I H P S

Institute for

Health Policy Studies

Contracts and Competition in the NHS

by

Martin Chalkley &
James M Malcomson

Research Paper

University of Southampton

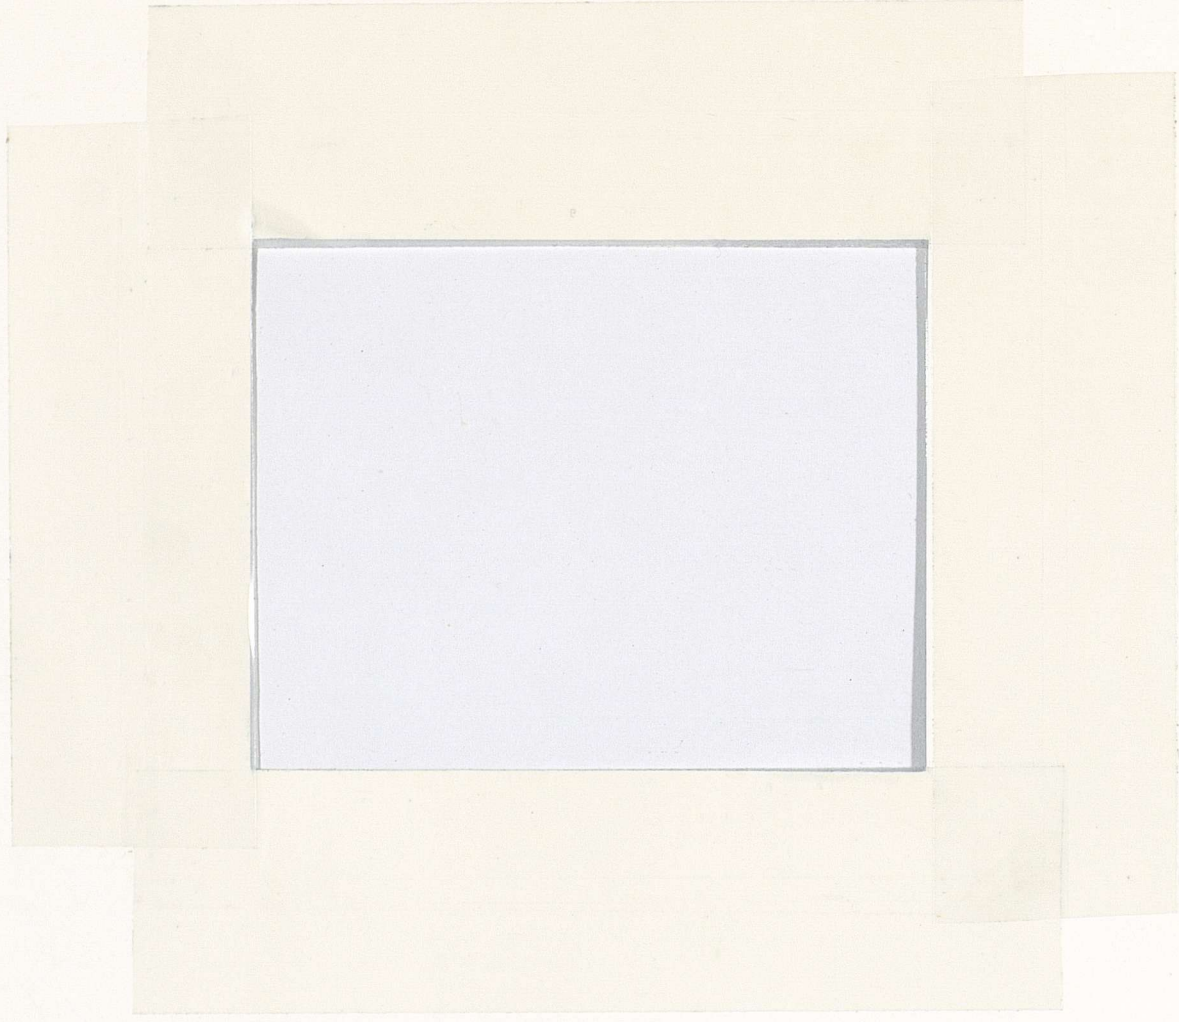

Contracts and Competition in the NHS

by

Martin Chalkley &
James M Malcomson

Contracts and Competition in the NHS[1]

Martin Chalkley and James M. Malcomson

University of Southampton

ISBN 085432 5719

October 1995

---

[1] This paper was originally prepared for presentation at the British Association Festival of Science 1994 held at Loughborough University of Technology, 5-8 September. The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged. The work was part of the ESRC Contracts and Competition Research Programme and was funded by ESRC award number L114251005. We are grateful to Tony Culyer, Ray Robinson and participants in both the British Association meeting and the ESRC Contracts and Competition programme workshop held at Robinson College, Cambridge, 19-20 September 1994, for helpful comments and discussions.

Contents

Foreword

Attitudes towards competition and contracting in the reformed NHS are in a state of considerable flux. On the one hand, there is a widespread view that the early emphasis on competition should be replaced by a more co-operative approach between purchasers and providers. An important element of this view is that longer-term relational contracts between purchasers and preferred providers should replace 'spot market' contracting. At the same time, however, the NHS Executive's recent guidance, The Operation of the Internal Market: Local Freedoms, National Responsibilities indicates that a pro-competition policy is desirable. In pursuit of this aim the guidance sets out in considerable detail the policy that will be adopted to prevent monopoly abuses arising from mergers between purchasers, mergers between providers and collusion. Meanwhile, the growth of standard fundholding, community fundholding and total purchasing all bring additional uncertainty about the appropriate forms of competition and contracting.

Much central policy and local practice in this area has developed pragmatically. Commonsense has prevailed. Although commonsense is not necessarily a poor guide in a complex world dominated by uncertainty, it can sometimes mislead. This is especially true when the complicated consequences of actions are counter-intuitive. In such cases, careful analysis is essential. With this aim in mind, IHPS is pleased to publish Contracts and Competition in the NHS by our colleagues from the Department of Economics, Dr Martin Chalkley and Professor James Malcomson. Their rigorous theoretical analysis of the potential role of competition between providers suggests that competition is not likely to be very important in reducing economic inefficiencies given the countervailing power of health authorities. Neither is it likely to reduce providers' influence over purchasers, although it can be expected to lead to waste in the use of resources. Increasing the responsiveness of providers to patients' wants is seen as more likely to develop through the growth of fundholding than through provider competition. Dr Chalkley and Professor Malcomson also show how contracting can be used to encourage the provision of quality services and how the ways in which this takes place will be determined by market structures and provider objectives. Overall their thoughtful analysis provides a powerful antidote to much of the superficial and glib rhetoric surrounding competition in the NHS.

Professor Ray Robinson
Director
Institute for Health Policy Studies

Abstract

One key element in the recent reforms of the NHS is the separation of *purchasers* from *providers*, with provision of services governed by *contracts* between them. A second key element is *competition* between providers. We discuss the potential roles for competition between providers and assess their significance and implications for the NHS. We also discuss how the amount and quality of services provided is influenced by the form of contract used and how competition interacts with this.

