Using Flexibility to Improve Value-for-Money in PFI Projects[1]

Richard de Neufville, Yun Shin Lee, Stefan Scholtes

June 2007

Abstract. This paper addresses PFI stakeholders – government and local clients on the public side, private construction companies, service providers and banks on the private side. We will argue that:

(a) PFI is cost focused rather than value focused. It doesn't deliver best value projects in the light of significant value uncertainty, driven by unpredictable changes in demand, technical or socio-political circumstances over the lifetime of the projects. What is needed is a value focus in PFI.

(b) Flexibility is key to dealing effectively with value uncertainty. We need flexible, "living" infrastructures rather than rigid monuments. In order to make a case for flexibility infrastructures, we need ways to articulate and communicate the value of a dynamic asset, in a simple and transparent manner. We will suggest some principles that should govern such dynamic valuations.

(c) To reap the value of flexible infrastructure systems, the PFI parties need to create a genuine long-term partnership. The private parties have to stay engaged during the project lifetime to help adjust the infrastructure to developing circumstances and thereby minimise value risks, maximise opportunities if and when they arise, and allow the public sector to reap as much social benefit from the infrastructure as possible.

Background: PFI in controlling cost and time

Since its introduction in 1997, the UK Private Finance Initiative (PFI) has delivered over 550 projects, primarily in the health, education and transportation sectors. Investments in 80 major PFI hospital projects alone exceed a total of £60 billion[2]. The background for the initiative is the appalling performance of traditional public procurement processes on core metrics, such as delivering projects *on time* and *on budget*.

The solution to politicians has been simple: Bring in the private sector. Private companies are assumed to be more disciplined with capital expenditure and to have a better understanding of and better access to the expertise required to carry out large infrastructure projects. Combine this with competition, via a bidding process, and performance should improve.

The PFI principle is equally simple: A private consortium finances and builds the infrastructure and then rents it to the government for a fixed period, typically 25-35

[1] The authors are grateful to Laing O'Rourke personnel for valuable discussions that helped shape their thinking on the issue addressed in the paper. Any claims in this paper are entirely the authors' responsibility.

[2] Department of Health (2007), Prioritised Capital Schemes approved to go ahead since May 1997, London.

years, after which ownership of the asset is transferred to the public sector. During the PFI-period, the government makes agreed annual rental payments to the private consortium.

PFI is a significant success with regard to delivery on time and on budget. Indeed, a recent study by McKee et al. (2006)[3] on a sample of conventionally procured and PFI projects gives clear evidence:

	On-Time	On-Budget
PFI projects	76%	79%
Conventionally procured projects	30%	27%

So, on these important project management metrics, PFI seems to work well, assuming that neither budgets nor time are inflated. Competition in the bidding process should keep total budgets down and the fact that positive cash flows from government rents only start when the infrastructure is operational gives the private sector an incentive to keep construction times low because they will otherwise loose part of their income stream.

Do current PFIs increase value?

The value proposition for any project has three main components: Construction Costs, Operational & Maintenance Costs, and Operational Benefits / Revenues. The McKee et al. study relates to construction costs and certifies the success of PFI in this respect.

How well do PFI projects do *after* construction? Might it be that the private consortia save on construction costs and time wherever possible and therefore produce infrastructure with significantly enlarged running costs, when compared to conventionally procured projects? This "moral hazard", as economists call it, was recognised early on in the design of the PFI process. The solution again is simple: PFI projects require the private consortium not only to build the infrastructure but also to maintain it during the lifetime of the contract. It is therefore in the private company's best interest to find a good balance point between spending money upfront for a robust infrastructure and spending money later on maintenance.

PFI has thus transferred two key risks, the risk of excessive construction costs and time and the risk of excessive running costs, from the public to the private sector.

This leaves us to look at the third value driver: the benefits or revenues generated from PFI procured projects. Here, the record of PFI is mixed. PFI projects can still be white elephants. PFI does not seem to address, systematically, the issue of uncertainty in revenues or social benefits in a systematic way. Outcomes can differ substantially from what was originally expected.

