximum weighting), we get

Result 5.7.6: Under collusion within SROs on quality only (regime r), relative welfares when $\rho = 1$ are

$$W_I^{rA} > W_{II}^{r} > W_I^{rS}$$

for all allowable parameter values.

That is, welfare is higher under entry if entry is of the member-stealing form, but lower under entry if entry is of the member-augmenting form. This result is especially interesting, in that the ranking of welfares is identical to the ranking of total industry profits. Thus, an industry which can choose the number of SROs will make the welfare-optimal decision, without any need for external intervention.

Turning to comparative statics, the reaction curves for this regime have following slopes

$$\frac{dk_1^r(k_2^r)}{dk_2^r} = -\frac{n_2(n_2 + 2)}{\lambda T^2 - (n_2 + 2)^2} < 0$$

$$\frac{dk_2^r(k_1^r)}{dk_1^r} = \frac{2(n_1 + 1)}{\lambda T^2 - 2(n_1 + 1)^2} > 0$$

These slopes are qualitatively the same as for full collusion.

Entry is also deterred in qualitatively the same way.

5.8 The model with Leland-style supply-side

An important problem with Leland's model (as described in chapter 2) is that its demand curve is ad hoc and somewhat unrealistic. My Hotelling model, on the other hand, is very useful for analysing the demand side of the industry, but makes rather ad hoc and unrealistic assumptions about the supply side of the industry. An obvious next step is to combine the supply side of Leland's model with the demand side of the Hotelling model.

The main difference between this new model and the simple Hotelling model is that total output and average quality are interdependent in this model. In the model above, SROs choose both quality levels and outputs. In the Leland framework, SROs have only one choice variable, their minimum quality levels, which determine both average quality levels and outputs.

A second important difference betwen the two set-ups is that, once Leland's supply-side framework is incorporated, the size of each SRO becomes endogenous, instead of being exogenous and fixed. Moreover, the entry question also becomes endogenous insofar as the incumbent SRO can choose to bar completely any potential entry into the industry by a second SRO, by choosing $L = 0$. Or, if each SRO faces an entry cost, then entry can be barred if the incumbent SRO sets $L = \underline{L}$, where $\underline{L}$ is such that an alternative SRO would just make negative profits.

When applied to simple one-SRO and two-SRO frameworks, this model gives qualitatively similar results to the simple Hotelling model. In this section, however, I use it to examine how a single SRO will react to the introduction of a non-SRO sector. This was impossible to model in the simple Hotelling model, because that model relies on quality collusion within sectors which seems unrealistic for a non-SRO sector.

The most interesting results are that

(1) It is ambiguous whether the SRO will raise or lower L in response to entry by a non-SRO sector, but that it will be more likely to raise it as transport costs and costs of quality become higher.

(2) The SRO will certainly make lower profits when there is entry.

(3) Consumer surplus (and, in fact, total welfare) will certainly be higher when there is entry, so that the SRO's preferences with respect to entry will be diametrically opposed to those of consumers.

In section 5.8.5, this model is also used to examine the optimal number of SROs (and whether these will emerge automatically in an industry). The analysis is relatively simple because total output must be one and SROs turn out not to change their L-choices in the face of further entry. So there is a simple trade-off between extra fixed costs and reduced transport costs. One conclusion in this section is that the optimal number of SROs will be increasing in the magnitude of transport costs and decreasing in that of fixed costs. A less obvious conclusion is that, for given transport costs, there will tend to be too many SROs in the industry for welfare-maximization when fixed costs are low, but too few SROs when fixed costs are high.

5.8.1 The model

The set of actual and potential sellers (the size of which is normalized to one) is assumed to consist of sellers with differing *innate* qualities θ, where θ is uniformly distributed on the support $[0,1]$. Thus, for an SRO with top quality $\bar{\theta}$ and minimum quality $\underline{\theta}$, supply $X = (\bar{\theta} - \underline{\theta})$ and expected quality $E(q)$ equals average quality $\bar{q} = \frac{1}{2}(\bar{\theta} + \underline{\theta})$.

It is assumed that each seller can make either zero or one unit of its good, and that a seller with innate quality θ will face a linear cost of quality function $\lambda\tau\theta$. Any sector or SRO also faces a fixed set-up cost Φ. Thus, the total costs for a sector/SRO with quality range $[\underline{\theta}, \bar{\theta}]$ will equal

$$\int_{\underline{\theta}}^{\bar{\theta}} \lambda\tau\theta \, d\theta + \Phi = \lambda\tau(\bar{\theta}^2 - \underline{\theta}^2)/2 + \Phi$$

These costs are assumed to be such that, at equilibrium, all potential and actual sellers could/can cover their costs. That is, there is no adverse selection problem whereby, as price falls, the highest quality sellers leave the industry. This simplifying assumption is not made by Leland.