1. Introduction

One key element in the recent reforms of the NHS is the separation of *purchasers* from *providers*. The health authorities that receive funds from the government to arrange for health services to be provided are now separate organisations from the hospitals, clinics, etc. who actually care for patients. The provision of services is governed by *contracts* between those purchasers and providers. A second key element is that there should be *competition* between providers so that no hospital can be sure that it will always get the contract to provide, for example, hip and joint replacements in a particular area of the country unless it provides a competitive service. These two key elements are seen as linked - competition is more genuine if none of the competing providers is part of the same organization as the purchaser.

The purpose of these reforms is to provide a better health service. Whetheror not they will do that is a broad question, to which we would expect an answer only as the result of experience. It nevertheless makes sense to ask the more limited questions of *what effects increased competition might have* and, if there is to be separation between purchasers and providers, *how the kinds of contracts used influences the health services delivered*. These are the questions with which this paper is concerned.

We first consider, in section 2, the potential roles for competition between providers, how significant these are likely to be for the NHS, and the implications of the recent reforms for these roles. Our general conclusion is that the role traditionally seen for competition in limiting the economic inefficiency associated with monopoly power is unlikely to be fundamental to the NHS, at least in its present form. Of more importance, as those employed in health authorities become more and more distanced from their former responsibility for the direct provision of services, is likely to be the role of competitive tendering in providing information both about the prices at which services can actually be delivered and about which providers can deliver those services at the lowest prices. There is of course much more to ensuring adequate health care provision than simply achieving the lowest overall cost to a health authority. With services like health care for which it is not easy to monitor all aspects of the *quality* of the service provided, there is a real danger that awarding a contract to the lowest bidder will result in degradation of quality. It is here that contracts between providers and purchasers have a vital role to play. Contracts have to be designed carefully to maintain quality of service at acceptable levels. We discuss ways to do this in section contracts. We argue there that the most appropriate

form of contract depends on the extent to which quality can be assessed by patients and on how much the particular provider actually cares about the quality of services delivered to patients. We relate our conclusions to the kinds of contracts that have been used in practice. When some potential providers care less than others, encouraging bidding for contracts raises a problem that we discuss in section 4. If contracts are awarded to the lowest bidder, there is a risk that the quality of services will be driven down to below the level that the health authority would wish to see provided. This raises a question regarding the possible alternative ways of ensuring quality.

One alternative to maintaining quality standards by contract is to rely on the reputations of providers. If those who provide good quality services are rewarded by having their contracts renewed whereas those suspected of skimping on quality do not, there can be a strong incentive to provide high quality. This will only work, however, if contracts have a relatively short duration - without that, the threat not to renew becomes less powerful. In section 5 we consider how the use of short term contracts gives rise to its own problems, in particular what are known as the *ratchet* and *hold-up* effects that result in higher costs and less investment in new methods and equipment.

Before we discuss these ideas in detail, it is important to make a few remarks about the framework we use. We do not discuss how the budgets of health authorities are, or should be, determined. We simply start from the undoubted fact that these budgets are not (and are never likely to be) enough for health authorities to purchase all the health services they would like to purchase. Nor do we discuss how health authorities should, or do, decide on what services are to have priority given their limited budgets. The issues of concern here are the implications of competition for the cost and quality of the services provided and how contracts and competition can be used to meet the priorities that have been set with a limited budget.

One criterion we discuss is *economic efficiency*. It is important to understand the precise sense in which this term is used. Economic efficiency is concerned with the health authority getting as far down the list of priorities as possible for *a given level of reward to the providing organisation and the people who work for it*. (*Reward* here is to be interpreted broadly to include nonmonetary, as well as monetary, rewards.) It would not be economically efficient if, for example, a health authority was prepared to pay more for some service than the *minimum* a hospital would accept to provide it but the service was not in fact purchased. Economic efficiency is *not* synonymous with reducing costs, driving wages down, or many of the other

senses in which the term efficiency has been used in debates on the NHS.

In this respect economic efficiency is different from *budgetary efficiency*, which is concerned solely with getting as far down the list of priorities as possible given a limited budget. Budgetary efficiency necessarily entails economic efficiency because an improvement in economic efficiency always enables the health authority to get further down its list of priorities for a given budget. But budgetary efficiency also entails driving down costs wherever possible by reducing the rewards to those providing services, something a health authority may be keen to do. We emphasize the difference between these concepts of efficiency because arguments based purely on economic efficiency are likely to be less controversial than those that make use of the additional criteria in budgetary efficiency.[2]

Our final introductory remark concerns terminology. For concreteness, we often refer to purchasers of health care as health authorities and to providers as hospitals. What we say about the latter, however, applies equally well to other types of providers, such as community health trusts, nursing homes, etc.

---

[2] In technical terminology, budgetary efficiency corresponds to the maximum of a purely utilitarian social welfare function when there is a deadweight loss to raising taxes for the health authority's budget.

2. Competition

There are different interpretations of the word competition. In everyday language, competition is often taken to mean sellers pursuing active strategies against each other - advertising, price cutting, sales promotions and the marketing of new and better products are seen as reflecting fierce competition. Here we use the word in its standard economic sense to refer to the number and type of actual and potential suppliers of a service. The limiting case is *perfect competition* where each supplier is individually powerless to influence the price at which services are provided because of the large number of other potential suppliers.

In general, we see four potential roles for competition: reducing the economic inefficiencies that arise from monopoly power; reducing the influence providers have over purchasers' decisions; ensuring customers receive the types of products and services they want; and providing information to purchasers about the prices at which suppliers *can* provide products and services so that they get what they want at the lowest cost. In this section, we discuss in turn the extent to which each of these applies to the reforms of the NHS.

2.1 *Competition, monopoly power and economic inefficiency*

There is a large literature that investigates the relationship between the *competitiveness* of supply (the number and types of suppliers) and *performance* (measured by quantity sold) compared to what would happen if supply were perfectly competitive. This literature is reviewed extensively in Tirole (1989). Perfect competition is used as a benchmark for these comparisons because, under certain conditions, it results in an economically efficient level of sales. Typically, economic models predict that greater competition (a greater number of competing suppliers) leads to performance closer to the perfectly competitive benchmark. At the other extreme a single (monopoly) supplier chooses a higher price than under perfect competition and, as a result, sales are inefficiently low. It is this prediction that is widely regarded as providing a case for increased competition and that forms the basis of much government competition policy.

The underlying argument is as follows. Perfect competition results in a price that reflects the *marginal cost* of an extra sale (the additional cost to the supplier of supplying one extra sale). This happens because it is in the interest of each individual supplier to continue to expand sales until the price received just covers the cost of supplying the last unit. In deciding how much to buy, customers consider whether an extra unit is worth the price. But when price is

equal to marginal cost, they are in effect considering whether an extra unit is worth the cost of supplying it. A health purchaser who, for example, decides on three extra hip replacements rather than an extra coronary bypass operation on the basis of what a hospital charges is then actually choosing on what it really entails in terms of resources to carry out the operations. Moreover, if the price of everything is equal to its marginal cost, it cannot be that a customer would prefer to have something else *unless it costs more to supply*. But if it costs more to supply, some other customer would, since resources are limited, necessarily be disadvantaged by the switch. One customer's interest can be advanced only by harming some other.