[3] M McKee, N. Edwards, R. Atun (2006), *Public private- partnerships for hospitals. Bulletin of the World Health Organisation*, 84, 890-895

Balmoral High School in Belfast is a notorious example. The school was build for 500 pupils. Due to demographic changes it could only attract some 150 pupils and therefore was uneconomical and had to be closed – leaving the local education authority to honour a 20-year contract of annual payments to the private sector, a cumulative cost of £7.4M, without any future benefit.

Another example is Coventry University Hospital, a £400M PFI project. The local health trust could not even afford the first annual payment after the construction period, and had to be bailed out through debt, underwritten by the primary care trust. To survive, the hospital trust will have to make significant savings, which will lead to shrinking services and quite possibly a distortion of clinical performance. This is hardly a case of good value for money.

The first PFI development of a museum, the Royal Armouries in Leeds, is a further example. Realised visitor volume in 1999 was only 400,000 against a projection of 750,000 that formed the basis of the PFI case. This led to the financial collapse of the museum.

Admittedly, these examples provide only anecdotal evidence – but they do point to an inherent weakness in the PFI process: Its preoccupation with cost control, rather than value delivery. If we define value as a benefit-cost ratio, then PFI works on minimizing cost – but does not specifically address the maximisation of benefits.

PFI can be in the way of good value for money

The existing PFI framework can be a significant obstacle to managing situations where the infrastructure does not create as much value as expected. The natural reaction to lack of demand, as in the case of Balmoral High School or Coventry University Hospital, is to change the infrastructure, or even sell it. Balmoral High School might be a good prospect for conversion into a cinema, gym, or other private sector facility. Within the existing PFI framework this is, however, easier said than done.

Firstly, the private consortium is naturally expected to act in the interest of its shareholders and therefore has no incentive to lower the contractually agreed rent if projected revenues do not materialise. In fact, it will have sold a significant proportion of the rental cash flows to the financial markets, e.g. to pension funds, to help finance the project in the first place. So the private consortium's base position in any negotiation is that the public partner should honour the contractually agreed payments. The negotiation space is reduced to potential savings in maintenance costs, typically a small proportion relative the total capital expenditure.

Secondly, there is a gridlock situation with regard to the right to alter the asset. On the one hand, the public sector has a right to usage as agreed in the contract, so the private consortium cannot unilaterally change the asset. On the other hand, the private sector owns the asset, so the public client cannot change it unilaterally to improve its value. Changes can only take place if all parties agree. To make matters worse, the membership of the private consortium may have changed, and the design and construction knowledge and expertise necessary for a creative and successful resolution of the fate of the failing asset may be lacking.

Even if flexibility is recognised in the design of the hospital, the PFI contract can be a stumbling block to its efficient use. An example in point is the Royal Victoria Infirmary PFI project in Newcastle-upon-Tyne, which includes an expansion provision for additional ward space. During the construction phase it became apparent that the Trust would rather use this space for offices, which was simpler and less costly than using the expansion space for wards. However, the potential use of the expansion space for offices was not explicitly stipulated in the contract and therefore the private consortium could not agree to the request of the Trust. The consortium argued that the change would require a potentially very costly re-rating of the bonds used to finance the PFI.

To summarise the challenge: PFI projects may well deliver a pre-specified operational service level at low cost – but still fail because during the contract lifetime it turns out that the installed infrastructure and pre-specified operational service level is actually not what is needed. To make the asset valuable, it needs to be changed. PFI is currently not set up to tackle this challenge and can indeed seriously strangle the *flexibility* to make changes when required.

