ΎΕ ∠



Alex Bates

Bonded Reimbursements Submissions

Alex Bates Associate Director GlaxoSmithKline

All opinions expressed by the author(s) do not necessarily reflect the opinions of the Office of Health Economics (OHE). The author(s) retain copyright over their submission but have granted non-exclusive rights to OHE to reproduce and publish the submission in whole or in part in any form at any time provided that the author(s) are duly acknowledged.

Alex Bates



FINALIST OHE Innovation Policy Prize 2022

1. INTRODUCTION

There are a number of characteristic failures in healthcare markets which create problems for regulators trying to design fair systems. This essay concerns what I believe is the biggest barrier to creating a system which balances access and innovation throughout the lifecycle of medicine, information asymmetry. Asymmetric information is when one party to a transaction has better knowledge about the value of the product being bought or sold (for example, how effective a drug might actually be in the real world). This is problematic for the less knowledgeable party, because it means they cannot be sure if they are getting good value for money. In extreme cases of asymmetric information, the healthcare market could collapse because of a breakdown of trust between manufacturers and regulators.

During the course of investigating a drug, a manufacturer will generate significant quantities of 'latent' knowledge (facts the manufacturer has reasonable grounds to believe are true but which cannot be systematically proved to regulators). For example, the manufacturer might have reasonable grounds to believe that dose modification could prevent an adverse event, but be prevented by their license from making a claim that this will actually happen. This 'latent' knowledge will eventually be converted into 'common' knowledge (meaning facts the regulators could in principle act upon). The two mechanisms by which latent knowledge is converted into common knowledge are – broadly:

- The manufacturer finding a way to translate their unsystematised best estimates into a form the regulator finds convincing, which resolves asymmetry between **the manufacturer and regulator**
- The regulator waits a long time to see what the 'all things considered' verdict on a drug is after it has been used for a while, which resolves asymmetry between **the regulator's knowledge now and their knowledge in the future.**

The first mechanism - translate latent knowledge into common knowledge prior to regulatory submission – is a source of some considerable tension between manufacturers and regulators. This is because manufacturers typically have many possible ways of presenting their data to regulators, and can combine analytical techniques in a myriad of different ways in order to make claims about the drug beyond what a direct reading of the data would suggest. For example, the choice between presenting a well-documented surrogate endpoint versus an exploratory but 'real' primary patient-reported outcome measure is complicated, and the manufacture can influence which endpoint is most suitable through trial design and approaches to statistical analysis. This has been poetically described as 'The Garden of Forking Paths' by Gelman & Loken (2013); they describe analysing a dataset as like wandering through a beautiful ornamental garden, with a path that splits and forks every time the researcher must make a decision. When regulators eventually see the result, they cannot know if the researcher took the most direct route through that garden or a meandering and torturous route which forces the data to come out one way or another.



Alex Bates

Gelman & Loken describe how it is possible to imagine a dishonest person systematically exploring every possible combination of analytical techniques, and then only reporting the combination that gives the best outcome for their drug. However, it is also possible to imagine an entirely honest researcher making reasonable choices in isolation which nevertheless would have been made differently had the data turned out differently. This is probably a better description of the asymmetric information problem in healthcare; individual employees of a pharmaceutical company are broadly honest people, but have incentives that lend themselves towards overoptimism on behalf of a drug. Given the choice between two reasonable techniques that could both honestly be selected for analysis, these employees will tend to select techniques which highlight strong features of the drug and minimise weaknesses.

Regulators therefore cannot know whether a result presented to them is genuinely the best description of the data, a reasonable-but-optimistic approach that needs exploring, or an outright 'fishing expedition' for the best result it is possible to credibly generate. The example box below highlights how patients can be harmed by regulators managing a situation of complex asymmetric information poorly.



Alex Bates

EXAMPLEMAB – A BAD DECISION IN THE GARDEN OF FORKING PATHS



Examplemab is a new drug for the treatment of Examploma, a rare type of blood cancer. Unlike all other therapies for Examploma (which are oral chemotherapies), Examplemab is an intravenous infusion. The clinical trial took place during the COVID 19 pandemic, meaning that the intervention group was exposed to COVID 19 at a significantly higher rate than the control group, because they spent longer in hospital. As a result, the death rate in the intervention arm is significantly higher than in the control arm.

The manufacturer knows that the objective trial outcome is misleading – they have a good enough understanding of the drug to know that it is impossible that the COVID-19 symptoms reported in the intervention group could possibly have been caused by Examplemab. Similarly, they might have other kinds of latent knowledge, for example unprompted feedback from doctors saying, "Examplemab was working great until the patient caught COVID". However, there is no precedent for a 'COVID 19 correction' to a pre-registered trial, so regulators must rely on the good faith of the manufacturer that whatever statistical adjustment they do make is fit for purpose.

In this hypothetical example, the manufacturer proposes a complex statistical correction to the regulators (for example, match-adjusting the intervention group exposed to COVID to the control group). The regulators lack the expertise to fully understand this correction, since it is justified with reference to a number of 'latent' assumptions about how the drug works (for example, which features of the patient population to match-adjust on and which to leave alone). As a result, the regulators reject the manufacturer's approach, and take a more conservative approach of using trial-reported outcomes only. This undermines the value proposition of the drug, and it is eventually rejected.

Here, patients are harmed by asymmetric information **in favour** of a drug – the manufacturer knows their drug is most likely safe, but cannot articulate that knowledge in a way the regulator can respond to it. Therefore, the regulator are forced to reject the drug, and patients die of Examploma unnecessarily in the years it takes the manufacturer to rerun the trial.



Alex Bates

2. INFORMATION ASYMMETRY

In order to understand why existing mechanisms for preventing bad outcomes due to asymmetric information do not resolve the problem, it is necessary to momentarily consider the economic logic behind asymmetric information in this particular market.

Figure 1 shows a generic model of market failure under conditions of asymmetric information. In this example, sellers know something about a product that buyers do not, and if buyers *did* know this information their demand curve would shift leftwards. In non-economic terms, this means they would switch from demanding a high amount of product at Q_A to a low amount of product at Q_S. This in turn means the seller is no longer able to charge P_A, but must instead charge the lower P_S. Given asymmetric information, this means consumers are buying too much product and paying too much for it, because they believe the product to be better than it actually is. This breaks a fundamental assumption of market economics; the benefit consumers are getting (MSB) is much less than the cost they are paying (MSC). We can quantify the extent to which consumers are being harmed – it is the red triangle highlighting the difference between MSC and MSB marked 'DWL' for 'deadweight loss'.