Critical to this argument is that individual suppliers cannot influence the price at which they can sell. If, however, there is a single supplier, that will not be the case. A single supplier can choose whether to sell a little at a high price or more at a lower price. The most profitable choice will be when an additional sale generates as much revenue *(marginal revenue)* as it adds to cost. But marginal revenue is lower than price because, to sell more, the supplier has to lower the price. The additional revenue from cutting the price to achieve an extra sale is thus the price at which that sale is made *less* the amount by which the revenue from all other sales is reduced because the price is lower. Since the most profitable choice of sales is where marginal revenue equals marginal cost and price is greater than marginal revenue, it follows that price is above marginal cost. Thus, when customers base their decisions on prices, these decisions do not necessarily reflect the actual costs of supply - a customer might prefer something different that actually costs no more to supply because prices are not equal to marginal costs. It will then in general be possible to advance one customer's interests without harming those of any others (and, indeed, even the interest of the monopolist), a classic case of economic inefficiency.

The essence of this analysis is that, in order to capture the profits made possible by monopoly power, a monopoly supplier sets price above marginal cost. Since customers base their decisions on prices, those decisions are not then based on the actual costs of supply. In many cases, including health services, the actual cost of supply is not known in advance, so a price set in advance cannot be set precisely equal to marginal cost. What perfect competition ensures in this case is that price is set at the lowest level at which the supplier is willing to undertake to supply in view of the uncertainty in costs. This price is termed the supplier's *reservation price*. But the substance of the argument continues to apply if we replace the term *marginal cost* by the term *reservation price.*

This argument for the economic inefficiency of monopoly hinges on customers simply accepting that the price quoted by the supplier is the price that has to be paid and basing their purchasing decisions on that. This is typically true for individual shoppers in a supermarket. It is certainly not true for health authorities which can, and do, bargain over what they will pay for services. It is well known, for example, that the central buying power of the NHS has enabled it to bargain for prices for drugs supplied to the NHS that are lower than in many other countries.

Where purchasers are active in negotiating prices, there are more efficient arrangements they can reach with a monopoly supplier than simply paying the monopoly price. A purchaser could, for example, offer to pay the reservation price plus a lump sum that ensures the supplier's profits are no less than at the monopoly price. (Some of the cost and volume contracts negotiated between health authorities and NHS hospitals involve what amounts to a lump sum in addition to a price for each treatment.) The supplier would then make as much profit as at the monopoly price and the purchaser would have the advantage of the additional treatments that it regards as cost effective at the supplier's reservation price but not at the monopoly price. The economic efficiency properties of perfect competition would then apply because price would be equal to the reservation price, despite there being only a single supplier. Budgetary efficiency would, in addition, involve reducing the lump sum as far as possible while still getting the service supplied.

Such mutually advantageous economically efficient arrangements are always *possible* when purchasers and providers negotiate over the terms of supply but that does not necessarily mean they will actually come about. The supplier would still make more profit supplying for a slightly lower lump sum than not supplying at all. The health authority, in pursuit of budgetary efficiency, would like to reduce the total outlay so that it can use the savings to purchase other health services and so will try to bargain down the amount of the lump sum. The supplier will bargain to resist this. Will the outcome of their bargaining be economically efficient and how much will the purchaser actually end up paying?

A standard way of answering these questions is with game theoretic models of bargaining. Gibbons (1992) and Binmore (1993) provide accessible discussions of these. A typical conclusion from this literature is that, provided the purchaser and the supplier know enough about each other, they will settle for an outcome that is as good as possible for the purchaser without harming the interest of the supplier. (There are some bargaining arrangements

that may result in other outcomes but for obvious reasons both purchaser and supplier have every interest in avoiding these.) The essential reasoning is as follows. Suppose the purchaser knows the lowest price (the reservation price) at which the provider is prepared to sell any given number and quality of treatments and the provider knows what the purchaser is prepared to pay for any given number and quality of treatments. If the purchaser were prepared to pay more for an additional treatment than the supplier's reservation price, the purchaser would gain without the supplier losing by having that treatment supplied at a price equal to the reservation price. If the supplier must be given part of the gain to be persuaded to agree, that can be achieved by an increase in the lump sum so that they *both* gain. Thus both their interests will be served by always settling for an outcome with the number and quality of treatments that the purchaser is prepared to pay for at the reservation price and bargaining over the lump sum to decide how the gains will be divided. But a change in the lump sum to the advantage of one party can only be at the expense of the other. Thus the bargained outcome will always have the property that the purchaser's position cannot be improved without the supplier being disadvantaged, precisely the criterion for economic efficiency in the provision of the service.

The conclusion we draw from this discussion is that, even with monopoly supply, it is not necessary to introduce competition to prevent the *economic* inefficiency associated with monopoly power provided (as in the NHS as currently constituted) the purchaser and supplier bargain over the prices of the services to be provided. That might change if, for example, a single hospital group were to buy up a chain of hospitals that formed a significant part of the UK market. It might also change as purchasers become more fragmented through the continued growth of GP fundholding. There are, however, other mechanisms for dealing with such issues. Indeed, a mechanism to counteract the fragmented purchasing power of GP fundholders is already developing in the form of purchasing consortia that negotiate with providers on behalf of groups of GP practices. Competition might, of course, still improve *budgetary* efficiency if it increases the bargaining power of purchasers relative to providers. That, however, is a different issue from the economic inefficiency traditionally associated with monopoly. We take up that issue in a later section.

Crucial to the argument just presented is that the health authority and the hospital know sufficient about, respectively, the price at which the hospital is prepared to provide treatments and the amount the health authority is prepared to pay. Without that, bargaining between them may not result in an efficient settlement. To see why, suppose the health authority does not know the price at which the provider is prepared to supply. By pretending that it requires a higher

price than it actually does, the hospital may induce the health authority to pay more than it needs to. So the hospital may have an interest in hiding the price at which it is prepared to supply treatments. Similarly, the health authority may have an interest in hiding how much it is really willing to pay. Unless bargaining reveals the truth, they may not settle on an outcome with the number and quality of treatments that the health authority would be prepared to pay for if priced at the reservation price simply because neither of them knows what that number and quality are. The outcome will not then be economically efficient.

This economic inefficiency is somewhat different from the inefficiency traditionally associated with monopoly power. In essence, it arises from lack of information. Lack of information about the reservation prices of suppliers may also prevent the health authority from achieving provision of services at the lowest cost (budgetary efficiency). We discuss the role competition may play in providing information later. But before we do that we discuss two further roles for competition, the first in reducing the ability of providers to influence purchasers' decisions, the second in helping to ensure that patients receive the services they really want.

2.2 Competition and provider influence over purchasers

The above analysis presumes that health authorities decide their priorities for services independently of the providers. But those individuals responsible for making purchasing decisions may be subject to pressures that affect their decisions. There are interest groups, including the providing institutions themselves, who may like to influence health care purchasing decisions in their private interests. Such groups may attempt to influence purchasers. If they are successful, the public interest may be perverted. Even if unsuccessful, they may use up valuable resources in the process of trying to obtain influence. This is an issue whenever the conduct of public policy, whether the purchasing of health care or the regulation of monopolised industries, is entrusted to individuals who may be susceptible to influence by interest groups. It has been analysed extensively in the case of regulatory bodies, in which it is known as *regulatory capture*. Laffont and Tirole (1993, chapter 11) discuss this issue in detail.