In times of significant uncertainty, be it commercial, technological, or socio-political in nature, rigid infrastructures are bound to under-perform. The challenge is to deliver high-value *flexible* infrastructures, "living systems", that can cope with this uncertainty by being suitably adapted as the future unfolds. Meeting this challenge within a PFI framework requires the involvement of the private sector in the *value* proposition of the project. This is not straight-forward. Furthermore, it calls for private sector partners who are bold enough to climb up the value chain and fully engage in the provision of long-term value to the public. This requires genuine public-private partnerships, not "fee for service" contracts as in the existing PFI framework. It amounts to nothing less than re-thinking public sector construction. It needs visionary construction companies and governments to meet this challenge.

Why should the private sector engage with this challenge at all? What is in it for the shareholders? Beside potentially large reputational gains, there are immediate financial opportunities: Where there is downside there is also upside. Being able to adapt infrastructure does not only allow the partners to cut losses in downside scenarios but also to exploit upside scenarios, to amplify gains if demand is higher than expected. Flexibility can, for example, come from a staged investment, with lower initial outlay of capital for the infrastructure as such but additional investment in a sound contingency plan for expansion if and when more is known about the size and scope of demand for the infrastructure.

Coping with value uncertainty: Flexibility is key

A key aspect of the management of sustained long-term value, as opposed to a short-term burst of investment costs, is the ability to deal with the significant level of uncertainty surrounding long-term demand for an installed infrastructure. Demand, e.g. for hospital-based health care, can change substantially in size and in scope, over a period of 30-40 years. It is not surprising that the private sector would like the tax-payer to carry this hot potato. However, it is the private sector that holds the key to managing the crucial value uncertainty: Private firms have the expertise and employ the best engineering designers who, together with the best minds of the public sector partners, can produce creative *flexible* infrastructures that can be adapted cost-efficiently to many different futures – and deliver value in many situations.

To make the vision of flexible infrastructures work, it is crucially important

(a) To develop a common language that allows the partners to communicate uncertainty and flexibility effectively and to articulate the value of alternative flexible designs systematically and
(b) to move from a contractor-client relationship to a genuine partnership, where the partners engage over the life-time of the project in a genuine joint operation, sharing costs, rewards and risks in an efficient way.

In the remainder of this paper we will expand on these themes.

Thinking flexible

Flexibility is often described as *the right but not obligation to a specific future action.* One way of thinking about any particular type of flexibility is to regard it as a system switch which is either in "on" or in "off" (default) setting. Switching to "on" will change the way the system operates. Building the switch in costs money – and switching from "off" (default setting) to "on" will often also cost money. In some situations one can switch on or off as often as one likes, sometimes one can only use the switch once. Of course a good flexible system will typically have several switches to allow a reaction to different circumstances.

There are two key issues with a system with switches:

(a) how can we lead designers to creatively think of switches to include in the design? and
(b) how can we compare the value of different designs with switches?

The prevalent discounted cash flow valuations are certainly inadequate as they do not account for switching.

Flexible design is nothing else than contingency planning. You have a spare tire in the boot of your car because that gives you the flexibility to replace a tire if you have a puncture on the way to the airport – so you won't miss your flight. If you are planning a windsurfing weekend on the beach, you might take along a good book because that gives you the flexibility to enjoy a good read on the beach if there is not enough wind for surfing. A very small additional investment - taking the book - can have great payoff – if the weather isn't as expected.

Flexibility helps reduce pain in some scenarios or increase gain in others. You will know that you had not built in enough flexibility if you moan "Had I only done X earlier, I could now do Y and would be in a better position". This should not be confused with hindsight argumentation, i.e. "Had I known that demand was so low, I would have built a smaller hospital". Flexible design is rather to be seen as the designer's reaction to Blaise Pascale's observation that "Chance favours the prepared".

How can we become prepared? The starting point is that we explore a range of possible futures. Is this done sufficiently in the design process for infrastructure that is meant to serve for 30-50 years? Or are we all too happy to design and build infrastructure for a single projection of the future. Everyone knows that forecasts, in particular long-term ones, are highly unreliable, in fact are "always wrong" in that what actually happens "never" conforms to predictions. Balmoral High School is not a singular case at all in this respect.