Lastly, it is assumed that there is asymmetric information so that consumers can only see each sector/SRO's quality range, rather than each individual seller's quality. When a consumer makes a purchase from a given SRO/sector, therefore, his *ex post* utility is a random variable (unless his 'most preferred' quality ω_i is lower than the SRO's minimum quality), for which he is assumed to know the distribution. If we assume that consumers are risk-neutral, then the expected utility of a consumer with 'most preferred quality' ω_i who buys from an SRO with quality range $[\underline{\theta}, \bar{\theta}]$ (where $\omega_i > \underline{\theta}$) is given by

$$EU = \underline{u} - p - \tau . \int_{\underline{\theta}}^{\omega_i} (\omega_i - \theta) . f(\theta | \theta \in [\underline{\theta}, \bar{\theta}]) . d\theta$$

where $f(\theta | \theta \in [\underline{\theta}, \bar{\theta}])$ is the conditional density of θ which, given the above assumptions on the distribution of θ, is given by

$$f(\theta | \theta \in [\underline{\theta}, \bar{\theta}]) = \frac{1}{(\bar{\theta} - \underline{\theta})}$$

By substituting this into the expression for EU, and calculating the integral, we get

$$EU(\omega_i) = \underline{u} - p - \frac{\tau(\omega_i - \underline{\theta})^2}{2(\bar{\theta} - \underline{\theta})} \qquad \text{if } \omega_i > \underline{\theta}$$

$$\text{and} \qquad = \underline{u} - p \qquad \text{if } \omega_i \leq \underline{\theta}$$

This result is interesting in that it gives a rationalization for the use of *ex ante* quadratic transport costs, as long as the initial use of linear transport costs *ex post* is reasonable. In much of the literature on Hotelling models, quadratic transport costs are used, yet there have been few attempts to justify the use of this assumption other than the fact that it is straightforward to work with.

In what follows, two different industry structures will be analysed and compared. These are:

Case I : One SRO and no non-SRO sector
Case II: One SRO plus a potential non-SRO sector

The effect of the potential non-SRO sector in Case II is to allow all sellers to function in the industry, whatever their θ. Thus, Case II is equivalent to a two SRO set-up, in which the second SRO has a minimum quality level of zero.

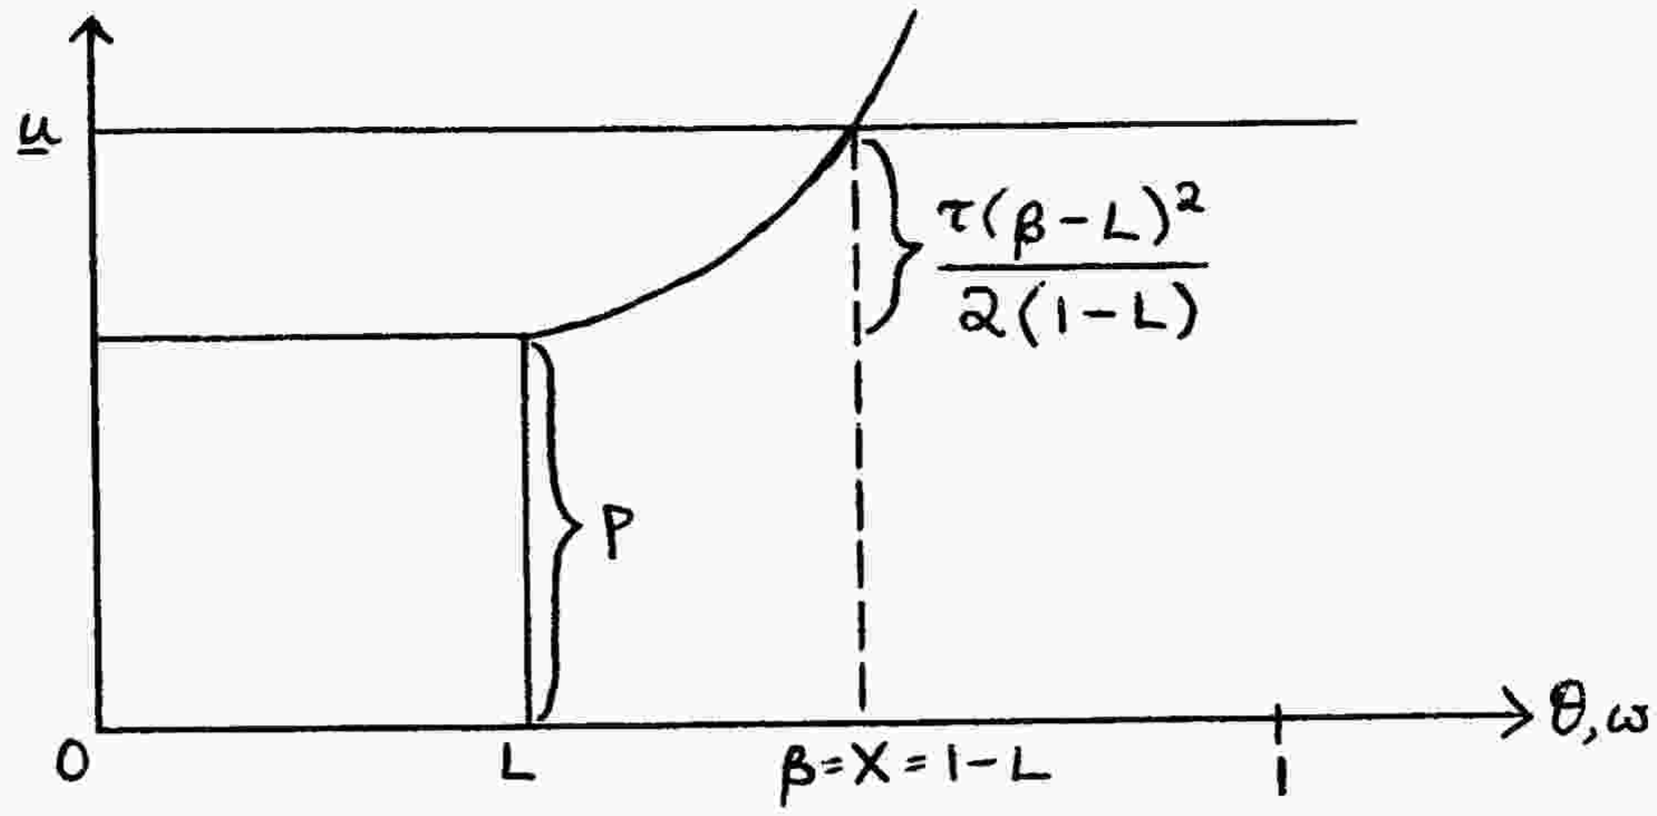
Note that sections 5.8.2 and 5.8.3 are relatively dry and mathematical, their aim being to provide the groundwork for the results in section 5.8.4.