Figure 1 – Archetypical example of deadweight loss due to asymmetric information, where the seller knows their product is worse than the buyer thinks

In reality, we can be fairly confident that manufacturer-regulator interactions do not function like a classic supply-demand market as depicted in Figure 1. For one thing, we have no *prima facie* reason to believe asymmetric information usually tends in the direction of hiding unfavourable



Alex Bates

information about a bad product as per the diagram above. In reality, it is equally possible that asymmetric information hides favourable information about a good product, which the regulators won't accept because the knowledge is 'latent' and the manufacturers can't find a way to demonstrate it to the regulator's satisfaction. In this situation, the regulators would be more likely to purchase the drug if they knew the full facts, which is not a situation that occurs in 'textbook' economic theory. However, the deadweight loss that occurs in this situation is symmetrical to the deadweight loss that occurs in Figure 1 - regulators purchase too little of the drug and patients miss out. A second big distinction between actual healthcare markets and the situation depicted in Figure 1 is that in many healthcare markets the regulators do not directly purchase the quantity of medicines the healthcare system will use, but rather they add the medicine to an 'approved' list of some sort with a certain probability (which rises in proportion to the social desirability of the medicine) and then allow other actors in the system to actually purchase the medicine. This is by no means universal (the USA's FDA puts drugs on an 'approved' list via regulatory approval, but then market mechanisms mostly allocate drugs afterwards), but it would alter the diagram above by replacing the x-axis with 'probability of reimbursement' or similar. Overall, there are perhaps more sophisticated models which capture the same insight with better granularity - for example iterated game theory models - but the basic model gets across the point that asymmetric information in healthcare systems will harm society by leading to suboptimal decisions by regulators.

How significant a problem asymmetric information is will depend on the particulars of the product being assessed and the healthcare system doing the assessing. For example, when knowledge of the withheld information wouldn't much change the conclusions of the regulator, the resulting aggregate demand curves will be close together and the potential loss of social value will be small. For example, a system which prioritised innovative clinical benefit over cost-reduction would be less worried about asymmetric information regarding costs as long as the primary clinical benefit was well demonstrated. On the other hand, there are other circumstances where even a tiny amount of asymmetric information could radically change the conclusions of regulators. For example, a system with extremely elastic demand curves due to the presence of many clinicallysimilar comparators would be especially prone to making the wrong decision in the presence of asymmetric information.

In general, certain therapy areas will be more prone to problematic trial data than others. For example:

- Areas where the drug modifies an intermediate endpoint with high certainty, but the relevance of that endpoint to patients is hard to establish through formal trial mechanisms (such as the recent concern that managing beta-amyloid plaques is insufficient to have a disease modifying effect on Alzheimer's Disease)
- Drugs with intensely expensive loading doses which are expected to offer benefit over a long period of time (such as CAR-T therapies), such that uncertainty about the exact level of benefit is magnified significantly
- Areas where conventional clinical trial execution is very difficult such as mental health conditions which affect the ability to consent to a trial (like depression and schizophrenia)



Alex Bates

- Rare diseases, which must necessarily conduct trials with fewer patients and therefore are especially reliant on 'latent knowledge' in the interpretation of noisy results
- Therapies with substantial qualitative benefits over existing standard of care, but which don't significantly modify objective disease response characteristics – that is, the importance of the qualitative benefit to patients is difficult to establish in the absence of extensive clinical experience with the drug.

Since manufacturers know that certain types of claim are higher-risk, this leads to 'adverse selection' (meaning manufacturers will vary the disease areas they investigate in order to minimise the risk of favourable information asymmetries being discarded by the regulators). It is possible to imagine this mechanism working in reverse too - where manufacturers know certain disease areas receive less scrutiny and so attempt to exploit their asymmetric information more fully than they would otherwise. An interesting 'real-world' example of this adverse selection phenomenon may be GLP-1 agonists. GLP-1 agonists, like semaglutide, control obesity in combination with diet and exercise. Obesity is an important medical condition which costs healthcare systems billions of dollars per year, but is a highly politicised space - it is not clear that it is a 'disease' in the conventional sense regulators are used to and pharmaceutical companies wishing to offer a GLP-1 agonist will need to convince regulators that they are addressing a genuine medical need rather than just offering a lifestyle intervention. The mechanism of mitochondrial uncoupling has been known about since the 1930s, but (potentially) because of the risks of innovating in this space no company has been prepared to develop this mechanism into a safe-in-humans medicinal product until very recently. Therefore, the potential 'chilling effect' of this adverse selection could reasonably be on the order of billions of dollars worth of value per year.

The diagram below demonstrates how this chilling effect might work in practice. Regulators would be prepared to pay a premium for innovative medicines, but they can't distinguish between innovative medicines and incrementally superior / 'me too' medicines, since every manufacturer claims their product is innovative (and manufactures with genuinely innovative products often only know their product is innovative because of latent knowledge, e.g. a deep understanding of the disease area which regulators lack). For example, a new immunology product which targets a unique interleukin signalling pathway might be meaningfully innovative or might be a 'me too' depending on the exact mechanism of action, and it is hard to issue a blanket statement about which pathways are innovative and which are not without a deep knowledge of the immunology space. Because of this, regulators clear the market assuming only one type of product, corresponding to the 'Aggregate Supply' line in Figure 2. This creates deadweight loss in two places; innovative products are less likely to be purchased than the socially optimal amount, and incrementally superior products are more likely to be purchased than the socially optimal amount. This is important for demonstrating firstly why manufacturers might preferentially target 'me too' drugs under conditions of information asymmetry (because they earn a premium on them relative to a fully free market) and secondly why asymmetric information is a **particular** problem for innovative medicines, rather than just a generic market failure.