Therefore, the first and arguably most important step towards more flexible infrastructure is to stop asking for accurate forecasts of the future! That is impossible. We ought to ask for a range of possible futures. Given such a range, the designer's task is to build an infrastructure that fairs well in many of these futures. The incorporation of flexibility into the system design is then a very natural step.

A case in point: a chief technical designer of one of the satellite telephone systems used to complain that one of his major problems was to get the company forecasters to give him a precise prediction of the future capacity required for the system. In the event, the company settled on one figure, and he and his team created a design based on that figure – which turned out to be off by a factor of more than 10, and was a prime cause of the bankruptcy of the company! They would have been far wiser to recognize in advance that it was impossible to predict the capacity required with any accuracy, that they needed to anticipate many different scenarios, and to create a design that was adaptable to this range of possibilities.

Examples of flexibility in hospital PFIs

There are many examples of flexibilities that can help hospitals to cope with uncertainty in scale and scope of demand.

The Royal Victoria Infirmary in Newcastle-upon-Tyne, a recent PFI project, includes the option to increase its height to make room for additional ward space, to increase ward space through a side expansion, or by converting car parking space into new buildings and building a multi-story car-park to cope with additional traffic. Also, parts of the hospital shell were included without specific functionality; the use of this space can be determined by the NHS trust at some future date, when more is known about needs. Some of this space includes footprints for specific facilities, such as operating theatres, but can also be used for other purposes. Designing storage space next to functional rooms enables the expansion of the latter, should the need arise.

The above flexibilities can be thought of as strategic (e.g. hospital expansion through the conversion of a car park) and tactical (e.g. expansion of an operating theatre). In addition, there are operational flexibilities, which enable fast small-scale changes, for example of the engineering services. To this end, hospitals are often designed to have mechanical and electrical distribution systems with spare capacity, typically an extra 25%.

Flexibilities, if added to rigid systems, might increase costs. Alternatively, they might reduce the cost, for example, if they allow us to start with a relatively smaller hospital initially but have the option to expand it. Given the pressure on capital expenditure minimization and value-for-money, it is crucial to have a systematic way of articulating the value of flexible designs.

Articulating the value of flexible design

The prevalent valuation methods, such as discounted cash flows, account for important value-drivers, e.g. economies of scale and time-value of money. There is, however, hitherto no standard systematic way of accounting for the value of flexibility when we compare projects and system designs.

What are the principles of flexibility valuation? The overriding principle is that flexibility can, by its very nature, only show its value if the valuation process considers more than one future. Indeed, in some futures we will not use the switch, in others we will. If there is only one future, only one projection, then the valuation cannot reflect the use of the switch – it will either stay off or be switched on for sure and in the latter case we may just as well "hard-wire" the switching and forget the switch itself.

We will use a fictional, but not far-fetched example of a hospital PFI to illustrate one simple process that allows the articulation and discussion of the value of flexible designs.

Suppose you are working on the design of a new hospital. A key input to the design is the growth of annual inpatient days over the next 25 years. Over the past years, inpatient days in the incumbent hospital have increased significantly. In fact that is the chief reasons for the new hospital project. Forecasts say that inpatient days will continue to rise. Therefore, the default plan is a large hospital of 1000 beds. Some analysts, however, warn that the trend could well reverse. Regulations are discussed that could give the private sector more access to patient care. Also, over the lifetime of the hospital, medical technology may drastically reduce length of stay. The majority of analysts, however, argue against this and quote the aging population, the deteriorating health of children, but also the increasingly successful recruitment of wealthy private patients from developing countries as drivers of demand.