One argument that has been made for increasing competition between suppliers to the NHS is that it breaks down the close personal relationships between employees of the purchaser and those of the provider, thus reducing the influence that providers have over the decisions of purchasers. We know of no formal argument that demonstrates this. However, even if it is

correct, there is another aspect to this issue. Suppliers who compete to influence a purchaser in their favour use up resources in doing so. There is a substantial literature on such *rent seeking* behaviour. See Nitzan (1994) for a survey. That literature suggests that competition between rent seekers is an extremely wasteful process. The resulting waste must be set against any gains from distancing providers from purchasers.

2.3 Competition and patient satisfaction

There is a common perception that where there is a lack of competition, customers have to accept the kinds of products and services that suppliers choose to offer. In contrast, competition between suppliers ensures that suppliers offer the kinds of products and services that customers really want. An obvious comparison is between the reputation for limited range, poor quality and shoddy service under Soviet central planning and the sometimes bewildering variety on offer in some western economies. At root, this would seem closely related to the use of monopoly power as discussed above. After all, as Hicks put it, the best of all monopoly profits is an easy life. However, the argument presented above implicitly took as given the nature of what was being supplied. It was not concerned with the effect of competition on the range on offer and on how that might change over time as one supplier tries to steal a competitive march on the others by designing a new and better product.

There has been a certain amount of discussion of these issues in the literature. See, for example, Lancaster (1979). It is not at all clear, however, that this is relevant in the context of the NHS. A characteristic of the NHS is that purchasers decide what services are to be made available to patients - it is not patients themselves who decide this. Under the recent reforms, purchasing authorities have been told to actively seek the views of local people before deciding their priorities for health care spending. But there is no inherent connection between consulting with users and the separation of purchasers from providers. Nor is it obvious why competition between suppliers should have the effect of meeting patients needs more closely when the decisions on purchasing are made by a health authority and the continued prosperity of the suppliers depends primarily on satisfying that health authority.

Health authorities are not, however, the only purchasers in the reformed NHS. Fundholding GPs also purchase services on behalf of patients. If GPs are more responsive to what their patients want than health authorities, or if patients are in a position to switch between GPs until they find one who is responsive to what they want, then the reforms may result in the

supply of services being more closely aligned to what patients want. Any such effect would, of course, have more to do with competition between purchasers for patients than with competition between suppliers. Whether or not such an effect will come about would seem to depend crucially on the terms under which patients can switch GPs and the response of GPs to patients who want to switch. These are issues on which others have more expertise than we do.

2.4 Competition and information

The argument we presented above for bargaining resulting in an outcome that is economically efficient depends on the parties having sufficient information about each other. In particular, it relies on the purchaser knowing the reservation price at which the supplier is prepared to provide a service and the supplier knowing how much the purchaser is really prepared to pay for that service. Differences in what the purchaser and supplier know may well increase over time in the NHS as the hitherto close relationships between provider and purchaser become more remote. Competition can help to provide the parties with information but, before we discuss how, it is helpful to explain precisely why the lack of information may result in the bargained outcome being economically inefficient. An example demonstrates this.

Suppose a health authority, in view of its limited budget, decides that it is prepared to pay a hospital up to £B for a particular treatment and that the hospital decides it really cannot afford to carry out the treatment for less than £C (so £C is its reservation price), where $B > C$ so that there is certainly a price for the treatment at which both would like it to go ahead. Suppose also that the hospital knows both B and C but that the health authority knows only B. The health authority wants the treatment carried out but it would also like to keep the cost down so that the savings can be used for other treatments. The hospital would like to keep the price high so that it can, for example, afford some new equipment for treating other patients. If the hospital is in a position to dictate the price, it will specify a price of £B (which it knows) and the treatment will be carried out. If the health authority is in a position to dictate price, it has to consider two things. If it offers a price of £B, the treatment will certainly be carried out but it will have no savings to use elsewhere. If it offers a lower price, the treatment will be carried out if the price is above £C and it will have some savings to use elsewhere. But if the price turns out to be below £C, the hospital will refuse to carry out the treatment. That outcome is economically inefficient - both parties would prefer the treatment to be carried out at some price between £C and £B but the lack of information prevents them reaching agreement on that.

In the example given, it is always possible to ensure the treatment is carried out by letting the hospital set the price, the procedure implicit in the NHS Management Executive guidelines. See Dawson (1994) and Ferguson and Palmer (1994) for discussions of those guidelines. Although the health authority may not be happy at how much it has to pay it is always *possible*, in this simple example, to achieve an efficient outcome. But even that may not be the case if the hospital does not know B. In particular, if it is not known for sure that $B > C$, there exists no bargaining procedure that ensures treatment will always be carried out promptly whenever $B > C$. (There may be a way of ensuring that, if they bargain long enough, the treatment will *eventually* be carried out whenever $B > C$, but this too is not economically efficient because it imposes additional suffering on patients awaiting treatment.) The classic statement of this result is Myerson and Satterthwaite (1983).

It is the lack of information that generates the inefficiency. But, if there are many hospitals capable of carrying out the treatment, competition between them can reveal information about the prices at which they would be prepared to carry out treatments. A common form of competition for public services is *competitive tendering* - any supplier who wishes can put in a bid for providing the service and the contract is awarded to the lowest bidder. Standard auction and bidding theory (see, for example, McAfee and McMillan (1987)) tells us a lot about the outcome of such competitive tendering. A standard result is that, if the suppliers have different reservation prices at which they are prepared to provide the service, the contract will go to the supplier with the lowest reservation price at a price that is, on average, just below the reservation price of the supplier with the second lowest reservation price. This is seen most easily in what is called an *English auction* in which bidding continues, with the bids being known by all participants, until all but one of the bidders has dropped out of the bidding. The supplier with the second lowest reservation price will drop out once the price has been bid down below its reservation price. That will leave the supplier with the lowest reservation price committed to providing the service at a price just slightly below the reservation price of the second lowest bidder. With a *sealed bid competitive tender* (in which suppliers make bids that are unknown to the other bidders and the purchaser takes the lowest bid), the winning bid may turn out higher or lower than with an English auction but on average it is exactly the same.

Bidding for contracts thus reveals (at least on average) the price at which the second lowest bidder is prepared to provide the service. As long as the purchaser is prepared to pay at least that for the service, it will be provided. Bidding therefore reduces the extent of the economic inefficiency associated with bargaining discussed above. Moreover increased

competition, in the form of more bidders, on average reduces the price that the purchaser has to pay, see McAfee and McMillan (1987, p. 711). In that way it also promotes budgetary efficiency. The essential intuition is that the more bidders there are, the lower on average is the price the second lowest bidder is prepared to accept. Thus increased competition has two advantages for the purchaser. First, it contributes to economic efficiency by reducing the likelihood of the service not being provided when the purchaser does not know the reservation price of a provider. Second, it contributes to budgetary efficiency by lowering the price that the purchaser has to pay.