Figure 2 – Example of asymmetric information in a two-product market leading to both types of product being consumed at socially non-optimal rates



3. EXISTING SOLUTIONS TO THE PROBLEM

Looking only at the theory of asymmetric information might give the impression that the relationship between pharmaceutical companies and regulators is a proverbial 'wild west', where pharmaceutical companies will make almost any outlandish claim provided it increases the chances of their drug being reimbursed. In fact, this is not the case in general, and many of the generic solutions to asymmetric information function in the pharmaceutical market too:

- Most pharmaceutical companies are major multinationals with well-developed pipelines. A company which developed a reputation for hoodwinking regulators would quickly find the benefit of the doubt drying up when it actually mattered.
- For many outcomes, the problem of asymmetric information is more or less solved. For example, there are industry standard methods of conducting trials (such as pre-registration of hypotheses, double-blinding investigators to which agent is being used on which patient, using intention-to-treat analysis, and so on). In many countries, regulators are experts in these methods due to having assessed so many submissions of this kind, and therefore have the skills to establish whether a company has deviated from these methods for good reasons or not. Companies deviating simply to exploit information asymmetry would be quickly identified as attempting to withhold information from a regulator.



Alex Bates

 Most pharmaceutical companies are staffed by basically honest people, who will refuse to go along with obviously unethical behaviour (especially behaviour which might harm patients). Even in the case of a hypothetical company staffed by entirely dishonest people, most countries have an industry body for the pharmaceutical industry which has an interest in generally protecting the reputation of pharma who would penalise this hypothetical company.

These mechanisms prevent a catastrophic breakdown of the market for healthcare. However, these solutions do not prevent the mechanisms described above from operating entirely. In particular, they are not scalable. For example, regulators are experts in standard methods of conducting trials. But as trials get more complex the healthcare system cannot continue to expect regulators to remain at the cutting edge of trial design theory. This means that specific and formal solutions are needed for the sorts of submissions which are most affected by information asymmetry (i.e. especially innovative medicines).

The current gold standard of asymmetric information resolution is 'pay-for-performance' (P4P), or 'payment by results' (PBR). Under a PBR scheme, the price paid for a technology is adjusted in some sense to reflect the value the technology brings to society. For example, a classic example of a PBR scheme would be a situation where a manufacturer offers the healthcare system their product for free during the 'induction' period of dosing, and then the healthcare system pays full price during the 'maintenance' period if a response is seen during induction. This should ensure that the healthcare system only pays for patients who are responding to the drug.

However, PBR is not a general solution to the problem of asymmetric information and suffers from a number of characteristic drawbacks:

- Even in the absolutely best case, a PBR scheme is burdensome for the healthcare system. For example, in the case above where induction is free but maintenance is not then the healthcare system must usually implement a layer of monitoring to ensure it has met its contractual obligations not to give any 'free' stock to patients who should be charged, or accidentally paid for stock that should have been 'free'. This burden increases significantly as the PBR scheme becomes more complex if the drug in question was a treatment for Alzheimer's Disease, then it might be necessary to give a battery of cognitive tests each time the drug was given to ensure that the performance of the drug was still at acceptable levels. This would be hugely burdensome for patients, and expensive and time-consuming for the healthcare system. Since generally more innovative drugs will require more complex PBR pricing, this mechanism tends to disincentivise innovative drugs coming to market.
- PBR schemes are inherently incentive-misaligned. Since payment is conditional on meeting some good outcome, but there is no penalty for failing to meet this good outcome, then there is no reason a manufacturer should not 'have a punt' on meeting an outcome if the alternative is not getting reimbursement at all. Depending on the details of the scheme, there may be a financial incentive to focus on maximising PBR payments rather than improving patient outcomes. For example, a manufacturer might preferentially target healthy patients who are more likely to meet the agreed endpoint



Alex Bates

than patients who will benefit from the drug more holistically, in a process called 'cherry picking'.

PBR schemes are not robust to product lifecycle considerations or considerations of clinical dynamics. If a manufacturer creates a product which might extend life by one year in combination with chemotherapy, then improvements in chemotherapy techniques will help the manufacturer reach this milestone even if the drug itself is not especially effective. Another example of this effect would be a manufacturer creating a superb drug which is significantly better than everything else on the market, which as a result is used mainly in very sick patients near end of life. Although this approach may be socially optimal (for example in societies which value health at end-of-life higher than other kinds of health) it will prevent the manufacturer from reaching endpoints agreed on the basis of a clinical trial against peer competitors, because their drug will only ever be used on patients who have inherently worse outcomes than for their competitors. Again, since innovative technologies are more likely to have unusual clinical dynamics or lifecycle considerations this effect tends to mean PBR selects against innovative technologies.

The example box below highlights how pay-for-performance can prevent good outcomes being reached.



Alex Bates

PANACEAMOL - A CATASTROPHIC FAILURE OF PBR



Panaceamol is a new antibiotic. It is universally understood to possess a revolutionary mechanism of action which cannot be evaded by MRSA, but regulators are uncertain how much better it is than ordinary antibiotics for day-to-day use. To manage the risk, the regulators propose a PBR scheme where the manufacturer is paid only if the patient does not require extra bed-days due to a hospital acquired infection.

Due to the high cost of Panaceamol compared to e.g. penicillin, doctors naturally withhold the new drug from most patients. In fact, most hospitals implement a cost-containment rule that Panaceamol should only be given to patients who are considered extremely high risk for MRSA – for example post-operative immunocompromised patients. While Panaceamol performs admirably in this group, they still acquire infections at a rate much higher than the regulators expected due to their underlying base rate of risk. As a result, Panaceamol is a significant loss-maker for the manufacturer and the manufacturer disinvests in their antibiotics R&D division.

In this example the manufacturer is probably at fault for failing to spot the obvious potential for problems the PBR scheme causes, but merely highlighting the problem is no guarantee there is any way for the healthcare system to route around it – in some healthcare systems there just isn't the depth of healthcare informatics to identify whether doctors are behaving in this way and what impact it has on reimbursable outcomes.



Alex Bates

Bonded reimbursement submissions are an incentive-aligned solution to the problem of asymmetric information

4. INTRODUCTION

This essay proposes a complete solution to the problem of information asymmetry described in the introduction. That solution is a **bonded reimbursement submission**. It is summarised in the diagram below:





This proposal assumes no changes to the existing reimbursement systems in any country, except that the manufacturer is offered a novel option for overcoming a negotiating deadlock; they can propose to make a 'bonded reimbursement submission'. This is conceptually an 'PBR-plus', in the sense that it is a risk sharing agreement like an PBR but with an additional step of creating a bond market in order to force it into incentive-compatibility. A bonded submission consists of the following four elements:

1. The manufacturer's definitive best estimate for the value of a parameter, which we assume logically must be the major sticking point preventing approval (it would be a waste of time and resources to undertake the following steps over an irrelevant parameter, but it would not be incentive-misaligned so no other harm would come of it).