Before we suggest our approach we paraphrase the traditional approach to design optimisation: Take the most likely forecast as starting point and optimise the hospital design for this demand growth projection. Let's assume this results in a 1000 bed hospital. Then stress-test the design. It turns out that this hospital will be economically viable even if the demand does not grow quite as rapidly as forecast. The hospital is built – but then demand under-performs dramatically and the hospital has to take drastic measures to survive, impacting clinical performance. The fate of Coventry University Hospital shows that this is not too fictional an example.

Have we just been incredibly unlucky in this situation? Stress-testing on demands is a standard procedure - but one of the problems is that stress tests are typically not bold enough in their assumptions on the range of possible futures, in particular in the light of the long life-time of many PFI projects. The figure below, which gives a sample of errors in traffic demand estimates, illustrates this point. It shows that (a) there has been a significant bias in estimating traffic demand as too low and (b) there has been a huge variation in the accuracy of traffic demand estimates. Stress-testing with a +/-10% variation on growth assumptions will not uncover the real risks[4].

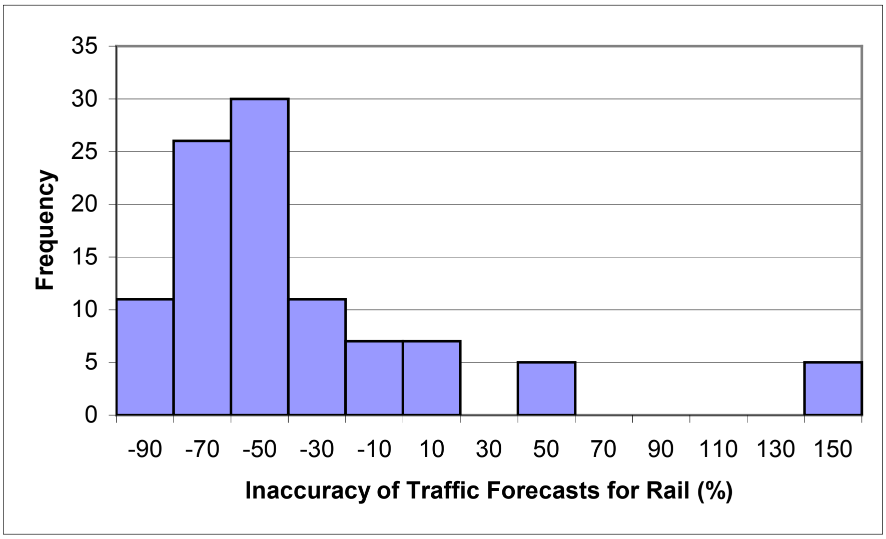

In our hospital example it is quite possible that if we took a wider, more realistic range of demand scenarios into account then there would not be a hospital that can perform well in most of these scenarios. A very uncomfortable situation – but precisely the situation where flexibility can help tremendously.

One possible flexible design might be to build a 700 bed hospital now, but design it in such a way that it can be expanded at a reasonable cost, say through the conversion of ground-level car parking space to additional hospital buildings and the building of multi-story car parks to cope with the additional traffic. This might give the hospital trust the flexibility to increase total capacity to up to 1200 beds if required in the future. Clearly on the basis of costs per bed the 700 bed hospital is more expensive. Even if it will be expanded to 1000 beds later, the fact that the capacity is not built immediately means that economies of scale savings fall below the rigid 1000 bed hospital. On the positive side, initial total expenditure is less for the 700 bed hospital. This is an important economic argument for this option. Nevertheless, if the demand is as expected, we will expand to 1000 beds later and the total phasing cost, including the cost of disruption during the expansion phase, may well exceed the savings due to phased capital outlay, in which case economic arguments seem to suggest the immediate construction of a 1000 bed hospital.

Intuitively, we know that we are overlooking something important in our valuation: If demand is low, we avoid a white elephant! Also, if demand is higher than expected,

[4] Figure adapted from Flyvjberg, B., Holm, M. and Buhl, S. (2005) "How (in)accurate are demand forecasts in public works projects? The case of transportation,". *J. American Planning Association,* 71(2), pp. 131-146 and Flyvjberg, B., Buzelius, N., and Rothengatter, W. (2003) Megaprojects and Risk: An Anatomy of Ambition, Cambridge University Press, Cambridge, UK.

we can actually grow the hospital to 1200 beds. How can we make an economic case for the value of this flexibility?