But getting the service provided at the lowest price may not be such a good idea if the lowest price service is also a lower quality service. Ensuring that this does not happen is a serious issue in bidding for health services, as indeed for many other public services. The reason is that specifying and monitoring the quality of the service to be provided is not straightforward. If competitive bidding is to lead to a *good*, as well as a cheap, service the contracts offered for tender must be designed carefully to maintain quality. We discuss what can be done in the next section.

3. Contracts

3.1 Contract enforcement

Contracts between purchasers and providers are used to govern the quantity and quality of the health services provided. But what can be achieved by contract is limited by what can be enforced, by going to court if necessary. Contract provisions may be unenforceable because they are not regarded as legally binding. That is the case with contracts between two NHS bodies (though not between, say, a health authority and a privately run nursing home) because the National Health Service and Community Care Act (1990) specifically states that these contracts are not subject to contract law. (There are other branches of the law under which *NHS contracts*, as these are known, might still be enforceable. See Barker (1993, 1995) for a discussion of this issue.) For NHS contracts there is, however, an arbitration procedure that replaces enforcement by contract law.

What the parties intend to happen under a contract may also not be enforced because a court (or arbitrator) does not have the information to enforce it. It might seem that the parties could deal with this by ensuring that the contract is clear. But it is not that simple. Even contracts that are drawn up carefully may remain obscure. The parties may not have envisaged some possibilities, so no provision is made for them. Even if they do envisage all the possibilities, it may just be too costly or time consuming to write all the relevant details into a contract. And even if they would like to write all the relevant details into a contract, they may be unable to do so in such a way that a court can enforce their intention. The first two of these are self-explanatory. An example may be useful to clarify the third. It may simply not be possible to describe sufficiently precisely in a contract the circumstances in which a health authority wants a particular medical procedure to be carried out in such a way that, if a dispute arises over whether that procedure should have been carried out, a court (or arbitrator) can verify months later whether those circumstances actually applied to a particular patient. Faced with a dispute, it is the job of the court or arbitrator to make a judgement one way or the other but the judgement may not be what the parties would have written into the contract had they been able to do so. When the information needed to enforce the parties' intentions is available to the court or arbitrator, that information is said to be *verifiable*.

In practice, lack of verifiability is a serious issue in contracting for health services. This is particularly true for those aspects of a contract that are concerned with the quality of service.

Conventions on medical malpractice set certain minimum standards but these may well be substantially below the standards a health authority would like. Somewhat higher standards may be enforced by careful monitoring and documentation to ensure the relevant information is verifiable. But, since monitoring the quality of services is costly and time consuming, a health authority cannot possibly hope to monitor the quality of service provided for each patient treated. The best it can hope to do in practice is check up on the services provided to a small number of patients by a quality assurance procedure. Where it is costly to make relevant information verifiable, it may make sense to design contracts that do not rely on it being verifiable. It is to the design of such contracts that we now turn. In discussing contract design, we use the term *quality* as a shorthand for those dimensions of a service that are of concern to purchasers and that the parties to the contract either cannot enforce by contract or for which they decide not to set up the costly mechanism to enable them to do so.

3.2 Contracting for quality of service

A major issue in contracting for health services is ensuring that an appropriate quality of service is provided while maintaining incentives for hospitals to keep costs under control. Central to this is how a hospital responds to a particular contract. That in turn depends partly on the extent to which patients both correctly perceive the quality of treatment that a hospital offers and are in a position to respond in choosing where they are treated.

If only one aspect of quality matters and patients respond quickly to that in their choice of where they seek treatment, the health authority can achieve precisely the level of quality it would like to achieve, while retaining economically efficient incentives to keep costs under control, by paying the hospital a fixed price per patient treated (a *cost per case contract*, or in the US terminology a *prospective payment* system), see Ma (1994). The intuition for this is that hospitals get paid more by attracting more patients and do this by increasing the quality of service they provide until the cost of providing additional quality offsets the price per treatment. By appropriate choice of price, the health authority can induce any level of quality it wishes. Moreover, with a fixed price per case, the hospital retains any cost savings it makes and so has an incentive to keep costs down.

In practice, there are many important aspects to quality such as quality of medical services, quality of nursing services, quality of hotel services, and many more. In Chalkley and Malcomson (1995a), we show that the same conclusion applies in this case but only if patients

perceive *every* aspect of the quality of service correctly and respond to all those aspects. Patient demand then keeps all aspects of quality in line even when patients do not themselves pay for the services they receive. This result emphasises the importance of providing information about quality to patients, or at least to the GPs who guide their choices of where to be treated. But it also indicates the limitation of using this mechanism to ensure quality. As has been emphasised certainly since Arrow's seminal (1963) article, health care is a service for which there are particular problems in ensuring that patients know as much as providers about the outcomes of the services they receive.

There are also other problems with relying on this mechanism. First, for hospitals that have long waiting lists for treatment, little incentive to provide quality in order to attract yet more patients is provided by a payment that is received only when the cases are actually treated. Second, patients may not be in a position to choose where they are treated on the basis of their perceptions of quality. This is particularly true for patients whose treatment is paid for by a health authority that has a contract with only one provider for that treatment. Third, some aspects of quality (*experience*, as opposed to *search*, aspects of quality in the jargon) can be assessed only while being treated. For many health services, an individual patient does not receive frequent treatment, so that patient's own demand is not affected by the quality received. The demand for services is then affected by quality only through the effects of reputation and only in the future, which may temper a hospital's incentive to provide those aspects of quality.

The first of these problems can be overcome by making part of the payment for each patient payable at the time a patient is added to the waiting list. By appropriate choice of payment, the health authority can then provide exactly the same incentives for the hospital to provide quality as when there is no waiting list. The second, however, seems fundamental to NHS services purchased by health authorities (in contrast to those purchased by GP fundholders). The third seems fundamental to any system of providing health services. It is thus important to consider ways in which contracts can ensure high quality services when patient demand will not do so.

In discussing this issue, it is important to distinguish between hospitals with different objectives. One possible hospital objective discussed in the literature is maximising profit. But it is important to remember that, even in the US, there are substantial numbers of hospitals that are non-profit institutions alongside hospitals that operate for profit. For understanding the contracting issues that arise when demand does not respond to quality, it is instructive to

consider three different types of objectives hospitals may have. The first is pure *self interest*. This is not synonymous with maximising profit. Non-profit institutions can be equally self interested (about, for example, the perks and the power of those who run them). But some hospitals are no doubt *benevolent* in the sense that they genuinely want to do what is best for their patients for its own sake, not just because of the contractual consequences. Given a budget to spend freely, such a hospital will spend it on its patients in exactly the way a health authority that has its patients' interests at heart would want the hospital to spend it. A hospital cannot, though, be expected to know the health authority's priorities for care purchased from other hospitals and so even a benevolent hospital may choose to spend all its budget on its *own* patients. (Since our brief here is not concerned with how the priorities of health authorities reflect those of patients, in what follows we discuss "benevolence" solely with respect to the health authority's priorities.) These cases are extremes - but extremes are useful for isolating analytical issues. Between them lies what we term a *partially benevolent* hospital that gives some weight to the goals of the health authority but also gives some weight to its own self interest. Even with the best will in the world, it is only natural that hospital employees have their own particular concerns.