- 2. A contractual stipulation that should a definitive trial find this parameter to be under the claimed level a rebate would be paid on every patient initiated on treatment. This wouldn't have to be a 100% rebate, but rather a reflection of the amount of uncertainty the bonded submission is attempting to bridge.
- 3. A description of what a 'definitive trial' would look like in reference to the above, potentially with some indicative costings to ensure it is feasible. It would be reasonable for this to be written in collaboration with the regulator, since it would be unfortunate if the regulator rejected the entire submission over a technicality in the trial description. However, importantly, the regulator does not *dictate* this step and the manufacturer can withdraw from the bonded submission without withdrawing from the medicinal product submission if the 'definitive trial' would be biased against their product in their view.
- 4. A legal undertaking to release a certain number of bonds to the public under a certain timetable. These bonds would pay out in the event that the 'definitive trial' found the parameter to be below the claimed level, and would be priced such it offers exactly zero expected value to bondholders given the manufacturer's estimate of the parameter of interest is accurate.

The regulator could then accept or deny this submission with a critical parameter 'in bond'. If the regulator accepted the submission, then all the steps above are actioned; the product is immediately made available to patients, and the first tranche of bonds become available for sale as per step 4. If the regulator denies the submission then the submission reverts to the regular pathway and whatever usually occurs in this situation will occur (e.g. the manufacturer could take the regulator to court if this was unreasonable, or offer a lower price, or withdraw from the market etc).

The key innovative element to this proposal is the use of conditional bonds, and especially the way these bonds can create a secondary market for information after the initial reimbursement decision. These bonds should be issued such that they offer zero expected value to bondholders, such that the only way the bondholder will make money on expectation is if they successfully predict the outcome of a 'definitive' clinical trial better than the manufacturer. In most cases this will be because the clinical situation on the ground changes – a new discovery the manufacturer could not have predicted causes the probable value of the bonded parameter to change. A secondary market in bonds allows us to track how the probability of the event changes over time, since the price at the margin for a bond on the secondary market is the most accurate assessment of the probability of an outcome we could wish for. As well as a price signal, the secondary market provides an incentive to translate academic research into clinical trials. For example, if a research group discover that the ABC pathway is implicated in some particular bad outcome, and a new drug targeting the ABC pathway has just been licensed under a bonded reimbursement then there is a strong incentive for the researchers' host institution to buy up a lot of the bond (increasing the price) and conduct the definitive trial necessary to make a profit. This rapid iteration of latent knowledge to common knowledge is of very high benefit to patients, and will happen automatically as a consequence of the primary bond issue.

The number of bonds being issued should result in equivalent income to the manufacturer as the amount they put 'at risk' from the PBR ('at risk' is slightly misleading because the only loss to the



Alex Bates

manufacturer is opportunity cost). In this way, the manufacturer can exactly cover their counterparty risk from the point of view of the regulator, even if the medicine at issue becomes an unexpected blockbuster. This could be achieved straightforwardly by issuing one appropriately-priced bond per patient initiated on treatment.

The goal of this system is threefold:

- The initial price bond is an incentive-aligned signal of the manufacturer's exact level of confidence in a particular claim. Insofar as a submission hinges on the truth of that claim, this allows regulators to overcome the information asymmetry barrier and make better decisions, which in turn removes the need for regulators to be conservative with respect to funding innovative medicines.
- 2. The *ex post* price of the bond in the secondary market could be used to manage uncertainties relating to the lifecycle of the product. For example, we would expect the price of a bond conditioned on the claim, "Examplemab is the most efficacious XYZ inhibitor" to change if a new XYZ inhibitor entered the market. This could be used as part of a pre-agreed schedule of incentive-aligned pricing in highly dynamic markets (without requiring a new PhIII trial every time a new product entered the market).
- 3. The *ex post* price of the bond also serves as an excellent signal of the social value of a definitive trial into certain claims.

The example box below shows how the bonded reimbursement submission could protect patients from unethical use of asymmetric information while still rewarding manufacturers with a fair price for innovative medicines.



Alex Bates



Earlier we discussed 'Examplemab' as an illustration of how significantly asymmetric information can harm patients. A **bonded reimbursement submission** could have protected patients while not penalising manufacturers for innovative analysis techniques.

Suppose the manufacturer believes there is a 60% chance that the excess deaths on Examplemab were caused by COVID, and therefore a 40% chance they were caused by Examplemab itself. They plan to sell Examplemab for \$10,000 per patient. Under the proposed scheme they would offer a bond to the market initially priced at \$4000 which pays out \$10,000 if a definitive trial concludes that Examplemab itself elevates death rates. The definitive trial, in this case, could be a relatively straightforward retrospective database analysis run in five years' time, assuming no further COVID outbreaks. Therefore, the manufacturers are claiming that in 60% of cases they will take in \$4000 and pay out nothing (expected value \$2400), and in 40% of cases they will take in \$4000 and pay out \$10,000 for a \$6000 loss (expected value negative \$2400). The manufacturer therefore has overall neutral expected value from the bond.

A particular worry of regulators is dishonest manufacturers who try every combination of analytical techniques until they find the one they like the best. An unethical manufacturer might therefore try to claim that it is almost certain the excess deaths were caused by COVID, rather than merely likely on balance. Suppose they claim that the deaths were to do with COVID with 95% probability, while knowing in reality that 60% is more realistic. Then they would be forced to offer the bond at just \$500, despite knowing that in 40% of cases they have to pay out \$10,000. Their expected value from the bond is negative \$3500, plus any additional rebates imposed by the regulator for failing to hit their PBR endpoint.

As demonstrated above, the scheme is incentive-aligned, since reporting the true best current estimate of a parameter is always in a company's best financial interest. In the first case, the manufacturer is rewarded for appropriate innovation (in this case, innovation in analysis methods – being the first to identify how to correct for COVID deaths in a trial). In the second case, the manufacturer is penalised for either deliberate dishonesty or simply misleading the regulator with their over-optimism. Given that this incentive alignment means that only honest manufacturers will enter the market, patients will benefit from access to the innovative Examplemab technology.

Patients also have an additional layer of protection from the risk that Examplemab is dangerous, which is that the price of the secondary market for the bonds will give up-to-the-minute information on the overall risk of mortality. If this price changes, the regulators can quickly re-assess their original decision.