As mentioned above – the starting point for the articulation of the economic value of a flexible design must be the recognition of many possible futures. To keep things simple let us assume that we work with a range of 10 possible futures with different assumptions, call them F1, F2,.., F10.

Next we develop a contingency plan for each design. Of course if the design is rigid as in the case of the 1000 bed hospital, then there is no contingency plan. But for the 700 bed hospital a contingency plan might be of the following form: "We will decide whether to expand or not in 5 years time. If total growth in demand over these 5 years exceeds 10% then we expand to a total of 1000 beds, if growth exceeds 15% then we will increase to 1200 beds. Otherwise we do not expand." With this contingency plan we can calculate the cash flow (or other cost-benefit metric) for each future F1,…,F10[5]. The results can be summarised in a bar-chart as follows:

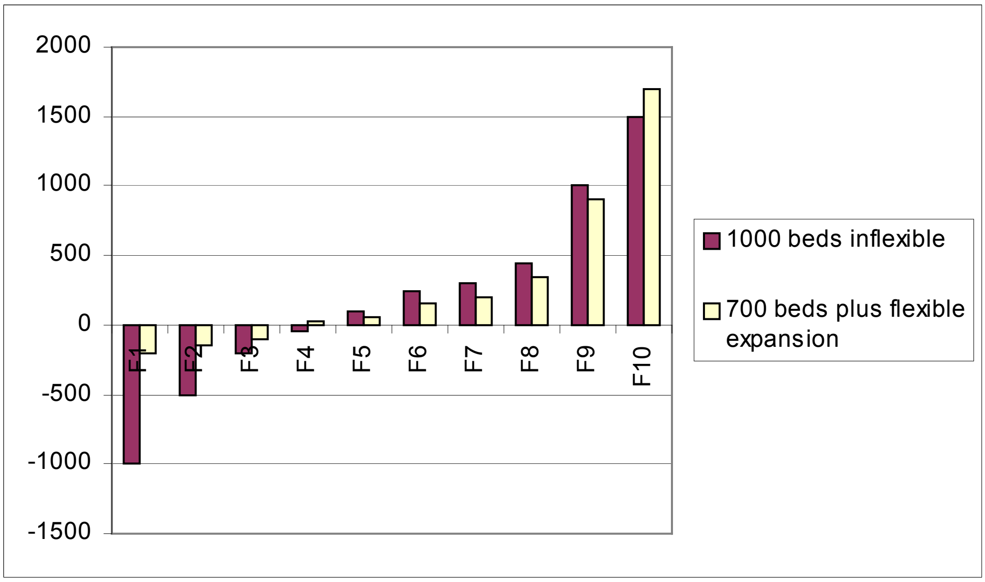

As you can see, in the low demand scenarios F1,.., F4 the flexible design performs much better because the 1000 bed hospital would be a white elephant. In the medium scenarios F5,…, F9 the flexible hospital is expanded from 700 to 1000 beds. The overall performance is worse – the saved capital costs due to delaying part of the investment were not sufficient to balance out economies of scale and the cost of operational disruption during the expansion. In the very high scenario the flexible design performs better again because the hospital can be expanded to 1200 beds.

It is not immediately clear from the bar chart which hospital design you should prefer. Of course if you were very worried about the low demand scenarios, the flexible

[5] In a professional analysis these scenario-by-scenario values would be calculated through a Monte Carlo Simulation.

system looks better. But you may have good reasons to believe that the chance of these low demand scenarios is low – and be prepared to take the gamble.