At first sight, it might seem that there would be no contracting problems with benevolent hospitals even if the demand for their services does not respond to the quality they provide - given a budget to spend freely, they would choose to do what the health authority would want them to do, so why not just give them the money and let them get on with it? But there is an important question of *how much* money to give them. If they are given a fixed amount to treat patients with a particular diagnosis (a *block contract* as it has come to be known) but the number of patients to be treated turns out to be larger than expected, the hospital will not have enough money to treat them all. It may end up closing wards towards the end of the financial year (a familiar characteristic in UK hospitals in recent years) because it has run out of money to treat non-emergency patients as the result of treating a lot of patients earlier in the year. If, on the other hand, the hospital is given enough money to treat the maximum number of patients to be treated under any circumstances and there turn out to be fewer than that, the health authority will have used up funds that it might have preferred to use elsewhere. With a benevolent hospital, a block contract will work well where the number of patients to be treated is known in advance - for example, non-emergency procedures for which capacity for treatments is limited and always fully utilized.

Where (as is obviously the case with emergency procedures) the number to be treated is not known in advance, ensuring that the hospital has just enough money to treat the appropriate number of patients can be achieved only if the amount it receives depends on the number to be treated, that is, the demand for its services. It might seem that an appropriate form of contract would then be to pay the hospital the cost of treating patients to the desired quality level for each patient treated. However, even with a benevolent hospital, that can distort the hospital's treatment decisions. The reason is slightly subtle. Because a benevolent hospital has the same objectives for its patients as the health authority, it will make the decisions that the health authority would want it to make about how to spend a *fixed* budget on them. But, if it receives a payment for *each* patient treated, it can increase its budget by treating more patients. Since it would always like to have more funds (to, for example, treat more non-emergency patients so as to reduce its waiting lists), it has an incentive to treat too many patients with a somewhat cheaper service in order that it can treat more patients in total. The detailed argument is given in Chalkley and Malcomson (1995b). One way to avoid this is to pay the hospital a fixed amount plus, for each patient treated, the cost of treatment at the desired quality level, but only up to the maximum number of patients the health authority would like to have treated given the length of the waiting list. This is one form of what is called a *cost and volume contract*. This discussion emphasizes the importance of putting a maximum on the number of patients to be treated in non-emergency categories (a maximum that may depend on the number wanting treatment) even with a benevolent hospital whose goals do not conflict with those of the health authority.

Now consider the other extreme, a self-interested hospital. If the demand for its services does not respond to the quality it provides, a purely self-interested hospital will always skimp on unverifiable quality unless its actual costs in providing high quality services are directly reimbursed. There are two difficulties with achieving quality by cost reimbursement. First, it must be possible to verify what the actual costs of a particular treatment are. Second, even if actual costs can be verified, it is not straightforward to establish whether those costs were actually necessary to achieve the desired quality. With a purely self interested hospital there is thus a dilemma. Either costs are reimbursed, in which case high quality services may be provided but with the risk that costs will soar, as with the cost reimbursement rules previously used for Medicare in the US. Or costs are not reimbursed, in which case there is a strong incentive for the hospital to keep costs down but quality will suffer. The only way to overcome this dilemma is by auditing procedures that make quality verifiable - there is nothing, short of specifying in detail what is to be done under each set of circumstances and monitoring it, that the design of a contract can do to help.

If the health authority decides to go for cost saving, the best that can be done is to pay the hospital a lump sum plus, for each case treated, a fixed amount that is independent of the number of cases treated. This takes us back to a *cost per case* contract or *prospective payment* system. The fixed amount per case treated should reflect the amount the health authority is prepared to pay for treating an additional patient at the quality level it can actually get the hospital to supply with the auditing procedures it has available. The economic intuition for this is as follows. Faced with a fixed payment per case, a self-interested hospital will always take full account of the costs of providing treatment. What it has to be induced by the contract to take account of is what the health authority regards as the benefit of treatment. Faced with the cost per case contract just described, it will treat patients up to the point at which the benefit from treating an additional patient (as reflected in the health authority's view about how much it is prepared to pay for the quality of service that will actually be delivered) just covers the additional cost. But unless quality can be enforced by effective auditing procedures, the quality of services will be lower than the health authority would choose if it could ensure that the chosen quality would be delivered.

We come now to the third type of hospital objective, a partially benevolent hospital. There are two forces pushing a hospital in the direction of benevolence. One is a natural concern for the welfare of patients, even if that concern is not so strong as to completely swamp self interest. That is a genuine case of partial benevolence. The second is not what one would normally think of as benevolence but has very much the same implications for the present analysis. Even if quality cannot be enforced by audit procedures, the hospital may be concerned that, if it delivers low quality services, word will get around, its reputation will suffer and, as a result, it will get fewer cases to treat in the future and possibly worse terms out of the health authority in any future contract negotiations. We discuss reputations in more detail in the next section. What is relevant for the moment is that reputation may not be sufficient to induce the hospital to act as if it were fully benevolent. The reason is that the effect of lower quality on reputation may well be uncertain and, in any case, loss of reputation has consequences only for the future and that may be of less concern to the hospital than the more immediate present.

But whatever the reason for partial benevolence, it enables a health authority to get better quality services without full audit than from a purely self-interested hospital even when demand does not respond to the quality provided. However, as explained in Chalkley and Malcomson (1995b), it will typically make sense for the health authority to include some degree of cost reimbursement in the contract as long as the monitoring of actual costs is not itself too costly.

The intuition is that having no cost reimbursement would provide the correct incentives to keep costs down but result in too low quality. At the margin, it is worth relaxing the incentive to keep costs down in order to achieve somewhat higher quality.

This is theory. Do hospitals actually respond to incentives in the way theory suggests? The change in the payment system for Medicare in the US in 1983 from cost reimbursement to prospective payment dramatically altered the incentives for hospitals to cut costs of treatment because with prospective payment additional costs are borne entirely by the hospital instead of being reimbursed from the Medicare budget.[3] Studies of hospital costs reviewed in Ellis and McGuire (1993) indicate that, before the change, the average cost per hospital discharge *increased* at about 2% per year. In the years immediately following the change, the cost per discharge *declined* by roughly 2% per year. The change in incentive system may not have been the only reason for the change in costs but the evidence is at least suggestive that it had an impact in the direction theory predicts. Hodgkin and McGuire (1994) have a fuller discussion of how the changes resulting from the introduction of prospective payment correspond to the predictions from a theoretical framework.

[3] Even under prospective payment there is some provision for reimbursing hospitals for exceptionally expensive cases but this applies to only a small proportion of cases.

4. Bidding for contracts

Suppose a health authority selects contracts according to the principles just discussed to protect quality of service. Two natural questions are: will it achieve that quality at the lowest price by having hospitals bid for contracts, so enhancing budgetary efficiency? and will greater competition (in the form of more bidders) enable it to fulfil its priorities more closely, so enhancing economic efficiency?