5.8.2 Case I: One SRO and no non-SRO sector

Suppose there is only one SRO, with quality range [L, 1], to which all of the firms in the industry must belong. From the above expressions for consumers' expected utilities, the consumer who is just indifferent between buying a good from this single SRO and not consuming anything will have a 'most preferred' quality β given by

$$p + \frac{\tau(\beta - L)^2}{2(1 - L)} = \underline{u}$$

where p is the price charged by the SRO for one unit of the good. This situation is pictured below



Clearly, β is equal to the demand facing the SRO. If we assume that, at equilibrium, demand will be equal to supply, β must equal (1 - L). Thus, we can derive the SRO's price and profit functions in terms of L.

$$p = \underline{u} - \frac{\tau(1 - 2L)^2}{2(1 - L)}$$

$$\Pi = \left[\underline{u} - \frac{\tau(1 - 2L)^2}{2(1 - L)}\right] \cdot (1 - L) - \frac{\lambda\tau(1 - L^2)}{2} - \Phi$$

The SRO chooses L to maximize Π. Rearranging the first-order Π-maximization condition gives

$$L_I = \frac{2\tau - \underline{u}}{\tau(4 - \lambda)} \qquad X_I = \frac{\tau(2 - \lambda) + \underline{u}}{\tau(4 - \lambda)}$$

$$p_I = \underline{u} - \frac{(2\underline{u} - \lambda\tau)^2}{2(4 - \lambda)(\tau(2 - \lambda) + \underline{u})}$$

$$\Pi_I = \frac{(\underline{u} - \lambda\tau)^2 + \tau(4\underline{u} - 3\lambda\tau)}{2\tau(4 - \lambda)}$$

Note that for profits to maximized (rather than minimized) when $d\Pi/dL = 0$, $d^2\Pi/dL^2$ must be negative, which requires that $\lambda < 4$.

Consumer surplus is found by integrating $EU(\omega_i)$, over all the $(1 - L)$ consumers which consume in the industry. Formally,

$$S_I = \int_0^L (\underline{u} - p)\, d\omega' + \int_L^{(1-L)} \left[\underline{u} - p - \frac{\tau(\omega' - L)^2}{2(1 - L)}\right] d\omega'$$

which comes out to be

$$S_I = \frac{\tau(1 - 2L)^2(2 - L)}{6(1 - L)} = \frac{(\underline{u} + 6\tau - 2\lambda\tau)(2\underline{u} - \lambda\tau)^2}{6\tau(\lambda - 4)^2(\underline{u} + 2\tau - \lambda\tau)}$$

5.8.3 Case II: One SRO (r) and a potential non-SRO sector (u)

Suppose there are two sectors in the industry, only one of which is self-regulated. If the SRO has a quality range [L, 1], then the second sector will have a quality range [0, L]. The consumer who is just indifferent between buying a unit from SRO 1 and not buying anything will have a 'most preferred' quality β, which is given, as before, by

$$p^r + \frac{\tau(\beta - L)^2}{2(1 - L)} = \underline{u}$$

The consumer who is just indifferent between buying a unit from the SRO and buying a unit from the unregulated sector will have a 'most preferred' quality γ given by