This brings us to the second useful ingredient of a valuation – the relative weight you assign to any of these futures. One way of thinking about a relative weight is terms of likelihood of occurrence of these futures. If you do not wish to rate one of the futures more likely then any other, then you are not adding new information to the bar chart above. If you do, however, give them different likelihood, then one way of incorporating this into the bar graph is to multiply every bar value with its probability of occurrence. This will then result in a bar graph as follows:

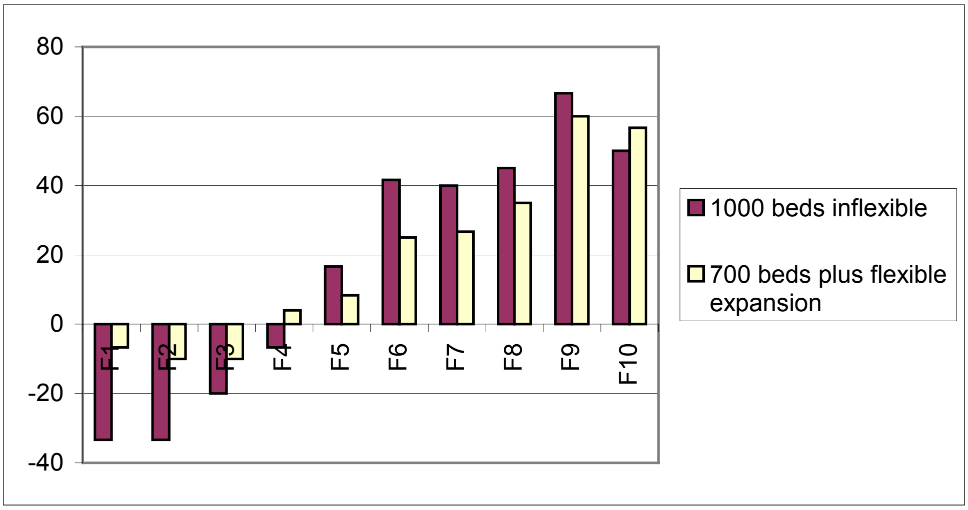

This graph has a different scale because you have multiplied the bars by your estimate of the probability of occurrence of the respective future. But what is more interesting is that the shape of the bar chart has changed. The inflexible design now looks significantly better in the medium demand scenarios because these were regarded as more likely.

One interesting metric is the sum of the bars, which gives you the expected value of the respective system, i.e., the outcome over all considered futures weighted with your judgement of their likelihood of occurrence. Here you get an expected value of £166 for the inflexible design and of £189 for the flexible design. In this case, flexibility has added value on average. The expected values, however, do not tell you the whole story. It is important to look at the variation of values across the futures to understand the full effect of flexibility.

Incorporating flexibility thinking in the PFI design process

Flexibility gets its value from uncertainty. The more uncertain the social, technological and commercial environment is, the more we should think about flexibilities that allow us to react when uncertainties unfold. There are, however, so many uncertainties and potential flexibilities that we may well feel snowed under very quickly. Common sense and expertise driven prioritisation is necessary. What are the most important uncertainties for the PFI project? This should be a key question at the outset of the planning process.

Who should be driving the exploration of possible futures? It should certainly be the future operators of the system, e.g. the NHS trust in the case of a hospital. They will know best what the future might hold. NHS managers will bring in commercial expertise, medical professionals the requisite technical expertise, outside health specialists and demographers the demographic expertise, etc. This team's main task is to provide a manageable range of possible futures. They need to be creative, not strangled by the desire to produce "one" forecast. In fact, it should be their remit to make sure that no important potential future is overlooked. This is a completely different set-up to a projection arrangement, where the end-product is a single future, plus a few stress-test variations. The outcome of an uncertainty brainstorming process is a reasonably comprehensive *range* of possible futures.