Alex Bates

5. MECHANISM OF ACTION

The key reason that bonded reimbursement submissions solve the problem of asymmetric information is that they offer an incentive-aligned reason for manufacturers to truthfully report their latent knowledge of a product, thus making knowledge symmetric, so regulators can confidently make decisions.

An important feature of this mechanism is that it is risk free to the manufacturers as long as they are conservative in their assumptions. If manufacturers are excessively conservative then their bonds won't sell, which exposes them to no risk and provides an unambiguous signal to regulators that the company is truthfully (under)-reporting its latent knowledge. However, if manufacturers are insufficiently conservative then they take on risk as the uptake for their bond increases. In some cases, this risk may be worth it to manufacturers – for example if the difference between convincing a regulator about a particular feature of the drug and not is the difference between launching a drug and not. In other cases, especially in cases where the clinical situation on the ground has changed, the manufacturer may no longer feel confident they can support the bonded submission claim. Under these circumstances the system should behave as though the definitive trial was negative, which would usually mean switching to a fallback price pre-agreed with the regulator.

An important element of this system is that *it does not matter if a 'definitive' trial is ever run*. In the same way that a company can have a non-zero value on the stock market even if it has never paid dividends, it is the 'threat' of a trial being run which forces manufacturers to honestly report their latent knowledge – not the actual trial itself. Indeed, insofar as we would expect the eventual trial to reflect informed clinical opinion of the product at the time the trial is run, the trial itself has zero social value in expectation.

In comparison to PBR, bonded reimbursement submissions:

- 'Frontload' all complexity, and load that complexity onto the manufacturers / regulators (who are well equipped to deal with it) rather than the broader healthcare system which is not set up to perform this kind of task
- Move the risk from the healthcare system (which pays for upside but not downside in PBR) onto the manufacturers who have the better end of the asymmetric knowledge tradeoff, and ensures that any risk the manufacturer takes on is balanced by a commensurate reward for their better knowledge
- Have a secondary market to provide ongoing information about the medicine even after approval. This is potentially the single biggest advantage of bonded submissions over PBR, since it allows sophisticated lifecycle management of products while rewarding innovation at the time it occurs

USE-CASES OF BONDED REIMBURSEMENT SUBMISSIONS

The principal use of bonded reimbursement submissions is to resolve the problem outlined in the introduction – systems which rely on an assessment of information to make decisions find it more difficult to reach fair prices and balances access / innovation in the presence of asymmetric information. This means that it is likely that bonded reimbursement submissions will find the most



Alex Bates

use in cost-effectiveness HTA countries, since these countries have healthcare systems which are the most susceptible to decision-altering asymmetric information. In principle, every piece of latent information is decision-altering in a cost-effectiveness HTA country since it will alter the price the manufacturer can charge and remain under the cost-effectiveness threshold.

However, it is not required that the country use a broadly cost-effectiveness HTA system in order for bonded reimbursement decisions to function. The example below considers a country which uses a broadly 'clinical benefit' HTA system. In all other respects it describes what I envisage as the primary use case of bonded reimbursement submissions, which is removing asymmetric information from a critical element of a reimbursement submission in order to make approvals of innovative products at a fair price more likely.



Alex Bates

HYPOTHETIMOL – BONDED REIMBURSEMENT SUBMISSIONS IN A 'CLINICAL BENEFIT' CONTEXT



Hypothetimol is a new small-molecule medicine for the treatment of diabetes. The medicine is expected to be a marginal improvement on existing treatments, but the real innovation is making the dose once-weekly instead of once-daily, which is expected to be a substantial quality of life improvement for some patients. The manufacturers, a large American company, have designed the pivotal trial to be non-inferiority, in order to rapidly enter key markets such as USA, China and Japan.

This approach causes problems for reimbursement in France, where a non-inferiority trial often means the drug will receive the lowest ASMR rating, especially when the qualitative importance of the innovation (convenience to patients) is hard to demonstrate. An ASMR V rating would mean that this drug cannot launch in France given the global pricing strategy, and French patients will miss out on the innovation of more convenient treatment for their diabetes.

Using a **bonded reimbursement submission** this issue can be resolved. The manufacturer proposes a (hypothetical) superiority trial for Hypothetimol versus a comparator as the 'definitive trial', and argue that if it were true that Hypothetimol is mildly superior then they should receive at least an ASMR IV rating. The regulators accept the design of the trial, and propose that as long as the company is at least 80% sure the improvement will be demonstrated in their trial then they can provisionally accept the ASMR IV rating. This is consistent with the company's pricing strategy, and so French patients benefit from the mildly improved convenience of treating their diabetes.

This is a useful demonstration of how the system can drive access to innovation even if the innovation is not 'game changing' for patients – in this case it is good that the market has found a fair way to compensate the company for their mildly innovative new treatment, even if the healthcare system of a particular country isn't traditionally set up to recognise such incremental qualitative improvements.



Alex Bates

A second use of bonded reimbursement submissions is to use them more generally to systematise latent knowledge. This usage of the technique would be especially appropriate in more market-oriented countries like the USA, where regulators tend to have less of a role in determining what medicines the healthcare system purchases and more of a role overseeing the mechanics of market competition.

The example below demonstrates how the bonded reimbursement submission can be used to systematise latent knowledge. The particular feature to draw attention to is that it is not possible to make this claim in the absence of a mechanism like a bonded reimbursement submission – it would simply not be economically feasible to run a ten-year trial to demonstrate a ten-year time-on-treatment, which means that patients are unable to learn about this important and innovative benefit of Illustra Citrate.



Alex Bates

ILLUSTRA CITRATE – USING BONDED REIMBURSEMENTS TO UNDERWRITE CLINICAL CLAIMS IN THE USA



Illustra Citrate is a new biologic treatment for Illustrative Arthritis The clinical particulars of Illustrative Arthritis are that patients develop a tolerance to all existing treatments after about five years, necessitating a disruptive and unpleasant treatment switch. The manufacturer is confident that Illustra Citrate will work longer than all other treatments due to its innovative mechanism of action, but weren't able to prove it in a clinical trial due to the cost of extending the trial that long.