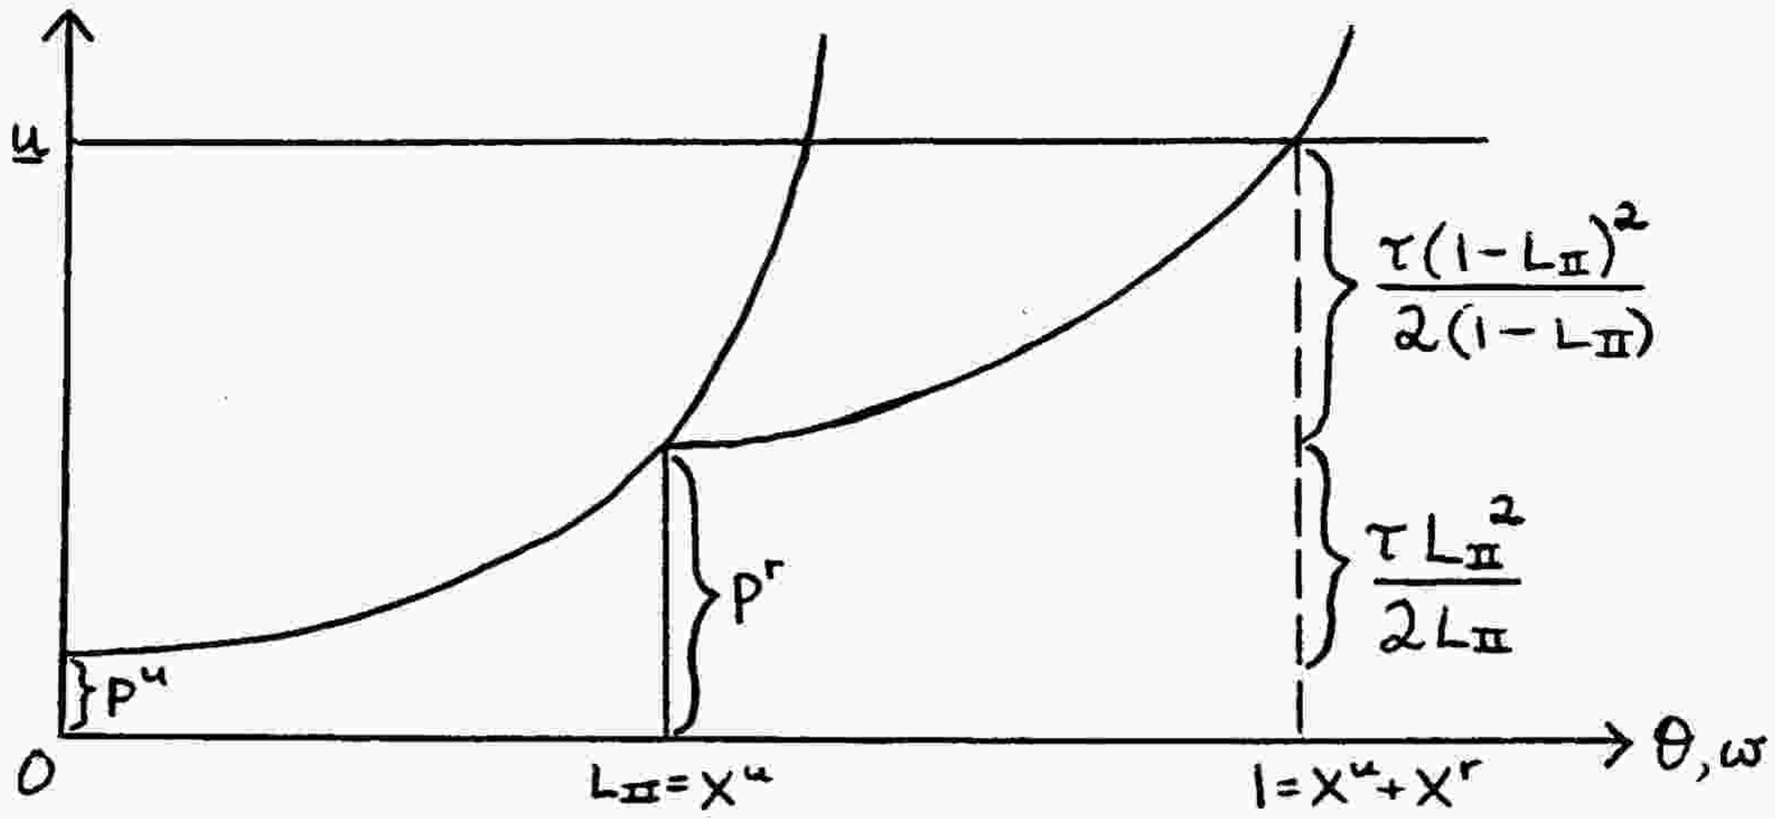
$$p^u + \frac{\tau\gamma^2}{2L} = p^r$$

The first thing to notice about this set-up is that total industry supply must equal one. Thus, when demand and supply for each sector are equal,

$$\beta = (X^r + X^u) = 1,$$

and $$\gamma = X^u = L$$

The situation with one SRO and an unregulated sector is pictured below.



Prices and profits are thus

$$p^r = \frac{2\underline{u} - \tau(1 - L)}{2} \qquad p^u = \frac{2\underline{u} - \tau}{2}$$

$$\Pi^r = \left[\frac{2\underline{u} - \tau(1 - L)}{2}\right].(1 - L) - \frac{\lambda\tau(1 - L^2)}{2} - \Phi$$

$$\Pi^u = \left[\frac{2\underline{u} - \tau}{2}\right].L - \frac{\lambda\tau L^2}{2} - \Phi$$

Note that, in order for the second sector to be profitable, it must be the case that $p_2 \geq 0$, which requires that $2\underline{u} \geq \tau$.