These are difficult questions to which the literature does not provide definitive answers. When the health authority does not know the price at which the potential bidders are able to provide any given level of service, it will enhance budgetary efficiency to alter the contracts suggested in the previous section to increase the provision for payment that depends on the costs the hospital actually incurs in providing the services (or add such a provision if it is not there), provided it is not too costly to make these costs verifiable. To see why, consider the case of a self-interested hospital with a contract that has a lump sum plus a fixed payment per patient treated. Such a contract provides a strong incentive for the hospital to keep costs to a minimum because it will retain any savings. But, if it is actually a very low cost hospital, it will make a lot of profit that the health authority would prefer (in line with budgetary efficiency) to use for treating other patients. Profits could be kept low by a contract that reimburses the hospital for its actual costs of treatment. This is what happens with private medical insurance and what used to happen with the cost reimbursement rules previously used for Medicare in the US. Such a contract can prevent the hospital making excessive profits but it also provides no incentive to keep costs low because any cost savings made by the hospital simply result in it being paid that much less for providing the service. The best the health authority can typically do under these circumstances is to offer a contract that specifies a maximum payment and a target cost, with the payment to the hospital reduced below the maximum by a proportion of the amount by which the actual cost falls below the target cost. With such a contract, the bidder still has an incentive to reduce costs because it receives a share of the cost savings but also, if costs turn out to be low, it does not make as large profits as with a contract that has payment independent of costs. Laffont and Tirole (1993, chapter 1) have a detailed (but technical) discussion of the precise form the contract should take.

Use of such a contract makes bidding more complicated. The reason is that the contract has several dimensions (the maximum payment, the target cost and the share of cost reductions retained by the bidder), so selecting the winning bid is not a simple matter of choosing the

lowest. Nevertheless, it is possible to design a bidding process that will select the lowest cost bidder. Moreover, the overall cost will on average fall as the number of bidders is increased, so increased competition enhances budgetary efficiency by reducing costs. The reason is essentially the same as in the case of simple bidding discussed above - the more bidders there are, the lower on average will be the reservation price of the second lowest bidder and, hence, the price of the winning bid. For further details, see Laffont and Tirole (1993, chapter 7).

That is not, however, the whole story because we are concerned here with what such bidding does to the quality of service provided. The conclusions just drawn apply to self interested hospitals when demand does not respond to quality and the contract provides incentives to keep costs down. In the case in which patients correctly assess quality and their demand for treatment responds rapidly to that, we conjecture that, because quality can then be ensured by patient demand, these conclusions continue to apply. These same conclusions almost certainly also apply to hospitals that are at least partially benevolent. But there is likely to be a real problem if demand does not respond to some aspects of quality and bidders consist of some self-interested hospitals and some benevolent hospitals, with the health authority not knowing which is which. For any given contract, benevolent hospitals will provide a quality of service closer to that which the health authority would like. However, a concern that has been expressed about many public services put out to tender in the UK is that, by awarding the contract to the lowest bidder, it will go to a self-interested bidder who has no concern for the quality of service or, at the very least, successful bidders who, although they would like to provide a higher quality service, have to settle for low quality in order to keep their prices down to the level of bidders who are unconcerned with quality. Whether this will in fact be the case and, if so, what a health authority can do about it is an issue for future research.

5. Quality, reputations and long term contracts

Contracts between health authorities and hospitals typically last for a year, after which time a new contract is negotiated. Each contract covers the services to be provided during that year. Short term, renewable contracts like these can give a powerful incentive for a hospital to provide a good quality service even if patient demand does not respond to quality as long as the health authority receives some feedback on the quality of service, however informal.

For quality to be enforced by contract, it needs to be verifiable. But, at the time a new contract is awarded, a health authority can make use of any information it has, even if not verifiable, to affect either the terms of the contract it awards or which provider it is awarded to. (Unverifiable information of this type is termed *observable* in the literature.) By offering less good terms, or threatening to switch providers, if the information indicates that quality is unsatisfactory, the health authority can make it costly in the future for a hospital to skimp on quality now and thus make the hospital behave as if it were at least partially benevolent. Indeed, provided sufficient information about quality is fed back to the authority, the hospital cares sufficiently about the future (its discount rate is not too high), and future custom is sufficiently valuable, the concern for reputation can ensure that the hospital provides precisely the quality of service the health authority would like. See Gibbons (1992) for an accessible discussion of this issue. The threat to award the contract to an alternative provider becomes more powerful if there is competition among providers.

The use of short term, renewable contracts can, however, give rise to problems. Two that have been widely discussed in the literature are the *ratchet effect* and the *hold-up effect*. The ratchet effect was first widely discussed in the literature on central planning in the Soviet Union. The essential issue is the following. Soviet factory managers were given an output quota to fulfil. To encourage production, the factory received a bonus if it overfulfilled its quota. But, by overfulfilling the quota, the managers would reveal to the central planners that they had the capacity to produce more than their current quota. So the planners would not unnaturally seek to increase the quota the following year, with the result that the factory would have to work harder to fulfil the higher quota. So canny managers were faced with a calculation. Overfulfilling quota would result in a bonus this year but harder work (or a lower bonus) next. The effect of the planners response to an overfulfilled quota was to reduce the incentive for higher output that the bonus was intended to provide.

The ratchet effect also arises with contract negotiation. A hospital that works hard to cut costs this year reveals that it can provide services at lower prices. A health authority that would like additional funds to spend elsewhere will then naturally bargain for lower prices next year, which reduces the incentive for the hospital to cut costs.

The hold-up effect can be illustrated as follows. Suppose a hospital can install equipment that would reduce the cost of providing services but that, once installed, could not be resold for anything like its original cost. The difference between the cost of installation and the resale value is a *sunk cost*. Under a fixed price contract, the hospital will get the cost savings that result from the installation while the contract lasts. But when it comes to negotiating the next contract, the health authority is in a position to bargain for services at lower prices because of the hospital's investment. In the extreme, if the health authority captures all the cost savings, the hospital will install equipment only if the cost savings under the *existing* contract will pay for the installation. But even if the health authority captures only a part of the cost savings, the hospital will invest too little in new equipment with consequences for both the cost and quality of service it is able to offer. The classic discussion of hold-up is in Williamson (1985). Hold-up is not just a theoretical possibility - one negotiator for an NHS trust asked us unprompted how the trust could prevent the health authority exploiting its investments in this way.

Both the ratchet effect and the hold-up effect can be alleviated by having a longer term contract.[4] In both cases, the longer it is before the contract comes up for renewal, the more time there is for the hospital to get the benefits of its investment before the health authority captures some of those benefits in a new contract. But long term contracts have the disadvantage of reducing the power of reputation effects because a hospital that delivers poor quality can do so for longer before losing its contract. There is thus a trade off between having longer term contracts to mitigate the ratchet and hold-up effects and having shorter term contracts to maintain quality. We know of no analysis that has investigated this trade off fully.

---

[4] We say *alleviated* and not *solved* advisedly. Even if the contract were to last indefinitely, the effects can occur because of circumstances in which it is in the interests of *both* parties to renegotiate the contract. A contract can always be renegotiated if both parties agree to this. The ratchet and hold-up effects can then occur because it is typically the case that renegotiation of an existing contract has some of the properties of negotiation of a new contract. Milgrom and Roberts (1992) provide an accessible discussion of this issue. We do not pursue it here.