When thinking about their launch in America, the FDA (the American regulators) don't really care about the additional time on treatment, since they are mostly concerned with whether the benefits outweigh the risks and the time on treatment is not necessary to demonstrate this (let us suppose). However, the manufacturer knows that there is significant unmet need for long-acting treatments, and would like to make some reference to this important and innovative benefit of Illustra Citrate when making claims about their product.

Using a **bonded reimbursement submission** this issue can be resolved. The manufacturers can approach the FDA with a design for a hypothetical trial that would allow them to make a claim, "Illustra Citrate lasts for ten years on average" and make reference to a bond market resolving that exact question. The FDA agree this claim can be made as long as the manufacturers are 90% certain of themselves, which is to say issue the bond at \$1.00 payout per \$0.10 investment if they are wrong. The manufacturer is happy with this risk, and consequently is able to gain a dominant market share in America on the basis of being able to make more patient-relevant efficacy claims than their competitors.



Alex Bates

A final use of bonded reimbursement submissions is to allow the secondary market for bonds to create a fair pricing environment throughout a product lifecycle. This is a more niche application than the other two uses, since manufacturers will often alter their prices in response to market dynamics in a way that makes the more complex bonded reimbursement submission redundant. However, where the application is valuable at all it is likely to be highly valuable, since the mechanism functions 'automatically' and so manufacturers and regulators can both trust that political and business considerations (respectively) will not cause the other party to defect and demand an unfair price for the drug in the future.

The example below shows a use-case for the secondary bond market, where the value of a drug increases over time and the company can capture this value if they are confident in their latent knowledge about this increase.



Alex Bates



EX-T is a revolutionary genetic therapy for the treatment of heart disease. A single dose has a high chance of reversing many common cardiac complications, and would be lifesaving for many patients currently on the waiting list for a heart operation. Unfortunately, EX-T can only be given to patients who are otherwise extremely healthy, because of the toll it takes on the patient during the activation phase.

The manufactures are confident that their trial is too conservative regarding the minimum standard required to be 'healthy', and that when clinicians gain experience with managing the side-effects of EX-T they will be able to treat significantly sicker patients. This is a classic case of 'latent knowledge' – the manufacturer learned something about the drug in the course of running the trial which it did not predict from benchtop research. However, the regulators are unwilling to risk expensive doses of EX-T on patients who might not benefit. Ironically, competitors to EX-T benefit from the clinical knowledge gained from EX-T prescribing and are able to be more aggressive in their trial design, effectively punishing the EX-T manufacturers for their innovation.

A **bonded reimbursement submission** can solve this problem. The manufacturers propose a hypothetical definitive trial that would effectively conclude that the benefits were worth the risk in some particular high-risk group (for example, a subgroup analysis of their original trial with enough power to detect the result they need). The manufacturer knows that these bonds will initially trade at a loss for them; at the moment EX-T is approved, the clinical community has no experience of the product and it would seem too risky to treat these patients. However, over time the manufacturer expects the clinical community to gain more familiarity with the treatment, to the point where they might even make a small profit on the bonds. The regulator agrees that when the bonds begin to trade below 10% then the drug can be reassessed for use in the high risk group.

In this case, the manufacturer is appropriately rewarded for their innovation and ambition for patients, but only when this reward makes sense in the context of the value of the drug over time; the value of EX-T increases as clinical experience of the drug increases. Because the manufacturers control their own exposure to risk, they can determine whether the risk of the bonded reimbursement submission justifies the increased speed to market (for example if competitors are hot on their heels) versus a slower and more certain route to market such as running a second trial with more a aggressive treatment regimen.



In the relatively limited space allotted for this essay it is not possible to explore every way in which bonded reimbursement submissions might be used. I have therefore tried to focus on areas where there is a clear need from manufacturers to be able to systematise their latent information (that is, where the cost / risk of making a bonded reimbursement submission is clearly financially incentivised for the manufacturer), and where the drug in question is sufficiently innovative that conventional mechanisms for reimbursement will not function well. The three above examples represent central use-cases for the scheme.

Implementation of bonded reimbursement submissions

Bonded reimbursement submissions are notable for their complexity. Complexity is inherently bad in the healthcare system, as it increases the workload on an already strained system. In general, however, it is reasonable to suggest that bonded reimbursement submissions and PRB are of approximately equal overall complexity. In contrast to a PBR scheme, the complexity is frontloaded onto the manufacturers and regulators (rather than spaced out over the whole healthcare system) and thereafter the mechanism functions more-or-less autonomously. It is still clear that bonded reimbursement is not a panacea, but given the importance of making the right decisions in assessing new medicinal products, it is likely to offer sufficient reward to justify the complexity in many cases.

Conceptually, there are two main areas where we might expect difficulties implementing the scheme because of the complexity; the negotiations over what constitutes a 'definitive trial' are likely to be challenging, and the implementation of the 'bond markets' requires the creation of a new type of financial instrument.

6. 'DEFINITIVE TRIALS'

The mechanism of the bonded reimbursement submission depends on there being some way to arbitrate the truth in principle, a 'definitive trial'. The intent behind a definitive trial is to settle a specific clinical question beyond any reasonable doubt, and for that clinical question to be sufficiently well specified that all parties agree that the trial investigates the critical issue for reimbursement.

The proposed mechanism by which a definitive trial is agreed is that the manufacturers include a detailed trial protocol for a hypothetical trial as part of their bonded reimbursement submission. This protocol should be detailed enough that it is clear what would and would not constitute a definitive investigation into a topic, and should be at least somewhat feasible (for example a trial with a clearly unethical comparator should not be proposed by manufacturers, although a trial which was merely very expensive might be appropriate). The expected level of detail would be – for example – a description of a trial protocol suitable for publication in a high-ranking clinical journal.



The figure below shows a possible 'definitive trial' schema for a hypothetical bonded reimbursement submission on immediate vs deferred radiotherapy for an incompletely excised tumour, taken from the NICE Clinical Guidance on the same topic. Professional bodies routinely take NICE Research Recommendations and translate them into meaningful trials, and therefore it is reasonable to point to NICE Research Recommendations as examples of good-practice trial design schema which could be adapted into this proposal's 'definitive trials'.