The SRO will choose L to maximize Π^r. Carrying out this first-order profit-maximization and rearranging gives, at equilibrium,

$$L_{II} = \frac{\tau - \underline{u}}{\tau(1 - \lambda)}$$

$$p_{II}{}^r = \frac{\underline{u}(1 - 2\lambda) + \tau\lambda}{2(1 - \lambda)} \qquad p_{II}{}^u = \frac{2\underline{u} - \tau}{2}$$

$$\Pi_{II}^{r} = \frac{(\underline{u} - \tau\lambda)^2}{2\tau(1 - \lambda)} \qquad \text{and} \qquad \Pi_{II}^{u} = \frac{(\tau - \underline{u})(\underline{u}(2 - \lambda) - \tau)}{2\tau(1 - \lambda)^2}$$

Note that, for L_{II} to maximize (rather than minimize) Π^{r}, it is necessary that $\lambda < 1$.

Lastly, consumer surplus is given by

$$S_{II} = \int_0^L \left[\underline{u} - p_{II}^{u} - \frac{\tau\omega^2}{2L} \right] d\omega + \int_L^1 \left[\underline{u} - p_{II}^{r} - \frac{\tau(\omega - L)^2}{2(1 - L)} \right] d\omega$$

which comes out to be

$$S_{II} = \frac{\tau(2 - L_{II} + L_{II}^2)}{6} = \frac{(\underline{u} - \tau\lambda)(\underline{u} - \tau) + 2\tau^2(1 - \lambda)^2}{6\tau(1 - \lambda)^2}$$

5.8.4 Results of allowing free entry: a comparison of cases I and II

We can use the above results to examine the effects of allowing non-SRO members to enter the industry. There are few clear-cut unambiguous conclusions to be made in this section, but it can be shown which factors will tend to affect the results. For simplicity, and without any loss of generality, let $\tau = \mu\underline{u}$. Note that, for the model to make sense, μ must lie in the interval $[1, (2 - \lambda)]$ and λ must lie in the interval $[0, 1]$.

First, let us compare L_I and L_{II}. Using simple algebra we get

Result 5.8.1: The incumbent SRO will raise its minimum quality level in response to entry into the industry by non-SRO members (i.e. $L_{II} > L_I$) if and only if

$$2\mu + \lambda\mu - 3 > 0$$

The sign of this term is ambiguous for the parameter constraints made so far. However, we can say that it will become more likely to be true as either μ or λ increases.

The intuition for this result is as follows. Low λ and μ mean that costs are low, so that the SRO cares mainly about total revenue. When there is no unregulated sector (case I), a reduction in the SRO's price will only result in a small increase in output. Thus, the SRO will tend to maximize revenue by choosing low X, which is equivalent to high L. When there is an unregulated sector in the industry, however, a reduction in price by the SRO will result in a large increase in its output. Thus, the SRO tends to maximize profits by choosing a higher X, which is equivalent to a lower L. Overall, then, when λ and μ are low, $L_I > L_{II}$.

As μ increases, transport costs become more important. This will raise L_I a little bit (since, at the old p and L, demand will now be less than supply). However, it will raise L_{II} by a lot since the SRO must serve the highest ω consumers, and an increased μ means that these consumers will be a lot less willing to purchase from a low L SRO. So, for given λ, L_{II} becomes more likely to be greater than L_I as μ rises.

As λ increases, meanwhile, the cost of quality rises. This will affect the SRO by more than a non-SRO (lower quality) sector. Thus, while L_I will increase slightly (so that price can increase to give a better return to the increased quality costs), L_{II} will increase by a lot, since the SRO is now much less able to compete with the non-SRO sector for mid-range ω consumers. Thus, for given μ, L_{II} becomes more likely to be greater than L_I as λ rises.

Turning to prices, the use of simple algebra gives us

Result 5.8.2: When $L_{II} \leq L_I$, then the SRO certainly lowers its price in response to entry into the industry by non-SRO members (i.e. $p_{II}^r < p_I$). When $L_{II} > L_I$, however, it *is* possible that the SRO's price will be higher when there is entry (i.e. $p_{II}^r > p_I$), although the conditions for this are

fairly stringent. More formally, $p_{II}^r > p_I$ only if $\lambda > 0.856$ and μ is pretty close to $(2 - \lambda)$.

The intuition behind this result is straightforward. There are two factors affecting the SRO's price in cases I and II: the SRO's quality (and so its quality costs) and the degree of competition. When there is a non-SRO sector, prices tend to be lower than when there is just a single SRO (the competition effect). However, if quality is a lot higher in case II ($L_{II} \gg L_I$), price may increased so much as to outweigh the competition effect. So, p_I will generally be greater than p_{II}^r, but it *will* be smaller than p_{II}^r for the highest values of $(L_{II} - L_I)$.

Turning to profits, since p_{II}^r can sometimes be greater than p_I, it seems possible that the incumbent SRO might actually like to allow entry by non-members. However, simple algebra gives us

Result 5.8.3: The SRO will make lower profits when there is entry by non-members (i.e. $\Pi_{II}^r < \Pi_I$), for all allowable parameter values.