It is crucially important to manage this brainstorming well – it can easily become a wonderland. We have no tested solution for such a process, yet, but it seems sensible to start from the macro-drivers of health futures, such as demography, nationally as well as locally, population health status development, possibly by population segment, potential technological advances, regulatory changes, etc. These macro-drivers will have impacts on operational metrics of the hospital, such as number of inpatients and outpatients, length of stay, or case-mix. A macro-micro matrix, highlighting the effects of the macro drivers on the operational (micro) metrics of the hospital might be a useful device to help steer the discussion. Such a matrix could look like this:

Uncertainty Matrix	**Operational Metrics**		
	No. of Inpatients	Length of Stay	No. of Outpatients
Immigrations of middle-aged population	+	-	+
Increase in population aged 65 and over	++	+	++
Development of new treatment	+	--	+

Here a plus means that an increase of the macro factor will lead to an increase of the micro factor, whilst a minus indicates a decrease of the micro factor as the macro factor increases. This matrix helps understand the relationships between the micro factor, which will eventually drive the operation and hence the value of the hospital. Starting from this, one might be able to create meaningful quantitative relationships between macro-level uncertainties and their impacts on the operation of a hospital.

Once the key uncertainties are understood and a range of futures is determined, the engineering design team will come in. They will work with the trust towards potential designs that will produce value for money in many of these futures. The challenge is to trade off optimal designs for any single future against miserable performance in other futures. Creative use of flexibility will be a natural in fact quite possibly the most important ingredient to the design discussion. The engineering designers will have to work actively with the professionals to understand what type of flexibility might help and how it can be incorporated into the overall design at reasonable costs. Again, it is easy to drown in the myriads of possibilities and prioritisation is key. A matrix structure might help to understand the relationships between the various operational uncertainties and the flexibilities and the flexibilities that might help to mitigate or exploit them.

Flexibility Matrix	**Operational Metrics**		
	No. of Inpatients	Length of Stay	No. of Outpatients
Expansion of hospital in height	+	+	+
Subletting hospital for secondary usages	-	-	-
'Shell' space	+/-	+/-	+/-

A plus now indicates that the respective flexibility helps if the corresponding operational metric increases, a minus indicates that the flexibility helps if the metric decreases. Some flexibilities might help in both cases (+/-). Again, this is just a pre-cursor to a more quantitative analysis. Notice that the two matrices, the macro-micro matrix and the uncertainty-flexibility matrix can be combined to understand the relationship between macro uncertainties and the flexibilities in the design[6].

We are not in a position, yet, to give advice on best practice in flexible design for PFI projects. Finding the right mix of simplicity and transparency on the one hand and rigour and consistency on the other is a crucial challenge for flexible infrastructure design.

Exploiting flexibility

We have illustrated earlier that the existing PFI framework can be in the way of exploiting flexibility. We mentioned the case where a contract allowed for the expansion of the ward space of a hospital but that it was not possible to use this space for an expansion of offices, which would have been easier and cheaper from a mere construction point of view. Flexibility has more value than one can reasonably foresee and put in a contract at the outset. There are unforeseen circumstances, such as the need for office space, when the flexible infrastructure will deliver more value than we had actually foreseen.

The existing relationship between the private and the public sector is essentially a sub-contractor relationship. A lot of what PFI does could also be achieved with modified traditional procurement schemes and suitable penalty clauses. What cannot be achieved without public private partnership is the creation and maintenance of high-value *flexible* systems, of "living infrastructures", where the operation, the reaping of value, is one part (the public sector's responsibility) and the regular

[6] For further details see Yun Shin Lee, *Flexible Design in Public Private Partnerships: A PFI Case Study in the National Health Service*, Master Thesis, Judge Business School, University of Cambridge 2007.

updating of the infrastructure, the exploitation of flexibilities is the other part (the private sector's responsibility). This might be coordinated, for example, by an on-site office with a remit to monitor unfolding circumstances, understanding their impact on the operation of the system and exercising flexibilities in the design as appropriate for the maximisation of the value reaped from the system. Only if the public and private sector form a partnership, aligned to achieve good long-term value for the public and good value for the shareholders, can the vision of a living infrastructure become reality.