Figure 4 – An example of a hypothetical trial design, taken from a NICE 'Recommendation for Research'

Population	Adults (18 years onwards) with an incompletely excised or inoperable grade I meningioma
Intervention	Immediate radiotherapy
Comparison	Deferred radiotherapy (given on clinical or radiological progression)
Outcomes	 Quality of life Neurocognitive decline Overall survival Progression-free survival Local control Radiation Therapy Oncology Group toxicity
Study design	Randomised controlled trial
Timeframe	10-year follow up

This protocol could be contentious in its own right. For example, if an issue at stake in a submission is whether a drug is effective in 'high risk' patients, the definition of 'high risk' is likely to be fought over by the manufacturer and regulator. The regulator would be encouraged to offer suggestions on how to improve the trial if it was narrowly rejected, but the regulators should avoid 'designing' the trial as far as possible, since the whole point is to allow the manufacturer to express latent knowledge that only they possess. In essence, the regulator tells the manufacturer what level of evidence would satisfy them on a contentious point, and the manufacturer determines whether it is possible to create a trial protocol satisfying this uncertainty. The manufacturer carries no risk at this stage – if they are unhappy with proposals made by the regulator they can withdraw from the bonded reimbursement submission and continue with the submission as usual, perhaps for example lowering their price to overcome the uncertainty rather than using the conditional bond instrument to overcome it.

It should be clearly specified and legally enforceable what the outcome will be if the 'definitive trial' returns a negative result from the point of view of the manufacturer. In most cases the appropriate action will be that the manufacturer immediately drops the price. It may also be appropriate to impose a PBR-style rebate in the event of a negative trial result, for example paying the difference between the premium the manufacturer obtained for their innovation and the cost of the next-most expensive drug. For some outcomes, such as those which directly relate to patient safety, it



Alex Bates

might be socially valuable to err on the side of caution and treat a high bond price on the secondary market as effectively being the same as a negative trial result. It is critical that the action all parties need to take is clearly specified in advance of the trial (and absolutely unambiguous) to prevent legal challenges being mounted on the enforcement of any particular bonded reimbursement scheme, which would undermine their value as a price signal.

In theory, the robustness of these preceding steps means that trial design and execution for the definitive trial should be completely automatic; the team of triallists who actually conduct it could blindly follow the protocol and reach the correct results. In practice, however, trial conduct is as much an art as a science and there must be flexibility to allow the triallists to avoid unethical or incoherent elements of the protocol. For example, the protocol may specify that patients can take 'concomitant steroids' but this may make no sense in a world where the standard of care adjunct medicine is no longer steroids. This specific example is not so problematic – we can stipulate that the regulators and manufacturer can amend the trial by mutual agreement. However more generally, we may worry that the team that actually conducts the trial might have a bias which could impact on results, which is undesirable.

To avoid this impacting outcomes, a few common-sense stipulations should be included such as preventing the company proposing the bonded reimbursement submission from funding / carrying out the definitive trial. No matter how carefully designed the protocol, there will always be unavoidable ambiguity in the result of the definitive trial and this ambiguity will be critical in some percentage of cases. Indeed, we'd expected the bonded reimbursement submission to be particularly valuable in cases where the manufacturer has a lot of latent knowledge – which is precisely the situation where the manufacturer would be most able to undetectably affect the outcome of the trial! It is likely therefore that the triallists for the definitive trials would be well-regarded neutral academic groups, who receive funding from a national study sponsor body (like the NIHR in the UK or NIH in the USA). A positive externality from the secondary bond market is that these national study sponsor bodies will find it much easier to identify trials with a positive social value to fund, since the market capitalisation of a particular bond market represents the value resolving the question has to society (i.e. markets which have a large number of bonds issued at a high value per bond, and which therefore represent a social consensus that the regulators have made the wrong decision).

One small technical point is that the bonded submission must be offered at a fixed price (reflecting the probability initially assigned by the manufacturer that a certain proposition is true) in order for the incentive-alignment mechanism to work. However, the manufacturer must sell these bonds onto an open market, where the price of the bond might have risen significantly higher than this fixed price. For example, if a bond was offered by the company at \$4000 but was trading at \$7000, then when the company offered the bond for \$4000 they would essentially be giving a 'free' \$3000 to the first buyer who happened to notice the bonds for sale. In principle this doesn't matter from the point of view of an incentive-aligned system, but in practice we might have reasons to prefer not to reward investors with huge payouts just for exploiting this predictable arbitrage opportunity. One strong suggestion for what to do with this overage would be to allow the company to offer the bonds at the market rate, but mandate that the company hold the difference between the agreed price and the spot price in an escrow account, payable to the host institution of the first group who



Alex Bates

conduct the relevant definitive trial. This would create a strong financial incentive to conduct trials where the social value of information was high (significant probability of overturning a regulatory decision and affecting a large number of patients), without altering the functioning of the mechanism.

7. 'BOND MARKETS'

The most unique feature of the proposal is the use of secondary bond markets to measure the probability of some critical outcome in the regulatory submission. Although complex, it is notable that in contrast to the concept of a 'definitive trial', the idea of 'a bond market' already exists, and bond markets are successfully used to trade trillions of dollars of value every day.

There are therefore doubtlessly many ways in which a bond market could be adapted to trade the specific financial instrument created during a bonded reimbursement submission. One of the simplest, and what I propose in this essay, is to repurpose existing 'prediction markets' to perform the market clearing function of a bond market. Prediction markets effectively trade bonds which are conditional on some real-world outcome (with the intention being to predict these real-world outcomes). The overlap with what a bonded reimbursement submission is trying to achieve is therefore excellent; prediction markets already exist, and are highly functional at resolving complex questions on the efficacy of medical interventions. For example, the screenshot below is taken from the prediction market Metaculus, and concerns the market's view of the probable efficacy of AstraZeneca's COVID vaccine. The Metaculus community predicted the efficacy would be 82.5%, and the trial eventually resolved with an efficacy of 76.0%. Importantly, the market reached a view that 76% was within the 95% confidence interval in December 2020 – more than a year before the official trial readout and based only on aggregating public knowledge about the mechanism of action of the drug (which would otherwise be 'latent' for the manufacturer).

Figure 5 – Screenshot from https://www.metaculus.com/questions/5800/astrazeneca-covid-vaccine-effectiveness/





A key advantage of bonded reimbursement submissions over conventional PBR is that they can address questions which are otherwise unresolvable with PBR mechanisms. For example, the screenshot below – also from Metaculus – shows the market's estimate for the probability that NAD+ boosters will extend human life expectancy. This is a highly innovative mechanism of action which it would be very challenging to demonstrate in a standard clinical trial, since the expected improvement in life expectancy would only be demonstrated after fifty to sixty years of taking the drug. Despite this challenge, the market below offers key operational insight to regulators; over time, the case for NAD+ seems to have weakened, and therefore regulators may wish to demand a higher risk premium from manufacturers of NAD+ if any make a reimbursement submission.