This result is intuitive.

Turning to consumer surplus, we get (using algebra and some numerical simulation) the following result.

Result 5.8.4: Consumer surplus is always higher when there is entry into the industry by non-SRO members, for all allowable parameter values.

We have shown that the SRO's profits will always be higher and consumer surplus will always be lower when there is a non-SRO sector. Thus, if the SRO has a choice about whether or not to allow a non-SRO sector, its optimal decision will not be in consumers' interests. We have yet to see, though, which decision would be welfare-superior. This is examined in the following results.

Result 5.8.5: When there is just a single SRO (case I), L_I will be set too high for welfare maximization (as in Shaked and Sutton's model).

Proof: When $d\Pi_I/dL_I = 0$, then $dW_I/dL_I = dS_I/dL_I$, which is given by

$$\frac{dS_I}{dL_I} = -\frac{\tau(1 - 2L_I)}{6(1 - L_I)^2}.(4L_I^2 - 10\ L_I + 7)$$

This is definitely negative. Thus, welfare is maximized at a lower L_I than a single SRO will choose. Note, however, that welfare maximization does not tend to require that $L_I = 0$, since this would involve high transport costs. Indeed, it is easy to show that welfare maximization requires $L_I = 0$ if and only if $\tau \leq 6\underline{u}/5$ (i.e. if and only if transport costs are not too high). □

This indicates that welfare might well be higher when there is a non-SRO sector. Indeed, using algebra and some numerical simulation, we get the following result.

Result 5.8.6: Even with the maximum weighting on profits ($\rho = 1$), total welfare is always higher when there is entry into the industry by non-SRO members, for all allowable parameter values.

Thus, while the SRO does best if entry by a non-SRO sector is not allowed, welfare is improved if it is. In this model, then, the choice about whether to allow entry by a non-SRO sector should not be left up to the SRO.

5.8.5 N SROs with a non-SRO sector

In this section, I look (very briefly) at the case where there are Z+1 sectors in the industry, only Z of which need be self-regulated so that the Z+1th sector has quality range $[0, L_Z]$. The first thing to notice about this set-up is that total industry supply must equal one. Thus, when demand and supply for each sector are equal,

$$\beta = (x_1 + x_2 + \ldots\ldots + x_{z+1}) = 1,$$

$$\gamma = (x_2 + x_3 + \ldots\ldots + x_{z+1}) = 1 - x_1 = 1 - (1 - L_1) = L_1,$$

and, by continuing the series,

$$\omega_i = (x_i + x_{i+1} + \ldots + x_{z+1}) = 1 - (1 - L_{i-1}) = L_{i-1}$$

Let us assume that the costs, for one SRO, of stealing all the firms from another SRO, will be high enough to prevent this from ever happening. That is, it will never be profitable for one SRO to set its minimum quality level at just higher than the most profitable SRO's minimum quality level, in the hope that it would so gain all of the latter's members and profits. This is a questionable assumption, in that it is somewhat *ad hoc*. However, it is perhaps rationalizable if we start from the reasnable assumption that each firm faces a small cost of joining any SRO (which doesn't go to the SRO). This cost will generally be unimportant to a firm deciding whether to join an SRO or remain unregulated, but it may become very important to a firm deciding whether to move form its current SRO to another, essentially identical, SRO.

Let us also assume that each SRO faces a fixed (but not sunk) cost of functioning Φ, which is identical for all SROs, whatever their size. The model says nothing about how SROs recoup this fixed cost from their member firms. It simply assumes that there is some method in place.

From the above, it is straightforward to work out the profits of the 1st and ith sectors

$$\Pi_1 = \left[\underline{u} - \frac{\tau(1 - L_1)}{2}\right].(1 - L_1) - \frac{\lambda\tau(1 - L_1^2)}{2} - \Phi$$

$$\Pi_i = \left[\underline{u} - \frac{\tau(1 - L_i)}{2}\right].(L_{i-1} - L_i) - \frac{(L_{i-1}^2 - L_i^2)}{2} - \Phi$$

Note that, in order for the Z+1th sector to be profitable, it must be the case that $p_{z+1} \geq 0$, which requires that $2\underline{u} \geq \tau$.

At Nash equilibrium,

$$\frac{\partial \Pi_1}{\partial L_1} = \mu(L_1) - \underline{u} + \tau(1 - L_1) = 0$$

$$\frac{\partial \Pi_i}{\partial L_i} = \mu(L_i) - \underline{u} + \frac{\tau(1 + L_{i-1} - 2L_i)}{2} = 0$$

The solution to this set of i equations is easy to find if you work from the equation for i = 1 upwards. That is, find L_1 from $\partial\Pi_1/\partial L_1 = 0$; substitute this into $\partial\Pi_2/\partial L_2 = 0$ to find L_2; substitute this into the condition for $\partial\Pi_3/\partial L_3 = 0$ to find L_3; and so on. The implication of these conditions and their solution is that, whenever there is a non-SRO sector with quality range $[0, L_2]$, then each SRO i will choose the same optimal L_i regardless of many other SROs there are in the industry.