6. Conclusion

Our discussion can be summarized as follows. We envisage four potential roles for competition between suppliers: reducing the economic inefficiencies that can arise from monopoly power; reducing the influence of providers over purchasers' decisions; ensuring that services respond to what patients want; and providing information to purchasers both about which supplier can provide services for the lowest prices and how low those prices can be. The first of these does not seem very important at present in view of the countervailing bargaining power that health authorities have. It might become a serious issue if, for example, one hospital group were to buy up a chain of hospitals that formed a significant part of the UK market or if purchasers were to become sufficiently fragmented, perhaps as a result of the growth of GP fundholding. There are, however, other mechanisms for dealing with these. Indeed, a mechanism to counteract the fragmented purchasing power of GP fundholders is already developing in the form of purchasing consortia that negotiate with providers on behalf of groups of GP practices. The second potential role is one on which we know of no formal analysis. However, there are formal analyses of competition between suppliers for influence over decision makers that suggest such competition is very wasteful and such waste must be set against any gains increased competition may have. The third may be partially addressed by increased use of GP fundholders if GPs are more closely in touch with the needs and wants of their patients than health authorities. Whether or not they are, however, is not an issue that theory can address and there is not as yet sufficient experience to provide an assessment of how this aspect of GP fundholding works in practice. The fourth role for competition may well become increasingly important as a consequence of the separation of purchasers from providers. As time goes by, those employed by health authorities are becoming more distanced from their former responsibility for the direct provision of services and may thus become less knowledgeable about the prices at which services could be provided. Increased competition between hospitals may then enhance economic efficiency and also, by enabling health authorities to get services provided at a lower cost, budgetary efficiency.

Getting hospitals to bid for contracts to provide health services is not, however, like conducting an auction for wheat because the quality (in the widest sense) of the services to be provided is not something that is easily specified by contract. Taking the lowest bid may simply result in the provision of poor quality services. It is here that the form of contract becomes important. Contracts must be designed carefully to ensure that the quantity and quality of services are kept as close as possible to what the health authority would like. An obvious

problem is that, when quality cannot be verified, high quality will be delivered only if the hospital chooses, or is induced by the contract, to provide it. The form of contract appropriate for this depends among other things on the extent to which patients can correctly assess the quality on offer and choose where to be treated on the basis of their assessments, and also on how much the provider really cares about the interests of patients, as opposed to just its own self interest. There is a danger in awarding a contract to the lowest bidder. Hospitals that care enough about their patients would provide the quality of service the health authority wants if given the funds to do so. But they will have to offer a low quality service to compete with bidders who are less concerned about their patients and who therefore offer lower quality.

The key problem that we return to again here is that there are many dimensions of quality that cannot be enforced by contract. One way to alleviate this problem is to use short term contracts that a health authority can refuse to renew if it has doubts about the quality of services that have been provided. Refusing to renew a contract does not need concrete evidence of the type that would be needed to demonstrate in court that a hospital had not fulfilled the quality terms in a contract. So a decision on renewal can make use of much more subjective assessments based on whatever information the health authority can glean. Where loss of reputation for quality results in the loss of valued future contracts, it can act as a powerful incentive to maintain standards. But the use of short term contracts has its own problems. In particular, it can give rise to the ratchet and hold-up effects which result in services being provided at a higher cost than with a long term contract. The trade off between using short term contracts to maintain quality and long term contracts to alleviate the ratchet and hold-up effects is not easily escaped.

Our overall assessment of the current situation is as follows. There are certainly potential roles for competition in the provision of health services, though we are not convinced that the traditional anti-monopoly role is the most important. But there are also potential dangers from competition that arise from the variability in the quality of health services and the fact that quality is not easily verified (and not, therefore, easily enforced by contract). Whether or not the recent reforms in the NHS lead to better health services depends at least in part on purchasers and providers understanding what these roles and dangers are and choosing contract forms that ensure standards of quality are high. In this paper, we have set out at least some of the issues involved.

References

Arrow, Kenneth J. (1963), 'Uncertainty and the welfare economics of medical care', *American Economic Review*, 53 (5), 941-973.

Barker, Kit (1993), 'NHS contracts, restitution and the internal market', *Modern Law Review*, 56 (6), 832-843.

Barker, Kit (1995), 'NHS contracting: shadows in the law of tort?', *Medical Law Review*, 3 (3), forthcoming.

Binmore, Ken (1992), *Fun and Games: A Text on Game Theory*, Lexington, MA: D.C. Heath and Company.

Chalkley, Martin and Malcomson, James M. (1995a), 'Contracting for health services with unmonitored quality', University of Southampton, Department of Economics.

Chalkley, Martin and Malcomson, James M. (1995b), 'Contracting for health services when demand does not reflect quality', University of Southampton, Department of Economics.

Dawson, Diane (1994), 'Costs and prices in the internal market: markets vs NHS Management Executive guidelines', University of York, Centre for Health Economics Discussion Paper 115.

Ellis, Randall P. and McGuire, Thomas G. (1993), 'Supply-side and demand-side cost sharing in health care', *Journal of Economic Perspectives*, 7 (4), 135-151.

Ferguson, Brian and Palmer, Stephen (1994), 'Markets and the NHSME guidelines: costs and prices in the NHS internal market', University of York, Centre for Health Economics Discussion Paper 120.

Gibbons, Robert (1992), *A Primer in Game Theory*, London: Harvester Wheatsheaf.

Hodgkin, Dominic and McGuire, Thomas G. (1994), 'Payment levels and hospital response to prospective payment', *Journal of Health Economics*, 13, 1-29.

Laffont, Jean-Jacques, and Tirole, Jean (1993), *A Theory of Incentives in Procurement and Regulation*, Cambridge, MA: MIT Press.

Lancaster, Kelvin J. (1979), *Variety, Equity and Efficiency*, New York: Columbia University Press.

Ma, Ching-to Albert (1994), 'Health care payment systems: cost and quality incentives',*Journal of Economics and Management Strategy*, 3 (1), 93-112.

McAfee, R. Preston and McMillan, John (1987), 'Auctions and bidding', *Journal of Economic Literature*, 25 (2), 699-738.

Milgrom, Paul and Roberts, John (1992), *Economics, Organization and Management*, Englewood Cliffs, NJ: Prentice-Hall.

Myerson, Roger and Satterthwaite, Mark (1983), 'Efficient mechanisms for bilateral trading.' *Journal of Economic Theory*, 28, 265-281.

Nitzan, Shmuel (1994), 'Modelling rent-seeking contests', *European Journal of Political Economy*, 10 (1), 41-60.

Tirole, Jean (1989), *The Theory of Industrial Organization*, Cambridge, MA: MIT Press.

Williamson, Oliver E. (1985), *The Economic Institutions of Capitalism*, New York: The Free Press.

Faculty of Social Sciences
University of Southampton
129 University Road
Highfield
Southampton SO17 1BJ

Telephone 01703 593896
Fax 01703 593177