Figure 6 – Screenshot from https://www.metaculus.com/questions/4290/will-nad-boosters-be-shown-in-a-systematic-review-to-increase-human-lifespan-by-5-by-2030/

Will NAD+ boosters be shown in a systematic review to increase human lifespan by >=5% by 2030?



Metaculus is used as an example of a prediction market because it is large and well-traded, but several other prediction markets exist. One slight barrier to their use is that the USA has a different regulatory framework for prediction markets, and so for most practical purposes markets which operate in the United States cannot operate elsewhere. This is not so problematic in the context for which the markets are intended to be used – a USA-specific question hosted on the market Kalshi does not necessarily compete with a Europe-specific question hosted on the market PredictIt, any more than the FTSE competes with the Dow Jones. However, it is more likely that the markets will be more liquid (and hence do a better job of resolving questions) if this regulatory mismatch was corrected before the deployment of the first bonded reimbursement submissions.

7. CONCLUSIONS

Asymmetric information is a particularly pernicious form of market failure in healthcare markets, because it disproportionately harms manufacturers with innovative assets. Because regulators have no way of distinguishing scrupulous manufacturers from manufacturers with a tendency to



Alex Bates

over-exaggerate the benefit of their drug, a natural reaction is for regulators to treat claims made about innovative medicines extremely sceptically. But conservative regulators are a barrier to a well-functioning healthcare system because they deny patients access to genuinely valuable medicines, as well as distort the market away from innovation and towards incrementally beneficial drugs.

There are relatively inelastic limits to how much asymmetric can distort the market. Regulators have come to expect a minimum standard of trial design to prevent genuinely dishonest manufacturers from outright fabricating claims, and since most manufacturers expect an ongoing working relationship with regulators there are limits on how optimistically a manufacturer will interpret a dataset in order to preserve their credibility in the long run. Regulators who lack sufficient expertise in a particular disease area can usually call upon the expertise of others (e.g. clinical experts) to vet claims which seem particularly fanciful, although this will not work in every case. Nevertheless, within the bounds of these relatively inelastic limits is more than enough information asymmetry for deadweight loss and other socially undesirable outcomes to occur. Furthermore, the direction of the modern healthcare market is tending more towards emphasising this effect than reducing it; innovative products (and associated complex trials) increase the relative edge information asymmetry gives manufactures over regualtors.

The current state-of-the-art in resolving the problem is 'pay-for-performance' (P4P), or 'payment by results' (PBR). This provides a disincentive for manufacturers to enter a market where they know their claims cannot be supported, since they are not paid if the drug does not perform well. As discussed in the body of the submission, this is directionally better than no attempt to align incentives, but there are characteristic flaws in P4P / PRB schemes which manufacturers could potentially exploit to the detriment of the healthcare system. The biggest (in the context of the OHE Question) is that P4P / PRB has no mechanisms which act sensibly throughout the lifecycle of medicines, and this is a problem which is particularly likely to affect innovative medicines where the market dynamics over time will be harder to predict for the regulators.

Bonded reimbursement submissions resolve this problem by creating a mechanism for converting 'latent' knowledge that only the manufacturer knows into 'common' knowledge which regulators can use as the basis for decision-making. A primary market for bonds issued at a fixed schedule by the manufacturer removes the hidden risk of latent knowledge, and a secondary market for bonds can manage lifecycle risk in a predictable and objective fashion. Bonded reimbursement submissions clearly carry significant administrative burden (incident on the manufacturer) and so would not be appropriate as a general tool for uncertainty reduction. However, for certain kinds of negotiation deadlock they represent a new mechanism for gaining access, with the significant advantage that no actor in the system must spend more money (on expectation) to resolve an otherwise intractable asymmetry.

The main expected use-case for bonded reimbursement submissions are in HTA-type systems (especially cost-effectiveness HTA systems). This is because HTA-type appraisals of a drug will more often hinge on the truth or falsity of claims which are difficult to demonstrate, such as the correct long-term extrapolation of Kaplan-Meier curves. However, the mechanism itself is both medicine- and system-agnostic; if there is value in resolving asymmetric information in a more



Alex Bates

free-market system like the USA then the bonded reimbursement submission can act effectively there too.

The most unique feature of bonded reimbursement submissions is the development of a secondary market in bonds, trading on new information about a product as it becomes available. The implications of being able to track up-to-date best clinical consensus on a medical claim long after its release into a market are potentially even greater than the implications of being able to resolve arbitrary information asymmetries, since this could greatly improve patient outcomes by providing more information to clinicians. Even in the relatively circumscribed use-case described in this essay, the secondary market in bonds would revolutionise the process for finding socially valuable trials for funding, since it would be possible to find trials with a high probability of overturning existing recommendations very easily.

More tools for access are desperately needed as manufacturers' submissions become more complex and the increasing expectations of regulatory bodies makes submissions more expensive and riskier to undertake. A major problem with many proposed tools for access is that they require one party to carry an unacceptable level of risk, or are badly incentive misaligned. A tool like the 'bonded reimbursement submission' - which is incentive-aligned and risk-neutral for all parties – could potentially result in significantly fairer prices, better access and more innovation across the entire lifecycle of a medicinal product.

All illustrations in the example boxes were generated by the Artificial Intelligence DALL-E



OHE Innovation Policy Prize

The OHE Innovation Policy Prize supports our charitable goal of improving the quantity and quality of debate on health economics, and has been designed to be a non-exclusive platform which:

- generates novel ideas and solutions
- facilitates sharing of perspectives across disciplines countries
- encourages more research into the economics of the life sciences sector

The OHE Innovation Policy Prize awards £40,000 for the entry that best fulfils our judging criteria; originality, empirical/theoretical foundations, global feasibility/implementation, potential impact, and clarity. All short-listed entries also have the opportunity to submit a paper to a special issue or section of PharmacoEconomics based on their entry.

The next OHE Innovation Policy Prize will launch in 2024.

Having been at the forefront of economics and policy of health and life sciences innovation for the past 60 years, OHE has a rich history of supporting and promoting original thinking which solves the problems of the day and the future.