Making welfare comparisons in this case is comparatively easy, since the number of consumers served by the industry is constant (= 1). This means that comparing having i+1 SROs with having i SROs simply involves weighing up the reduction in transport costs against the increase in fixed costs Φ (as long as equal weight is given in the welfare function to consumer surplus and SRO profits). Moreover, the analysis is made even more straightforward by the fact that the i SROs will not change their L_i choices in the face of entry.

Suppose there are currently i-1 SROs in the industry (with $L_{i-1} > 0$) and that we are considering the welfare consequences of allowing an ith SRO to enter. If $L_i = 0$, there would be an unambiguous welfare loss to entry, since there will be an increase in fixed costs (= Φ), but no reduction in transport costs. Suppose, instead, that $L_i > 0$. In this case, it is straightforward to show that the reduction in transport costs as a result of entry, $\Delta(TC)$, is given by

$$\int_0^{L_{i-1}} \frac{\tau\omega^2}{2L_{i-1}}\,d\omega - \int_0^{L_i} \frac{\tau\omega^2}{2L_i}\,d\omega - \int_{L_i}^{L_{i-1}} \frac{\tau(\omega - L_i)^2}{2(L_{i-1} - L_i)}\,d\omega + \frac{\tau L_i}{2}(L_{i-1} - L_i)$$

where $L_0 = 1$ (i.e. zero SROs). This simplifies out to

$$\Delta(TC) = \frac{5\tau L_i}{6}(L_{i-1} - L_i)$$

which will always be positive, but which will clearly be falling as i increases, since both L_i and $(L_{i-1} - L_i)$ will fall. Overall, then, we get the following result.

Result 5.8.7: Entry by the ith SRO will be beneficial to welfare if and only if

$$\frac{5\tau L_i}{6}(L_{i-1} - L_i) > \Phi$$

This result can now be used to determine whether there will tend to be too many or too few SROs in an industry, from a welfarist point of view.

Result 5.8.8: Informally, when Φ is low there will be too many SROs in the industry, and when Φ is high there will be too few.

Proof: It is clear that an SRO will only enter the industry if it will make non-zero profits. From the term for Π^i above, there will be levels of Φ at which entry by an ith seller will be beneficial to welfare, but not profitable, if and only if

$$\frac{5\tau L_i}{6} > \underline{u} - \frac{\tau(1 - L_i)}{2} - \frac{\lambda\tau}{2}(L_{i-1} + L_i)$$

Substituting the Nash equilibrium results into this inequality gives

$$L_i > \hat{L} = \frac{3(2\underline{u} - \tau)(1 - \lambda)}{\tau(2 + 9\lambda - 6\lambda^2)}$$

The term on the right hand side of this expression (which is defined as $\hat{L}$) is definitely non-negative and will be positive when $(2\underline{u} - \tau)$ is strictly positive. This means that if Φ is very low, so that even very low levels of L_i ($L_i < \hat{L}$) are profitable, then there will be too many SROs in the industry. However, if Φ is high, so that only higher levels of L_i ($L_i > \hat{L}$) are profitable, then there will be too few SROs in the industry. □

REFERENCES

AKERLOF, G. 1970. "The market for lemons: Qualitative uncertainty and the market mechanism", *Quarterly Journal of Economics* 84: 488 - 500.

ALLEN, F. 1984. "Reputation and product quality", *Rand Journal of Economics* 15: 311 - 327.

AVERCH, H. and L. JOHNSON. 1962. "Behaviour of the firm under regulatory constraint", *American Economic Review* 52: 1052 - 69.

BARON, D. and D. BESANKO. 1984. "Regulation and information in a continuing relationship", *Information Economics and Policy* 1: 267 - 302.

BARON, D. and R. MYERSON. 1982. "Regulating a monopolist with unknown costs", *Econometrica* 50: 911 - 930.

BESANKO, D. and D. SPULBER. 1989. "Delegated law enforcement and non-cooperative behaviour", *Journal of Law, Economics and Organization* 5: 25 - 52.

BLAIR, R. and S. RUBIN. 1980. *Regulating the Professions: A Public Policy Symposium.* Lexington Books.

CALCUTT, D. 1993. *Review of Press Self-Regulation.* HMSO.

DEMSKI, J. and D. SAPPINGTON. 1987. "Hierarchical regulatory control", *Rand Journal of Economics* 18: 369 - 383.

DEPARTMENT OF TRADE AND INDUSTRY. 1989. *Opening Markets: New Policy on Restrictive Trade Practices Policy.* Cm. 727. HMSO.

DIXIT, A. 1980. "The role of investment in entry deterrence", *Economic Journal* 90: 95 - 106.

DIXIT, A. and J. STIGLITZ. 1977. "Monopolistic competition and optimum product diversity", *American Economic Review* 67: 297 - 308.

FARRELL, J. 1986. "Moral hazard as an entry barrier", *Rand Journal of Economics* 17: 440 - 449.

FLETCHER, A. 1991. *Theories of Self-regulation.* MPhil Thesis.

FRANKS, J. and C. MAYER. 1989. *Risk, Regulation, and Investor Protection.* Clarendon Press.

GILBERT, R. and D. NEWBERY. 1988. "Regulation games", *CEPR Discussion Paper no. 267.*

GOODHART, C. and A. SELDON. 1988. *Financial Regulation - Or Over-Regulation.* Institute of Economic Affairs.

GROSSMAN, G. and C. SHAPIRO. 1986. "Research joint ventures: An antitrust analysis", *Journal of Law, Economics and Organization* 2: 315 - 337.

HOTELLING, H. 1929. "Stability in competition", *Economic Journal* 39: 41 - 57.

KATZ, M. 1986. "An analysis of cooperative research and development", *Rand Journal of Economics* 17: 527 - 543.

KAY, J. and J. VICKERS. 1988. "Regulatory reform in Britain", *Economic Policy* 7: 285 - 351.

KLEIN, B. and K. LEFFLER. 1981. "The role of market forces in assuring contractual performance", *Journal of Political Economy* 89: 615 - 641.

KREPS, D. and R. WILSON. 1982. "Reputation and imperfect information", *Journal of Economic Theory* 27: 253 - 279.

LAFFONT, J-J. and J. TIROLE. 1988. "The dynamics of incentive contracts", *Econometrica* 56: 1153 - 75.

LAFFONT, J-J. and J. TIROLE. 1990. "The politics of goverment decision making: Regulatory institutions", *Journal of Law, Economics and Organization* 6: 1 - 32.

LAFFONT, J-J. and J. TIROLE. 1993. *A Theory of Incentives in Procurement and Regulation.* MIT Press.

LELAND, H. 1979. "Quacks, lemons, and licensing: A theory of minimum quality standards", *Journal of Political Economy* 87: 1328 - 46.

LEWIS, T. and D. SAPPINGTON. 1988. "Regulating a monopoly with unknown demand", *American Economic Review* 78: 986 - 998.

LEWIS, T. and D. SAPPINGTON. 1991. "Incentives for monitoring quality", *Rand Journal of Economics* 22: 370 - 384.

MANKIW, G. and M. WHINSTON. 1986. "Free entry and social inefficiency", *Rand Journal of Economics* 17: 48 - 58.

MARRIS, R. 1964. *The Economic Theory of Managerial Capitalism.* Macmillan.

MAYER, C. and D. NEVEN. 1991. "European financial regulation: A framework for policy analysis", in *European Financial Regulation,* eds. A. Giovannini and C. Mayer. Cambridge University Press.

MILGROM, P., D. NORTH, and B. WEINGAST. 1990. "The role of institutions in the revival of trade: The law merchant, private judges, and the Champagne fairs", *Economics and Politics* 2: 1 - 24.

MILGROM, P. and J. ROBERTS. 1986. "Prices and advertising signals of product quality", *Journal of Political Economy* 94: 796 - 821.

NELSON, P. 1974. "Advertising as information", *Journal of Political Economy* 81: 729 - 754.

OECD. 1985. *Competition Policy and the Professions: A Report by the Committee of Experts on Restrictive Business Practices.* OECD.

ORDOVER, J. and R. WILLIG. 1985. "Antitrust for high-technology industries: Assessing research joint ventures and mergers", *Journal of Law and Economics* 28: 311 - 333.

OSWALD, A. 1985. "The economic theory of trade unions: An introductory survey", *Scandinavian Journal of Economics* 87: 160 -193.

RONNEN, U. 1991. "Minimum quality standards, fixed costs, and competition", *Rand Journal of Economics* 22: 490 - 504.

SHAKED, A. and J. SUTTON. 1981. "The self-regulating profession", *Review of Economic Studies* 48: 217 - 234.

SHAKED, A. and J. SUTTON. 1982. "Imperfect information, perceived quality, and the formation of professional groups", *Journal of Economic Theory* 27: 170 - 181.

SHAPIRO, C. 1983. "Premiums for high quality products as returns to reputations", *Quarterly Journal of Economics* 98: 659 - 679.

SHAPIRO, C. 1986. "Investment, moral hazard, and occupational licensing", *Review of Economic Studies* 53: 843 - 862.

SPENCE, M. 1973. "Job market signalling", *Quarterly Journal of Economics* 87: 355 - 374.

SPENCE, M. 1975. "Monopoly, quality, and regulation", *Bell Journal of Economics* 6: 417 - 429.

SPENCE, M. 1977. "Entry, capacity, investment and oligopolistic pricing", *Bell Journal of Economics* 8: 534 - 544.

SPILLER, P. 1990. "Politicians, interest groups, and regulators: A multiple-principals agency theory of regulation (or 'let them be bribed')", *Journal of Law and Economics* 33: 65 - 101.

SPULBER, D. 1988. "Bargaining and regulation with asymmetric information about demand", *Journal of Economic Theory* 44: 251 - 268.

TIROLE, J. 1989. *The Theory of Industrial Organization.* MIT Press.

TOWNSEND, R. 1981. "Models of money with spatially separated agents", in *Models of Monetary Economics*, eds. J. Karekan and N. Wallace. Federal Reserve Bank, Minneapolis.

WOLFRAM, S. 1991. *Mathematica; A System for doing Mathematics by Computer*. Addison Wesley.